Formalizing Negotiation in Engineering Design

Thesis by
Michael J. Scott

In Partial Fulfillment of the Requirements
for the Degree of
Doctor of Philosophy

California Institute of Technology
Pasadena, California

1999

(Defended September 25, 1998)

Acknowledgements

I offer my first thanks to my advisor Erik Antonsson, and to those other engineering faculty at Caltech who were willing to take a chance on a math teacher with an undergraduate degree in philosophy. Without their initial confidence this thesis would have been impossible, and I hope that this effort will allow them to feel pleased with their decision.

The evolution of these ideas owes a debt to more people here at Caltech, and graduate students who have moved on, than I would care to try to list. Suffice it to say that my education here has come from many directions; I hope that I have offered half as much as I have received. Conversations with research engineers at Ford Motor Company, Alcoa, and Volkswagen Wolfsburg were of great value to a student exploring a theoretical topic. Several individuals helped directly with parts of this thesis: the hardware testing and solid modelling described in Chapter 8 were the work of two undergraduates, Zhee Khoo and Juan Nuño, and the aeroshell analysis was done in conjunction with two graduate students, Rolf Kaiser and Matt Dilligan, and with Robert J. Glaser of the Mechanical Systems Engineering and Research Division of Caltech's Jet Propulsion Laboratory.

The National Science Foundation and the Jet Propulsion Laboratory provided financial support, through my advisor, for work reflected in this thesis. The views expressed here are of course mine, not theirs.

Finally, I thank the companion of my days, Joy Clendenning, for moving across a continent for the sake of my studies, and for making my life as wonderful as it has been these past years. While it would be almost customary for me to thank her as well for enduring the difficulties that grad school imposes on families, I need not do that: on top of everything else, my advisor, Professor Erik Antonsson, made graduate school a pleasant experience.

Formalizing Negotiation in Engineering Design

by

Michael J. Scott

In Partial Fulfillment of the

Requirements for the Degree of

Doctor of Philosophy

Abstract

Negotiations are common in engineering design, especially on large projects, and are typically conducted informally. Often, negotiation is used to handle the imprecision or uncertainty that is inherent in the design process. Performance targets, initially specified as hard numerical constraints, are adjusted throughout the design process in negotiations between engineers and managers. Crucial unmeasured or unmeasurable aspects of performances, such as aesthetic concerns, are commonly negotiated. Negotiations settle conflicts between engineering groups over values of shared design variables and distribution of limited design resources.

In this thesis, a formal description of negotiation in engineering design is presented. This formal model builds on earlier work at Caltech in the modelling of imprecision in engineering design. Negotiation is modelled mathematically as the aggregation of preferences. A complete characterization of the aggregation problem and of the aggregation operators suitable for engineering design is given. This class of operators spans a range of rational decisions and allows for different possible levels of compensation among goals. Furthermore, the entire range of aggregation operators is necessary to capture all possible engineering design decisions. Techniques are presented for determining which aggregation operators are appropriate for particular problems.

As the aggregation of preference is also a component of other fields, notably decision theory applied to economics and social choice, various concerns raised in those fields about the legitimacy of preference aggregation are treated. A more comprehensive justification is

presented here of the approach to modelling imprecision known as the *Method of Imprecision*, or M_oI, than has previously been offered.

Although the decision model presented here is statically one of choice among given alternatives, refinement and redesign are crucial in the engineering design process. The consideration of information about entire *sets* of designs not only accrues computational benefits, but is a more natural model for how designers reason, and can be of significantly more use to designers in refinement and redesign, than information about individual *point* designs. Conditions under which the negotiation model can support set-based information and still yield consistent answers are here explored and presented.

Finally, an example of the application of these ideas to a preliminary vehicle structure design is presented. This example was undertaken as a demonstration of the method for research engineers at Volkswagen Wolfsburg, and serves to help introduce many of the ideas in a more concrete manner.

Contents

1 Introduction **1**

1.1 Contributions of the thesis and dependence of chapters 5

2 Research Context **7**

2.1 Design theory and methodology 7

2.2 Economics and decision theory 9

2.3 Fuzzy sets 9

2.3.1 Theory 10

2.3.2 Application 11

2.4 Artificial intelligence 11

2.5 Summary 12

3 The Method of Imprecision **13**

4 Individual and Overall Preference **19**

4.1 Arrow's Impossibility Theorem and implications for the aggregation of preference 21

4.1.1 The motivating paradox 21

4.1.2 Axioms for the social choice problem 22

4.1.3 The resulting contradiction 24

4.1.4 Ways around the contradiction 24

4.1.5 Weights in the social choice problem 25

4.2 Decision with multiple criteria 29

4.3 Individual preferences 34

4.3.1 Utility 34

4.3.2 M$_o$I preference 36

4.4 Summary 37

5 Aggregation of Preference **39**

5.1 Aggregation in the M$_o$I 39

5.2 The axioms of the M$_o$I 41

5.3 Fuzzy multi-attribute decision making 46

5.4 Weighted means 47

5.5 Supercompensating functions 56

5.6 Example 59

5.7 Modelling negotiation: which decisions can be captured? 66

5.7.1 Theoretical possibilities 67

5.7.2 Managed negotiation 70

5.7.3 Negotiation with predetermined weights 72

5.7.4 Misrepresentation of preference as a negotiation strategy 75

5.8 Hierarchical negotiation 75

5.8.1 Direct aggregation of more than two preferences 75

5.8.2 Hierarchical aggregation of more than two preferences 76

5.9 Examples of negotiation in engineering design 77

5.9.1 Unreachable target performance values 77

5.9.2 Trade-offs between facets of performance 78

5.9.3 Conflicts between design and manufacturing 79

5.9.4 Conflicts between engineering groups 79

5.9.5 The incorporation of unquantifiable performances 79

5.10 Summary and implications for MADM 80

6 Convexity and Set-based Design **82**

6.1 Convexity of preference in the M$_o$I 84

6.2 Proofs 92

7 Computation Methods Comparison **96**

7.1 Optimization strategies . . . 97

7.1.1 Exhaustive search . . . 98

7.1.2 Classical optimization . . . 98

7.1.3 Genetic algorithms . . . 99

7.2 Optimizing for sets of designs . . . 99

7.3 Approximation techniques . . . 100

7.4 Example . . . 100

7.5 Results . . . 103

7.5.1 Classical optimization . . . 104

7.5.2 Genetic algorithm . . . 105

7.5.3 Set-based search with minimum assumption . . . 107

7.5.4 Approximations . . . 110

7.6 Summary . . . 110

8 Example: Passenger Vehicle Structure Design **113**

8.1 Preliminary vehicle structure design . . . 113

8.2 Applying the M_oI to include imprecise information . . . 119

8.2.1 Performance preferences . . . 120

8.2.2 Design preferences . . . 120

8.2.3 Weights and strategies . . . 123

8.3 Results . . . 127

8.3.1 Approximations . . . 128

8.4 Discussion . . . 131

8.5 Summary . . . 132

9 Conclusions **133**

A Aeroshell Design and Analysis **136**

A.1 Aeroshell design . . . 137

A.2 Problem scope . . . 139

A.3 Formal treatment of the problem 142
A.4 Application to the example 144
A.4.1 Aggregation of data from disparate sources 150
A.5 Summary 155

List of Figures

3.1 Example imprecise specification 14

4.1 Single-peaked functions 26

5.1 Functions between *min* and $\mathcal{P}_{\Pi}$ 51

5.2 Functions that exceed $\mathcal{P}_{\Pi}$ 57

5.3 Decision space with optimal region 61

5.4 Decision surface with minimum operator 63

5.5 Decision surface with product of powers operator 64

5.6 Optimal points varying with parameter s 65

5.7 Hierarchy of preferences for VW example 77

6.1 Functions convex with respect to μ, one dimension 86

6.2 Proximity of designs 87

6.3 μ-convex μ_1, μ_2 that combine to non-μ-convex $\bar{\mu}$ 90

7.1 Finite element model in wireframe 101

7.2 Representative design preference 102

7.3 Performance preference 103

7.4 Comparison of three optimization schemes 109

7.5 Ranges of design variables, three optimization schemes 111

8.1 1980 VW Rabbit in stiffness testing 114

8.2 Load test, bending stiffness 115

8.3 Load test, torsional stiffness 116

8.4 Geometric model of body-in-white in SDRC I-DEAS 117

8.5 Finite element model of body-in-white 118
8.6 Imprecise performance requirements 121
8.7 Designer preferences 122
8.8 Hierarchy of preferences for VW example 124
8.9 Graphical user interface for preference display 129
8.10 3-D graphical user interface for preference display 130

A.1 Launch frequency and payload size of NASA space missions, 1955–2015 (Used by permission, John Peterson, JPL, May 15, 1997) 137
A.2 Schematic of aeroshell 140
A.3 Example quality curve 146
A.4 Empirical data: C_a for three cone angles θ 149
A.5 Newtonian flow model: C_a for various θ 151
A.6 Free-molecular flow model: C_a for two θ 152
A.7 Fuzzy set expressing applicability of regime 154
A.8 Interpolated experimental results with $B = 0$, $F = 0$, $\theta = 15°$ 155
A.9 Results from free molecular code with $B = 0$, $F = 0$, $\theta = 15°$ 155

List of Tables

4.1 Weak orders of three voters 27

5.1 Axioms of the M_oI for aggregation operators 42

5.2 Axioms of the M_oI for aggregation functions 44

5.3 Properties of the weighted mean 48

5.4 Undominated points in the decision space 60

8.1 Designer preferences 123

8.2 Peak performance points, various strategies 128

Chapter 1

Introduction

> There is no new thing under the sun.
>
> *Solomon (Ecclesiastes 1:9, King James Version)*

> Everything that can be invented has been invented.
>
> *Charles Duell, Director of U.S. Patent Office, 1899*

Design is a fundamental human activity. Whatever the field of endeavor, design involves the conception of something new to satisfy a need. All design involves creativity (the generation of alternative solutions) and decision (choice among those alternatives). It is a common view that design must be intuitive, and as no course of study can substitute for human genius, both creativity and decision are ineffable and mysterious. Yet past experience and methods are considered the proper study of the designer in all fields. *Engineering* design distinguishes itself from other fields of design by its use of calculation and analysis. Indeed, analysis of alternatives should be inserted into the taxonomy of design proposed above: for engineering design, there is generation of alternatives, analysis of alternatives, refinement of alternatives, and decision among alternatives. Generation of alternative solutions may remain an informal, even mysterious process in engineering design, but the use of calculation formalizes the evaluation of alternatives, from which it follows that decision can potentially be made formal as well.

This view of design is, however, deceptive. Solomon's wisdom and Duell's comical short-sightedness notwithstanding, design *is* the creation of the new. Furthermore, the process is iterative. The neat division of design into stages of concept generation and concept selection is artificial. The engineered solution to a problem is most often an alternative that was not considered at first. The designer must go "back to the drawing board," and not simply to create a larger sample from which to choose. The designer learns from the analysis and the decision, and thus generates better alternatives. Creativity is informed by evaluation and experience; there is feedback in the system. If engineering design were nothing more than trial and error, then a simple model of decision among alternatives would suffice; as design is considerably more complex, such an approach is simplistic at best.

Formalizing design creativity may be, for now, unattainable, but a formal approach to design decisions should as much as possible recognize the connection between decisions and creativity. The formalism discussed in this thesis attempts to do so by supporting the consideration of whole sets of designs at once. An algorithm to find a single optimum among a collection of alternative designs is of little, if any, use to a designer in the generation of further alternatives. Information about the behavior of entire sets of designs may be of greater value.

A formal theory of engineering design seeks to model the activity of design well enough that it can be computed. Some engineering design tasks, like the sizing of fasteners, bearings, or motors, rely mainly on formal calculation. Others, like the selection of fastener type (rivet, press fit, snap fit, bolt, weld, Velcro®) rely more on informal intuition. The informal makes for a large part of design, as can be clearly seen by how few design tasks can be automated. Research in engineering design seeks formal structure to explain what humans accomplish informally, and thus tends to increase the number of design tasks that can be computed.

Why compute design tasks? Routine tasks can be accomplished more quickly and efficiently through computation, and the results can be as good as those that the most experienced and talented individuals can deliver; the designer's time is better spent on questions that merit creative thinking. Although there is interaction between creativity and analysis, and enumerating more alternatives is not the only way to solve design problems, it is generally advantageous to consider a larger set of alternatives, to "increase the size of the design space." When design can be understood well enough to be computed, that understanding is useful to the novice and expert, to the inspired and the mundane designer alike, in the productive inclusion of more alternatives in the decision process. Thus the goal of a theory of engineering design is to automate the tedious, analyze the intuitive, and communicate the experiential.

When design is conducted by more than one person, communication becomes as fundamental as creativity and decision. In such collaborative design, decisions must be made jointly. The research presented here was inspired by conversations with vehicle structure engineers from industry, who commented that many of their design decisions took the form of *negotiations*. Negotiation, defined broadly (paraphrasing Webster's) as conference with others to arrive at a settlement of some matter, need not be adversarial. Yet a divided design task creates a situation in which individuals or groups can readily identify the sub-goals associated with their assigned task, perhaps more readily than the common goal to fulfill a given need. When a corporate culture separates personal self-interest from the fulfillment of the common goal, the negotiations that arise between different groups may be adversarial negotiations, when they ought to be planning sessions among partners striving together towards a monolithic end. Even when the participants identify their own self-interest with a single unified goal, the communication among groups is a negotiation. Thus negotiation plays an important role in many design decisions, and is almost always conducted informally. This thesis presents a formal description of negotiation in engineering design, and thus the groundwork for the development of a methodology to conduct negotiations. The development of a formal model for negotiation shows that the negotiation problem can be equated with the problem of aggregation of several preferences. The aggregation model serves in many situations that would not obviously be thought of as "negotiations." A formal understanding of the negotiation process can shed light on the question of how to structure that process so that participants in the negotiation behave to further a unified, rather than a selfish, self-interest.

Engineering design research seeks both mathematical rigor and real-world applicability. The two are often connected: a sound mathematical foundation helps guarantee the success of an application, while the desire to solve particular problems uncovers needs of the theory. The two components of engineering design research, rigor and application, are sometimes traded-off, in a metaphorical negotiation, one against the other. In engineering design decision making, one possible formal approach is to design by enumeration of alternatives, and this approach, indeed, underpins most if not all formal decision methods. Certainly, there is some theoretical justification for viewing design as a choice among alternatives. However, this assumes that all possible alternatives are actually considered, an assumption that seems completely unjustified in light of the iterative nature of real design.

There are actually two difficulties here, one perhaps deeper than the other. If it is accepted that the design decision problem is to choose among some *given* alternatives, there is the simple problem of the inherent computational difficulty of exhaustive search. An understanding of the structure of the design problem can make much of that search unnecessary. Elegant methods use structural knowledge to avoid unnecessary computation. The deeper problem lies in the assumption that the decision is restricted to the given alternatives. An overarching formalism for the consideration of unspecified alternatives is well outside the scope of this thesis. The methods presented here are restricted to considering a given collection of alternatives, but, by considering them as sets rather than points, provide information that may be useful to a designer in later iterations.

Preliminary design is inherently imprecise, and has enormous economic importance. Much of the cost of a design is determined by preliminary decisions [99], which are often informal and rely on imprecise information. This thesis extends earlier work in the modelling of imprecision in engineering design. It contributes to the understanding of decision making, and to the mathematics of fuzzy sets. It examines economic theory's possible extension to engineering design decisions. Related research is done in engineering, in design theory, in multi-objective optimization, and in multi-criteria decision making. Researchers in artificial intelligence have also addressed the problem of negotiations. There is a considerable body of work, which will be considered below, on related questions in the application of fuzzy sets to various decision problems.

While negotiation is a focal point of the thesis, to arrive at the simple formal model requires detailed consideration of preference, aggregation, and convexity. These ideas are central to the development of the thesis; they are also intertwined. The sequence in which the three are presented is to some extent arbitrary: all are needed to paint a complete picture.

An undercurrent throughout the thesis is an interest in set-based design, the notion that designers reason not with individual, unique design instances, but with whole sets of possible designs. Formal descriptions of engineering design decisions must be rational: irrational decisions are inadmissible. It is not so clear that a formal description must be set-based, but there are several advantages to making it so. Perhaps the most important advantage is that it provides a natural representation for reasoning in design, though it turns out that there are computational advantages as well.

In addition to formally modelling actual negotiations, the mathematical description can be applied to other situations that are not so obviously negotiations. The aeroshell example presented in Appendix A is an extension of the mathematics to the realm of analysis, for example. Even when communication between groups is not significant, some aggregations are "negotiations" between a designer and him or herself.

1.1 Contributions of the thesis and dependence of chapters

The broad contribution of this thesis is the development of a mathematical model for negotiated decisions in engineering design. While negotiation in engineering design has been studied previously (previous work is reviewed in Chapter 2), the explicit mathematical representation presented here is new.

The formal negotiation problem is a problem of the aggregation of preferences. A set of axioms for preference aggregation in engineering design has been previously proposed by researchers at Caltech [56], and was used in the investigations described in this thesis. These axioms differ from the axioms that describe economic and social choice decision problems, and the preference aggregation results from those fields are not applicable to the engineering design problem. One contribution of this thesis is the explicit distinction that it draws between the engineering design decision problem and these other decision problems. This is especially important because it allows an accurate assessment of the potential application of economic decision theory to engineering design problems. In particular, the implications of one well-known result from social choice theory, Kenneth J. Arrow's Impossibility Theorem, are shown not to present any difficulties for the engineering design decision problem.

A chief contribution of this thesis is a complete characterization of the operators that are appropriate for decision making in engineering design. The range of acceptable operators includes both of the aggregation operators that were in prior use; it is also demonstrated here that the entire range is necessary to formalize all design decisions. In addition to their relevance to the decision problem, the results about preference aggregation operators are relevant to the problem of the combination of fuzzy sets in general contexts.

The characterization of design-appropriate aggregation operators includes not only proof that the entire collection of operators is necessary to capture all decisions, but also methods for the determination of the correct operator for a particular decision. Thus this thesis demonstrates that the decision methods presently in use in engineering design in industry are liable to favor incorrect choices, and offers alternatives that are demonstrably accurate.

A further contribution of this thesis explores the potential use of the negotiation models and methods presented here to support set-based design, in which designers reason with entire sets of candidate designs, rather than with individual designs. The thesis presents the concept of convexity of preference, and uses that concept to derive sufficient conditions for the use of the negotiation methods for set-based design.

The dependence of chapters is as follows: Chapter 2 provides general research context. Chapter 3 is a review of earlier work at Caltech, and should be read before all following chapters. Chapters 4 and 5 discuss preference and aggregation, and should be read next. The discussion of convexity in Chapter 6 relies on Chapters 4 and 5, but may also clarify some of the ideas of those earlier chapters, and it may be useful to revisit them after reading Chapter 6. The discussion of computation methods in Chapter 7 depends only on Chapter 3, as does the presentation of aeroshell analysis in Appendix A. The vehicle structure example in Chapter 8 refers to all earlier chapters, but can profitably be read immediately after Chapter 3; an early visit to the example may be helpful in reading, especially, Chapters 4–6.

Chapter 2

Research Context

The formalization of negotiation in engineering design draws upon several fields of research. It is a part of a larger and more general effort in design theory and methodology. It is related to decision theory, and to a lesser extent, to some economic decision models. As some of the models discussed here rely on the mathematics of fuzzy sets, this work fits also into a tradition, both theoretical and practical, of decision making with fuzzy sets. Finally, there are some commonalities between this work and work that has been done in artificial intelligence, though the two approaches to negotiation are distinct. In this chapter, the varied research background and context for the work of the thesis is presented.

2.1 Design theory and methodology

Perhaps the most comprehensive survey of research in mechanical engineering design to date is Finger and Dixon's two-part study of 1989 [30, 31]. They consider six distinct areas of engineering design research:

- descriptive models,
- prescriptive models,
- computer-based models,
- languages, representations, and environments for design,
- analysis in support of design,

- design for manufacture and the life-cycle.

The work presented in this thesis is placed by Finger and Dixon in the arena of representations and environments for design. However, they describe engineering design research as an emerging field in a pre-theory state, without a wide consensus as to either what are the interesting, outstanding questions, or what research methodologies should be employed. Since the field continues to define itself, no taxonomy of design research will be strict, and other areas of design research are directly relevant to the work discussed here.

Engineering design research can also be classified with respect to its formality, though Finger and Dixon did not choose to do so. Formal models and methods are those which allow codification if not computation, and thus can potentially be automated. The work presented in this thesis is on the formal side of design research.

Even among the community of design researchers whose work can be described as formal, there is little consensus as to problems and methodologies. Prescriptive models of the design process were developed earlier in Europe than in the United States. The work of Pahl and Beitz [61], which essentially became the German VDI (*Verein Deutscher Ingenieure*) standard for design [91], and Hubka and Eder [40], might be described as pre-formal codification of the engineering design process. Yoshikawa's General Design Theory, or GDT [105, 88] (see also Reich's review [73] of GDT) was an axiomatic, ambitious approach to the formalization of design; Yoshikawa was also early in asking the question of what is well enough understood in design that it can be automated [106]. Stiny's groundbreaking work on shape grammars in architecture [83] has proven to be the departure point for some interesting formal work in engineering design [3, 78]. Ward *et al.* have developed interval analysis methods for set-based

design [96, 97, 98], as well as conducting studies of their application in the automobile industry [95]; their work is directly related to this thesis.

A number of semi-formal decision methods have seen application in industry: Quality Function Deployment (QFD) [38], Pugh charts [67], and Saaty's Analytic Hierarchy Process, or AHP [75] are examples. These methods are semi-formal in that they provide rules or formulae for the organization of information, but offer no rigorous justification for the heuristic calculations. Somewhat more formal are Taguchi methods [16], and experimental design for quality control and process planning [52, 64]. A line of research into formal modelling of imprecision in preliminary design has been undertaken here at Caltech. This work, collectively referred to as the Method of Imprecision, or M_oI, is of direct relevance to this thesis and is covered in greater detail in Chapter 3.

2.2 Economics and decision theory

There is a rich literature on probabilistic decision making going back at least to Bayes's work in 1763 [11]. This decision making theory has been relatively well-developed, especially as applied to economics. The original formulations of utility theory and game theory by von Neumann and Morgenstern [93] were taken up by Keeney and Raiffa [42] and Luce and Raiffa [49], among others. (See French [32] for a good overview.) The foundational works are of relevance to the formalization of decision decision making presented here. However, the development of economic decision theory has been directed from the outset to the problem of decision making under probabilistic uncertainty, which is a different portion of the engineering design decision problem than is considered here.

The design decision problem is one of decision with multiple criteria, though it is sometimes confused with two other problems in decision theory, decision under probabilistic uncertainty and group decision making. At the worst this can lead to the inappropriate application of results without a verification of conditions. More constructively, the success of formal models in economics have led to attempts to extend those models to engineering contexts. A few examples of utility models in engineering are Siddall [82], Bradley and Agogino [14], and Thurston [87]. The distinction between probabilistic decision and decision in engineering design will be discussed further in Chapter 4.

2.3 Fuzzy sets

The starting point for the Method of Imprecision, the modelling of imprecision in preliminary design, was the recognition of the intrinsic imprecision in uncompleted designs. Initial exploration led to the consideration of fuzzy sets [101]. Fuzzy sets research is relevant to this thesis in two ways. On the theoretical side, the negotiation problem treated here can be considered as a problem in the aggregation of fuzzy sets, and the results on aggregation will be of general interest to the fuzzy sets community. On the side of implementation, many researchers are applying fuzzy set theory to decision problems in engineering design and related fields.

2.3.1 Theory

Fuzzy sets were first proposed in 1965 by Zadeh [107] as a generalization of classical set theory. In classical set theory, for any set, every entity is either in the set (has a membership of 1) or not in the set (has a membership of 0). In a fuzzy set, elements can take on varying degrees of membership between 0 (not at all in the set) and 1 (completely in the set); a *membership function* specifies the membership level of all elements in the set. This assignment of degrees of membership to the set could be used, for example, to capture linguistic notions such as the set of "tall men." Any function that assigns numbers in $[0, 1]$ to all the elements of a set induces a fuzzy set. The preferences that will be discussed in Chapter 4 can be viewed and manipulated as fuzzy sets.

The negotiation model discussed in this thesis involves the aggregation of preferences, and thus the problem of aggregation of fuzzy sets is directly relevant to the negotiation model presented here. Many different aggregation functions for fuzzy sets have been proposed and studied. Much fuzzy set research has focused on *t-norms* and *t-conorms* [26], implicitly or explicitly equating fuzzy set aggregation with the extension of classical binary logic to fuzzy sets. (This can be seen as a natural consequence of the initial interest in fuzzy sets to capture linguistic expressions.) T-norms are bounded above by the *min* function, and are the appropriate model for extensions of the logical AND to fuzzy sets, while t-conorms are bounded below by the *max* function and are an extension of the logical OR. Both classes of functions have been studied in detail, and research is ongoing. Multi Attribute Decision Making (MADM) schemes have applied a wide range of t-norms and other operators to decision problems. While averaging operators, which fall between *min* and *max*, have been acknowledged for some time [26, 27, 103, 104], comparatively little attention has been devoted to these connectives. Averaging operators are not appropriate for binary logic, but they are well suited to engineering design decisions: indeed, the axioms for engineering design preference aggregation presented in Section 5.2 require the aggregation functions to fall between *min* and *max*. The treatment of averaging operators offered here is more thorough than has been presented previously, and so the results are of relevance to more general research in fuzzy set theory.

2.3.2 Application

There is significant recent interest in fuzzy methods for decision making. Recent contributions include: Dhingra on multiobjective fuzzy optimization techniques for engineering design [20]; Diaz on fuzzy optimization methods [21, 22]; Djouad on chemical process synthesis [23]; Dubois on fuzzy constraint propagation applied to manufacturing [25]; Fargier's fuzzy scheduling [29]; an application of fuzzy methods to windturbine design by Gerhart [33]; Grabot's multiobjective scheduling [35]; Hamburg and Hamburg on management of uncertain knowledge in engineering design [36]; Hsu *et al.* on engineering design optimization [39]; Otto on imprecise calculations in engineering design [17]; Knosala and Pedrycz on evaluation of design alternatives [44]; Müller and Thärigen's fuzzy MADM methods in system design [54]; Posthoff on fuzzy evaluations [66]; multiobjective fuzzy optimization techniques for engineering design, by Rao and colleagues [68, 69, 70, 71, 72]; Sakawa and Kato on multiobjective fuzzy optimization [76, 81]; Schleiffer's fuzzy design with evolutionary algorithms [77], with reference to the M_oI; Thurston and Carnahan on fuzzy ratings and utility analysis in preliminary design evaluation of multiple attributes [85, 86]; scheduling system design by Turksen [90]; and work by Zimmermann and Sebastian on fuzzy-multi-criteria decision making [110, 111, 112, 113].

2.4 Artificial intelligence

Researchers in artificial intelligence have noted that conflict is an integral part of the design process. A few are mentioned here: Oh and Sharpe [55] present a bibliography of current research and a thoughtful list of potential sources of conflict in addition to their own work on a design support environment called Schemebuilder. Bahler *et al.* [9] have approached conflict from a utility theory point of view; their work is perhaps the most comparable to the research direction developed here, but they have focused on a computer implementation of the decision and have limited themselves to a linear weighted sum model and fairly restrictive representations for goals. The work accomplished previously with the M_oI, and the results presented in this thesis, offer more possibilities for the modelling of the design but less automation of the decisions. Some interesting work has been done on a design support system using Pareto optimality by Petrie *et al.* [63]. The system does not calculate optimal solutions, but rather tracks a history of design decisions and automatically notifies agents when it seems that a better design might be overlooked.

These approaches to managing conflict in design, and others from the artificial intelligence community [15, 37], have concentrated on environments that model the design process itself, with the idea that such a model will be applicable in any design situation, thus approaching the design problem from above. The act of negotiation is seen as an entity to be modelled. This top-down approach stands in contrast to the approach taken in this thesis, where the design model, rather than the negotiation model, is primary.

2.5 Summary

This chapter presents some of the prior and ongoing work that is relevant to the thesis. This relevant work is seen to span a wide range of fields, including engineering, decision theory, economics, fuzzy set theory, and artificial intelligence. Finger and Dixon's description in 1989 of design theory as a "pre-theory" field still holds for the negotiation research presented in this thesis. While the prior work in artificial intelligence comes closest to the work presented here, artificial intelligence has considered the negotiation problem from above, attempting to model the act of negotiation for universal application. The research discussed here, in contrast, approaches the design problem from below, where the crucial problem is to model the imprecision inherent in design information, and to use that model of design imprecision to guide negotiation.

Chapter 3

The Method of Imprecision

The investigations of this thesis were conducted within the framework of the Method of Imprecision, or $\mathrm{M_oI}$ [46, 56, 100, 102], a method developed at Caltech for incorporating imprecise information into a design process. In this chapter, the necessary background and notation are presented. Since this thesis focuses on the use of the formalisms of the $\mathrm{M_oI}$ to represent negotiation in engineering design, certain concepts presented briefly here will be defined in greater detail later, while previous developments that are not directly relevant to the negotiation problem will be treated lightly or not at all.

The original work on the $\mathrm{M_oI}$ formulated the design problem as a decision problem: given a set of candidate designs, identified by vectors $\mathbf{d}$ of *design variables* in a Design Variable Space $\mathbb{D}$ (sometimes called the *DVS*), a set of performances, described by vectors $\mathbf{p}$ of *performance variables* in a Performance Variable Space $\mathbb{P}$ (sometimes called the *PVS*), and a mapping $f : \mathbf{d} \mapsto \mathbf{p}$, choose the candidate design $\mathbf{d}^*$ which maps to the "best" possible performance $\mathbf{p}^* = f(\mathbf{d}^*)$. So stated, this model of the design problem is unacceptably abstract. Several considerations must be explicitly incorporated into the decision problem:

- requirements are imprecise;
- some preferences are not modelled by f;
- there is no obvious unique way to compare different performance variables which are usually not even expressed in the same units.

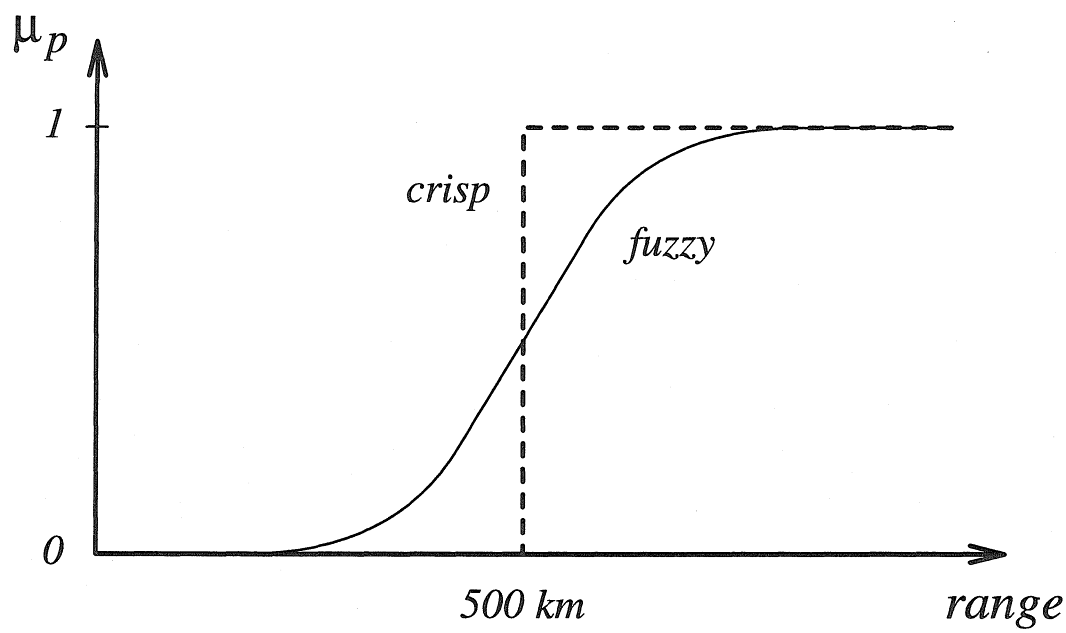

Figure 3.1: Example imprecise specification

In addition, the selection of a *single* best design is an oversimplification of the problem: more generally, the selection of a suitable set of (not yet completed) designs is desired.

The need to include imprecision in engineering design can be illustrated by a simple example. Figure 3.1 shows a specification for one performance variable (p_j). As specifications are commonly written, $p_j \geq 500$ km would be represented by the dashed line (the sharp-edged rectangular step), where $\mu_p = 1$ in the acceptable region. However, this crisp specification (or requirement) indicates that two different designs, one with $d_j = 500 - \epsilon$, and another with $d_j = 500 + \epsilon$, would have completely different acceptabilities, no matter how small ϵ becomes. Thus two designs, indistinguishably different in d_j (as $\epsilon \rightarrow 0$), have completely different preferences: one is completely acceptable and one is unacceptable. This is clearly an inadequate model.

Alternatively, the solid line shown in Figure 3.1 indicates a transition of acceptability of performances from unacceptable ($\mu_p = 0$) to most desired ($\mu_p = 1$), and thus reflects a more realistic specification. The range over which the transition from unacceptable performance to most desired performance takes place will depend on the particular design problem, as will the shape of the curve. The curve may be non-differentiable or even discontinuous (a government regulation, for example).

Thus the MoI introduces the notion of *preferences*, mappings (denoted by μ) which take values on the closed unit interval $[0, 1]$, both to represent the imprecision inherent in the preliminary design problem, and to provide a basis for comparison between different attributes. *Performance preferences* $\mu_{\mathbf{P}} : \mathbf{P} \rightarrow [0, 1]$ express the customer's requirements for potential performance values more completely than crisp targets. In addition, engineers express *design preferences* on design variables ($\mu_{\mathbf{D}} : \mathbf{D} \rightarrow [0, 1]$), allowing the incorporation of performance aspects that are not explicitly calculated by f.

Designs are often judged by criteria that are not calculated in an engineer's analysis. Some of these criteria can be crucial to the success of the design: style, for example, is close to paramount for many consumer goods, and manufacturability is a prerequisite. When the engineer does not have a calculable model for these criteria, the incorporation of these unmodelled aspects of performance (unmodelled by f, that is) is conducted informally, often in negotiations with other groups such as stylists or manufacturing engineers, or with managers. The specification of design preferences allows for the explicit representation of otherwise unmodelled concerns in the design decision problem.

To summarize and clarify the notation presented so far:

$\mathbf{X}$	any set, a generic parameterized set
$\mathbf{x}_1, \mathbf{x}_2, \ldots, \mathbf{x}_n \in \mathbf{X}$	vector elements of $\mathbf{X}$
$\mathbf{x} = (x_1, \ldots, x_m)$	components of the vector $\mathbf{x}$
$\mathbf{D}$, the *DVS*	the set of possible designs
$\mathbf{d} \in \mathbf{D}$	a candidate design
$\mathbf{P}$, the *PVS*	the set of possible performances
$\mathbf{p} \in \mathbf{P}$	a design performance
$\mathbf{d}_i, \mathbf{p}_j$	particular designs and performances (as with $\mathbf{x}$)
d_i, p_j	components of designs and performances (as with $\mathbf{x}$)
$f : \mathbf{d} \mapsto \mathbf{p}$	a mapping that determines the performance of each design
$\mu : \mathbf{X} \to [0, 1]$	preference on a generic set
$\mu_{\mathbf{D}} : \mathbf{D} \to [0, 1]$	design preference on the set of design variables
$\mu_{\mathbf{P}} : \mathbf{P} \to [0, 1]$	performance preference on the set of performance variables
$\alpha, \beta \in [0, 1]$	elements of the closed unit interval

The design preferences $\mu_{\mathbf{D}}(\mathbf{d})$, which are specified on the *DVS*, can be mapped onto the *PVS* by use of the *extension principle* [108]. This induces a map $\hat{\mu}_{\mathbf{D}} : \mathbf{P} \to [0, 1]$, defined by

$$\hat{\mu}_{\mathbf{D}}(\mathbf{p}) = \sup_{\mathbf{d} | f(\mathbf{d}) = \mathbf{p}} \mu_{\mathbf{D}}(\mathbf{d})$$

The performance preferences give rise to an induced map on the *DVS* which is defined by simple composition:

$$\hat{\mu}_{\mathbf{P}}(\mathbf{d}) = \mu_{\mathbf{P}}(f(\mathbf{d}))$$

Since it is clear from the context when the induced map is invoked, the $\hat{\mu}$ notation will not be used, and both the original and induced maps will be denoted μ. The preference maps take values on the closed unit interval $[0, 1]$; these values will usually be denoted by lowercase Greek letters α and β, but in some cases where the context is clear an abuse of notation may identify μ with $\mu(\mathbf{x}) \in [0, 1]$.

A design is typically judged on the basis of more than one preference; the $\text{M}_\text{O}\text{I}$ employs an explicit *aggregation* of all preferences to compare and combine the different aspects of performance, modelled and unmodelled, on which the design is judged. Several preferences are combined with an *aggregation function* $\mathcal{P}$. (In previous work on the $\mathrm{M_oI}$, aggregation functions were considered to act on preference values, not preference functions; a more precise definition of aggregation will be given in Chapter 5.) At first, the $\mathrm{M_oI}$ made use of two different aggregation functions [57], the non-compensating $\mathcal{P}_{\min}(\alpha_1, \alpha_2) = \min(\alpha_1, \alpha_2)$ for situations where the overall performance is dictated by the lowest-performing attribute, and the compensating $\mathcal{P}_{\Pi}(\alpha_1, \alpha_2) = \sqrt{\alpha_1 \alpha_2}$, when high performance on one attribute is deemed to partly compensate for lower performance on another. Each candidate design $\mathbf{d}$ thus has an associated overall preference:

$$\bar{\mu}(\mathbf{d}) = \mathcal{P}\left(\mu_{\mathbf{D}}(\mathbf{d}), \mu_{\mathbf{P}}(f(\mathbf{d}))\right)$$

(where $\mu_{\mathbf{D}}(\mathbf{d})$ and $\mu_{\mathbf{P}}(\mathbf{p})$ are themselves aggregations of their constituent preferences). Candidate designs can be compared on the basis of this overall preference.

Definition 3.1 *The* design decision problem *is to identify a subset of* $\mathbf{D}$ *which yields particular values of* $\bar{\mu}$. *This may be a search for a single design* $\mathbf{d}$ *to maximize* $\bar{\mu}$, *or for the set of designs that exceed a particular* $\bar{\mu}$.

An additional concern in the design decision problem is often that of computation cost: the answer must be found with a limited number of calculations. In Chapter 5 a thorough analysis of possible aggregation operators and justification for selecting a particular operator will be presented.

The foregoing is sufficient background to develop the ideas presented in this thesis, but it should be noted that the $\mathrm{M_oI}$ has also been developed along other avenues. Earlier research developed techniques for including noise [60] and adjustments or *tuning parameters* [59] in the imprecision calculations, and placed an axiomatic framework on the calculations [56]. Implementation of the $\mathrm{M_oI}$ continued with the development of a computational tool [47]. The applicability of the method was seen to be limited by large computational requirements, so the inclusion of Design of Experiments (DOE) approximations [48] and other computational innovations [46] followed. A number of overview articles and book chapters on the $\mathrm{M_oI}$ have been recently published [4, 5, 6, 80]; they contain details, and comparison to other methods, that are not germane here.

Chapter 4

Individual and Overall Preference

As a prerequisite to the formal representation of negotiation in engineering design, the aggregation of separate or individual preferences into a single, overarching preference, must be addressed. This aggregation, in turn, depends on the formalization of the notion of (individual) preferences. Chapter 3 presented the M_oI background to this problem, while Chapter 2 described economic decision theories and multi-criteria decision-making systems. This chapter treats the question of an assignment of a numerical scale for preferences. Such a numerical scale can be specified in a number of different ways, and the choice of a scale depends upon its intended use.

Any formal method for decision making must represent the comparable acceptability of different alternatives. A point of notation: generic alternatives will be denoted A, B, C, *etc*. Parameterized design alternatives are still denoted $\mathbf{d}$. The most basic concept in the ranking of alternatives is simple comparison. In such comparison there is no association of numbers with alternatives, but only the idea that one alternative A is preferred to another alternative B. A ranking that depends only on simple comparison is called a *weak order*:

Definition 4.1 *A* weak order *on a set of alternatives* $\mathbf{X} = \{\mathrm{A}, \mathrm{B}, \mathrm{C}, \ldots\}$ *is a transitive binary relation* $\succeq$ *such that for any two elements* A *and* B, *either* $\mathrm{A} \succeq \mathrm{B}$ *(*A *is at least as preferable as* B*), or* $\mathrm{B} \succeq \mathrm{A}$ *(*B *is at least as preferable as* A*). Indifference is possible: if* $\mathrm{A} \succeq \mathrm{B}$ *and* $\mathrm{B} \succeq \mathrm{A}$*, then one writes* $\mathrm{A} \sim \mathrm{B}$ *(*A *is indifferent to* B*). If* $\mathrm{A} \succeq \mathrm{B}$ *but* $\mathrm{B} \not\succeq \mathrm{A}$*, then* A *is (strictly) preferred to* B*, written* $\mathrm{A} \succ \mathrm{B}$.

A weak order is an ordinal ranking: it orders the alternatives without assigning numerical values. Any *computational* method for decision making requires the further structure of a numerical scale that ranks alternatives. Such a numerical scale will be called a *value function*. The familiar $>$ and $\geq$ on the real numbers of the value function correspond to the preference relations $\succ$ and $\succeq$ among alternatives:

Definition 4.2 *A* value function *is an assignment of real numbers to alternatives that preserves a weak order of acceptability of those alternatives. A value function maps a set together with a weak order* $\{\mathbf{X}, \succeq\}$ *to the real numbers with its usual ordering* $\{\mathbb{R}, \geq\}$*. For a value function v, $v(\mathrm{A}) \geq v(\mathrm{B})$ iff.* $\mathrm{A} \succeq \mathrm{B}$*, with equality for indifference.*

While it is always possible to construct a value function from a weak order [45], there is nothing inherent in the definition of a value function that permits a measure of *degree* of acceptability. In other words, there is no interpretation of relative value beyond the weak ordering of alternatives. A value function can be given additional structure that allows the interpretation of the numerical value. For computation to be meaningful, some such extra structure is required. This chapter discusses the question of choosing a structure for numerical preference.

Three distinct decision problems are decision with multiple criteria, group decision making, and decision under uncertainty [32]. The aggregation of individual preferences in engineering design, and the particular paradigm of resolution by negotiation, is principally a problem in decision with multiple criteria, though it bears superficial resemblance to the other two decision problems. In the theory of group decision making, a well-known objection to the validity of combining separate weak orders into a single ("social") order at all, is Kenneth J. Arrow's so-called Impossibility Theorem [7, 8]. The computable negotiation method for engineering design presented in this thesis is also a decision method that depends on the aggregation of several weak orders into a single order. Because of its apparent similarity to group decision making, it is appropriate to address here the concerns raised by the Impossibility Theorem.

All three sorts of decision problems ultimately rely on a weak ordering among alternatives. Multiple criteria analysis and decision under uncertainty usually overlay the weak ordering of alternatives with a value function (and an interpretation of the numerical scale).

The special character of the multiple criteria decision problem, and its particular value function, turn out to be crucial in the resolution of the difficulties raised by Arrow's theorem.

In this chapter, the objections raised by Arrow's theorem will first be treated. Then, the concept of individual preference, and possible interpretations of a numerical scale on the underlying weak ordering of alternatives, will be discussed. The value functions that are associated with multiple criteria analysis and decision under uncertainty both overcome the objections of Arrow's theorem, but there are still important differences between the two.

4.1 Arrow's Impossibility Theorem and implications for the aggregation of preference

Kenneth J. Arrow's Impossibility Theorem is an important and powerful result in the theory of social choice. For that reason, and because a thorough understanding of that result will facilitate a comparison between the social choice and multi-criteria decision problems, the Impossibility Theorem will be presented here. The treatment here refers mainly to two books by Arrow [7, 8].

4.1.1 The motivating paradox

The backdrop for Arrow's work is politico-economic. Political scientists are interested in determining a "fair" method of reconciling the potentially conflicting interests of individuals in a society. Economists seek the most "satisfactory" distribution of a set of commodities throughout a society. The similarities between the two problems are evident, and indeed both can be formalized in the same way; the notions of "fair" and "satisfactory" are explored through this formalization.

The *majority method of decision making* is one possible answer to the loosely formulated question of fair social choice, and one that is sufficiently obvious that a contradiction that arises from its employment motivates the Impossibility Theorem. For an odd number of people and two options to choose among, a simple vote is guaranteed to satisfy the most people (degree of satisfaction, which will be discussed later, is assumed not to be an issue here). But when there are three choices, a paradox arises: say that there are three voters, one of whom prefers option A to option B to option C, another who prefers B to C to A, and a third who prefers C to A to B. All three voters have rationally ordered preferences, and yet a pairwise vote shows that as a group, these three prefer A to B, and prefer B to C, yet also prefer C to A. The resulting social order is not rational, and provides no basis on which to make a decision. This paradox is called the *failure to ensure the transitivity of the majority method*, or the *paradox of voting* [7, p. 2]. It is in the context of this paradox that economists and political scientists explore the limits of "fair" and "satisfactory" social choice: is there any procedure for aggregating social preferences that can avoid this paradox? Arrow's Impossibility Theorem shows that, given a particular set of axioms that define fair and satisfactory, there is no procedure that can (always) fulfill them all. The formal proof proceeds from the description of the problem with a set of axioms.

4.1.2 Axioms for the social choice problem

By introducing axioms to define any decision problem, two ends are accomplished. Primarily, the problem is modelled so that conclusions about the problem can be derived mathematically. Results are certain with respect to the axiomatic model; their certainty with respect to real problems depends on the validity of the axioms. An additional end of axiomatization is the casting of vague descriptions such as "fair" and "satisfactory" in precise terms.

The axioms which model the social choice problem formalize a number of assumptions about the nature of the problem. When the multi-criteria decision problem is discussed below, the applicability of these assumptions will also be considered:

1. The social choice problem is the aggregation of many weak orderings into a single "social" weak ordering.

2. Every individual ordering carries equal weight. This assumption is natural when each voter must be accorded the same fundamental worth (as a human being) in order to maintain a notion of fairness.

3. There is no interpersonal comparison of utility.[1] In other words, there is no standard

[1]Note that Arrow uses "utility" for any value function; these are not necessarily von Neumann–Morgenstern

for comparison of strength of feeling or value to individuals, so there is no mechanism for choosing A over B because one voter "strongly" prefers A to B while the other two have only a "mild" preference for B over A.

4. There are at least three alternatives, since the problem of two alternatives is trivial. There are an odd and finite number of alternatives, a condition that avoids technicalities without excluding any interesting cases.

In other words, the social choice problem considers decision cases where all options are known, mutually exclusive, and ordered by individuals, and where the task is to produce a single social order yielding the greatest overall benefit while respecting the (equal) worth of each individual. To formalize this decision situation, Arrow introduces five axioms:

Axiom 4.1 (unrestricted domain) *Each individual is free to order the alternatives in any way.*

Restricting Axiom 4.1 is one way to address the paradox, and methods that guarantee the transitivity of the majority method can be ranked by how severely they restrict this freedom. It is not at all obvious that this is reasonable for design decisions, as will be discussed below.

Axiom 4.2 (positive response) *If a set of orders ranks A before B, and a second set of orders is identical except that individuals who ranked B before A are entitled to switch, then A is before B in the second set of orders.*

Axiom 4.2 is an ordinal version of monotonicity.

Axiom 4.3 (independence of irrelevant alternatives) *If A is before B in a social order, then A is still before B if a third alternative C is ignored or added.*

Note that Axiom 4.3 is violated in the motivating paradox, where the relative rankings of A and B are influenced by the addition of the alternative C.

Axiom 4.4 (not imposed) *An order is called* imposed *if some A is before some B in all possible social orders. The social choice problem must not be imposed.*

Axiom 4.5 (not dictatorial) *An order is called* dictatorial *if there is one individual whose decisions dictate the social order. This is likewise not allowed.*

utilities.

4.1.3 The resulting contradiction

The General Possibility Theorem, now commonly known as Arrow's Impossibility Theorem, shows that a social choice function satisfying all five conditions is an impossibility:

Theorem 4.1 *Any social choice function satisfying Axioms 4.1–4.3 must be either imposed or dictatorial.*

The proof is fairly straightforward. The reader is referred to one of [7, 8] for details, but the basic line of reasoning is as follows: A *decisive* set of individuals for A over B is a set who guarantee that A will be preferred to B whenever they unanimously agree so; any decisive set must contain a smaller decisive set; there is always a decisive set; any set that is decisive for A over B is decisive for A over anything else and for anything else over B, and thus for all A over all B; thus there must be a dictator. The only way to avoid this dictatorship is to impose some preferences, violating Axiom 4.4.

Thus the paradox of the intransitivity of the majority method is a manifestation of a difficulty so deeply embedded in the social choice problem that it cannot be resolved without compromising the defining axioms.

4.1.4 Ways around the contradiction

Arrow's Theorem shows only that there is no method of aggregating social choice that is guaranteed to satisfy all five axioms, not that all instances of social choice will violate one axiom or another. Are all socio-political systems then fundamentally irrational, or are there systems that diverge from the given axioms?

Arrow and others have attempted to resolve the paradox by weakening the first condition, arguing that in real political, economic, and even moral[2] systems participants tacitly agree to structure their choices in a "logical" way, *i.e.*, in a way that keeps contradictions from arising. Thus Arrow introduces the notion of *single-peakedness* as a way around the dilemma—the set of alternatives is ordered on some (one-dimensional) external scale, so that each individual is free to choose his favorite, but then must hold descending regard for the other alternatives to the two sides of his first choice. This notion of single-peakedness

[2]Arrow goes so far as to quote Kant [7, pp. 81–82]

(see Figure 4.1) is closely related to the convexity that will be discussed later when the engineering design decision problem is considered, though the limitation to one dimension does not hold there. The example of a political spectrum is given: each voter has a preferred, or ideal party, and each step away from the ideal party, whether to the left or to the right, is an ever less desirable alternative. This condition says nothing about comparison between parties to the left and parties to the right of ideal. If a condition of single-peakedness is substituted for the axiom of unrestricted domain, then the Impossibility Theorem no longer holds. A structure of this sort may be what holds parliamentary systems together. Of course, in a two-party system there is no contradiction, as the two-alternative situation is not paradoxical.

In general, the difficulty of the Impossibility Theorem can be overcome by restricting the freedom of individuals participating in the process by structuring their preferences in some way. Ranking all alternatives on an external scale as discussed above is one form of structure; allowing limited veto power is another. Various options for structuring preferences will be discussed below.

4.1.5 Weights in the social choice problem

Before turning to the multi-criteria decision problem, there is one non-solution to the problem that bears discussion. It may be tempting to resolve the dilemma with weights, or with some measure of strength of feeling of the individuals involved. This amounts to the assignment of a numerical ordering in the place of the weak order of alternatives by individuals. For instance, in the motivating paradox, instead of a simple weak order, let each individual rank each option on a scale from 0 to 1 and use a weighted sum or some other aggregation. In many cases, the resulting order will be transitive. For another example, consider votes by elected representatives on measures brought before them, where strength of feeling comes into play when a member is willing to vote for a measure he weakly opposes in order to gain support for another measure he strongly supports.

Neither of these examples resolves the basic difficulty of the Impossibility Theorem; each rather recasts the problem to a particular instance where the paradox is hidden. Indeed, in each case the problem can be expanded to one where the Impossibility Theorem applies

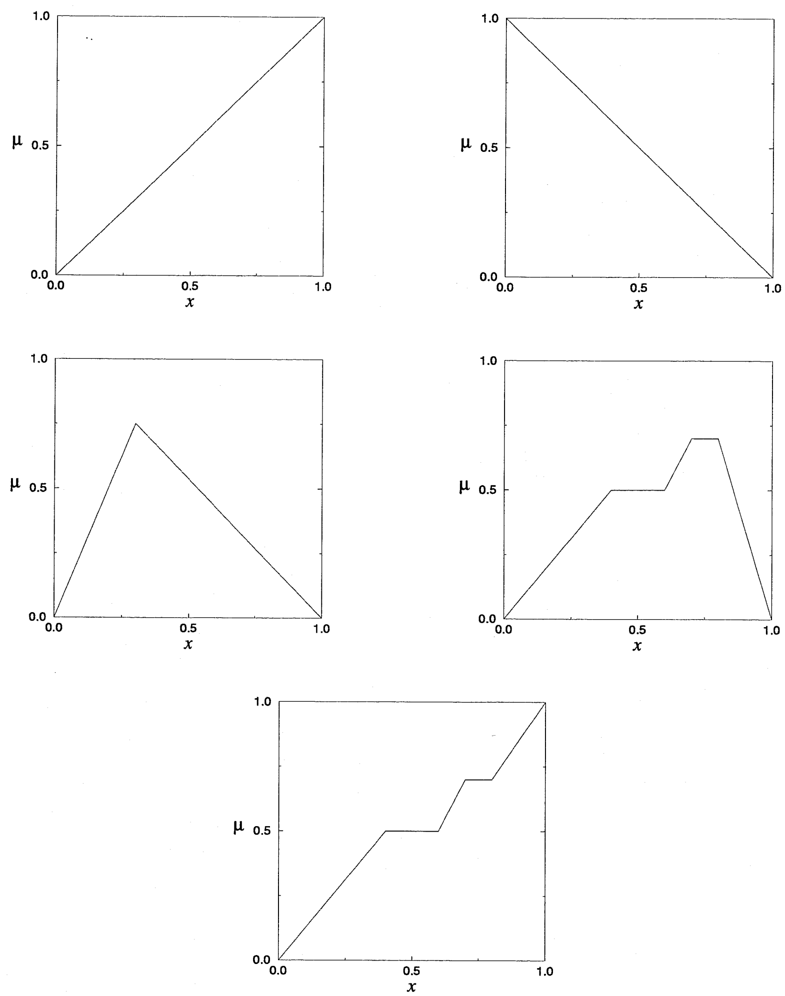

Figure 4.1: Single-peaked functions

	1st	2nd	3rd
Voter 1	A	B	C
Voter 2	B	C	A
Voter 3	C	A	B

Table 4.1: Weak orders of three voters

directly. In the case of the elected official, one must consider the outcomes on all votes at once (though they are separated in time) to get a weak order for each representative. In the case of summed weighted rankings of the alternatives, a little arithmetic reformulates the problem as one with more decision makers with simple weak orderings. Consider the motivating paradox of Section 4.1.1, with the individual voters' weak orders shown again in Table 4.1. A majority pairwise vote to combine these three weak orders leads to an intransitive, and thus untenable answer. Weights and a measure of strength of feeling can be added to the problem: let us suppose that each voter is given 10 points to distribute among the three alternatives, and that Voter 1 is assigned double the weight of the other two voters. A representative combination with a weighted sum is shown here:

	A	B	C	weight
Voter 1	6	3	1	2
Voter 2	2	5	3	1
Voter 3	2	0	8	1
Total	16	11	13	

In this case, there is no ambiguity: A is clearly preferred to B and C, and B is clearly preferred to C.

This solution, however, is accidental. If Voter 3, for instance, holds the slight different preferences shown here, preferences which are still consistent with the weak order in Table 4.1, there is no longer a clear choice between B and C:

	A	B	C	weight
Voter 1	6	3	1	2
Voter 2	2	5	3	1
Voter 3	2	1	7	1
Total	16	12	12	

Indeed, there are numerical preferences consistent with the weak orders in Table 4.1 that end in indifference among all three alternatives.

Any of these sets of numerical preferences can be recast as weak orders held by more voters. To express 10 preference points requires 20 individual weak orders: for example, if alternative A received 1 of the 10 possible points, that is expressed by one weak order $A \succ B \succ C$, and one $A \succ C \succ B$. (If there are more than three alternatives, then more weak orders are required.) Voter 2's numerical preferences, for instance, are equivalent to the following 20 weak orders:

1st	2nd	3rd	instances of this order
A	B	C	2
A	C	B	2
B	C	A	5
B	A	C	5
C	A	B	3
C	B	A	3

Voter 1 carries double the weight and will require 40 orders to capture their numerical preferences. The entire weighted sum aggregation is equivalent to the majority method applied to 80 individual voters. It is not surprising that the paradoxical situation of three voters is resolved by some of the many ways that the situation can be rewritten with 80 voters. Nevertheless, Arrow's Theorem still applies: the weighted sum method does not guarantee transitivity. Nor is the difficulty overcome by aggregating by other than a weighted sum: the arithmetic to recast the problem may be more complicated, but it is not deeper.

Arrow assumes, appropriately for the social choice problem, that interpersonal comparisons of criteria are meaningless, so that one alternative cannot be chosen because one individual would derive such great satisfaction from it that their "happiness" (or utility) would be greater than the loss felt by the other two who (only somewhat weakly) prefer a second alternative. Clearly, in many decision environments, including engineering design, strength of feeling or preference *is* important. However, the foregoing discussion shows that simply incorporating a value function as a numerical measure of strength of feeling is not sufficient to resolve the dilemma highlighted by the Impossibility Theorem. The difficulties raised by the Impossibility Theorem will be addressed in the following sections, and the solution will include the specification of value functions in place of weak orders. It is not the mere assignment of numbers for preference, however, that resolves the difficulties raised by the Impossibility Theorem, but the careful definition of the engineering design decision problem and the interpretation of the preference scale in that context.

Arrow's work raises several issues for design decision theory. First, it shows that it is hard to construct a provably consistent decision method: the assumptions are few, and yet contradictions arise. Second, it shows that consistency can be reached with additional structure. An axiomatic basis for design decision making must make explicit the structural peculiarities of design decision problems.

4.2 Decision with multiple criteria

The problem of decision with multiple criteria is to rank a number of alternatives, each of which is ranked separately by several ranking criteria. It is sometimes called the Multi Objective Decision Making (MODM) or Multi Attribute Decision Making (MADM) problem. The ranking of alternatives on the basis of each objective is assumed to be given, and the problem is to define an overall ranking based on some combination of the individual criteria.

Though this problem appears superficially similar to the social choice problem, since it seeks to combine several individual rankings into one, it is a distinct problem. Two differences are:

- In the social choice problem, all orderings are accorded equal worth. In the multiple criteria problem, it is desirable to be able to assign importance weightings to criteria. While it is natural to accord all human voters equal worth, there is no obvious reason to require equal weighting of the different criteria that describe a design.

This difference, as mentioned above, diminishes or even disappears if the weighted problem is defined as an unweighted problem with more individuals.

- The social choice problem permits no interpersonal comparison of utilities (preferences), and is thus limited to the discussion of weak orders. The heart of the multiple criteria problem is the inter-attribute comparison of preferences. When considering many design goals, it is crucial to understand their relative importance and the way in which they interact. Again, what is natural to require when modelling the sovereignty of individual citizens is not necessarily applicable to separate design criteria.

  This difference makes decision with multiple criteria structurally different from social choice, and has deep implications for the applicability of the Impossibility Theorem to the former problem.

Even the informal motivating paradox for the Impossibility Theorem (where a majority vote ranked A before B before C before A) loses much of its power if cast in the framework of multi-criteria decision making. Consider the analogous example of a design or a product that is to be judged on the basis of three criteria: X, Y, and Z. It is certainly plausible to assume that the designer may be faced with a choice of three candidate designs A, B, and C such that A is better than B is better than C with respect to criterion X, B is better than C is better than A with respect to Y, and C is better than A is better than B with respect to Z. The analogous "paradox" here is that giving X, Y, and Z one vote each as a method to determine the best design yields no obvious answer. In other words, if all that is known about a design is one weak order among alternatives for each of the three criteria X, Y, and Z, then there is not enough information to decide upon an overall best design. This "paradox" is resolved in the M_oI by more careful consideration of preferences for X, Y, and Z, and the knowledge of how those preferences will be aggregated; this is presented in more detail in Chapter 5.

In general, in a real design situation, there *will* be a (rational) weak order among A, B, and C. The "paradox" is merely that additional information beyond the weak orders on X, Y, and Z is required to recognize the overall order. The question asked here is whether and when it is possible to find consistent, rational techniques to discover the ranking among A, B, and C. If so, it will necessarily be with a slightly different set of axioms from the ones of the Impossibility Theorem, axioms more appropriate for engineering decision making than for social choice.

A careful examination of the axioms is necessary before considering the Impossibility Theorem in the context of the design decision problem. When combining engineering criteria (the MADM problem) rather than individual orderings (the social choice problem), the axioms of positive response, independence of irrelevant alternatives, and inadmissibility of dictatorial solutions still seem to hold. However, it is not so obvious that domains must be unrestricted or that orders must not be imposed. Consider two motivating examples:

1. Preferences for engineering requirements are commonly single-peaked (see Figure 4.1) around an ideal target; weight is an example, as is stiffness. Indeed, nearly all engineering requirements are of one of three forms: less is better, more is better, or closer to a particular target is better [16]. All three of these forms are single-peaked.

2. Designs have constraints. A maximum stress indicates the point at which a design breaks and fails; government regulations must be fulfilled or a design is not allowed on the market. The positions of alternatives that fail on the basis of a single criterion are thus imposed (to be last) in any aggregated order.

Not all evaluation criteria will behave in a single-peaked manner. A design preference for availability of a particular material stock may be one criterion for a design, and it may change over time and take on any order. The preference for the frequency of the first acoustic mode is often to *avoid* a particular unpleasant range. However, single-peaked criteria are common, and designers often can and do restrict criteria that are not globally single-peaked to regions of local single-peakedness. The vehicle structure designer seeking to avoid a particular range of frequencies of the first acoustic mode, for example, chooses to target either higher or lower frequencies, thus considering only a range over which the criterion is single-peaked. Thus while the completely generic design decision problem should obey Arrow's axiom of unrestricted domain, designers strive to avoid the generic problem, and rather to cast each problem so that domains, rather than being unrestricted, are single-peaked along the obvious external scales provided by the design parameterization. Indeed, in terms of the decision problem, the parameterization of a design serves to restrict domains. For the multi-criteria decision problem, the axiom of unrestricted domain is replaced by an exhortation to the designer to verify that criteria are single-peaked, or restrict the problem until they are. It is understood that this may not always be possible, but when it is not, the designer realizes that the design problem is not completely well-conditioned. These problems are difficult for formal methods and informal methods alike.

Constraints in engineering design, if translated into social choice terms, are a sort of veto that individual criteria may exercise over the entire design, and thus violate the axiom of no imposed orders. However, when the design is acceptable on the basis of each individual criterion, there is no reason to abandon the axiom of no imposed orders. For the multi-criteria engineering decision problem, the axiom of no imposed orders is weakened:

Axiom 4.4a (limited imposed orders) Axiom 4.4 holds, with the exception that some alternatives may be declared unacceptable, and thus last in any combined ("social") order, on the basis of an unacceptable ranking on a single criterion. All unacceptable alternatives are equally unacceptable.

These two differences between the axioms for social choice and the axioms for multi-criteria decision making are sufficient to make the results of Arrow's Theorem inapplicable to the engineering design decision-making problem.

Decision with multiple criteria differs empirically from social choice in an important way. In the former, there is always a well-defined aggregated order that takes into account all dimensions, and which is available in principle to anyone with the time and resources to query the decision maker directly about all possible combinations; in the latter, Arrow's theorem calls into question the very existence of a well-defined aggregated order. A direct specification of preference in many dimensions in the multi-criteria problem presents no more theoretical difficulties than a direct specification of preferences in one dimension; the practical implementation, however, can present great difficulty. The general problem of preference in many dimensions can be approached in several ways. Utility theory assumes the aggregated order is primary (and there is a single omniscient decision maker), and seeks conditions under which the aggregated order can be simply expressed in terms of orders on the individual criteria [43]. The M_oI and other systems (such as QFD [38]) address the case when the overall order is not directly available, but must be constructed from information about individual criteria. The negotiation problem presented in this thesis is on intermediate ground between the case of an omniscient decision maker and the case of antagonistic parties who share information only strategically.

The disturbing conclusion of the Impossibility Theorem for the social choice problem is twofold: first, it is impossible to construct a method to arrive at a fair and rational social order, and second, such an order may not exist. Clearly, an aggregated order exists in the MADM problem. Does Arrow's Impossibility Theorem shed any light on the search for a method that can construct the aggregated order from orders on the individual attributes?

The axioms of the MADM problem are not identical to those of the social choice problem, and some of the differences are known to invalidate the Impossibility Theorem. There are clear structural differences between the two cases. The social choice problem does not admit interpersonal comparison; the M_oI problem would be meaningless without it. The social choice problem must respect individuals by imposing no structures on the orders; in a design situation, cultural, customer, or managerial structure is almost always imposed. For instance, if three candidate vehicle structure designs (A, B, and C) have bending stiffnesses of 3000, 3200, and 3400 N/mm respectively, Axiom 4.4 states that any individual is free to prefer C over both A and B, and to prefer A over B. A vehicle structures group, however, which proposed this order to management, would be taken to task for "irrational" preferences *over bending stiffness*. This order, ranking C before A before B, is transitive, and any transitive order must be considered rational in the social choice problem; it would be an acceptable final order of candidate designs. With respect to the particular evaluation criterion of maximizing bending stiffness, however, it is not rational; many such transitive orders would be considered irrational in an engineering context. No individual is given veto power in the social choice context;[3] almost any attribute of an engineered design has a level so unacceptable as to veto the entire design.

Axioms 4.1 and 4.4 are crucial to the social choice problem. In the MADM problem,

---

[3] Some social systems, such as consensus, rely on (judiciously exercised) individual veto to prick the group conscience.

they disappear, or appear only in a modified form. It is this difference in axioms that allows MADM methods to operate, at least on some large classes of problems, without violating the conclusions of the Impossibility Theorem.

4.3 Individual preferences

There is more than one way to assign a value function when a weak order among alternatives is given. The assignment of a value function, and the interpretation of the numerical scale, are determined by the intended use. In this section, two distinct approaches to the assignment of a value function will be discussed: the *preferences* of the M_OI, and the *utilities* (or *von Neumann–Morgenstern utilities*) of economic theory. Both depend ultimately on weak orderings among alternatives, but the underlying assumptions are different, as are the interpretations.

4.3.1 Utility

The specification of utility depends on a weak order among alternatives, and on the mathematics of expectation. In order to determine a utility function, the so-called *lottery question* must have an answer: "Given that A is preferred to B, and B is preferred to C, at which probability p is there indifference between the two choices 'B with probability 1' and 'a lottery that yields A with probability p and C with probability $(1 - p)$'?" (Note that the question need not have a *direct* answer; see [94], for instance, for a discussion of the elicitation of von Neumann–Morgenstern utilities with little or no probability information.) Von Neumann and Morgenstern [93] show that, given the assumption that utilities combine with the mathematics of expectation, the numerical utility scale is determined up to an affine transformation.

The assumption of the use of mathematical expectation arises because utility theory is intended to treat questions of decision making under probabilistic uncertainty, such as those that are germane to gambling. This makes the specification of *relative* utilities (relative, that is, to each other) with probabilities natural.

However, the development of utility specifically excluded the notion of interpersonal comparison of utility as too difficult to address:

> We do not undertake to fix an absolute zero and an absolute unit of utility. [93, footnote, p. 25]

Utility theory is intended for use in decision making under uncertainty or risk, rather than as a solution of the multi-criteria decision making problem. Certainly the specification of individual utilities does not answer, by itself, the negotiation problem. Even in the case of a single decision maker, where utilities can be assumed to be comparable, utility theory derives conditions under which the decision maker's overall utility can be determined to be a simple function of individual, independent utility functions [42].

The success of the von Neumann–Morgenstern utility paradigm, and the ease of its application in terms of quantified risk, have led to the identification of many decision problems with problems of (economic) decision making under uncertainty. The lottery question seems natural, and so it is assumed that the lottery question is the right way to impose a numerical scale on preferences. The need for this particular numerical scale to make decisions under uncertainty is hidden, and other scales are not considered. Nevertheless, design may *not* be best classified as decision making under risk and uncertainty. Utility theory is one paradigm for decision making, appropriate for a particular set of problems, those where the "estimation of expectations for each option" is the most pertinent information. When design reaches the manufacturing stage, and probability distributions over manufacturing tolerances are the most relevant uncertainties, the design decision problem is much closer to the problem addressed by utility theory. Earlier in the design process, where uncertainty will be resolved by refinement of a design alternative, rather than by random selection from a perceived distribution among alternatives, a utility model is less appropriate.

From the point of view of classical utility theory, the design decision problem described in Definition 3.1 would be a case of decision making under certainty. The construction of the utility function and the choice of a "best" solution with limited computation would be uninteresting problems.

4.3.2 M_oI preference

Preferences specified in the M_oI also depend on a weak order among alternatives. Preferences, unlike utilities, are expressed on an absolute scale, where a preference of $\mu = 1$ indicates a completely acceptable value, and $\mu = 0$ a completely unacceptable value. The negotiated design decision combines many individual preferences into a single, overall preference. The individual preference orders may be generated by or associated with different people or groups involved in the design, but they are distinguished by the attributes that form the rationale for each preference. Since "interpersonal" (actually, inter-attribute) comparison of preference is required, all preference values must have consistent meaning. A preference level of α must mean the same thing regardless of which attribute it is based on, and regardless of which person or group provides the number.

Empirical studies into sensation (a heading that can be reasonably supposed to include preference) have shown that human beings are, in general, capable of sorting into 7 ± 2 categories [50]. Direct specification of an entire preference curve would then be unrealistic, though it is mathematically convenient to assume its existence at every point. Since a limited number of points (around seven) can thus be fixed directly, it is often most convenient to assume that preferences are piecewise linear between those points. A lottery method cannot fix points on an absolute scale, but might be used to fill in between the fixed points.

As preference will play a large role in following chapters, it is helpful to give a formal definition:

Definition 4.3 *A* preference *(an* M_oI preference*) is a map* $\mu : \mathbf{X} \rightarrow [0, 1]$*, where* $\mathbf{X}$ *is a set and will usually be identified with the design space* $\mathbf{D}$ *or the performance space* $\mathbf{P}$*.*

Preferences specified on the design space will be $\mu_{\mathbf{D}}$, preferences specified on the performance space will be $\mu_{\mathbf{P}}$. As discussed in Chapter 3, when there is a map $f : \mathbf{D} \rightarrow \mathbf{P}$ that specifies a preference for each design, there are also induced maps $\mu_{\mathbf{P}} : \mathbf{D} \rightarrow [0, 1]$ defined by $\mu_{\mathbf{P}}(\mathbf{d}) = \mu_{\mathbf{P}}(f(\mathbf{d}))$, and $\mu_{\mathbf{D}} : \mathbf{P} \rightarrow [0, 1]$, defined by the extension principle:

$$\mu_{\mathbf{D}}(\mathbf{p}) = \sup_{\mathbf{d} \mid f(\mathbf{d}) = \mathbf{p}} \mu_{\mathbf{D}}(\mathbf{d})$$

When there are coordinates for $\mathbf{X}$ and preferences are independent with respect to some of the coordinates, those preferences can be specified with respect to subsets of $\mathbf{X}$. In particular, when $\mathbf{X} = \mathbb{R}^n \times \mathbb{Z}^m$, μ may map one copy of $\mathbb{R}$ or $\mathbb{Z}$ to $[0, 1]$, representing the expression of preference on a single variable.

4.4 Summary

The aim of this chapter was to introduce the formal notion of preference, and to establish the legitimacy of comparing preferences in the engineering design problem, even though such comparison is not permissible in the social choice problem. The next chapter delves deeper into the question of legitimate aggregations of preference to effect such comparison.

The formalism for preference employed in this thesis is the absolute scale on the interval $[0, 1]$ used previously in the MoI. Because this formalism differs from the well-known formalism of preference through utility functions, some differences between the two were discussed. Utilities are specified so as to be useful in the selection of alternatives when the outcomes of those alternatives are not certain, but the probability distribution of consequences is known. On the other hand, utilities are not meant for inter-attribute comparison. The preferences of the MoI are not intended for use in expectation calculations, but their use of an absolute scale suits them for inter-attribute comparison and for the natural inclusion of constraints.

Both the social choice problem and the engineering design decision problem have another dimension that has not been dealt with here, that of *manipulability*. Can individuals benefit by misrepresenting their preferences? Game theory [93] considers the problem of strategic sharing of information to maximize personal gain. The negotiation problem considered in this thesis occupies an intriguing middle ground between cooperative and competitive games: the participants should have a common goal, but may not always perceive that. Indeed, it is a managerial challenge in any large design effort to prevent the individuals involved in the design from developing a set of preferences that is contrary to the group effort. The use of the model of negotiations presented here does not, by itself, force all participants to act in the common self-interest. It does, however, provide a more thorough structural analysis of the decision problem, and can thus serve as a tool for those who wish to enforce honesty.

Chapter 5

Aggregation of Preference

At the heart of the negotiation problem lies the aggregation of preference. Individual preference functions, each representing a particular performance measure or point of view, must be reconciled into a single, overall preference function. This chapter considers the problem of choosing an aggregation function for multiple criteria decision making in general, and for design decision making in particular. The results are presented within the context of the M_oI, but they have general applicability in MADM and MODM approaches. The axioms of the M_oI and the reasons for their use in modelling engineering design decisions will be presented. In the context of these axioms, potential aggregation functions will be discussed. The M_oI formerly used two different aggregation functions to model two different design trade-off strategies. Here the range of possible aggregation functions suitable for such decision making will be explored. A parameterized family of functions that satisfies all the axioms for design and that models a continuum of strategies between the two existing strategies of the M_oI will be presented. The possibility of suitable aggregation functions outside this range will be discussed, and an example will be given. Having found a complete range of design-appropriate aggregations, the question of which function to choose will also be addressed.

5.1 Aggregation in the M_oI

An engineering decision was described earlier as a choice among alternatives, or more generally as a search for the "best performing" (sets of) alternative(s). Definition 4.3 made formal the notion of preference as a map $\mu : \mathbf{X} \rightarrow [0,1]$. The interpretation of preference is that $\mathbf{x}_1$ is preferred to $\mathbf{x}_2$ if $\mu(\mathbf{x}_1) > \mu(\mathbf{x}_2)$. However, there will typically be several different preferences associated with a design, and in general it is impossible to maximize all preferences simultaneously. Thus it is necessary to formalize the notion of an aggregated preference that takes all individual preferences into account. Thus the following definition:

Definition 5.1 *For a set of preferences* $\mu_i : \mathbf{X} \rightarrow [0,1]$, *the* aggregation *of the* μ_i *is itself a preference* $\mu_{\text{agg}} : \mathbf{X} \rightarrow [0,1]$, *defined (pointwise) by a functional* aggregation operator $\mathcal{P} : PF[\mathbf{X}]^n \times \mathbb{R}^{+n} \setminus \{0\} \rightarrow PF[\mathbf{X}]$

$$\mu_{\text{agg}}(\mathbf{x}) = \mathcal{P}(\mu_1, \ldots, \mu_n; \omega_1, \ldots, \omega_n)(\mathbf{x})$$

where the parameters $\omega_i \in \mathbb{R}^+$ *are weights, and* $PF[\mathbf{X}]$ *is the space of preference functions on* $\mathbf{X}$, *the set* $\{\mu | \mu : \mathbf{X} \rightarrow [0,1]\}$. *An aggregation operator must satisfy:*

$$\mathcal{P}(\mu_1, \ldots, \mu_n; \omega_1, \ldots, \omega_n)(\mathbf{x}_1) = \mathcal{P}({\mu_1}', \ldots, {\mu_n}'; \omega_1, \ldots, \omega_n)(\mathbf{x}_2)$$

whenever $\mu_i(\mathbf{x}_1) = {\mu_i}'(\mathbf{x}_2)$ *for all* i.

When the aggregated preference μ_{agg} is the combined overall preference that considers all individual preferences describing the performance of a design, it is referred to as the *overall preference* for that design, and is denoted $\bar{\mu}$.

The aggregation of preference is defined for an arbitrary (finite) number of individual preference functions. The following discussion and results shall consider the case of aggregation of exactly two individual preferences. Aggregation of more than two preferences is accomplished hierarchically. The operators that are used for aggregation extend simply to several arguments; nevertheless, in the case where several preferences are assigned hierarchically using different aggregation operators, the overall preference is in general not a simple extension of a constituent operator. Note that the final condition of the definition, that the aggregation be identical whenever the constituent preference values are the same, will allow the identification of an aggregation operator with a function on the space of preference values.

The M$_o$I has previously used two different aggregation functions to model two different situations in decision making in design. When the overall preference for the performance of a design is limited by the attribute with the lowest performance, the decision-making problem is said to be *non-compensating*, and the aggregation function used has been the simple minimum:

$$\mu_{\mathrm{agg}}(\mathbf{x}) = \min\Big(\mu_1(\mathbf{x}), \mu_2(\mathbf{x})\Big)$$

(In this case, all non-zero weights are equivalent.) When good performance on one attribute is perceived to partially compensate for lower performance on another, the problem is called *compensating*, and the weighted geometric mean or product of powers has been used:[1]

$$\mu_{\mathrm{agg}}(\mathbf{x}) = \Big(\mu_1(\mathbf{x})^{\omega_1}\mu_2(\mathbf{x})^{\omega_2}\Big)^{\frac{1}{\omega_1+\omega_2}}$$

The definition of a functional aggregation operator $\mathcal{P}$ places no *a priori* restrictions on the behavior of $\mathcal{P}$ beyond the specification of its domain and range. Nevertheless, not all functional operators are appropriate for engineering design decisions. In the following sections, axioms for engineering design aggregation will be presented, and those axioms will then be solved (in the sense that a characterization of necessary and sufficient conditions for appropriate functions will be given).

5.2 The axioms of the M$_o$I

In any multi-criteria decision system, it is desirable that the aggregation functions used be justifiable models for decision-making behavior. The choice of an aggregation function may be justified in several ways. Empirical tests, such as those conducted by other researchers [84], can help determine which aggregation functions best model human decision making in various contexts. Computational simplicity is often used as a basis for the choice of an aggregation function, for practical rather than fundamental reasons. As the M$_o$I is a *formal* theory, its development has been to appeal to intuitive notions of rational human be-

[1]Both of these aggregations are analogous to Pareto-optimal solutions in game theory [49, 93], and the product of powers is analogous to a Nash solution. However, neither correspondence is mathematically precise, since preferences are not equivalent to utilities.

At each point x the following hold:	
AO.1	Monotonicity: $\mathcal{P}(\mu_1, \mu_2; \omega_1, \omega_2)(\mathbf{x}) \leq \mathcal{P}(\mu_1, \mu_2'; \omega_1, \omega_2)(\mathbf{x}) \;\; \forall\, \mu_2(\mathbf{x}) \leq \mu_2'(\mathbf{x})$ $\mathcal{P}(\mu_1, \mu_2; \omega_1, \omega_2)(\mathbf{x}) \leq \mathcal{P}(\mu_1, \mu_2; \omega_1, \omega_2')(\mathbf{x}) \;\; \forall\, \omega_2 \leq \omega_2'; \;\; \mu_1(\mathbf{x}) < \mu_2(\mathbf{x})$
AO.2	Symmetry: $\mathcal{P}(\mu_1, \mu_2; \omega_1, \omega_2)(\mathbf{x}) = \mathcal{P}(\mu_2, \mu_1; \omega_2, \omega_1)(\mathbf{x})$
AO.3	Continuity: $\mathcal{P}(\mu_1, \mu_2; \omega_1, \omega_2)(\mathbf{x}) = \lim_{\mu_2'(\mathbf{x}) \to \mu_2(\mathbf{x})} \mathcal{P}(\mu_1, \mu_2'; \omega_1, \omega_2)(\mathbf{x})$ $\mathcal{P}(\mu_1, \mu_2; \omega_1, \omega_2)(\mathbf{x}) = \lim_{\omega_2' \to \omega_2} \mathcal{P}(\mu_1, \mu_2; \omega_1, \omega_2')(\mathbf{x})$
AO.4	Idempotency: $\mathcal{P}(\mu, \mu; \omega_1, \omega_2)(\mathbf{x}) = \mu(\mathbf{x}) \;\; \forall\, \omega_1 + \omega_2 > 0$
AO.5	Annihilation: $\mathcal{P}(\mu, 0; \omega_1, \omega_2)(\mathbf{x}) = 0 \;\; \forall\, \omega_2 \neq 0$
AO.6	Self-scaling weights: $\mathcal{P}(\mu_1, \mu_2; \omega_1 t, \omega_2 t)(\mathbf{x}) = \mathcal{P}(\mu_1, \mu_2; \omega_1, \omega_2)(\mathbf{x}) \;\; \forall\, \omega_1 + \omega_2, t > 0$
AO.7	Zero weights: $\mathcal{P}(\mu_1, \mu_2; \omega_1, 0)(\mathbf{x}) = \mu_1(\mathbf{x}) \;\; \forall\, \omega_1 \neq 0$

Table 5.1: Axioms of the M_oI for aggregation operators

havior, and to formalize this rationality in a set of axioms that its aggregation functions must follow. The axioms of the M_oI (see Table 5.1) [56] are a formal description of restrictions on any preference aggregation operator for (rational) engineering design.

There is no universally accepted definition of rationality. A formal decision-making system presents its definition of rationality in its formal axioms. This thesis follows Tribus [89] in holding rational behavior to be behavior which is consistent with the pursuit of the stated objectives. The monotonicity axiom is a clear example of rational behavior: the designer who wishes to maximize preferences would be irrational to prefer one alternative over another if the second offered higher preferences on all attributes.

Following Definition 5.1, a preference aggregation operator $\mathcal{P}$ takes as its inputs two preference functions μ_1 and μ_2, and returns an aggregated preference function μ_{agg}. There are also two parameters, ω_1 and ω_2, which represent weights. The axioms in Table 5.1 are expressed pointwise. The aggregated preference at a point $\mathbf{x}$ is often calculated by composition:

$$\mathcal{P}(\mu_1, \mu_2; \omega_1, \omega_2)(\mathbf{x}) = \hat{\mathcal{P}}(\mu_1(\mathbf{x}), \mu_2(\mathbf{x}); \omega_1, \omega_2)$$

and it is convenient to consider the map $\hat{\mathcal{P}} : [0,1]^2 \times \mathbb{R}^{+2} \setminus \{0\} \rightarrow [0,1]$ which operates directly on preference values, rather than the aggregation operator $\mathcal{P}$. The requirement that aggregated preferences be consistent for all $\mathbf{x}$ implies that $\mathcal{P}$ and $\hat{\mathcal{P}}$ are equivalent. Call $\hat{\mathcal{P}}$ the *aggregation function* corresponding to the *aggregation operator* $\mathcal{P}$.

Proposition 5.1 *For any aggregation operator $\mathcal{P}$, there is a unique $\hat{\mathcal{P}}$ such that*

$$\mathcal{P}(\mu_1, \mu_2; \omega_1, \omega_2)(\mathbf{x}) = \hat{\mathcal{P}}\left(\mu_1(\mathbf{x}), \mu_2(\mathbf{x}); \omega_1, \omega_2\right)$$

Proof of Proposition 5.1 *The result follows immediately from the consistency requirement in Definition 5.1.*

All the axioms in Table 5.1 can be rewritten simply in terms of $\hat{\mathcal{P}}$ (see Table 5.2). Since an aggregation function can be identified with an aggregation operator, where context is clear the aggregation function may sometimes be written $\mathcal{P}$.

The axioms of monotonicity, symmetry and continuity are common to many multi-attribute decision-making schemes. Monotonicity formalizes the notion that an increase in preference for one aspect of a design should never result in a decrease of overall preference. Symmetry, like the consistency condition imposed in the definition of an aggregation operator, formalizes the idea that the overall preference should depend only on the individual preferences assigned, not on the reasons that those preferences were specified. Continuity of aggregation states that continuous inputs should yield continuous outputs: there may be discontinuities in an aggregated preference, but they should arise from discontinuities in individual preferences, not from the aggregation itself.

The axioms that refer to weights are conventionally accepted. An attribute with a weight

AF.1	Monotonicity: $\hat{\mathcal{P}}(\alpha_1, \alpha_2; \omega_1, \omega_2) \leq \hat{\mathcal{P}}(\alpha_1, \alpha_2'; \omega_1, \omega_2) \ \forall\, \alpha_2 \leq \alpha_2'$ $\hat{\mathcal{P}}(\alpha_1, \alpha_2; \omega_1, \omega_2) \leq \hat{\mathcal{P}}(\alpha_1, \alpha_2; \omega_1, \omega_2') \ \forall\, \omega_2 \leq \omega_2';\ \ \alpha_1 < \alpha_2$
AF.2	Symmetry: $\hat{\mathcal{P}}(\alpha_1, \alpha_2; \omega_1, \omega_2) = \hat{\mathcal{P}}(\alpha_2, \alpha_1; \omega_2, \omega_1)$
AF.3	Continuity: $\hat{\mathcal{P}}(\alpha_1, \alpha_2; \omega_1, \omega_2) = \lim_{\alpha_2' \to \alpha_2} \hat{\mathcal{P}}(\alpha_1, \alpha_2'; \omega_1, \omega_2)$ $\hat{\mathcal{P}}(\alpha_1, \alpha_2; \omega_1, \omega_2) = \lim_{\omega_2' \to \omega_2} \hat{\mathcal{P}}(\alpha_1, \alpha_2; \omega_1, \omega_2')$
AF.4	Idempotency: $\hat{\mathcal{P}}(\mu, \mu; \omega_1, \omega_2) = \mu \ \forall\, \omega_1 + \omega_2 > 0$
AF.5	Annihilation: $\hat{\mathcal{P}}(\mu, 0; \omega_1, \omega_2) = 0 \ \forall\, \omega_2 \neq 0$
AF.6	Self-scaling weights: $\hat{\mathcal{P}}(\alpha_1, \alpha_2; \omega_1 t, \omega_2 t) = \hat{\mathcal{P}}(\alpha_1, \alpha_2; \omega_1, \omega_2) \ \ \forall\, \omega_1 + \omega_2, t > 0$
AF.7	Zero weights: $\hat{\mathcal{P}}(\alpha_1, \alpha_2; \omega_1, 0) = \alpha_1 \ \forall\, \omega_1 \neq 0$

Table 5.2: Axioms of the M_OI for aggregation functions

of zero should contribute nothing to the calculation of the overall performance. Self-scaling weights are convenient for the hierarchical combination of an arbitrary number of attributes; alternatively, one could create an axiomatic system that relied on normalized weights. However, two of the axioms, idempotency (A0.4 and AF.4) and annihilation (AO.5 and AF.5), are fundamental to engineering design decision making, and thus to the M_oI.

The idempotency axiom appeals to a notion of rational behavior. It states that if several identical individual preferences are combined, the overall preference must be the same as the (identical) preferences on the individual variables. Idempotency reflects the constraint that the overall preference for a design should never exceed the preference of the highest-ranked attribute, nor fall below the preference of the lowest-ranked attribute. Idempotency and monotonicity together lead to the requirement that $min \leq \hat{\mathcal{P}} \leq max$. Idempotency has significant implications for the specification of preferences: it enforces the ability to compare different attributes.

The annihilation axiom is also specific to engineering design, and others have argued its validity [13, 58, 92]. It states that if the preference for any one attribute of the design sinks to zero (unacceptable) then the overall preference for the design is zero. For example, the tensile strength limit for a material cannot be exceeded no matter how great the reduction in the design's cost or weight. This is in contrast to a decision-making situation in which all performances can be converted into monetary units; in the latter case, two goals can always be traded, or bought, off. The annihilation axiom must also be considered carefully in the specification of individual preferences: an attribute may have no value that descends to a preference level of zero.

One axiom that is *not* necessary for design-appropriate aggregation functions is an axiom of strict monotonicity, and such a requirement would be incompatible with annihilation. Thus the M_oI allows for some indifference in aggregation. The non-compensating function *min* is an example of a function that fulfills all of the axioms of the M_oI and is not strictly monotonic.

The axioms of idempotency and annihilation set the M_oI apart from other multi-attribute decision-making systems. It should be noted that these axioms simultaneously define the preference aggregation and provide interpretation for the specification of preference. These axioms should not, for instance, be applied to the aggregation of utilities, which are specified with regard to their potential use in situations of decision making with quasi-monetary outcomes and probabilistic uncertainty. Here the axioms of the M_oI are posited as rational for design. There may be decision-making problems, outside of the field of engineering design, for which these axioms do not supply a rational model. Nevertheless, the fundamental results on aggregation discussed here are relevant to any other (even non-design) MADM schemes that obey similar axioms.

Definition 5.2 *An aggregation operator (function) is termed* design-appropriate *if it satisfies the axioms in Table 5.1 (Table 5.2).*

Having defined design-appropriate functions, it is natural to seek a complete functional characterization of such functions. This is undertaken in Section 5.4, after a digression to discuss some other MADM systems.

5.3 Fuzzy multi-attribute decision making

M_oI preferences and membership functions for fuzzy sets are both maps $\mathbf{X} \rightarrow [0,1]$. A designer's preference for particular values of a design variable defines a fuzzy set that might be called "Values of design variable $\mathbf{d}$ preferred by the designer." Fuzzy sets have been used in the M_oI since its inception, to represent imprecise quantities. The problem of aggregation of preference in the M_oI is thus a multi-attribute decision-making problem with fuzzy sets, sometimes called the *fuzzy MADM* problem. Other researchers have investigated this problem, and some of the literature on fuzzy aggregation methods is presented here.

In their recent book on the subject, Chen and Hwang [18] identify 18 fuzzy MADM methods, which they systematically classify into eight categories: simple additive weighting methods, the Analytic Hierarchy Process (AHP), the Conjunction/Disjunction method, MAUF, the General MADM method, the outranking method, maximin, and their own proposed MADM method. In their survey, Chen and Hwang do not draw the distinction observed by Zimmermann [109] between continuous Multi Objective Decision Making (MODM) problems and discrete Multi Attribute Decision Making (MADM) problems.

This thesis shall follow Chen and Hwang and use MADM to refer to the general problem, whether continuous or discrete.

Several of the methods surveyed by Chen and Hwang are similar to the M_oI. In addition, the application of utility theory [43] to decision problems bears some similarity to the M_oI and to the methods listed above. The possible application of utility theory to engineering design has been considered previously, and shown to be problematic [58]. Matrix methods such as QFD [38] and Pugh charts [67] also support decision making by simple additive aggregation over several requirements.

Aggregation operators are important in all MADM methods, from the most formal to the most casual. The arithmetic mean or weighted sum is popular in matrix methods and elsewhere, as it is simple to calculate. The *min* enjoys considerable popularity as well. Chen and Hwang provide an overview of commonly used aggregation operators; the original *min* and the product operators of the M_oI appear in their list, as do weighted sums. However, the general weighted means that shall be shown to solve the functional axioms do not.

5.4 Weighted means

Fuzzy set researchers have productively applied the study of functional equations [1] to explore t-norms and t-conorms [26]. This section applies the same general approach to design-appropriate aggregation functions: the intuitively reasonable set of axioms is translated into a set of functional equations, and these equations are then solved.

A promising class of functions is the class of weighted means. The properties that define weighted means are listed in Table 5.3. While weighted means are defined here as functions of two arguments, they can be extended to several arguments.

The properties of the weighted mean include all of the properties of design-appropriate aggregation functions with the exception of annihilation; a comparison of these properties with the axioms of the M_oI shows that **any weighted mean that satisfies annihilation is design-appropriate.** The properties of the weighted mean also include conditions that are not explicitly design axioms; nevertheless, these conditions are consistent with the axioms of the M_oI. The bisymmetry condition is a surrogate for commutativity and as-

WM.1	Idempotency: $\mathcal{P}(\mu,\mu;\omega_1,\omega_2)=\mu \quad \forall\, \mu,\omega_1,\omega_2$
WM.2	Internality: $\exists\, \mu_a<\mu_b$ such that $\forall\, \omega_1,\omega_2>0$ $\quad \mu_a=\mathcal{P}(\mu_a,\mu_b;1,0)<\mathcal{P}(\mu_a,\mu_b;\omega_1,\omega_2)<\mathcal{P}(\mu_a,\mu_b;0,1)=\mu_b$
WM.3	Homogeneity of weights: $\mathcal{P}(\mu_a,\mu_b;\omega_1 t,\omega_2 t)=\mathcal{P}(\mu_a,\mu_b;\omega_1,\omega_2)$ $\quad \forall\, \omega_1,\omega_2\geq 0;\ \omega_1+\omega_2,t>0$
WM.4	Bisymmetry: $\mathcal{P}\left(\mathcal{P}(\mu_1,\mu_2;\omega_1,\omega_2),\mathcal{P}(\mu_3,\mu_4;\omega_3,\omega_4);\omega_1+\omega_2,\omega_3+\omega_4\right)$ $\quad =\mathcal{P}\left(\mathcal{P}(\mu_1,\mu_3;\omega_1,\omega_3),\mathcal{P}(\mu_2,\mu_4;\omega_2,\omega_4);\omega_1+\omega_3,\omega_2+\omega_4\right)$
WM.5	Increasing in weights: $\mathcal{P}(\mu_a,\mu_b;\omega_1,\omega_2)<\mathcal{P}(\mu_a,\mu_b;\omega_1,\omega_3)$ $\quad$ for $\omega_2<\omega_3 \quad (\mu_a<\mu_b)$
WM.6	Increasing in variables: $\mathcal{P}(\mu_1,\mu_2;\omega_1,\omega_2)<\mathcal{P}(\mu_1,\mu_3;\omega_1,\omega_2)$ $\quad$ for $\mu_2<\mu_3,\ \ \omega_2\neq 0$

Table 5.3: Properties of the weighted mean

sociativity, and assures that $\mathcal{P}$ can be consistently defined for more than two arguments. Weighted means are strictly monotonic, which is a stronger condition than the monotonicity of the design axioms, but it can be verified that the other properties of the weighted mean are satisfied by arbitrary design-appropriate functions. **Thus any strictly monotonic design-appropriate aggregation function must be a weighted mean.** There are design-appropriate functions that are not weighted means, since they are monotonic but fail to satisfy strict monotonicity. Such operators are often conditional rather than algebraic; the *min* is but one example. The class of weighted means does not encompass any of these aggregation functions that are only weakly monotonic. However, we shall see that the weak monotonic operators presently used for design can be approximated arbitrarily closely by strictly monotonic operators.

The structure of the class of weighted means is described completely in the following theorem, proven in [1]. Note that the notational conventions of functional equations have not been preserved. Thus the generating function is here g, rather than the customary f, which is the performance measure of the MoI. Note as well that here μ represents a real number rather than a function:

Theorem 5.1 *The properties of the weighted mean are necessary and sufficient for the function $\mathcal{P}(\mu_1, \mu_2; \omega_1, \omega_2)$ to be of the form*

$$\mathcal{P}(\mu_1, \mu_2; \omega_1, \omega_2) = g\left(\frac{\omega_1 g^{-1}(\mu_1) + \omega_2 g^{-1}(\mu_2)}{\omega_1 + \omega_2}\right)$$

where $\mu_a \leq \mu_1, \mu_2 \leq \mu_b$; $\omega_1, \omega_2 \geq 0$; $\omega_1 + \omega_2 > 0$ and g is a strictly monotonic, continuous function with inverse g^{-1}.

It follows that the space of design-appropriate aggregation functions can be identified with the space of strictly increasing homeomorphisms on $\mathbb{R}$, and that any strictly monotonic design-appropriate function must have such a generating function g. For example, $g(t) = e^t$ (with $\mu_a = 1$ and $\mu_b = e$) generates the familiar weighted product of powers, denoted $\mathcal{P}_\Pi$ and also known as the geometric mean:

$$\mathcal{P}_\Pi(\mu_1, \mu_2; \omega_1, \omega_2) = (\mu_1{}^{\omega_1} \mu_2{}^{\omega_2})^{\frac{1}{\omega_1 + \omega_2}}$$

This function only satisfies the properties of the weighted mean for $\mu_1, \mu_2 > 0$, but it satisfies all of the design axioms, including annihilation, on the closed interval $[0, 1]$.

The "parameterization" of the set of all design-appropriate functions by strictly increasing homeomorphisms g is broader than is desirable. A parameterized family of equations of particular interest for design is generated by the functions $g(t) = t^{\frac{1}{s}}$, where s is a real number. The aggregation function so generated is

$$\mathcal{P}_s(\mu_1, \mu_2; \omega_1, \omega_2) = \left(\frac{\omega_1 {\mu_1}^s + \omega_2 {\mu_2}^s}{\omega_1 + \omega_2} \right)^{\frac{1}{s}}$$

Note that a weighted mean satisfies annihilation if and only if $g^{-1}(0)$ is unbounded for that function. For $s < 0$, $g^{-1}(t) = t^s$ is unbounded at $t = 0$ and $\mathcal{P}_s$ satisfies annihilation. Similarly, $\mathcal{P}_\Pi$ satisfies annihilation, as $g^{-1}(t) = \ln(t)$ is unbounded at $t = 0$. Figure 5.1 shows the behavior of $\mathcal{P}_s$ for several negative values of the parameter s, and for equal weights ($\omega_1 = \omega_2$). In this graph, $\mu_2 = 0.5$ is fixed and μ_1 varies from 0 to 1 along the x-axis. It is graphically evident, and easily shown analytically, that $\mathcal{P}_0$ is identical to the weighted product of powers $\mathcal{P}_\Pi$. Furthermore, as s tends to $-\infty$, $\mathcal{P}_s(\mu_1, \mu_2; \omega_1, \omega_2)$ tends to $\min(\mu_1, \mu_2)$, regardless of the weights.

Proposition 5.2 *$\mathcal{P}_0$ is identical to the weighted product of powers $\mathcal{P}_\Pi$:*

$$\lim_{s \to 0} \mathcal{P}_s(\mu_1, \mu_2; \omega_1, \omega_2) = \mathcal{P}_\Pi(\mu_1, \mu_2; \omega_1, \omega_2)$$

Proof of Proposition 5.2

$$\begin{aligned}
&\lim_{s \to 0} \mathcal{P}_s(\mu_1, \mu_2; \omega_1, \omega_2) \\
&= \lim_{s \to 0} \left(\frac{\omega_1 {\mu_1}^s + \omega_2 {\mu_2}^s}{\omega_1 + \omega_2} \right)^{\frac{1}{s}} \\
&= \exp^{\lim_{s \to 0} \ln \left(\frac{\omega_1 {\mu_1}^s + \omega_2 {\mu_2}^s}{\omega_1 + \omega_2} \right)^{\frac{1}{s}}} \\
&= \exp^{\lim_{s \to 0} \frac{1}{s} \ln \left(\frac{\omega_1 {\mu_1}^s + \omega_2 {\mu_2}^s}{\omega_1 + \omega_2} \right)} \qquad (5.1)
\end{aligned}$$

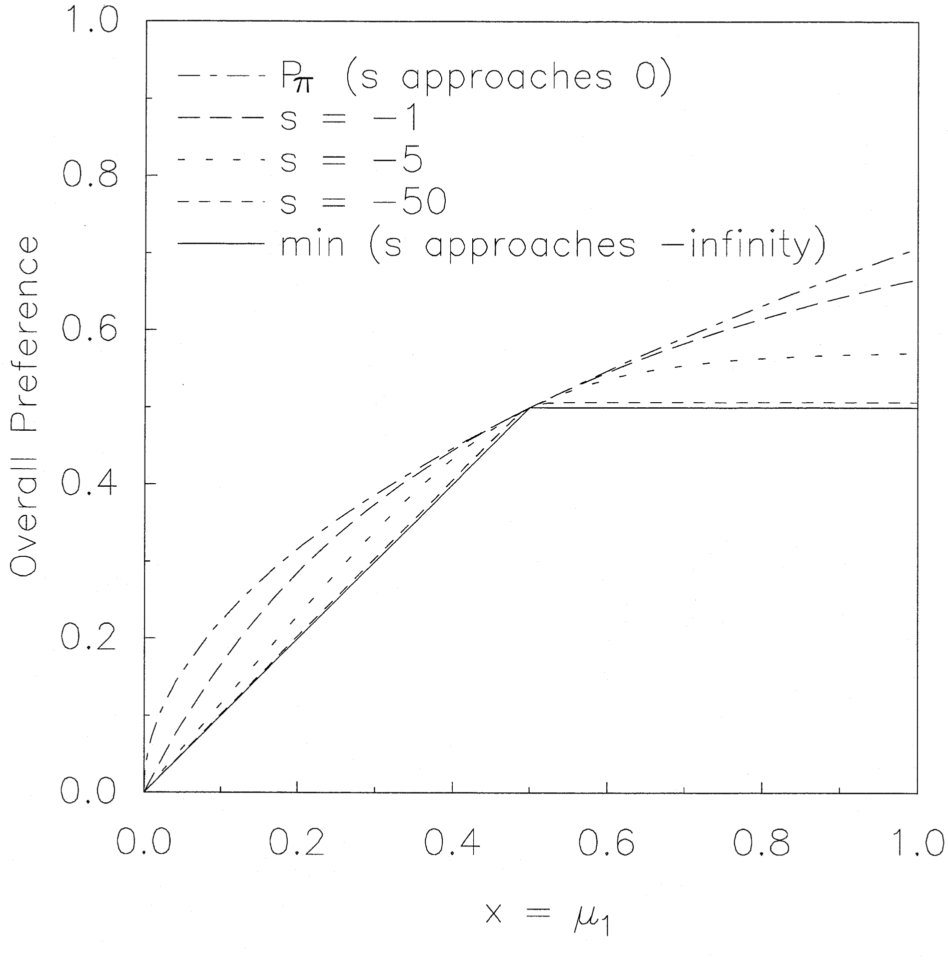

Figure 5.1: Functions between *min* and $\mathcal{P}_\Pi$

Note that, by the definition of the derivative,

$$\begin{aligned}
&\lim_{s \to 0} \frac{1}{s} \ln \left(\frac{\omega_1 {\mu_1}^s + \omega_2 {\mu_2}^s}{\omega_1 + \omega_2} \right) \\
&= \frac{d}{ds} \left[\ln \left(\frac{\omega_1 {\mu_1}^s + \omega_2 {\mu_2}^s}{\omega_1 + \omega_2} \right) \right]_{s=0} \\
&= \frac{\omega_1 \ln \mu_1 + \omega_2 \ln \mu_2}{\omega_1 + \omega_2}
\end{aligned}$$

Thus, proceeding from (5.1), it follows that

$$\begin{aligned}
&\lim_{s \to 0} \mathcal{P}_s(\mu_1, \mu_2; \omega_1, \omega_2) \\
&= \exp^{\frac{\omega_1 \ln \mu_1 + \omega_2 \ln \mu_2}{\omega_1 + \omega_2}} \\
&= ({\mu_1}^{\omega_1} {\mu_2}^{\omega_2})^{\frac{1}{\omega_1 + \omega_2}} \\
&= \mathcal{P}_\Pi(\mu_1, \mu_2; \omega_1, \omega_2)
\end{aligned}$$

which proves the proposition. ■

In light of the preceding proposition, the geometric mean $\mathcal{P}_\Pi$ will also be referred to as $\mathcal{P}_0$.

Proposition 5.3 *As* $s \to -\infty$, $\mathcal{P}_s$ *tends to min:*

$$\lim_{s \to -\infty} \mathcal{P}_s(\mu_1, \mu_2; \omega_1, \omega_2) = \min(\mu_1, \mu_2)$$

Proof of Proposition 5.3

$$\begin{aligned}
&\lim_{s \to -\infty} \left(\frac{\omega_1 {\mu_1}^s + \omega_2 {\mu_2}^s}{\omega_1 + \omega_2} \right)^{\frac{1}{s}} \\
&= \lim_{s \to -\infty} \left(\frac{\omega_1}{\omega_1 + \omega_2} \right)^{\frac{1}{s}} \left({\mu_1}^s + \frac{\omega_2}{\omega_1} {\mu_2}^s \right)^{\frac{1}{s}} \\
&= \lim_{t \to +\infty} \left(\frac{1}{{\mu_1}^t} + \frac{\omega_2}{\omega_1 {\mu_2}^t} \right)^{\frac{1}{-t}}
\end{aligned}$$

$$= \lim_{t \to +\infty} \left(\frac{{\mu_2}^t}{\left(\frac{\mu_2}{\mu_1}\right)^t + \frac{\omega_2}{\omega_1}} \right)^{\frac{1}{t}}$$

$$= \lim_{t \to +\infty} \frac{\mu_2}{\left(\left(\frac{\mu_2}{\mu_1}\right)^t + \frac{\omega_2}{\omega_1}\right)^{\frac{1}{t}}}$$

Note that if $\mu_2 \leq \mu_1$, then the denominator tends to one and the limit is μ_2; if $\mu_1 < \mu_2$, then the denominator tends to $\frac{\mu_2}{\mu_1}$ and the limit is μ_1. This proves the claim. ■

Thus the two functions originally proposed as aggregation functions for the M_0I, the *min* and the product of powers, turn out to be two limiting cases of a parameterized family of weighted means. Each $\mathcal{P}_s$, $s \leq 0$, models a point on a continuum of trade-off strategies between the original non-compensating and compensating functions.

The notion of a level of compensation can be made more precise. First, the family of functions $\mathcal{P}_s$ increases in level of compensation as s increases, in that the overall preference increases with s:

Proposition 5.4 *$\mathcal{P}_s(\alpha_1, \alpha_2; \omega_1, \omega_2)$ is non-decreasing as a function of s:*

$$\frac{d}{ds}\mathcal{P}_s(\alpha_1, \alpha_2; \omega_1, \omega_2) \geq 0$$

Proof of Proposition 5.4 *Assume, for computational simplicity, that $\omega_1 + \omega_2 = 1$. This assumption can be made without loss of generality by invoking Axiom AF.6 or AO.6. Then*

$$\mathcal{P}_s(\alpha_1, \alpha_2; \omega_1, \omega_2) = (\omega_1 {\alpha_1}^s + \omega_2 {\alpha_2}^s)^{\frac{1}{s}}$$

and

$$\frac{d}{ds}\mathcal{P}_s(\alpha_1, \alpha_2; \omega_1, \omega_2) = \frac{1}{s}\mathcal{P}_s(\alpha_1, \alpha_2; \omega_1, \omega_2) \Big((\omega_1 {\alpha_1}^s + \omega_2 {\alpha_2}^s)^{-1}(\omega_1 {\alpha_1}^s \log \alpha_1 + \omega_2 {\alpha_2}^s \log \alpha_2) - \frac{1}{s}\log(\omega_1 {\alpha_1}^s + \omega_2 {\alpha_2}^s) \Big)$$

In the case that $s < 0$, *we wish to establish that*

$$(\omega_1{\alpha_1}^s + \omega_2{\alpha_2}^s)^{-1}(\omega_1{\alpha_1}^s \log\alpha_1 + \omega_2{\alpha_2}^s \log\alpha_2) - \frac{1}{s}\log(\omega_1{\alpha_1}^s + \omega_2{\alpha_2}^s) \leq 0 \quad (5.2)$$

Multiplying both sides by $s(\omega_1{\alpha_1}^s + \omega_2{\alpha_2}^s)$ *(which is non-positive), we now require*

$$s(\omega_1{\alpha_1}^s \log\alpha_1 + \omega_2{\alpha_2}^s \log\alpha_2) \geq (\omega_1{\alpha_1}^s + \omega_2{\alpha_2}^s)\log(\omega_1{\alpha_1}^s + \omega_2{\alpha_2}^s) \quad (5.3)$$

Bringing s *inside on the left side of the inequality, we now need to show*

$$\omega_1{\alpha_1}^s \log({\alpha_1}^s) + \omega_2{\alpha_2}^s \log({\alpha_2}^s) \geq (\omega_1{\alpha_1}^s + \omega_2{\alpha_2}^s)\log(\omega_1{\alpha_1}^s + \omega_2{\alpha_2}^s) \quad (5.4)$$

Since $x\log x$ *is a concave function for* $x > 0$, *as* $\frac{d^2}{dx^s}x\log x = \frac{1}{x} > 0$, *this is the case. Let* $g(x) = x\log x$, $x_1 = {\alpha_1}^s$, $x_2 = {\alpha_2}^s$. *Then* $g(\omega_1 x_1 + \omega_2 x_2) \leq \omega_1 g(x_1) + \omega_2 g(x_2)$, *proving (5.4).*

If $s > 0$, *the inequality must be reversed in (5.2), but since* $s(\omega_1{\alpha_1}^s + \omega_2{\alpha_2}^s)$ *is now non-negative, the sign of the inequality is recovered in (5.3).*

This proves the proposition. It should be noted that if α_1 *and* α_2 *are non-zero, and not equal, then strict inequalities hold throughout the proof.* ■

The absolute overall preference is not as important as the relative preferences among alternatives. If the overall preferences for all alternatives increase uniformly, then the level of compensation has not changed. Another measure of compensation is thus more compelling. The idea is that if there are two points (α_1, α_2) and (β_1, β_2), with $\alpha_1 > \beta_1, \beta_2 > \alpha_2$, then larger s always favor (α_1, α_2) over (β_1, β_2). More precisely, the ratio $\frac{\mathcal{P}_s(\alpha)}{\mathcal{P}_s(\beta)}$ is an increasing function of s.

Proposition 5.5 *If* $(\alpha_1, \alpha_2), (\beta_1, \beta_2)$ *are two preference points, with* $\alpha_1 > \beta_1, \beta_2 > \alpha_2$, *then*

$$\frac{d}{ds}\frac{\mathcal{P}_s(\alpha_1, \alpha_2; \omega_1, \omega_2)}{\mathcal{P}_s(\beta_1, \beta_2; \omega_1, \omega_2)} \geq 0$$

as long as $\omega_1, \omega_2 > 0$.

The proof requires a lemma:

Lemma 5.1 *Let*

$$h(x, y; v, w) = \frac{vx \log(x) + wy \log(y) - (vx + wy) \log(vx + wy)}{vx + wy}$$

If $x, y, v, w > 0$*, then*

$$\frac{d}{dx} h(x, y; v, w) \begin{cases} < 0 & \text{if } x < y \\ > 0 & \text{if } x > y \end{cases}$$

Proof of Lemma 5.1

$$\frac{d}{dx} h(x, y; v, w) = \frac{vwy \log\left(\frac{x}{y}\right)}{(vx + wy)^2}$$

Since all other terms are positive, the sign depends on the sign of $\log\left(\frac{x}{y}\right)$*, which is indeed positive for* $x > y$ *and negative for* $x < y$*.* ∎

It is a consequence of Lemma 5.1 that if the interval (x_1, x_2) is wholly contained inside the interval (y_1, y_2), then $h(x_1, x_2; v, w) < h(y_1, y_2; v, w)$.

Proof of Proposition 5.5 *With a little algebra,* $\frac{d}{ds} \frac{\mathcal{P}_s(\alpha_1, \alpha_2; \omega_1, \omega_2)}{\mathcal{P}_s(\beta_1, \beta_2; \omega_1, \omega_2)}$ *is seen to be:*

$$\frac{\mathcal{P}_s(\alpha_1, \alpha_2; \omega_1, \omega_2)}{s^2 \mathcal{P}_s(\beta_1, \beta_2; \omega_1, \omega_2)} \left(\frac{\omega_1 {\alpha_1}^s \log({\alpha_1}^s) + \omega_2 {\alpha_2}^s \log({\alpha_2}^s) - (\omega_1 {\alpha_1}^s + \omega_2 {\alpha_2}^s) \log(\omega_1 {\alpha_1}^s + \omega_2 {\alpha_2}^s)}{\omega_1 {\alpha_1}^s + \omega_2 {\alpha_2}^s} \right.$$

$$\left. - \frac{\omega_1 {\beta_1}^s \log({\beta_1}^s) + \omega_2 {\beta_2}^s \log({\beta_2}^s) - (\omega_1 {\beta_1}^s + \omega_2 {\beta_2}^s) \log(\omega_1 {\beta_1}^s + \omega_2 {\beta_2}^s)}{\omega_1 {\beta_1}^s + \omega_2 {\beta_2}^s} \right)$$

Since the first factor is positive, the derivative is nonnegative whenever

$$\frac{\omega_1 {\alpha_1}^s \log({\alpha_1}^s) + \omega_2 {\alpha_2}^s \log({\alpha_2}^s) - (\omega_1 {\alpha_1}^s + \omega_2 {\alpha_2}^s) \log(\omega_1 {\alpha_1}^s + \omega_2 {\alpha_2}^s)}{\omega_1 {\alpha_1}^s + \omega_2 {\alpha_2}^s}$$

$$\geq \frac{\omega_1 {\beta_1}^s \log({\beta_1}^s) + \omega_2 {\beta_2}^s \log({\beta_2}^s) - (\omega_1 {\beta_1}^s + \omega_2 {\beta_2}^s) \log(\omega_1 {\beta_1}^s + \omega_2 {\beta_2}^s)}{\omega_1 {\beta_1}^s + \omega_2 {\beta_2}^s}$$

Setting $x_i = {\alpha_i}^s, y_i = {\beta_i}^s$*, and* $\omega_1 = v, \omega_2 = w$*, and applying Lemma 5.1, this is always the case.* ∎

Consider two alternatives A and B, each ranked on the basis of two separate criteria, the preferences for A being (α_1, α_2), and those for B being (β_1, β_2). If $\alpha_1 > \beta_1, \beta_2 > \alpha_2$, then the least compensating function, the *min* ($s = -\infty$), dictates a choice of B over A (regardless of weights, provided they are non-zero). The most compensating function, the max ($s = -\infty$), indicates that A is preferred to B. Proposition 5.5 states that as the parameter s increases, the relative preference for A rather than B increases as well. Thus the parameter s is indeed a measure of level of compensation among goals.

5.5 Supercompensating functions

This section justifies the use of functions that exceed the geometric mean $\mathcal{P}_0$.

In multi-attribute decision making, aggregation functions should provide a useful, justifiable model of the design decision process. The parameterized family $\mathcal{P}_s$ provides a continuum of weighted means between the two existing functions of the M$_o$I. While this family of functions is useful for design decision making, it is not exhaustive. Other generating functions give rise to other aggregation functions that satisfy all the axioms of the M$_o$I but behave differently from any of the $\mathcal{P}_s$.

The *min* is the least compensating possible design-appropriate function. No design aggregation function can take on values less than the *min* at any point in a design space. The *min* function defines a boundary not only of a certain family of weighted means, but also of design-appropriate functions in general. The product of powers $\mathcal{P}_\Pi$ is a pivotal example among weighted means, but it is not so clear that it is maximal among design-appropriate functions. Indeed, a maximal design-appropriate function is problematic: such a function $\mathcal{P}_{\max}$ would satisfy

$$\mathcal{P}_{\max}(0, \mu; \omega_1, \omega_2) = 0 \quad \forall\, \mu$$

but also satisfy

$$\mathcal{P}_{\max}(\epsilon, \mu; \omega_1, \omega_2) = \mu \quad \forall\, \mu > \epsilon > 0$$

for all non-zero weights ω_1, ω_2. $\mathcal{P}_{\max}$ so defined takes on the largest possible value while satisfying both idempotency and annihilation, but it fails another design axiom: it is discontinuous at zero. It is clear that there is no design function that provides an upper bound

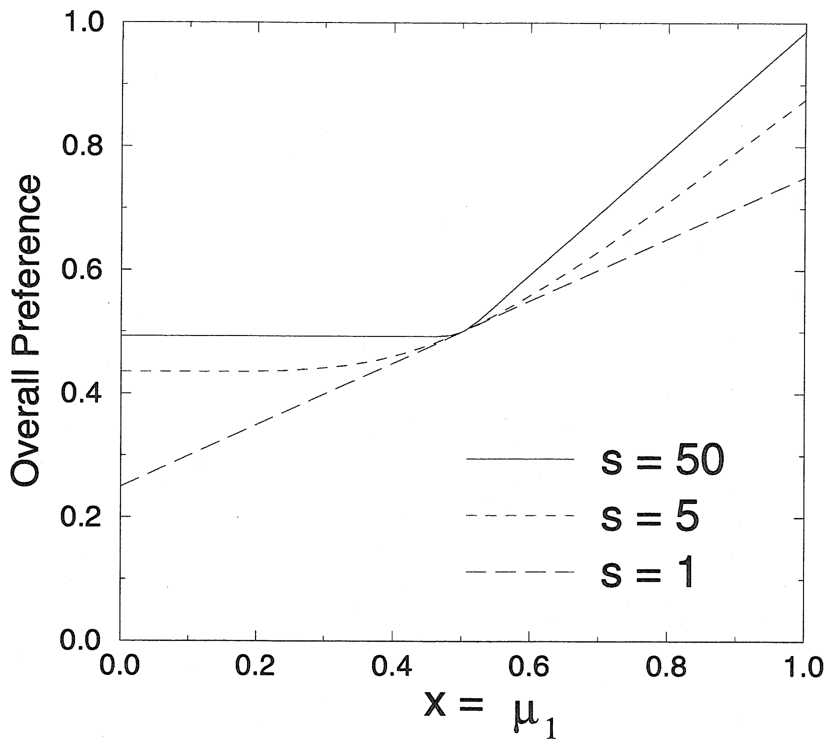

Figure 5.2: Functions that exceed $\mathcal{P}_\Pi$

in the same way that the *min* is the lower bound.

The family of weighted means that formed a continuum between the *min* and $\mathcal{P}_\Pi$ was found by varying the parameter s in $\mathcal{P}_s$ between $-\infty$ and 0. It is interesting to examine what happens if the parameter is varied between 0 and $+\infty$. In this case a family of aggregation functions between $\mathcal{P}_0 = \mathcal{P}_\Pi$ and $\mathcal{P}_\infty = max$ is generated. These functions do not satisfy annihilation, so they do not appear to be appropriate for design. As long as no preference approaches zero, however, these functions satisfy all the axioms of the $\mathrm{M_oI}$. Figure 5.2 shows $\mathcal{P}_s$ plotted for several positive values of s. As $s \to \infty$, $\mathcal{P}_s \to max$.

It was noted above that the arithmetic mean is an aggregation function commonly used in multi-attribute decision making. Yet the arithmetic mean does not satisfy all of the axioms of the $\mathrm{M_oI}$, as it fails annihilation. The arithmetic mean is the aggregation function $\mathcal{P}_1$, and allows goals to compensate more strongly than the geometric mean. Indeed, $\mathcal{P}_s$, for $s > 0$, always fails annihilation, and the level of compensation between goals increases with s all the way to $\mathcal{P}_{+\infty} = max$. If the arithmetic mean is chosen only for computational simplicity, then its use must be further justified.

In practice, the MoI is implemented using discrete functions. This is partly an artifact of computer implementation, but more fundamentally arises from the fact that some changes in preference are too small to be distinguished. (Consider Miller's research [50] indicating that humans can distinguish 7 ± 2 categories.) A designer does not actually specify a continuous preference function on the interval $[0, 1]$, but rather gives the values for each variable on several different α-cuts [4]. For example, for a particular design variable, the designer may specify which values correspond to $\mu = 0$, $\mu = 0.25$, $\mu = 0.5$, $\mu = 0.75$, and $\mu = 1$.

The discontinuous manner in which preferences are specified provides some justification for allowing the use of $\mathcal{P}_s$ with $s > 0$, or even the *max* operator, as an aggregation function when all of the preferences achieve some level. At each $\mu > 0$, $\mathcal{P}_s$ can be calculated, but where $\mu = 0$, annihilation governs. This seems to model the design process, as well: when all attributes are performing to some acceptable standard, a designer may choose to allow the highest preferences great importance. Unacceptable performance on a single objective, however, can scuttle the design, so annihilation is still satisfied.

Theorem 5.1 provides a technical justification for the use of the family of operators between $\mathcal{P}_0$ and $\mathcal{P}_{+\infty}$. If it can be assumed that preferences less than some small ϵ, say 0.1, are not relevant to the designer, then the theorem indicates that there is a continuous aggregation function that satisfies $\mathcal{P}(0, \alpha_2; \omega_1, \omega_2) = 0$ and

$$\mathcal{P}(\mu_1, \mu_2; \omega_1, \omega_2) = \mathcal{P}_s(\mu_1, \mu_2; \omega_1, \omega_2) \text{ for } \mu_1, \mu_2 \geq \epsilon$$

The theorem guarantees that there is a formal operator that models this level of compensation without violating continuity or any other axiom of the MoI.

Thus trade-offs for all cases of compensation, from none ($s = -\infty$) to fully compensating ($s = 0$) to supercompensating ($s > 0$) are accommodated by the weighted means in Theorem 5.1. Note that the parameterized family of supercompensating functions obeys the same results that were shown for the compensating functions in Section 5.4.

5.6 Example

To illustrate the family of aggregation operators outlined here and their importance, consider example 12-10 from Prof. Zimmermann's textbook [109]. This example is originally presented in the text as a fuzzy linear programming problem with continuous variables. In the expression of the problem, however, the variables are quantities of two products to be produced, and it would be natural to assume that they can only take on integer values. Thus the problem can be thought of as a discrete MADM problem.

The example involves a company that produces two products, which yield different returns in profit and balance of trade. (Product 1 yields \$2 profit but requires \$1 in imports; product 2 can be exported for \$2 revenue but makes only \$1 profit.) The problem is to decide on a "best" production schedule to achieve high profits and a favorable balance of trade. The production schedule is subject to capacity constraints and is modelled as follows:

$$\text{“maximize”} \qquad \mathbf{z}(\mathbf{x}) = \begin{pmatrix} -1 & 2 \\ 2 & 1 \end{pmatrix} \begin{pmatrix} x_1 \\ x_2 \end{pmatrix}$$

where

$$\begin{aligned} x_1 &= \text{number of Product 1 manufactured} \\ x_2 &= \text{number of Product 2 manufactured} \\ z_1 &= \text{balance of trade} \\ z_2 &= \text{profit} \end{aligned}$$

subject to:

$$\begin{aligned} \text{C1:} \quad -x_1 + 3x_2 &\leq 21 \\ \text{C2:} \quad x_1 + 3x_2 &\leq 27 \\ \text{C3:} \quad 4x_1 + 3x_2 &\leq 45 \\ \text{C4:} \quad 3x_1 + x_2 &\leq 30 \\ \text{C5:} \quad x_1 &\geq 0 \\ \text{C6:} \quad x_2 &\geq 0 \end{aligned}$$

(x_1, x_2)	z_1	z_2
(0,7)	14	7
(3,8)	13	14
(4,7)	10	15
(5,7)	9	17
(6,7)	8	19
(8,4)	0	20
(9,3)	-3	21

Table 5.4: Undominated points in the decision space

Prof. Zimmermann shows a plot of the decision space with a region of optimal values, similar to the one shown in Figure 5.3.

When decisions are based on the aggregation of performance of more than one attribute, there is the notion of an *undominated solution* to the decision problem.

Definition 5.3 *The alternative A* dominates *the alternative B if A performs no worse than B on all attributes, and better than B on at least one attribute. In this case, regardless of the weights or the strategy, it is always better to choose A over B. A feasible solution is* undominated *if there is no other feasible solution which dominates it.*

For a given problem, there is usually a (possibly infinite) set of undominated solutions (called the *undominated set*). The choice among the undominated solutions is accomplished by negotiation. There are seven undominated points, shown in Table 5.4, in the decision space for this problem. The constraint inequalities are also shown on the plot.

It is clear that this problem is not simply an exercise in mathematical programming. Three questions must be answered to determine an objective function which is a scalar function of the original variables. First, the decision maker must specify what it means to achieve the two individual goals (high profit, favorable balance of trade). Second, the relative importance of the goals must be addressed. Third, to what extent should high performance with respect to one goal be allowed to compensate for low performance elsewhere?

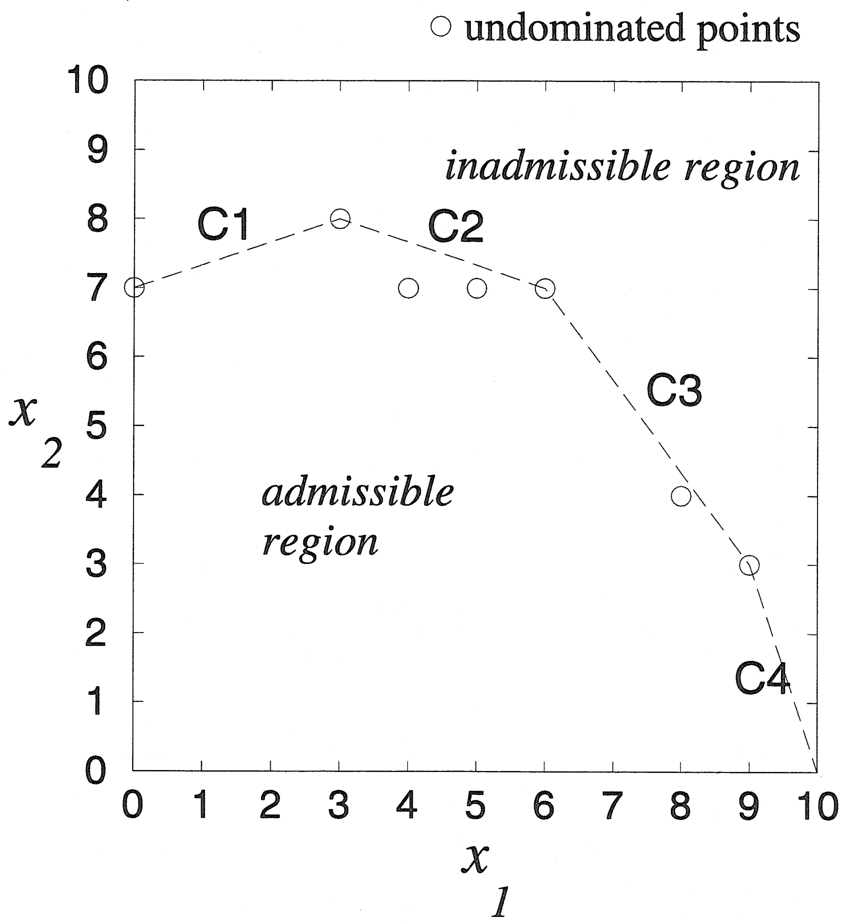

Figure 5.3: Decision space with optimal region

The example is solved in the textbook by an application of fuzzy sets that is substantively similar to that used by the $\text{M}_\text{o}\text{I}$. The first step is to determine a level of satisfaction for each of the two goals, in essence to create the two fuzzy sets "Decisions that satisfy the profit goal" and "Decisions that satisfy the balance of trade goal." What the $\text{M}_\text{o}\text{I}$ would call *preference* the textbook refers to as *level of satisfaction.* The preference or satisfaction for the performance on balance of trade increases linearly from $\mu_1(\mathbf{x}) = 0$ at $z_1(\mathbf{x}) = -3$ to $\mu_1(\mathbf{x}) = 1$ at $z_1(\mathbf{x}) = 14$. The preference for profit increases linearly from $\mu_2(\mathbf{x}) = 0$ at $z_2(\mathbf{x}) = 7$ to $\mu_2(\mathbf{x}) = 1$ at $z_2(\mathbf{x}) = 21$. These preferences are generated in the textbook with reference to the values listed in Table 5.4: $\mathbf{z}(0, 7) = (14, 7)$, while $\mathbf{z}(9, 3) = (-3, 21)$. Note that by annihilation, $\mathbf{x} = (0, 7)$ and $\mathbf{x} = (9, 3)$ are now dominated points. In the general application of the $\text{M}_\text{o}\text{I}$, the memberships μ_i may be specified using a different process. In any event, the given fuzzy sets "Decisions that satisfy the objectives," with memberships ranging from 0 to 1 throughout the decision space, are valid preference functions for the $\text{M}_\text{o}\text{I}$.

Since the textbook does not address the issue of importance weights for the two goals, it shall be assumed that they are equally weighted. However, the choice of an aggregation function, which encapsulates the decision of how much high performance on one attribute is to compensate for low performance on the other, remains. If (unequal) weights were specified, the choice of an aggregation function would still remain, and the conclusions reached below would be qualitatively the same.

If the decision problem is treated as a fuzzy linear programming problem, as in the text, the aggregation function used is the *min*. The *min* is "natural" for ease of computation, but not necessarily natural for the decision. When the problem is solved using the *min* operator, as shown in Figure 5.4, the maximum degree of "overall satisfaction" is given by the point $\mathbf{x}_A = (5.03, 7.32)$, with $\bar{\mu} = 0.74$. Among the integer choices available, $\mathbf{x}_B = (5, 7)$ is the best, with $\bar{\mu} = 0.71$. This corresponds to the original $\text{M}_\text{o}\text{I}$ solution for a non-compensating problem, *i.e.*, when the overall performance is limited by the lowest performance of all attributes.

It was noted earlier that in many situations, the overall performance of a design, or the general attractiveness of a decision, is not limited by the lowest performance among the at-

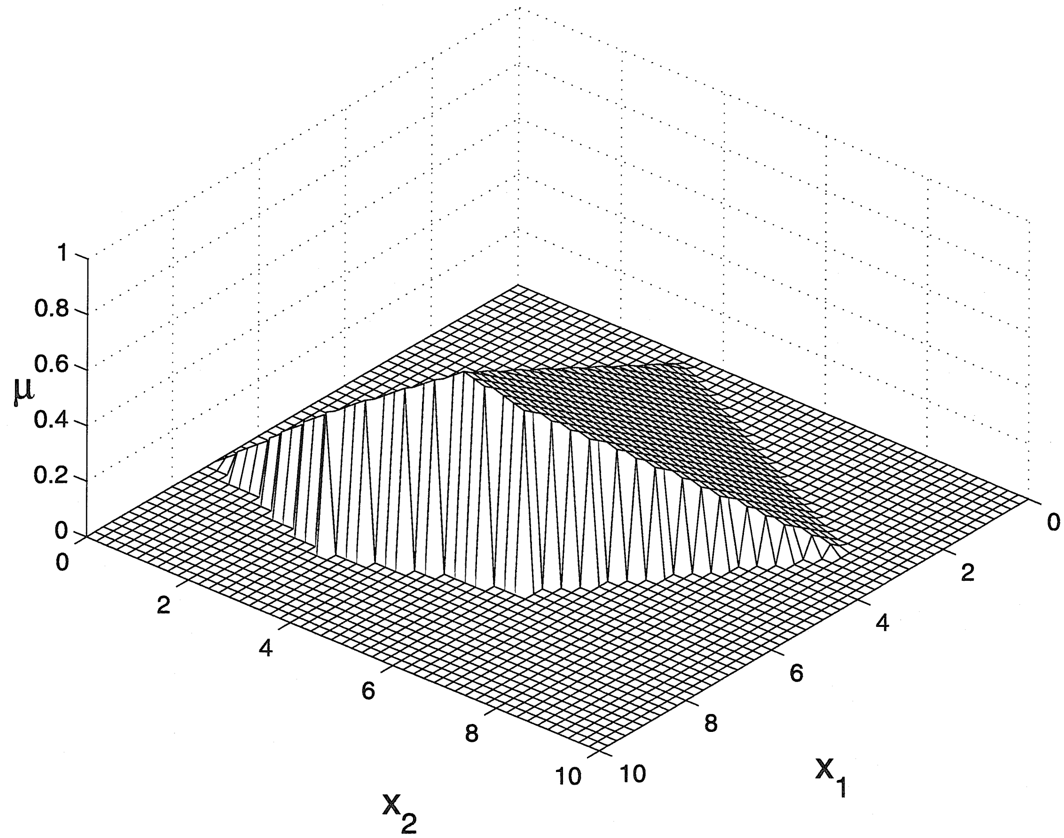

Figure 5.4: Decision surface with minimum operator

tributes. The MoI originally used a weighted product of powers for all "compensating" problems. Using this function, the point with highest combined preference is $\mathbf{x}_C = (5.70, 7.10)$ with $\bar{\mu} = 0.75$, as shown in Figure 5.5. The point of highest preference is not far from the point of highest preference achieved with the *min* operator. However, if one considers only the discrete points with integer values (if the company cannot manufacture 5.70 of a product), the highest performing point is $\mathbf{x}_D = (6, 7)$ (with $\bar{\mu} = 0.74$). A different level of compensation among goals leads to a different decision.

The application of the entire family of aggregation functions $\mathcal{P}_s$ to this example problem shows that there are at least three optimal points, each suitable over a range of values of the parameter s, and thus over a range of levels of compensation of goals. In addition, the point $(1, 7)$, though dominated by the point $(3, 8)$, approaches it asymptotically in preference

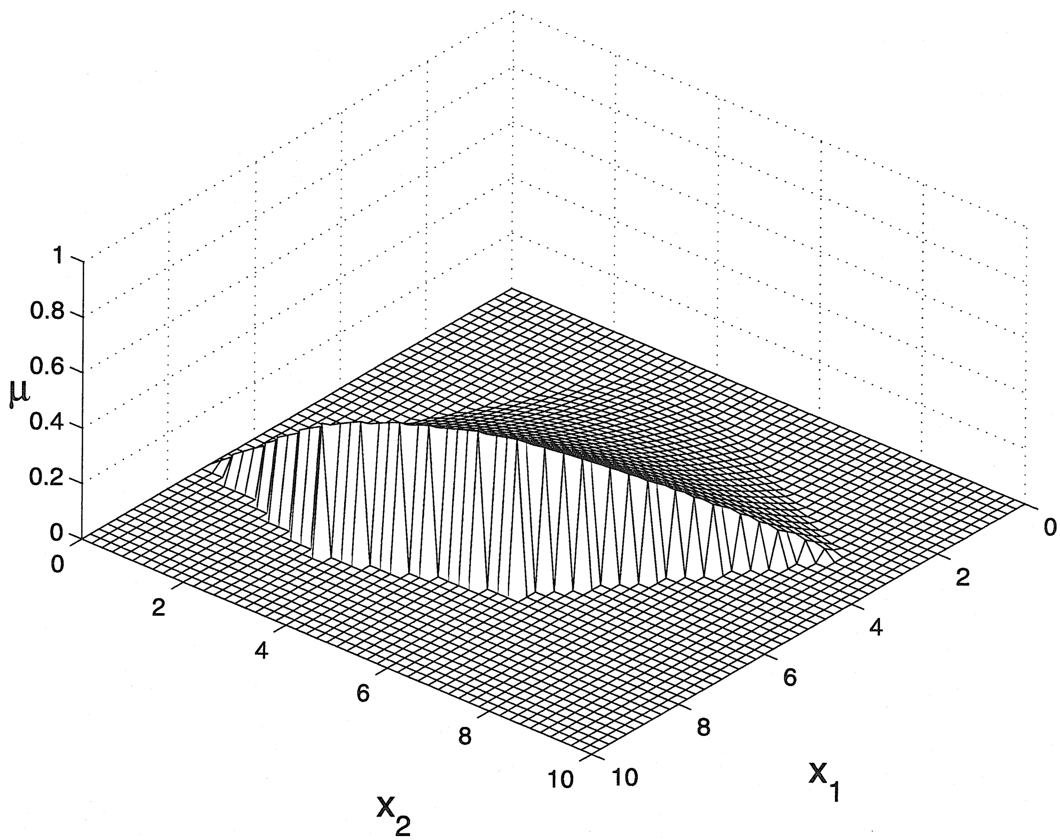

Figure 5.5: Decision surface with product of powers operator

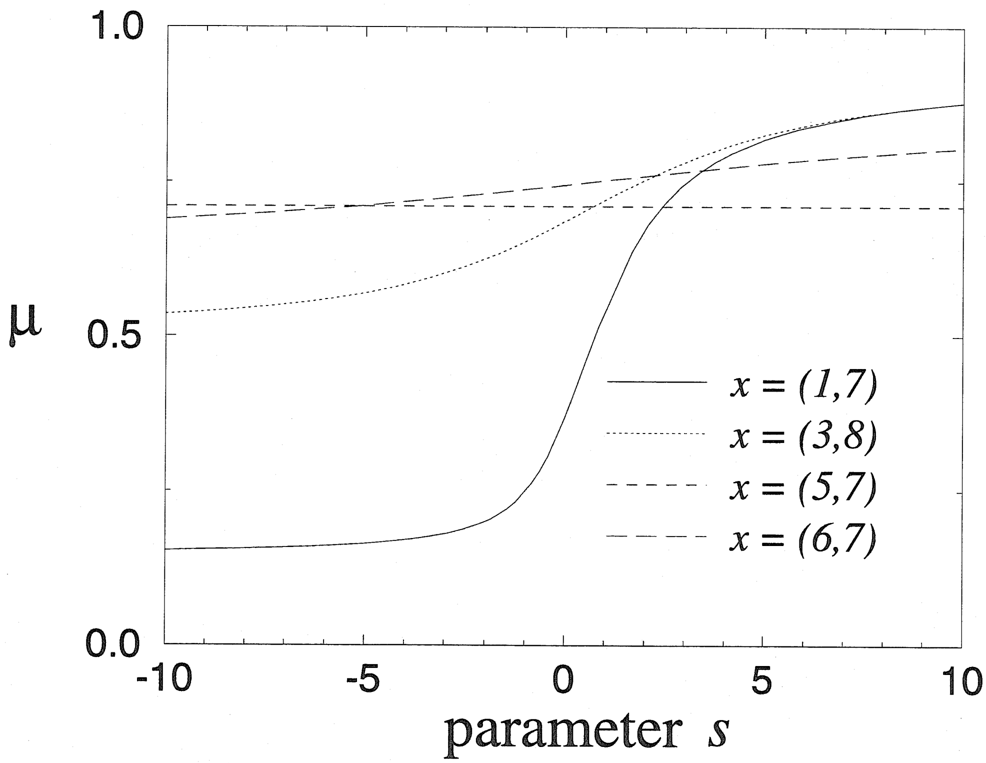

Figure 5.6: Optimal points varying with parameter s

as the level of compensation approaches the maximum. The *a priori* assumption of equal weights prevents the other undominated points from Table 5.4 from being optimal choices for any strategy: this point will be discussed further in the following section.

Figure 5.6 shows the overall preferences for all of these "optimal points" calculated using $\mathcal{P}_s$, with the parameter s ranging from -10 to 10. When $\mathcal{P}_s = \min$ ($s \to -\infty$), the point with the highest preference is $\mathbf{x} = (5, 7)$, the textbook answer. As s grows and $\mathcal{P}_s$ becomes more compensating, the preference for other points grows stronger. As s crosses -5.25, $\mathbf{x} = (6, 7)$ overtakes $\mathbf{x} = (5, 7)$ and remains the most preferred point through approximately $s = 2.25$. This region includes two common aggregation functions: the weighted product of powers or geometric mean ($s = 0$), and the classical weighted sum or arithmetic mean ($s = 1$). Values of s greater than 2.25, corresponding to even stronger levels of compensation, lead to an overall preference for the point $(3, 8)$, and as $s \to +\infty$, the preference for the dominated point $\mathbf{x} = (1, 7)$ approaches the preference for $\mathbf{x} = (3, 8)$ asymptotically. The limits $s = 10$ and $s = -10$ were chosen to show nearly the full range of behavior of this problem; values of s outside this range are certainly permissible.

So, even for this simple textbook problem, a careful examination shows that the answer depends on the aggregation functions employed. Different aggregation functions lead to different decisions; it is important to be able to model any tenable decision, and not to have the decision restricted *a priori* by a limited selection of aggregations. The functions discussed here provide models for a continuum of trade-offs ranging from the non-compensating *min* to the compensating $\mathcal{P}_\Pi$ all the way to the *max* operator. Instead of two aggregation functions, there is a parameterized family of functions ranging from the *min* to $\mathcal{P}_\Pi$, and another from $\mathcal{P}_\Pi$ to the *max*.

5.7 Modelling negotiation: which decisions can be captured?

There are three distinct inputs to a multi-attribute design decision described here: the individual preferences, the relative weightings of those preferences, and the selection of a strategy. It seems that the true level of individual preference should not be affected by negotiation; the preference for one aspect of a design should not depend on other aspects.

Another matter altogether is the willful misrepresentation of preference as a negotiation tactic, which will be discussed below. A natural definition of "negotiation" implies that two or more parties come to the table with their preferences set. The weightings of the preferences, however, and the choice of a particular strategy, are decisions that separate parties must consider, or negotiate, together. The relative weighting of the attributes and the choice of a particular strategy both affect the level of preference that is achieved for each and every attribute. Thus, in the broadest sense a negotiation considers both weights and strategies.

Definition 5.4 *A* negotiation *is the selection of weights and a strategy. In particular, it is the choice of the weights* ω_i *and of the compensation parameter* s*.*

The common usage of the word "negotiate" connotes a situation where two or more interested parties discuss a matter until a conclusion is reached. The definition of negotiation used here is somewhat broad, in that it encompasses the selection of strategy and weights whether or not those decisions are arrived at by adversarial parties in consultation. Here, negotiation encompasses the situation where a single external decision maker, a manager perhaps, imposes strategies and weights without consulting the parties who provided the individual preferences. It encompasses the situation in which weights are fixed by some external means, and strategies are all that can be considered. While various techniques for arriving at particular strategies and weights will be presented, the emphasis is on the mathematical model of the result, rather than on the actions taken by competing parties.

5.7.1 Theoretical possibilities

In this section it is shown that Definition 5.4 does not *a priori* exclude any undominated points from selection in a full negotiation. Here, a "full" negotiation is one in which both weights and strategies are to be chosen.

It is desirable that the result of a negotiation not be predetermined. In principle, it should be possible for the result of a negotiation to be any one of the undominated solutions. In other words, a negotiation model must permit the selection of any one of the feasible undominated points. The negotiation model presented in this thesis fulfills this condition, in that for any individual in a set of undominated solutions, there exist a strategy s and a ratio of weights $\omega = \frac{\omega_2}{\omega_1}$ that select that individual as the "best" overall solution. The result is shown directly on preferences:

Proposition 5.6 *For any pair (α_1, α_2) in an undominated finite set M of preference pairs, there exist s^* and ${\omega_1}^*, {\omega_2}^*$ such that*

$$\mathcal{P}_{s^*}(\alpha_1, \alpha_2; {\omega_1}^*, {\omega_2}^*) = \max_{(\alpha_i, \alpha_j) \in M} \mathcal{P}_{s^*}(\alpha_i, \alpha_j; {\omega_1}^*, {\omega_2}^*)$$

The proof uses a lemma:

Lemma 5.2 *If M is an undominated finite set of preference pairs and $(\alpha_1, \alpha_2) \in M$, then there exists $\epsilon > 0$ such that every other element of M is dominated by either $(\alpha_1 - \epsilon, 1)$ or by $(1, \alpha_2 - \epsilon)$.*

Proof of Lemma 5.2 *Consider $(\beta_1, \beta_2) \in M$. Since (β_1, β_2) does not dominate (α_1, α_2), either $\alpha_1 > \beta_1$, or $\alpha_2 > \beta_2$. If $\alpha_1 > \beta_1$, let $\delta = \frac{\alpha_1 - \beta_1}{2}$. Then $(\alpha_1 - \delta, 1)$ dominates (β_1, β_2). If $\alpha_2 > \beta_2$, let $\delta = \frac{\alpha_2 - \beta_2}{2}$. Then $(1, \alpha_2 - \delta)$ dominates (β_1, β_2). Since M is finite, this can be repeated for all remaining elements of M; let ϵ be the smallest such δ, and every element of M except (α_1, α_2) is dominated either by $(\alpha_1 - \epsilon, 1)$, or by $(1, \alpha_2 - \epsilon)$.* ∎

Proof of Proposition 5.6 *Let $\omega = \frac{\omega_1}{\omega_2}$. By Lemma 5.2, it suffices to show that for any ϵ, the following pair of equations can be solved for s and ω:*

$$\begin{aligned} \mathcal{P}_s(\alpha_1, \alpha_2; 1, \omega) &= \mathcal{P}_s(1, \alpha_2 - \epsilon; 1, \omega) \\ \mathcal{P}_s(\alpha_1, \alpha_2; 1, \omega) &= \mathcal{P}_s(\alpha_1 - \epsilon, \alpha_2; 1, \omega) \end{aligned} \tag{5.5}$$

If $s = 0$ is a solution, then the proposition is proved. If $s \neq 0$, then the following equations must be solved:

$${\alpha_1}^s + \omega {\alpha_2}^s = (\alpha_1 - \epsilon)^s + \omega \tag{5.6}$$

$${\alpha_1}^s + \omega {\alpha_2}^s = 1 + \omega(\alpha_2 - \epsilon)^s \tag{5.7}$$

Solving (5.7) for ω yields:

$$\omega = \frac{1 - {\alpha_1}^s}{{\alpha_2}^s - (\alpha_2 - \epsilon)^s}$$

and plugging this into (5.6), and rearranging terms as below, gives the following that must be solved for s:

$$\begin{aligned}(1-\alpha_1{}^s)(1-\alpha_2{}^s)-(\alpha_1{}^s-(\alpha_1-\epsilon)^s)\,(\alpha_2{}^s-(\alpha_2-\epsilon)^s) &= 0\\ 1-\alpha_1{}^s-\alpha_2{}^s-(\alpha_1-\epsilon)^s(\alpha_2-\epsilon)^s+\alpha_1{}^s(\alpha_2-\epsilon)^s+\alpha_2{}^s(\alpha_1-\epsilon)^s &= 0\end{aligned}$$

which can be rewritten as:

$$\begin{aligned}1-e^{s\log\alpha_1}-e^{s\log\alpha_2}-e^{s\log((\alpha_1-\epsilon)(\alpha_2-\epsilon))}\\ +e^{s\log(\alpha_1(\alpha_2-\epsilon))}+e^{s\log(\alpha_2(\alpha_1-\epsilon))} &= 0\end{aligned} \tag{5.8}$$

Now, the first derivative of (5.8) with respect to s, evaluated at $s=0$, is 0; the second derivative is positive at $s=0$. It thus suffices to show that the left-hand side of (5.8) is negative for some $s<0$. Since $\alpha_1, \alpha_1-\epsilon\in(0,1)$ and $\epsilon>0$, we can set $A_1=\log\alpha_1$. Since $\log(\alpha_1-\epsilon)<\log\alpha_1<0$, there is some $\delta_1>0$ such that $A_1-\delta_1=\log(\alpha_1-\epsilon)$. Similarly, we can define A_2,δ_2 with $A_2=\log\alpha_2$ and $A_2-\delta_2=\log(\alpha_2-\epsilon)$. Then consider the limit:

$$\lim_{s\to-\infty} 1-e^{sA_1}-e^{sA_2}-e^{s(A_1-\delta_1+A_2-\delta_2)}+e^{s(A_1+A_2-\delta_2)}+e^{s(A_2+A_1-\delta_1)}$$

which is the same as:

$$\lim_{t\to\infty} 1-e^{t|A_1|}-e^{t|A_2|}-e^{t(|A_1|+|A_2|+\delta_1+\delta_2)}+e^{t(|A_1|+|A_2|+\delta_2)}+e^{t(|A_1|+|A_2|+\delta_1)}$$

Since this expression is dominated as $t\to\infty$ by the term $e^{t(|A_1|+|A_2|+\delta_1+\delta_2)}$, it must take on negative values for some $t>0$, i.e., *for some $s<0$. Therefore, the set of equations (5.5) has a solution.* ■

Proposition 5.6 says that it is mathematically possible to choose a weight and a strategy to select any particular undominated point, at least from a finite set of possible solutions. In practice, the (s,ω) pair to select one particular point from an undominated set is often a heavily weighted almost-*min*. Indeed, the proof shows that there is always a solution with

$s < 0$. It is neither surprising nor detrimental that a decision problem that compensates strongly, or that has any predetermined strategy, will exclude some undominated points from the feasible set of solutions. Furthermore, the method discussed below for determining a proper compensation parameter will always find $s \leq 1$. Likewise, when weights are predetermined, not all apparently undominated points are actually eligible to be the highest overall.

5.7.2 Managed negotiation

Perhaps the simplest case of negotiation is the case where both strategies and weights are determined by an outside decision maker, ostensibly separate from any of the interested parties. This case can be thought of as the situation where a manager is allowed to impose the "negotiation," and might be called the case of imposed cooperation. The underlying mathematics of aggregation are of course the same. In this section this simplest case of selection of (s, ω) is presented.

When a single individual has complete authority over the negotiation, strategies and weights can be considered simultaneously, and their values can be calculated from *indifference points*. Two points are considered indifferent if they have the same preference; it is not necessary that the numerical preferences be known. The procedure is as follows:

1. Determine α_1 and α_2 such that $\mathcal{P}_s(\alpha_1, 1) = \mathcal{P}_s(1, \alpha_2) = 0.5$. In other words, at which value α_1 is there indifference between a design that achieves preferences of α_1 on the first attribute and 1 on the second attribute, and a design that achieves preferences of 0.5 on both attributes (and thus, by idempotency, has a combined preference of 0.5)? A similar question is asked for α_2. Sometimes it is easier to ask for values of $\mathbf{x}$ and calculate α_i; sometimes it is easier to seek α_i directly. Either approach to determining the indifference points is acceptable.

2. Let $b = \frac{\omega_{\alpha_2}}{\omega_{\alpha_1}}$.

3. If $\alpha_1 = \alpha_2$, then $b = 1$:

   (a) If $\alpha_1 = 0.5$, then $s = -\infty$.

(b) If $\alpha_1 = 0.25$, then $s = 0$.

(c) If $\alpha_1 > 0.25$, then $s \in (-\infty, 0)$. Solve ${\alpha_1}^s + 1 = 2(0.5)^s$ numerically.

(d) If $\alpha_1 < 0.25$, then $s \in (0, \infty)$. Solve ${\alpha_1}^s + 1 = 2(0.5)^s$ numerically.

4. If $\alpha_1 \neq \alpha_2$, then $b \neq 1$. Note that if $s = 0$,

$$ {\alpha_1}^m = 0.5 = {\alpha_2}^{1-m} \Rightarrow {\alpha_2}^{1-\log_{\alpha_1} 0.5} = 0.5 $$

Thus:

(a) If ${\alpha_2}^{1-\log_{\alpha_1} 0.5} = 0.5$, then $s = 0$, and $b = \frac{1-\log_{\alpha_1} 0.5}{\log_{\alpha_1} 0.5}$

(b) If ${\alpha_2}^{1-\log_{\alpha_1} 0.5} > 0.5$, then $s < 0$.

If ${\alpha_2}^{1-\log_{\alpha_1} 0.5} < 0.5$, then $s > 0$.

Solve numerically for s from

$$ \left(\frac{1 + b{\alpha_2}^s}{1+b}\right)^{\frac{1}{s}} = \left(\frac{{\alpha_1}^s + b}{1+b}\right)^{\frac{1}{s}} = 0.5 $$

which reduces to

$$ \left({\alpha_1}^s - 0.5^s\right)\left({\alpha_2}^s - 0.5^s\right) = (1 - 0.5^s)^2 $$

Once this is solved numerically for s, then b can also be determined.

Some remarks on this procedure: first, the procedure will never return an answer $s > 1$. Since the supercompensating functions can conflict with the annihilation axiom, this offers no difficulty and is in fact an advantage. It should also be noted that if either α_1 or α_2 is close to 0, the (s, b) pair is quite sensitive to small differences in α_1 and α_2. In these cases, it might be preferable to elicit other indifference points to determine s and b. In the procedure described above, points that are equivalent to $(0.5, 0.5)$ are chosen; the procedure could easily be modified to consider indifference to some other reference point. Indeed, if the procedure is applied more than once with different reference points, the redundant information serves as a check on the accuracy of the specification. Alternatively, Saaty's Analytic Hierarchy Process (AHP) [75], a method for normalizing directly specified weights, could be used in pairwise comparison at the start, as a check. However, the procedure given above is somewhat richer than the AHP, so such a comparison may not be useful. There is presently no comparable normalization scheme for strategies s.

5.7.3 Negotiation with predetermined weights

The definition of negotiation used here includes the choice of both strategies and weights, without prejudice as to the order of selection. It is common in formal decision making to separate the specification of weights from the aggregation of preferences. (Of course, it is common to choose a particular strategy at random, which is not to say that it is correct.) Methods such as Saaty's Analytic Hierarchy Process (AHP) [75] have been developed for the determination of weights.

When weights are determined prior to the negotiation stage, then it is no longer always possible to select any undominated point. For instance, in the example presented in Section 5.6, the weights of the two goals were determined to be equal before a strategy was chosen. Of seven undominated points in the decision space, only four were possible "best" points for particular choices of the compensation parameter s. The equally weighted problem excluded three undominated points from consideration.

Because preferences in two or more dimensions are ordered, some points which appear to be undominated can become dominated for some assigned weights. Consider an example: if weights are unrestricted, then the two points $\mathbf{x}_1 = (0.5, 0.9)$ and $\mathbf{x}_2 = (0.6, 0.1)$ do not dominate each other, and there exist strategy-weight pairs to select either one over the other. If the weights are declared equal, then the order of elements is irrelevant, and $\mathbf{x}_1$ dominates $\mathbf{x}_2$. Indeed, whenever the second element has a higher weight than the first, then $\mathbf{x}_1$ dominates $\mathbf{x}_2$. This situation only arises, however, when reversing the order of one preference makes one preference dominate the other. Whenever two points do not dominate each other, even if the elements of one point are reversed, then for any weight there is a strategy that will achieve equality. When reversing the order of elements allows one point to dominate the other, then there are always some weights which preclude equality for all strategies.

Proposition 5.7 *If (α_1, α_2) and (β_1, β_2) are two preference points, neither of which dominates the other (assume without loss of generality that $\alpha_1 > \beta_1$ and $\beta_2 > \alpha_2$), then the following are true:*

1. *For every strategy $s \in (-\infty, +\infty)$, there is a weight ω such that* $\mathcal{P}_s(\alpha_1, \alpha_2; 1, \omega) = \mathcal{P}_s(\beta_1, \beta_2; 1, \omega)$.

2. *If $\alpha_1 > \beta_1, \beta_2 > \alpha_2$, then for all $\omega > 0$ there exists a strategy s such that* $\mathcal{P}_s(\alpha_1, \alpha_2; 1, \omega) = \mathcal{P}_s(\beta_1, \beta_2; 1, \omega)$.

3. *If $\alpha_1 > \beta_2 > \alpha_2 > \beta_1$, then there exists a weight $\omega > 0$ such that there is no strategy s such that $\mathcal{P}_s(\alpha_1, \alpha_2; 1, \omega) = \mathcal{P}_s(\beta_1, \beta_2; 1, \omega)$.*

Proof of Proposition 5.7 *In each case, equality of overall preference is established by finding s and ω such that:*

$$\left(\frac{{\alpha_1}^s + \omega{\alpha_2}^s}{1+\omega}\right)^{\frac{1}{s}} = \left(\frac{{\beta_1}^s + \omega{\beta_2}^s}{1+\omega}\right)^{\frac{1}{s}}$$

Although this is not defined at $s = 0$, it is continuous through that point, and it suffices to consider

$${\alpha_1}^s + \omega{\alpha_2}^s = {\beta_1}^s + \omega{\beta_2}^s$$

Solving for ω gives:

$$\omega = \frac{{\alpha_1}^s - \omega{\beta_1}^s}{{\beta_2}^s - \omega{\alpha_2}^s} \tag{5.9}$$

1. For $s \neq 0$, (5.9) has a solution. This solution is positive since $\alpha_1 > \beta_1$ and $\beta_2 > \alpha_2$. In addition, in the limit as $s \longrightarrow 0$,

$$\omega = \frac{\log\alpha_1 - \log\beta_1}{\log\beta_2 - \log\alpha_2}$$

which is likewise a positive quantity.

2. If $\alpha_1 > \beta_1, \beta_2 > \alpha_2$, then consider the following two limits:

$$\lim_{s\longrightarrow-\infty} \frac{{\alpha_1}^s - \omega{\beta_1}^s}{{\beta_2}^s - \omega{\alpha_2}^s} \quad = \quad 0$$

$$\lim_{s\to\infty} \frac{{\alpha_1}^s - \omega{\beta_1}^s}{{\beta_2}^s - \omega{\alpha_2}^s} = \infty$$

Thus the range of the right-hand side of (5.9) is all positive ω, and for every ω there is a strategy s.

3. If $\alpha_1 > \beta_2 > \alpha_2 > \beta_1$, then

$$\mathcal{P}_s(\alpha_1, \alpha_2; 1, 1) > \mathcal{P}_s(\beta_1, \beta_2; 1, 1)$$

Furthermore, for any $\omega \in (0, 1)$,

$$\mathcal{P}_s(\alpha_1, \alpha_2; 1, \omega) > \mathcal{P}_s(\alpha_1, \alpha_2; 1, 1) > \mathcal{P}_s(\beta_1, \beta_2; 1, 1) > \mathcal{P}_s(\beta_1, \beta_2; 1, \omega)$$

Thus it is impossible to find a strategy s to achieve equality for all possible weights ω.

Furthermore,

$$\lim_{s\to-\infty} \frac{{\alpha_1}^s - \omega{\beta_1}^s}{{\beta_2}^s - \omega{\alpha_2}^s} = \infty$$
$$\lim_{s\to\infty} \frac{{\alpha_1}^s - \omega{\beta_1}^s}{{\beta_2}^s - \omega{\alpha_2}^s} = \infty$$

and there is a minimal ω which is reached when

$$\frac{d}{ds}\frac{{\alpha_1}^s - \omega{\beta_1}^s}{{\beta_2}^s - \omega{\alpha_2}^s} = \frac{1}{({\beta_2}^s - {\alpha_2}^s)^2} \times \Big(({\beta_2}^s - {\alpha_2}^s)({\alpha_1}^s \log\alpha_1 - {\beta_1}^s \log\beta_1) - ({\alpha_1}^s - {\beta_1}^s)({\beta_2}^s \log\beta_2 - {\alpha_2}^s \log\alpha_2) \Big)$$

is equal to 0. ■

Thus, if the strategy is chosen in advance, and is not the *min* or the *max*, then it is always possible to choose one point of two undominated points, by appropriate choice of weights. However, it is no longer always possible to choose any one of an undominated set. Methods which rely exclusively on the arithmetic mean, for example, do not truly consider all undominated points as potential solutions. There is, of course, nothing intrinsically wrong with restricting the set of possible points by selecting weights first, or strategies first, rather than simultaneously. In general, however, it will narrow the set of possible best alternatives.

5.7.4 Misrepresentation of preference as a negotiation strategy

The formal representation of a negotiation depends on an accurate portrayal of individual preferences. If the situation is indeed an adversarial negotiation in the ordinary sense of the word, then the results of the negotiation calculations can be manipulated through willful misrepresentation of preference.

The treatment in this thesis does not address the problem of manipulability, though it is undoubtedly important. Implementation of the model as a decision-making procedure presently relies on enforced honesty, and how to enforce it is left as a challenge to managers.

5.8 Hierarchical negotiation

The results of the preceding sections have been presented with respect to the aggregation of two preferences. In general, a design decision problem will require the aggregation of more, possibly many more, than two preferences. When more than two preferences must be combined, the aggregation is effected hierarchically: individual preferences are compared in pairs, and then groups are compared to each other. The choice of an aggregation hierarchy will depend on the problem, and on corporate structure.

5.8.1 Direct aggregation of more than two preferences

There is one special situation in which more than two preferences may be directly compared. If three or more preferences combine pairwise with the same strategy s, then all can be aggregated at once, regardless of their weights. To see this, consider three preferences α_1, α_2, and α_3, with respective weights ω_1, ω_2, and ω_3. Aggregating α_1 and α_2 using the strategy s gives the combined preference:

$$\alpha_{12} = \left(\frac{\omega_1 {\alpha_1}^s + \omega_2 {\alpha_2}^s}{\omega_1 + \omega_2} \right)^{\frac{1}{s}}$$

Assigning this new preference α_{12} the combined weight $\omega_1 + \omega_2$ and aggregating it with α_3 using the same strategy s gives:

$$\left(\frac{(\omega_1 + \omega_2){\alpha_{12}}^s + \omega_3{\alpha_3}^s}{(\omega_1 + \omega_2) + \omega_3}\right)^{\frac{1}{s}}$$

which reduces almost immediately to:

$$\left(\frac{\omega_1{\alpha_1}^s + \omega_2{\alpha_2}^s + \omega_3{\alpha_3}^s}{\omega_1 + \omega_2 + \omega_3}\right)^{\frac{1}{s}}$$

This is the same result as would be achieved by directly aggregating all three preferences. The consistency of the result depends upon the use of the same strategy s at each stage of the aggregation, and on the addition of the constituent weights for any intermediate preferences. The calculations can be continued to as many preferences as are needed, and a similar calculation holds for the special case when $s = 0$.

Direct aggregation of more than two preferences may be implemented whenever those preferences are determined to trade-off pairwise using the same strategy s. Although this situation may arise fortuitously by successive application of the method of Section 5.7.2, it applies more commonly when a decision maker determines that several attributes can sensibly be held to interact in an equivalent manner. For example, manufacturing preferences specified on several different sheet metal thicknesses may all combine in the same way.

5.8.2 Hierarchical aggregation of more than two preferences

When preferences do not all aggregate using the same strategy s, many preferences are still combined hierarchically. In the more general case, pairs of attributes are compared and then grouped together; the groups are then paired and then grouped together, and so on. The division into groups reflects the problem and the structure of the engineering team. It may be natural, for instance, to aggregate all preferences within each working group and then to combine those groups, or to aggregate design and performance preferences separately. It is helpful, but not necessary, for the person responsible for determining strategies and weights to be able to specify the indifference points of Section 5.7.2 using preference values. In the

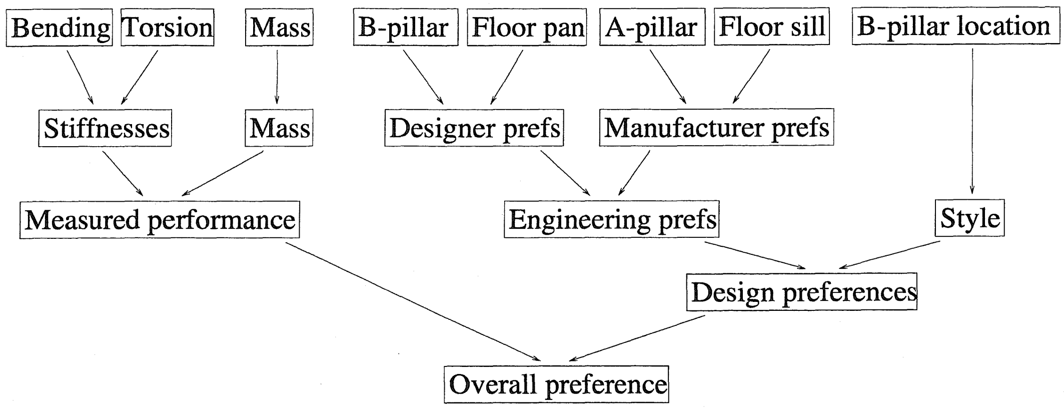

Figure 5.7: Hierarchy of preferences for VW example

example presented in Chapter 8, eight preferences are aggregated in a hierarchy, reproduced here in Figure 5.7; a complete discussion of this hierarchy is deferred to Chapter 8.

Although the separate aggregations in a hierarchy may all use different strategies, weights are propagated through the hierarchy in the same manner as when all strategies are identical. The individual weights assigned to the individual preferences are added to find the weight of the combined preference at each aggregation step. It is not necessary to specify weights at the outset, however. In a typical hierarchy, each aggregation will be assigned a strategy and a ratio of weights, and from these the weights of individual preferences can be calculated. The example in Chapter 8 uses this method to determine weights.

5.9 Examples of negotiation in engineering design

The following examples are not exhaustive, but they indicate a wide range of design negotiation situations. In addition, there are situations that are formally modelled as negotiations, though it is not so intuitively obvious that they are so.

5.9.1 Unreachable target performance values

One example of design negotiation occurs when an engineer or engineering group is given the task of designing a product to a target performance specification. When the product is a newer model of an existing product, the target is often an incremental improvement over last year's model. As an example, consider an automobile chassis, where an existing model has a torsional stiffness of 5000 N-m/degree and this year's requirement is to exceed that by 10%. If the engineers are unable to reach the fixed target easily, they will return to the manager who set the task and begin a negotiation process. Indeed, this meeting may be scheduled long before any potential problems are known. The engineers may ask for more resources, for a relaxation of other targets, or for a compromise on the original target. Targets are almost never immovable, and managers are commonly willing to negotiate.

Here, negotiation serves to address an inadequacy in the original description of the problem: the ostensibly exact (or *crisp*) requirement is in fact fuzzy, and through negotiation the two groups (in this case, chassis designers and their managers) explore the nature of the "constraints." To formalize this negotiation and reduce pre-distorted bargaining positions, it is necessary to represent the inherent fuzziness in the constraint.

5.9.2 Trade-offs between facets of performance

An additional layer of complication is added when several target performances are considered at once; here, negotiation can occur even when all specifications are met. In fact, there are usually at least two specifications, since cost of engineering and production resources is almost always a factor. In the example of the chassis design, the designers' position may be to offer the manager a choice between a 6% improvement at a production cost slightly lower than the present model's, or an 11% improvement at a substantially higher cost. To this, the managers may well counter that the new target is 8% improvement, as cheaply as possible.

The trade-off between cost and performance is one of many conflicts that are resolved through negotiation. A typical project will have an array of performance targets. The chassis example mentioned above will also have bending stiffness, weight, noise, and vibrational targets in addition to the torsional stiffness. The overall performance of the design depends on the individual performances, but the exact nature of the dependence varies greatly with the particular problem. The negotiation process is a means by which the true measure (and compromise) of overall performance is uncovered. A method to formalize negotiation may provide quicker and more complete information about the overall performance relationship.

5.9.3 Conflicts between design and manufacturing

The problem of design for manufacturability has been addressed by others (see [31]), but it has not previously been noted that conflicts between design and manufacturing are often resolved through a negotiation process. Sometimes the issue is the rejection of an unmanufacturable design by the production engineer. In many cases, a production engineer suggests changes that will make manufacture simpler, and negotiates with the design engineer for a compromise that will give the most satisfactory overall performance when production cost and reliability are taken into account. In the most optimistic case, a manufacturing group may suggest changes that improve the overall design performance. Much more often, there is either a degradation of performance or a need for the designer to expend more resources changing the design. A formalism for negotiation can help to facilitate the resolution of these conflicts, and can at the same time provide an unambiguous record of decisions that are made at each step of the design process.

5.9.4 Conflicts between engineering groups

When different working groups have responsibility for different subsystems, or for different aspects, of a design, the requirements of one group may conflict with the requirements of another. Stiffeners added to improve the structural rigidity of a frame might eliminate space that the fuel system group was counting on for the fuel tank. While in mature designs a structural part may well be described by a volume envelope and a few immutable points of contact, there are many situations in which the interaction between parts is not so rigidly described. Even when constraints are imposed in an attempt to avoid conflict between working groups, points of intersection between subsystems are often negotiated.

5.9.5 The incorporation of unquantifiable performances

Many design problems include performance criteria that are difficult, if not impossible, to measure, yet these criteria can be so important as to drive a design. Aesthetic and emotional concerns are certainly of great importance in the auto industry [19], and they also play a surprisingly significant role in other fields, from heavy machinery to military aircraft.

Style, beauty, appearance of solidity, color, and image are all examples of immeasurable attributes that can play a substantial role in the desirability of an engineered object. The fact that they are not easily quantified can lead to either underestimation or overestimation of their role in a design. An engineer designing for more concrete performance specifications may ignore them altogether, yet that same engineer may need to work within the strict geometrical constraints dictated by a stylist's aesthetic vision.

Immeasurable performances present the greatest challenge in the formalization of design negotiation as presented in this thesis. Still, steps can be taken to formalize this part of the design process, by allowing the engineer to interview other parties, such as stylists, and attempt to map their preferences onto the engineering performance model. The example presented in Chapter 8 includes preferences specified through such methods. The formalization can lead to a clearer picture of true overall design requirements.

5.10 Summary and implications for MADM

This chapter of the thesis treated the problem of selection of an aggregation function for Multi Attribute Decision Making. The problem was investigated within the context of the M_oI. The M_oI casts the preliminary design decision problem as a MADM problem, and uses different aggregation functions to formally model different trade-off strategies. The class of functions appropriate for the aggregation of these M_oI preferences has been explored. The results of this chapter are directly applicable to decision making in engineering design, and are also relevant to other MADM schemes.

In this chapter, a complete characterization of aggregation functions that satisfy the axioms of the M_oI was presented. The class of functions known as quasi-linear weighted means was shown to be crucial. It was demonstrated that any strictly monotonic design-appropriate aggregation function is generated by a generating function as detailed in Theorem 5.1. The conditional operators in use, while not weighted means, were shown to be limits of sequences of such functions. A parameterized family of functions was detailed, spanning two continua of possible design strategies, one between the non-compensating *min* and the compensating $\mathcal{P}_\Pi$, the two original aggregation functions of the M_oI, and one between $\mathcal{P}_\Pi$ and the (supercompensating) *max*.

Once aggregation is understood and characterized, negotiation can be formally defined. Since preferences arise prior to any negotiation, the specification of preference is not a feature of negotiation. It is the strategies and weights that are decided in concert that form the object of negotiations. The interconnection between weights and strategies, and a technique for the determination of the two, were presented and discussed.

Many MADM systems use aggregation functions, such as the arithmetic mean, that compensate between goals more aggressively than the existing functions of the MoI. These highly compensating functions may seem to be in conflict with the axioms of the MoI. This chapter has assessed the possibility of using these common aggregation functions in design decision-making problems.

There are an infinite number of aggregation operators that are suitable for engineering design. In particular, the parameterized family $\mathcal{P}_s$ is a range of functions that models a broad spectrum of design decision-making situations, with the parameter s indicating the degree of compensation permitted among performance criteria. The appropriate choice of the parameter s is problem-specific, and it was seen that a full negotiation should consider not only the compensation parameter s, but also the weights assigned to the individual preferences.

The use of an aggregation function in any MADM system may be justified on empirical as well as theoretical grounds. The development here has focused on determining a rational basis for the choice of an aggregation function. More empirical studies are needed to confirm in practice the results presented here.

Chapter 6

Convexity and Set-based Design

In the application of the M_OI to the MADM problem, three distinct questions can be asked. The first question is raised by the consideration of Arrow's Impossibility Theorem: Does the method give well-defined answers? The second also arises from the Impossibility Theorem: Are these answers guaranteed to be rational, and what does that mean in the context of engineering design? The third question goes beyond the Impossibility Theorem: What uncertainty is there in the decision problem, and how can the method handle it? This chapter attempts to address these questions.

In the consideration of decision making in Chapter 4, one requirement was that decisions be "rational." No formal definition of rationality was offered there, and the informal requirements to guarantee rationality were purposefully weak, so as to make the set of rational decisions as inclusive as possible. Indeed, the only example given of *irrational* decision was intransitivity of preference. Chapter 5 refined and formalized somewhat the notion of rational design. Violation of the aggregation axioms can be thought of as irrationalities in aggregation. There are other desirable features of engineering design decisions. It is not immediately obvious that violating these other requirements implies irrational design, but unquestionable advantages accrue when they are satisfied. In Section 5.7, for example, it was shown that the negotiation model presented here satisfies one such desirable condition: the result of a negotiation can choose any undominated point in the design space. This stands in marked contrast to other decision-making methods, which are less flexible in their choice of strategy.

Engineering design is naturally set-based [95]. The question at all stages of design is not, "What is the single best alternative?" but rather "What sets of designs are the best to consider at this stage?" Even a final design description contains set-based information, in the manufacturing tolerances. The reason for set-based description is twofold. First, it is assumed that design is always conducted with a model of the actual artifact. Even a physical prototype is a model, since it cannot be exactly identical to all instances of the final artifact. The use of a model assumes that a number of distinct instances of the artifact can be considered together as a set. Put another way, there is some resolution at which a set of designs can be identified with a single instance. Second, and perhaps more profound for the engineering design process, the design engineer thinks in terms of sets of designs at all stages of the process. Recognizing that different members of a set of designs have measurable, perhaps even gross, differences in parameters and performance, the engineer still considers the set as a single entity. Such set-based thinking has two obvious advantages: one, there is simply not enough time to consider all individual designs separately, especially in preliminary design; two, knowledge about sets of designs is much more useful than knowledge about individual "point" designs if the engineer must return to the drawing board to expand or alter the set of possible designs.

While there is nothing fundamentally irrational about a design decision/negotiation model that does not allow set-based descriptions, its value and applicability are certainly enhanced by a set-based approach. In this chapter, the notion of convexity of preference and its importance for set-based design are discussed. It will become clear that this convexity is different from convexity in linear programming (even multi-attribute linear programming), and also from Arrow's "convexity" that was discussed in Section 4.1. set-based design is also a natural solution to a sort of uncertainty that typically arises in preliminary design, the uncertainty of measured performance either because the description of the design is imprecise, or because "exact" answers, though available, are expensive.

6.1 Convexity of preference in the M_oI

The first question raised at the beginning of this chapter asks whether the method gives well-defined answers. Is the overall preference function of the M_oI a well-defined entity? There is no need to appeal to set-based reasoning to establish this point: it is assumed at the outset that there is an ordering among all possible alternatives that could be determined directly, if tediously, by exhaustive pairwise comparison. Also, the aggregation of independent preferences into a single preference function by the M_oI is a well-defined operation, so each alternative is assigned a single overall preference value. In Chapter 4 it was seen that Arrow's Impossibility Theorem casts doubt on the (well-defined) existence of a combined preference function *in the social choice problem.* That presents no difficulties to the M_oI, however, as the individual preferences specified in the engineering design problem do not obey Axioms 4.1 (unrestricted domain) and 4.4 (no imposed orders).

The second question asks what rationality means in the context of engineering design. This question requires the consideration of set-based reasoning. In the motivating paradox that began the discussion of the Impossibility Theorem, an intuitive notion of rationality was offended by an intransitive ordering. In the engineering design decision problem, especially at its preliminary stages, transitivity is only a first requirement for a rational order. An ability to represent *sets* of designs is another important feature; the designer is usually interested in determining which sets of designs are worth pursuing.

The sets in set-based reasoning are not arbitrary lists of individual designs. Designers have an intuitive notion of sets of designs as consisting of designs that are "close" to one another, so that it is sensible to consider them together. In addition, there is an intuitive notion that sets of designs ought to be "connected," so that a set can be reasonably supposed not to ignore any such "nearby" designs. These two ideas, of proximity and connectedness of sets of design alternatives, allow for the meaningful consideration of sets of design. Here we use a third concept, that of convexity, to give more precise definitions of proximity and connectedness of sets. set-based design will be seen to depend on convexity of design preference.

One approach to set-based design is interval analysis [53, 96], which imposes a strong connectedness condition on the sets by assuming that they are hypercubes in $\mathbb{R}^n$ described by their endpoints. This condition may be imposed on grounds of convenience, but it is clearly not necessary for rationality. The MoI uses a less restrictive connectedness condition.

Interval analysis makes use of a crucial feature of engineering design, that the design variables are ordered on external scales, so that the design variable space $\mathbf{D}$ can be treated as a subset of $\mathbb{R}^n$. In other words, design variables have dimensions. Sometimes the design variables are discrete, and $\mathbf{D}$ is contained in $\mathbb{R}^n \times \mathbb{Z}^m$; since the external scale still has meaning as a continuous variable, for purposes of interval analysis $\mathbf{D} \subset \mathbb{R}^{n+m}$. Arrow made use of an external scale to avoid the problems of the Impossibility Theorem, and included the further assumption that all preference orders were single-peaked along the single scale; the relevance of that assumption in the engineering design context will be discussed later. Sets in the MoI are ordered on the same external scale as in interval analysis, but the sets need not be hypercubes.

The identification of $\mathbf{D}$ with $\mathbb{R}^n$ allows the use of the metric space structure of the latter:

Definition 6.1 *A set of designs is called* connected *in the MoI if it is connected in the underlying space* $\mathbb{R}^n$.

While connectedness of sets of designs can be naturally identified with connectedness of sets in $\mathbb{R}^n$, the measurement of proximity of designs is not so clear. Separate dimensions are often not comparable, and a difference of k units may have different meanings at two different places along any one dimension.

A natural convexity condition arises from the identification of the *DVS* with $\mathbb{R}^n$. Since $\mu(\mathbf{x}) \in [0, 1]$, by adding a preference axis to the *DVS*, the overall preference function becomes an n-dimensional surface in $\mathbb{R}^{n+1}$. An important concept here, borrowed from fuzzy sets, is that of the *α-cut*:

Definition 6.2 *Given a preference function μ on a set $\mathbf{X}$, the set of all elements $\mathbf{x} \in \mathbf{X}$ such that $\mu(\mathbf{x}) \geq \alpha_1$ is called the* α-cut *of $\mathbf{X}$ at α_1.*

Note that an α-cut need not be connected. Connected α-cuts lend themselves naturally to set-based reasoning with respect to the preference μ; it is rational to use such an α-cut for set-based reasoning. Such a set has a simple geometric interpretation: the slice of the

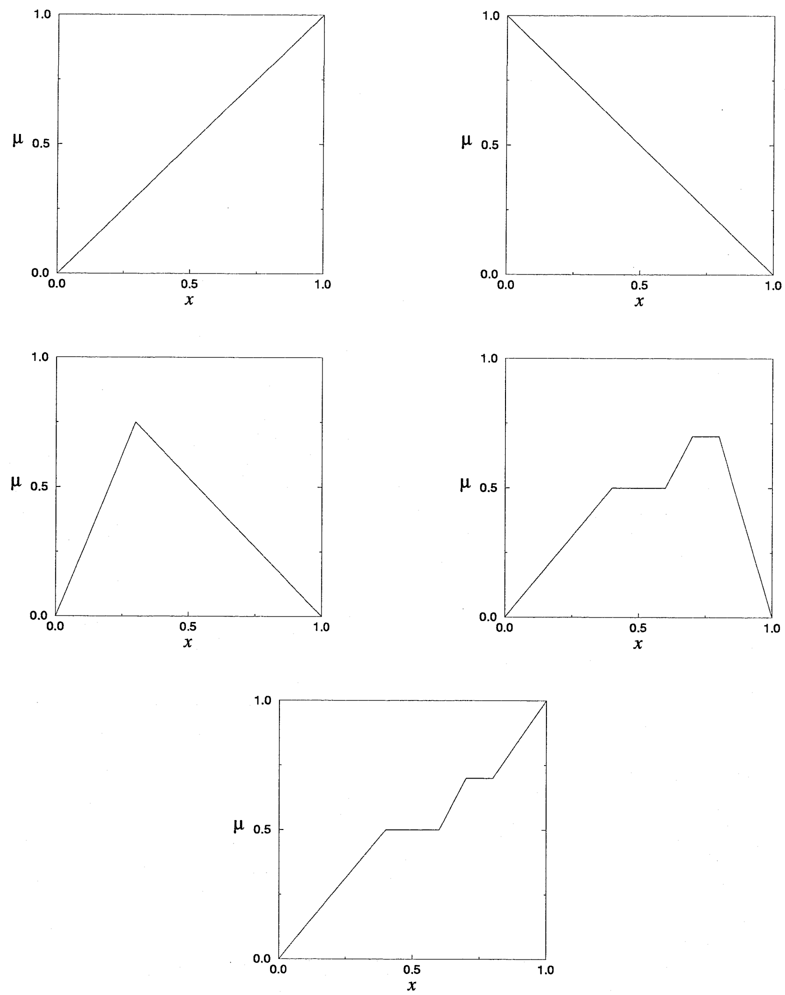

Figure 6.1: Functions convex with respect to μ, one dimension

preference surface, taken perpendicular to the μ-axis at α, must be connected. When all possible α-cuts are connected, the function is termed *μ-convex*:

Definition 6.3 *A preference function μ_i is called* convex with respect to μ, *or* μ-convex, *if every α-cut it induces is connected.*

The idea is easily represented (see Figure 6.1) when $n = 1$ (the case of a single design variable). For any of the μ-convex functions in Figure 6.1, any line parallel to the x-axis has only one connected section which is less than or equal to the function plotted. Note that the notion of μ-convexity is weaker than that of convexity in the usual sense. For example,

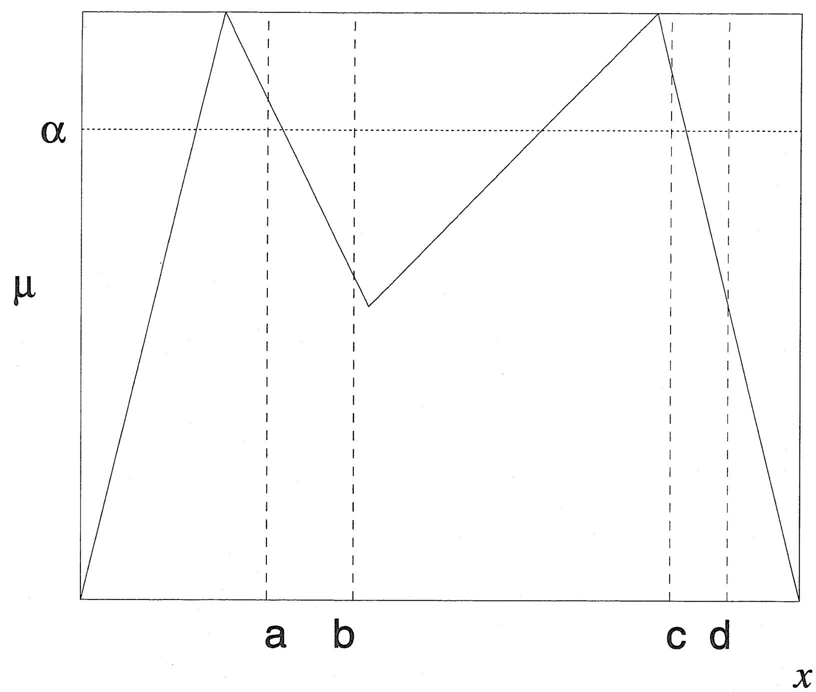

Figure 6.2: Proximity of designs

for two of the graphs in Figure 6.1 it is possible to draw a straight line, not parallel to the x-axis, that has more than one connected region under the function plotted.

At this point it is possible to discuss proximity of designs. Consider the preference function shown in Figure 6.2. Designs a and b, and designs c and d, appear "near" along the x-axis, yet a is much closer to c than to b in terms of preference μ, and indeed both are members of the (disjoint) α-cut at α that does not include b. Since the function graphed is not μ-convex, it is not clear which designs should be considered together as a set. For μ-convex functions, on the other hand, proximity in preference provides an unambiguous measure of proximity which can be used to unify the different scales of the other axes. While μ-convexity is not a necessary condition for the rationality of engineering design preferences, it is both common and desirable. Certainly it is possible to have two or more disjoint sets of "good" designs, as in Figure 6.2, but these sets are then often considered separately, effectively reasoning with the μ-convex portions of a divided set. Preferences may be rational without being μ-convex, but all μ-convex preferences are rational. Furthermore, μ-convexity is common. By determining the conditions under which μ-convexity holds, we can formalize a large subset of engineering design problems. For those (less common) problems which violate μ-convexity, rationality may still be assured if the problem can be divided into regions, perhaps disjoint, in which μ-convexity holds.

Thus by identifying a parameterized design space with $\mathbb{R}^n$, and by considering the additional axis of preference, it is seen that nearby designs are those that belong to the same μ-convex set, and are close in preference. If preference changes little over a wide range of a design variable d_i, two designs can be "near" to each other with large differences in d_i. If preference is sensitive to changes in another variable d_j, then small changes in d_j can make for designs that are "far apart." One must be careful, however. If there are two disjoint sets that achieve equal preference, those two sets cannot be considered to be near each other.

The μ-convexity of preferences indicates that the results of the M_oI will be useful in set-based reasoning. When can μ-convexity be guaranteed? A few results that begin to answer that question will now be presented. The design decision problem involves the combination of preferences over the n-dimensional *DVS*. Nevertheless, much can be learned from the case of a single design variable ($n = 1$). The proofs here (all of which are deferred to a later section) about the combination of two preferences on a one-dimensional *DVS* generalize to multiple preferences in higher dimensions. For instance, of all the family $\mathcal{P}_s$ of design-appropriate aggregation functions used in the M_oI, only the *min* guarantees that the combination of several μ-convex sets will also be μ-convex:

Proposition 6.1 *Assume that preferences are continuous. When the DVS is a connected subset of* $\mathbb{R}$, *if*

$$\bar{\mu}(x) = \mathcal{P}(\mu_1, \ldots, \mu_n; \omega_1, \ldots, \omega_n)(x) = \min\Big(\mu_1(x), \ldots, \mu_n(x)\Big),$$

then μ*-convexity of* μ_i *for all* i *is sufficient to guarantee* μ*-convexity of* $\bar{\mu}$*. If* μ_i *is not* μ*-convex for some* i*, then there exist* μ*-convex* $\mu_1, \ldots, \mu_n$ *(not including* μ_i*) that combine to a* $\bar{\mu}$ *that is not* μ*-convex.*

Note that the same result does not hold for continuous functions that exceed the *min*. In fact, it does not hold for arbitrary t-norms either. (A t-norm is bounded above by *min*.) Consider the following counterexample: Let $\mathcal{P}(\mu_1, \mu_2; \omega_1, \omega_2)$ obey the axioms, and let $\hat{\mathcal{P}}(\alpha, \beta; \omega_1, \omega_2) > \min(\alpha, \beta)$ for the corresponding $\hat{\mathcal{P}}$. We may assume $\alpha > \beta$ without loss of generality. Then construct the two preferences μ_1, μ_2 (shown in Figure 6.3):

$$\begin{aligned}
\mu_1(x) &= \alpha, x \leq x_1 \\
& \quad \beta, x \geq x_2 \\
& \quad \text{with linear interpolation from } x_1 - x_2 \\
\mu_2(x) &= \beta, x \leq x_2 \\
& \quad \alpha, x \geq x_3 \\
& \quad \text{with linear interpolation from } x_2 - x_3
\end{aligned}$$

Then by idempotency:

$$\mathcal{P}(\mu_1, \mu_2; \omega_1, \omega_2)(x_2) = \beta$$

but by definition of $\mathcal{P}$:

$$\begin{aligned}
\mathcal{P}(\mu_1, \mu_2; \omega_1, \omega_2)(x_1) &> \beta \\
\mathcal{P}(\mu_1, \mu_2; \omega_1, \omega_2)(x_3) &> \beta
\end{aligned}$$

so $\bar{\mu}$ defined by $\mathcal{P}$ is not μ-convex.

Proposition 6.1 shows that a requirement that all individual preferences be μ-convex is insufficient to guarantee that the overall preference is μ-convex. Necessary and sufficient conditions for μ-convexity of the overall preference function are not known, but a sufficient condition is presented here. While μ-convexity of the overall preference is intuitively desirable, a slightly stronger condition, that of ordinary convexity, can be imposed, and will assure rational combinations. Recall that a function is convex whenever the chord connecting two points on the graph of the function lies under or on the graph:

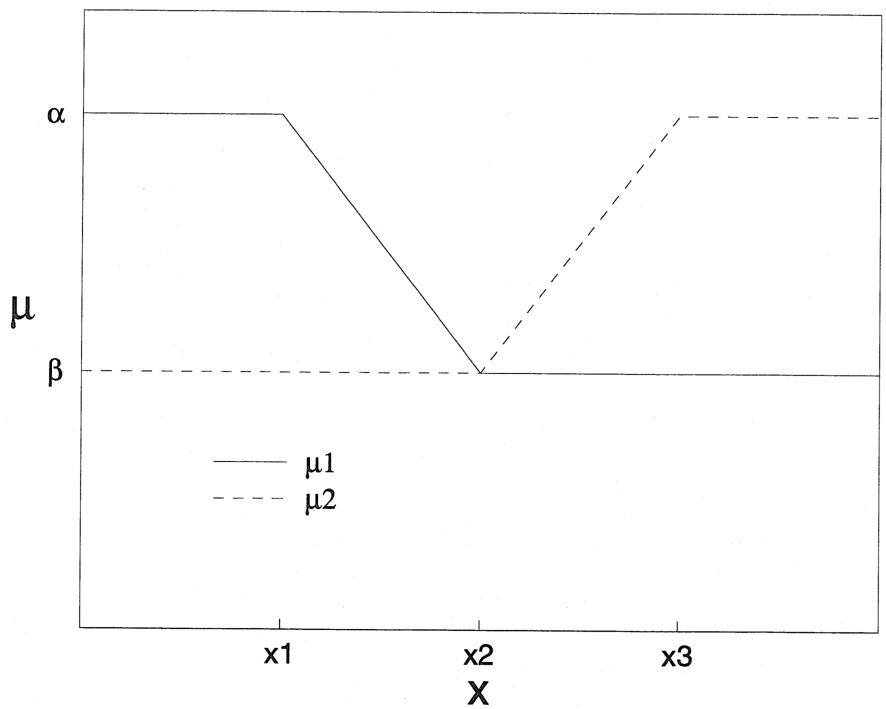

Figure 6.3: μ-convex μ_1, μ_2 that combine to non-μ-convex $\bar{\mu}$

Definition 6.4 *A function f on* $\mathbb{R}$ *is* convex *if*

$$f(x_2) \geq \frac{x_2 - x_1}{x_3 - x_1} f(x_1) + \frac{x_3 - x_2}{x_3 - x_1} f(x_3)$$

whenever $x_1 < x_2 < x_3$. *A function that is not convex is called* concave.

The combination of convex (in the ordinary sense) preference functions is again convex for the entire family of MoI aggregation functions up to and including the arithmetic mean. As was demonstrated above, when individual preferences are μ-convex but concave, then only the *min* can guarantee μ-convexity.

Proposition 6.2 *Assume that preferences are continuous. When the* DVS *is a connected subset of* $\mathbb{R}$, *if* $s \leq 1$, *and the preference functions* μ_1 *and* μ_2 *are convex, then the overall preference function*

$$\bar{\mu}(x) = \mathcal{P}_s(\mu_1, \mu_2; \omega_1, \omega_2)(x)$$

is also convex.

Thus the arithmetic mean is in some sense a bounding function of this family of design-appropriate functions. If any concavity (in the ordinary sense) is allowed in the individual preference functions, then μ-convexity of $\bar{\mu}$, and hence the connectedness of an arbitrary α-cut, can only be guaranteed by using the *min* as an aggregation function.

Proposition 6.2 also shows that for functions more compensating than the arithmetic mean (*i.e.*, for $s > 1$), not even convexity of all individual preferences suffices to guarantee μ-convexity of the overall preference. Functions that compensate to a higher degree than the arithmetic mean fail this particular test of rationality. This complements a limitation on s that was discovered earlier: in the procedure presented in Section 5.7, it was seen that functions with $s > 1$ lie outside the bounds of elicitation with indifference curves.

Propositions 6.1 and 6.2 are also applicable in arbitrary dimensions; the same proofs hold for any two-dimensional vertical slice in any direction, *i.e.*, for any one-dimensional subspace of $\mathbb{R}^n$ and its Cartesian product with the preference space $[0, 1]$.

Convexity is a sufficient but not a necessary condition for rationality (in its broad new meaning). It has been applied not just for mathematical simplicity, but because it is a common feature of engineering design problems. It is beyond the scope of this thesis to determine if the prevalence of μ-convexity is a structural feature imposed on problems by designers, or if it is unavoidable in engineering design.

6.2 Proofs

Proof of Proposition 6.1 *Since* $\mathcal{P} = \min$ *is associative, it suffices to prove the claim for* $n = 2$. *Say that* $\mu_1, \mu_2 : DVS \mapsto [0,1]$ *are both* μ*-convex. Consider any three points* $x_1 < x_2 < x_3 \in DVS$. *Then* μ*-convexity of* μ_1, μ_2 *implies:*

$$
\begin{aligned}
\mu_1(x_2) &\geq \min\Big(\mu_1(x_1), \mu_1(x_3)\Big) \\
\mu_2(x_2) &\geq \min\Big(\mu_2(x_1), \mu_2(x_3)\Big)
\end{aligned}
$$

Otherwise, if $\mu_1(x_2) < \mu_1(x_1) \leq \mu_1(x_3)$ *for instance, then the* α*-cut at* $\mu_1(x_1)$ *would not be connected. Therefore*

$$
\min\Big(\mu_1(x_2), \mu_2(x_2)\Big) \geq \min\Big(\mu_1(x_1), \mu_2(x_1), \mu_1(x_3), \mu_2(x_3)\Big) \tag{6.1}
$$

Now, by definition of $\mathcal{P}$*:*

$$
\begin{aligned}
\bar{\mu}(x_1) &= \min\Big(\mu_1(x_1), \mu_2(x_1)\Big) \\
\bar{\mu}(x_2) &= \min\Big(\mu_1(x_2), \mu_2(x_2)\Big) \\
\bar{\mu}(x_3) &= \min\Big(\mu_1(x_3), \mu_2(x_3)\Big)
\end{aligned} \tag{6.2}
$$

and it follows that

$$
\min\Big(\bar{\mu}(x_1), \bar{\mu}(x_3)\Big) = \min\Big(\mu_1(x_1), \mu_2(x_1), \mu_1(x_3), \mu_2(x_3)\Big) \tag{6.3}
$$

Taking (6.1), (6.2), and (6.3) together shows that:

$$
\bar{\mu}(x_2) \geq \min\Big(\bar{\mu}(x_1), \bar{\mu}(x_3)\Big)
$$

proving that $\bar{\mu}$ is μ-convex.

If any μ_i is not μ-convex, then setting all other μ_j identically equal to 1 makes $\bar{\mu} = \mu_i$, and therefore not μ-convex. ■

Proof of Proposition 6.2 *If a preference μ is convex, then for any three points $x_1 < x_2 < x_3$:*

$$\mu(x_2) \geq \frac{x_2 - x_1}{x_3 - x_1}\mu(x_1) + \frac{x_3 - x_2}{x_3 - x_1}\mu(x_3)$$

If two preference functions μ_1 and μ_2 are convex, but $\mathcal{P}_s(\mu_1, \mu_2; \omega_1, \omega_2)$ is not, then there exist $x_1 < x_2 < x_3$ such that:

$$\begin{aligned}\mathcal{P}_s\left(\mu_1(x_2), \mu_2(x_2); \omega_1, \omega_2\right) < \\ \frac{x_2 - x_1}{x_3 - x_1}\mathcal{P}_s\left(\mu_1(x_1), \mu_2(x_1); \omega_1, \omega_2\right) + \frac{x_3 - x_2}{x_3 - x_1}\mathcal{P}_s\left(\mu_1(x_3), \mu_2(x_3); \omega_1, \omega_2\right)\end{aligned} \tag{6.4}$$

Since μ_1 is convex,

$$\mu_1(x_2) \geq \frac{x_2 - x_1}{x_3 - x_1}\mu_1(x_1) + \frac{x_3 - x_2}{x_3 - x_1}\mu_1(x_3)$$

and thus, by monotonicity of $\mathcal{P}_s$ and (6.4):

$$\begin{aligned}\mathcal{P}_s\left(\frac{x_2 - x_1}{x_3 - x_1}\mu(x_1) + \frac{x_3 - x_2}{x_3 - x_1}\mu(x_3), \mu_2(x_2); \omega_1, \omega_2\right) < \\ \frac{x_2 - x_1}{x_3 - x_1}\mathcal{P}_s\left(\mu_1(x_1), \mu_2(x_1); \omega_1, \omega_2\right) + \frac{x_3 - x_2}{x_3 - x_1}\mathcal{P}_s\left(\mu_1(x_3), \mu_2(x_3); \omega_1, \omega_2\right)\end{aligned} \tag{6.5}$$

Likewise, using convexity of μ_2 and monotonicity of $\mathcal{P}_s$ together with (6.5):

$$\begin{aligned}\mathcal{P}_s\left(\frac{x_2 - x_1}{x_3 - x_1}\mu_1(x_1) + \frac{x_3 - x_2}{x_3 - x_1}\mu_1(x_3), \frac{x_2 - x_1}{x_3 - x_1}\mu_2(x_1) + \frac{x_3 - x_2}{x_3 - x_1}\mu_2(x_3); \omega_1, \omega_2\right) < \\ \frac{x_2 - x_1}{x_3 - x_1}\mathcal{P}_s\left(\mu_1(x_1), \mu_2(x_1); \omega_1, \omega_2\right) + \frac{x_3 - x_2}{x_3 - x_1}\mathcal{P}_s\left(\mu_1(x_3), \mu_2(x_3); \omega_1, \omega_2\right)\end{aligned}$$

Therefore, if two convex preferences combine to a non-convex preference, then two linear preferences (which interpolate the original preference functions at x_1 and x_3) also combine to a non-concave preference. It thus suffices to show that $\mathcal{P}_s(\mu_1, \mu_2; \omega_1, \omega_2)$ is convex whenever $s \leq 1$ and μ_1 and μ_2 are linear.

Consider two linear preference functions μ_1 and μ_2:

$$\begin{aligned}\mu_1(x) &= ax+b\\ \mu_2(x) &= cx+d\end{aligned}$$

Assume without loss of generality that weights are normalized so that $\omega_1+\omega_2=1$. Then the aggregated preference is:

$$\mathcal{P}_s(\mu_1,\mu_2;\omega_1,\omega_2)(x) = \Big(\omega_1(ax+b)^s+\omega_2(cx+d)^s\Big)^{\frac{1}{s}}$$

Differentiating twice with respect to x yields:

$$\begin{aligned}&\frac{d^2}{dx^2}\mathcal{P}_s(\mu_1,\mu_2;\omega_1,\omega_2) = \qquad (6.6)\\ &\left(\frac{1}{s}-1\right)\Big(\omega_1a(ax+b)^{s-1}+\omega_2c(cx+d)^{s-1}\Big)^2\Big(\omega_1(ax+b)^s+\omega_2(cx+d)^s\Big)^{\frac{1}{s}-2}+\\ &(s-1)\Big(\omega_1a^2(ax+b)^{s-2}+\omega_2c^2(cx+d)^{s-2}\Big)\Big(\omega_1(ax+b)^s+\omega_2(cx+d)^s\Big)^{\frac{1}{s}-1}\end{aligned}$$

Consider three cases:

1. $s\le 1$, $s\neq 0$.

   In this case, both terms on the right-hand side of (6.6) are non-positive. Note that:

$$\begin{aligned}ax+b, cx+d &\in [0,1]\\ \Big(\omega_1(ax+b)^s+\omega_2(cx+d)^s\Big)^{\frac{1}{s}} &\in [0,1]\end{aligned}$$

   as they are preferences. Therefore, $\frac{d^2}{dx^2}\mathcal{P}_s(\mu_1,\mu_2;\omega_1,\omega_2)\le 0$, which is equivalent to convexity.

2. $s=0$.

   If $s=0$, then (6.6) cannot be used. In this case,

$$\mathcal{P}_0(\mu_1,\mu_2;\omega_1,\omega_2) = (ax+b)^{\omega_1}(cx+d)^{\omega_2}$$

and

$$\frac{d^2}{dx^2}\mathcal{P}_s(\mu_1, \mu_2; \omega_1, \omega_2) =$$
$$a^2\omega_1(\omega_1 - 1)(ax + b)^{\omega_1 - 2} + c^2\omega_2(\omega_2 - 1)(cx + d)^{\omega_2 - 2}$$

which is also clearly non-positive, so the combined preference is again convex.

3. $s > 1$.

   When $s > 1$, *it is no longer true that convex preferences combine to convex preferences. Consider the two simple linear preference functions, defined for* $x \in [0, 1]$*:*

$$\begin{aligned} \mu_1(x) &= x \\ \mu_2(x) &= 1 - x \end{aligned}$$

   Then

$$\mathcal{P}_s\left(\mu_1(0.5), \mu_2(0.5); 0.5, 0.5\right) = 0.5$$

   but as long as $s > 1$,

$$\begin{aligned} \mathcal{P}_s\left(\mu_1(0), \mu_2(0); 0.5, 0.5\right) &> 0.5 \\ \mathcal{P}_s\left(\mu_1(0), \mu_2(0); 0.5, 0.5\right) &> 0.5 \end{aligned}$$

   showing that $\mathcal{P}_s(\mu_1, \mu_2; 0.5, 0.5)$ *is not convex.*

This completes the proof. ■

Chapter 7

Computation Methods Comparison

Every design problem is a "maximization" problem, in the sense that the problem includes a search for a "best" design. When a computable measure of design performance is available, engineers can apply a number of mathematical optimization methods to maximize that performance, and indeed the design engineer usually has numerical targets that a design should achieve. These measurable performance targets are only a part of the complete performance of a design, and designs that achieve the highest levels of measured performance will not necessarily be chosen, as the computed performance is compared with other, less easily computed measures, such as aesthetics, market preferences, or manufacturability. The importance of uncalculated performances such as styling and manufacturing has been an obstacle to the direct application of optimization methods to the complete design problem.

The M_oI, as discussed above, takes an engineering analysis problem and expands it to include (imprecise) information that is relevant to the decision but presently only incorporated informally. The result is the formal calculation of an overall preference $\bar{\mu} \in [0, 1]$ for each candidate design. Since $\bar{\mu}$ is a scalar function of the design variables, the design problem becomes a computable maximization problem: maximize overall preference for the design.

One possible approach to this maximization problem is to apply an optimization algorithm directly to the scalar function $\bar{\mu}$. The M_oI has incorporated optimization internally, but has not treated the overall preference function $\bar{\mu}$ as an objective function to be maximized. One reason is that optimization routines usually offer point-by-point information, but more information is available to the designer when considering sets of designs or ranges of design variable values to pursue than in a single overall optimum. In addition, the preference function $\bar{\mu}$ is not a black box, but has some known structure that comes from the specification of the imprecise information.

Several different optimization methods were applied to a vehicle structure design problem with 16 design variables. In this chapter, data from the implementation of these methods are presented and compared. The methods were not fine-tuned to the problem, as the goal was not to conduct a competition between methods, but to look for insight into which sorts of optimization might be useful to use with or incorporate into the design problem.

7.1 Optimization strategies

A typical design problem can have from a handful to many tens of design variables. When f is not an analytic function, as is usually the case, exhaustive enumeration and evaluation of all possible candidate designs is at least enormously expensive and often impossible. Various optimization techniques are available to search for good solutions without resorting to exhaustive enumeration, including classical, random, and hybrid methods.

The MoI has previously incorporated classical optimization into its internal computations (*e.g.*, for finding internal extremal values of non-monotonic f's [48]), but not to identify the point in the *DVS* with the best value of overall preference. Since the focus of the MoI is to perform trade-offs to determine the performance and preference for *sets* of designs, and the computation of a single point $\bar{\mu}$ may be computationally expensive if f is expensive (*e.g.*, a finite element model), we have chosen not to incur the computational costs of searching for the "best" design point. Additionally, $\bar{\mu}$ is often relatively well-behaved, because the functions f are often more or less unimodal with respect to at least some of the d_i (*e.g.*, stiffness tends to increase with thickness), and because the preference structure is rational.

The optimization strategies outlined below have been implemented without sophisticated efforts to tailor the method to any particular problem. The study was not meant to be exhaustive; simulated annealing, tunneling algorithms, and hybrid methods were not considered. The results of this study could be used to guide the further application and refinement of optimization techniques to the preliminary design problem as modelled by the $M_O I$.

7.1.1 Exhaustive search

The simplest, most naïve search technique to find a global maximum of a scalar function is an exhaustive search of a suitably fine grid of the entire search space. The well-known "curse of dimensionality" [12] is a reminder that such a technique will require hopelessly many function evaluations whenever there are more than a few variables or a need for a relatively fine grid. In a case such as the example considered in Section 7.4, where there are 16 design variables, and where it is reasonable to require a discretization of 10 steps in each dimension, even at a million evaluations a second the search would require more than 300 years; the situation is even more dire when each evaluation takes 5 seconds (which is the case in the example). The considerably trimmed example in Chapter 8, with only five design variables but a more complicated analysis model, would take 70 days to explore ten points in each dimension.

7.1.2 Classical optimization

Classical multi-variable optimization schemes [2] fall into two categories: calculus-based gradient methods, and search methods. Since the former require derivative information, which is rarely available from the functions f encountered in engineering design analysis, a multi-variable method known as a *pattern search* was employed here. A pattern search starts with a relatively coarse step size, proceeds in ever larger steps in the apparently best direction, until progress ceases, when it stops at the last best guess and starts off again with the original step size. When it converges to a point at the given step size, a new, smaller step size is chosen, and exploration continues. A slightly more sophisticated method known as Powell's method was previously implemented in the internal computations of the $M_O I$ [46].

A pattern search is not expected to handle multiple optima gracefully, but can be quite efficient when a function is reasonably well-behaved.

7.1.3 Genetic algorithms

Genetic algorithms [34] are structured random explorations of a space, which proceed by analogy to the mechanism of natural selection in biology. A simple genetic algorithm describes the candidate designs, not by the values of the design variables d_i, but by a binary encoding of those values. Thus each design, or *individual*, is a string of 1's and 0's. Here, each d_i was simply discretized into 2^n possibilities, and the binary representations were concatenated into a single string of length nm, where m represents the dimension of the *DVS*. A random population of individuals is generated as a starting point. For each generation, the *fitness* (here, measured by overall preference) of each individual is calculated, and a random group, weighted towards the fitter individuals, is selected to populate the next generation. These individuals are *crossed* (a pair of individual strings is broken at a random point, and the parts are recombined to make two new individuals), and at rare intervals mutated by swapping a single bit. The random element of the genetic search is meant to prevent the algorithm from stopping at local optima.

Genetic algorithms require more computations than search methods, but should be more robust to multiple optima. In addition, a genetic algorithm tracks many individuals at once, so it can in theory simultaneously converge to multiple optima. This makes it an attractive candidate to provide set-based information.

Many refinements to the simple genetic algorithm are possible; here the only departure from a simple evaluate–cross–mutate scheme was the enforced survival of the single fittest individual in each generation to the subsequent generation. The scheme was also tried without this addition.

7.2 Optimizing for sets of designs

Casting the design problem as a scalar optimization problem obscures a key feature of the original problem: in general, the designer is not as interested in a single most-preferred or best-performing design, as in *sets* of designs that are most promising [95]. Particularly at the preliminary design stage, the designer wants a design direction to pursue, and that is much better described by ranges on the design variables d_i rather than a single best point.

If, for example, the peak preference of $\bar{\mu}^* = 0.7$ occurs at $d_1 = 53$, the designer may want to know over which range a preference of $\bar{\mu} = 0.6$ can be expected. The design decisions may be quite different depending on the size of this set.

The implementation of the M_oI makes use of the natural set-based information that is available from the specification of preferences to search for sets of candidate designs. Any traditional optimization scheme runs the risk of omitting that information.

7.3 Approximation techniques

Computation cost is a limiting factor for much of design analysis. Preliminary design decisions often rely on intuition and judgement because computation is expensive, and the information describing each design alternative is imprecise. The computation of overall preference $\bar{\mu}$ becomes more costly as the number of design variables increases. For any number of variables, the calculations and aggregations of preferences are often of negligible computational cost compared to the cost of evaluating f. A typical finite element model, for instance, can be expected to take several minutes on the fastest available computer, while the preference calculations take microseconds. While the curse of dimensionality must ultimately be addressed for any function, the analysis function f is a clear candidate for possible approximation. Previous work exploring parsimonious selection of points at which to compute f, and approximation between these points, can be found in [46].

7.4 Example

A finite element model for stiffness evaluation of a passenger vehicle, where sheet-metal thickness of the various panels can be varied, was used as a test case. This problem was chosen for several reasons: it was a black box, so the answer was not known in advance; it had 16 design variables and could be computed in about 5 seconds, so exhaustive search was all but impossible, but searches of several thousands or even tens of thousands of runs were feasible. However, the model is not an accurate test for any particular automobile; the design variables are for evaluating and comparing methods, not designing a car, and the preferences on those variables are likewise for testing the methods. This example, which

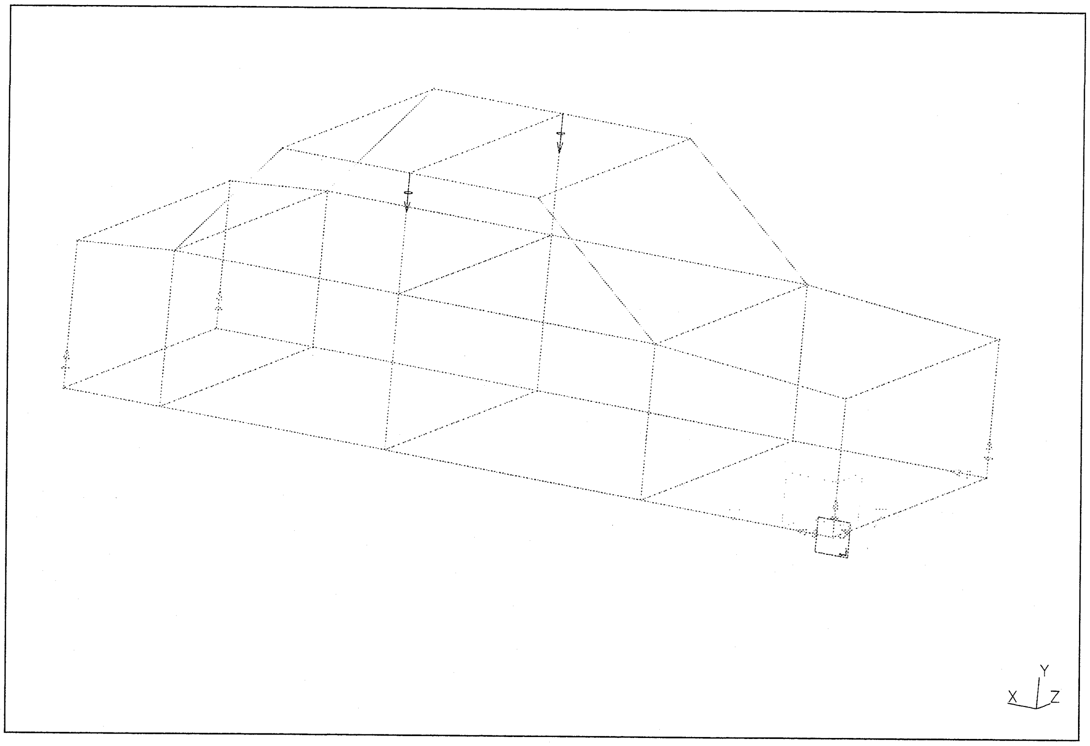

Figure 7.1: Finite element model in wireframe

focuses on computational issues, complements the example in Chapter 8, which focuses on decision issues.

A simple finite element model (see Figured 7.1) for calculating the bending stiffness of a passenger car body was used to test the optimization methods. The model was kept extremely simple so that the stiffness calculation could be performed in approximately 5 seconds using available computation (MSC Nastran on a Sun UltraSparc 170MHz processor workstation). The methodology would be the same for more complicated analysis models. The thickness of each body panel shown in Figure 7.1, and the cross section of two of the three pillars, were used as design variables. (This set of design variables was chosen for convenience in testing the model and the tools.) The panel thicknesses were allowed to vary between 5 and 20 millimeters, and the pillars were taken to have square cross sections with sides between 50 and 150 millimeters. Designer preferences, representing additional information that would ordinarily be considered informally, were specified on each of the design variables. The design preference on the first design variable, $\mu_{\mathrm{D}}(d_1)$ (Figure 7.2),

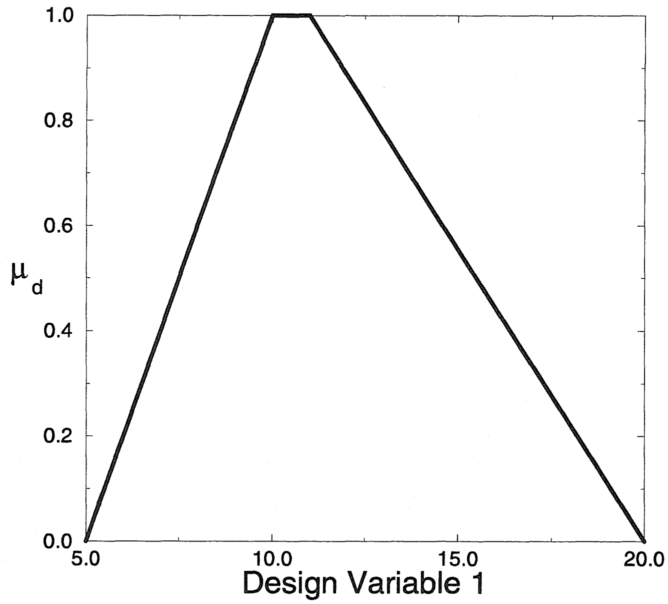

Figure 7.2: Representative design preference

is representative. The preference of $\mu_{\mathbf{D}} = 0$ at the endpoints $d_1 = 5$ and $d_1 = 20$ represents unacceptable values of d_1, while the preference $\mu_{\mathbf{D}} = 1$ for $d_1 \in [10, 11]$ indicates that these values are the most preferred, with intermediate levels of preference between. These preferences may come from manufacturing engineers, for example, and mean that they have found that sheet metal in that thickness range is most easily formed into that body panel. Preferences can also capture styling concerns, information about availability, and simple design intuition (a thick pillar in this simple model, for instance, may mean a more complicated and expensive part; the engineer uses this simple model with preferences attached rather than detailing the potential complication of a "thicker" pillar). Other design preferences were similarly specified.

For this simplified evaluation of methods, the finite element model has a single calculated engineering performance: bending stiffness. The preference for bending stiffness (Figure 7.3) is an imprecise performance target. The preference for the bending stiffness achieved by a design is combined with and traded-off against the design preferences on the individual variables. In this example, the individual design preferences were all weighted

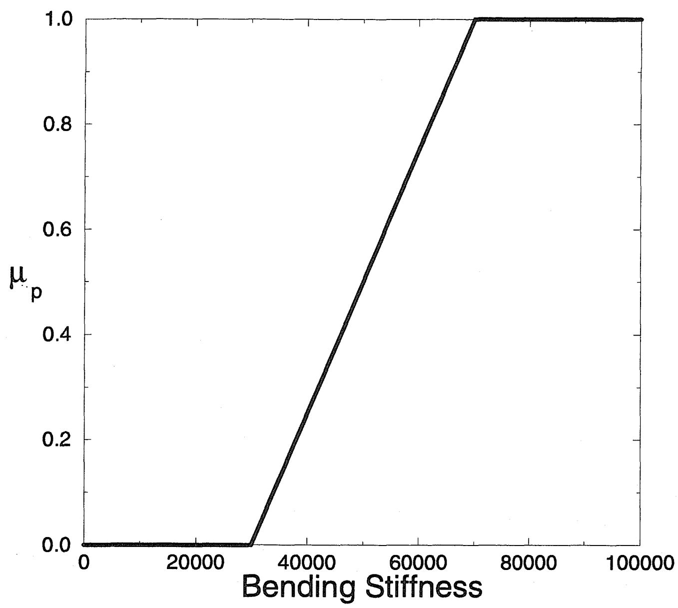

Figure 7.3: Performance preference

equally, and were allowed to trade-off in a non-compensating manner [57]. The aggregated design preference traded-off in a compensating manner with the performance preference [57]. In a traditional design process, the vehicle structure engineer would seek to maximize the bending stiffness, and then make trade-offs from the stiffest design in response to the needs of other groups working on the same part. Here, these negotiated changes are incorporated from the beginning.

7.5 Results

At 5 seconds for each finite element analysis, an enumerative search at a resolution of 10 steps in each design dimension would take 10^{16} evaluations, or on the order of 10^{11} years. Since five years is a more typical figure for the entire automobile design process, and preliminary or conceptual design is usually limited to a few months, the enumerative option was not pursued. Since an enumerative option was unrealistic, and the function f was not known precisely, the maximum achievable $\bar{\mu}$ is not certain. This is of course the most common case in a real design situation: it is not often that a designer can guarantee optimality or precision.

All of the methods achieved similar maximum overall preferences, at slightly less than $\bar{\mu} = 0.30$. Results from the various approaches are detailed below:

7.5.1 Classical optimization

The pattern search was started with an initial step size of 20% of the design variable range, and converged to a minimum step size of 0.15% of the range. The pattern search made fast progress early in the search, and then slowed down:

Evaluations	max $\bar{\mu}$
34	0.1686
67	0.2447
133	0.2678
297	0.2717
395	0.2792
525	0.2792
623	0.2803
721	0.2803

The best-performing individual design had design variable values:

d_1	17.7500
d_2	12.5000
d_3	16.2500
d_4	20.0000
d_5	17.9375
d_6	20.0000
d_7	13.7187
d_8	15.6875
d_9	18.6875
d_{10}	18.6875
d_{11}	18.6875
d_{12}	18.6875
d_{13}	16.2500
d_{14}	20.0000
d_{15}	162.5000
d_{16}	137.5000
$\bar{\mu}$	0.2803

7.5.2 Genetic algorithm

The performance of the genetic algorithm was less predictable than that of the pattern search. Indeed, since mutation (at a low but nonzero probability) is employed to avoid local minima, the algorithm does not necessarily make progress at each step:

Evaluations	max $\bar{\mu}$	average $\bar{\mu}$
63	0.0828	0.0039
105	0.1240	0.0844
252	0.1313	0.1045
525	0.1546	0.1013
630	0.2105	0.1515
1659	0.2226	0.1662
3507	0.2206	0.1717
3528	0.2502	0.1989
5124	0.2618	0.2239
6237	0.2622	0.2177
7896	0.2663	0.2476
11382	0.2686	0.2065
12936	0.2644	0.2077

The general trend is to improved performance, but in 6000 function evaluations the genetic algorithm had not yet achieved the value that the pattern search reached in 133 evaluations. Above 5000 function evaluations, the routine makes little progress. The highest average fitness was achieved after 7896 evaluations (376 generations), the highest maximum after 11382 evaluations, and the program was stopped after 12936 evaluations.

The population size was set at 21 for this search, and it is interesting to see the range of each variable. In the generation that included the highest maximum performance, seven individuals survived with $\bar{\mu} > 0.26$:

	Individuals						
	A	B	C	D	E	F	G
d_1	18.10	18.10	17.86	17.86	17.86	17.86	17.86
d_2	12.62	12.62	12.62	12.86	12.62	12.62	12.62
d_3	16.43	17.38	16.43	17.38	8.81	16.43	16.43
d_4	20.00	20.00	19.52	20.00	19.52	20.00	20.00
d_5	14.05	15.71	14.29	16.19	14.29	13.33	15.48
d_6	12.86	12.86	12.86	12.86	12.86	12.86	13.33
d_7	12.38	12.38	12.38	12.38	11.43	12.38	12.14
d_8	11.19	11.19	11.19	11.19	11.19	11.19	11.19
d_9	15.00	15.00	14.52	15.00	15.95	14.52	14.52
d_{10}	18.81	18.33	16.90	18.81	18.81	18.81	16.90
d_{11}	18.81	18.81	18.81	18.81	18.81	18.81	18.81
d_{12}	13.33	13.33	12.86	13.33	13.33	13.33	13.33
d_{13}	16.43	16.43	16.43	16.43	16.43	16.43	16.43
d_{14}	20.00	19.52	20.00	19.52	20.00	19.52	15.95
d_{15}	94.44	94.44	69.05	94.44	94.44	97.62	94.44
d_{16}	83.33	83.33	83.33	83.33	81.75	83.33	78.57
$\bar{\mu}$	0.2681	0.2680	0.2668	0.2680	0.2649	0.2628	0.2686

While these figures do not indicate that all designs with variables in those ranges above will achieve $\bar{\mu} = 0.26$, they give the designer more information than a single point showing the maximum performance found.

7.5.3 Set-based search with minimum assumption

The M$_\text{O}$I has historically employed a preference propagation mechanism based on the extension principle from fuzzy set theory [24, 108]. This mechanism offers computational efficiency at the expense of two assumptions about design problem: the analysis function f must be well-behaved (indeed, monotonic), and the preference combination must be non-compensating. Since many problems (including the example) come close to meeting these conditions, this parsimonious analysis can shed light on the design problem, without yielding "exact" answers, while requiring little computation.

The M_oI took only 18 function evaluations to conclude that the maximum achievable $\bar{\mu}$ is approximately 0.30, and that a minimum value of $\bar{\mu} = 0.24$ is achieved with the following ranges:

	Ranges
d_1	$[16.75, 17.87]$
d_2	$[11.64, 12.58]$
d_3	$[14.71, 16.44]$
d_4	$[20.00, 20.00]$
d_5	$[17.10, 18.10]$
d_6	$[19.84, 20.00]$
d_7	$[12.84, 13.81]$
d_8	$[15.08, 15.81]$
d_9	$[17.84, 18.81]$
d_{10}	$[17.84, 18.81]$
d_{11}	$[17.84, 18.81]$
d_{12}	$[17.84, 18.81]$
d_{13}	$[14.71, 16.44]$
d_{14}	$[14.71, 16.44]$
d_{15}	$[147.09, 164.43]$
d_{16}	$[128.36, 138.14]$

These ranges are in almost exact agreement with the results of the pattern search; only thickness d_{14} was different. There were more differences between the M_oI's ranges and the set generated by the genetic algorithm, with differences in d_5, d_6, d_8, d_9, d_{12}, d_{14}, d_{15}, and d_{16}. These differences are not as great as they might appear, for when a sensitivity analysis was conducted as part of the approximations described below, these variables were among the least significant.

It must be emphasized how small a number of function evaluations 18 is. Even the linearization scheme described in the next section required 129 evaluations to achieve its

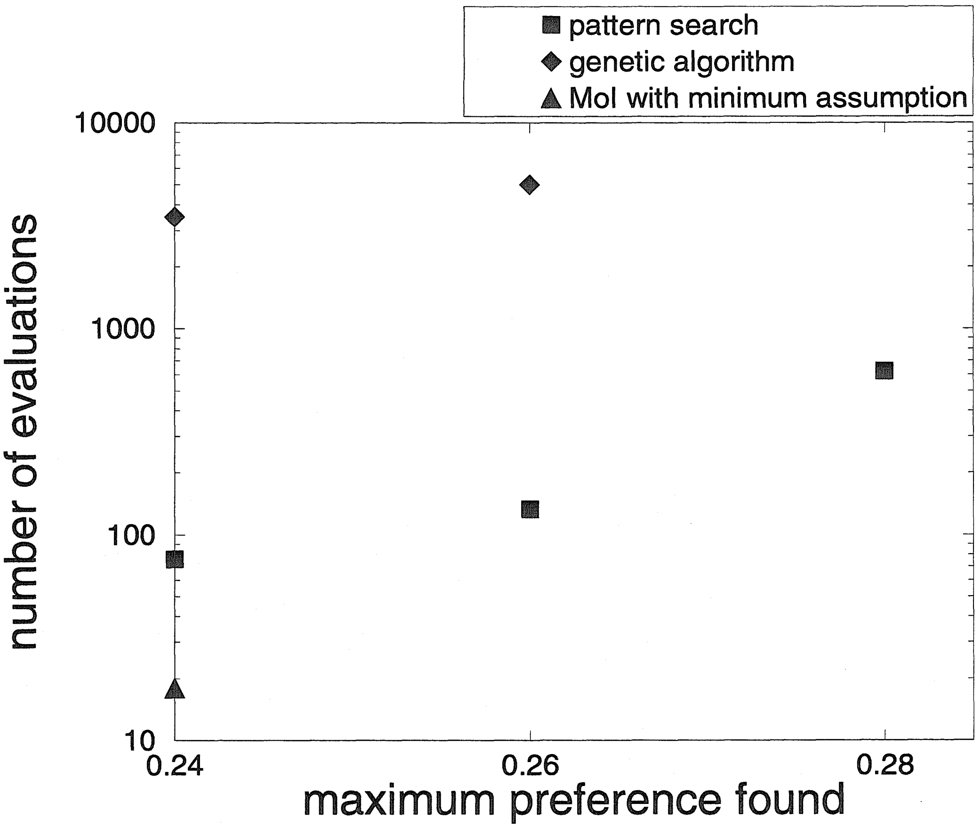

Figure 7.4: Comparison of three optimization schemes

approximate results, and the pattern search and genetic algorithm were significantly more expensive. The M_oI is thus the only method discussed here that would be realistic to apply in problems on the order of 100 variables with non-trivial analysis functions. As with any of the methods discussed here, results will not be exact, but they can provide useful guidance to designers.

A comparison of the number of function evaluations required by each method to find designs with different levels of overall preference is shown in Figure 7.4. The set-based approach is the fastest, but does not achieve results of the higher preferences. Pattern search takes fewer function evaluations than the genetic algorithm to arrive at the same levels, and reaches higher levels of overall preference as well. Figure 7.5 shows the ranges of design variables provided by each scheme. The M_oI with minimum assumption, represented by solid lines, returns a set of designs that exceed $\bar{\mu} = 0.24$. The genetic algorithm is shown with dashed lines, and those sets of designs exceed $\bar{\mu} = 0.26$. Finally, the single black square on each plot is the maximum at $\bar{\mu} = 0.28$ found by the pattern search, which is not necessarily the global maximum. The disparities between the genetic algorithm and the other methods may indicate that the genetic algorithm is missing some optimal regions in the design space, or that the other two methods have converged prematurely.

7.5.4 Approximations

All of the optimization techniques discussed above executed the finite element model directly. The implementation of the M_oI includes the possibility of using a method based on Design of Experiments (DOE) [10] to find, where possible, a linear approximation to the function f. The function f, rather than the overall preference $\bar{\mu}$, is chosen for approximation because it is orders-of-magnitude more expensive to compute, but more important, because it is unlikely to change, while the preference structure may be modified. Changes in the preference structure are inexpensive, so long as f does not need to be recalculated. This is a good reason not to optimize directly on overall preference, but to use optimization routines on pieces of the problem.

A resolution IV DOE approximation required 129 evaluations to construct a linear approximation in 9 of the 16 variables; the other seven variables were deemed non-linear. This can reduce the problem to one with a seven-dimensional *DVS*, and thus significantly reduce the computation cost for any of the methods discussed above.

7.6 Summary

This chapter summarizes an investigation into the possible application of various optimization schemes to the scalar optimization problem formalized by the M_oI. An example problem was constructed using a finite element model with 16 design variables as the analysis function, and classical, evolutionary, and M_oI-specific optimization schemes were considered. The results of this chapter are preliminary: the methods tested were not optimized for performance. The results can only be preliminary in another way as well: the negotiation

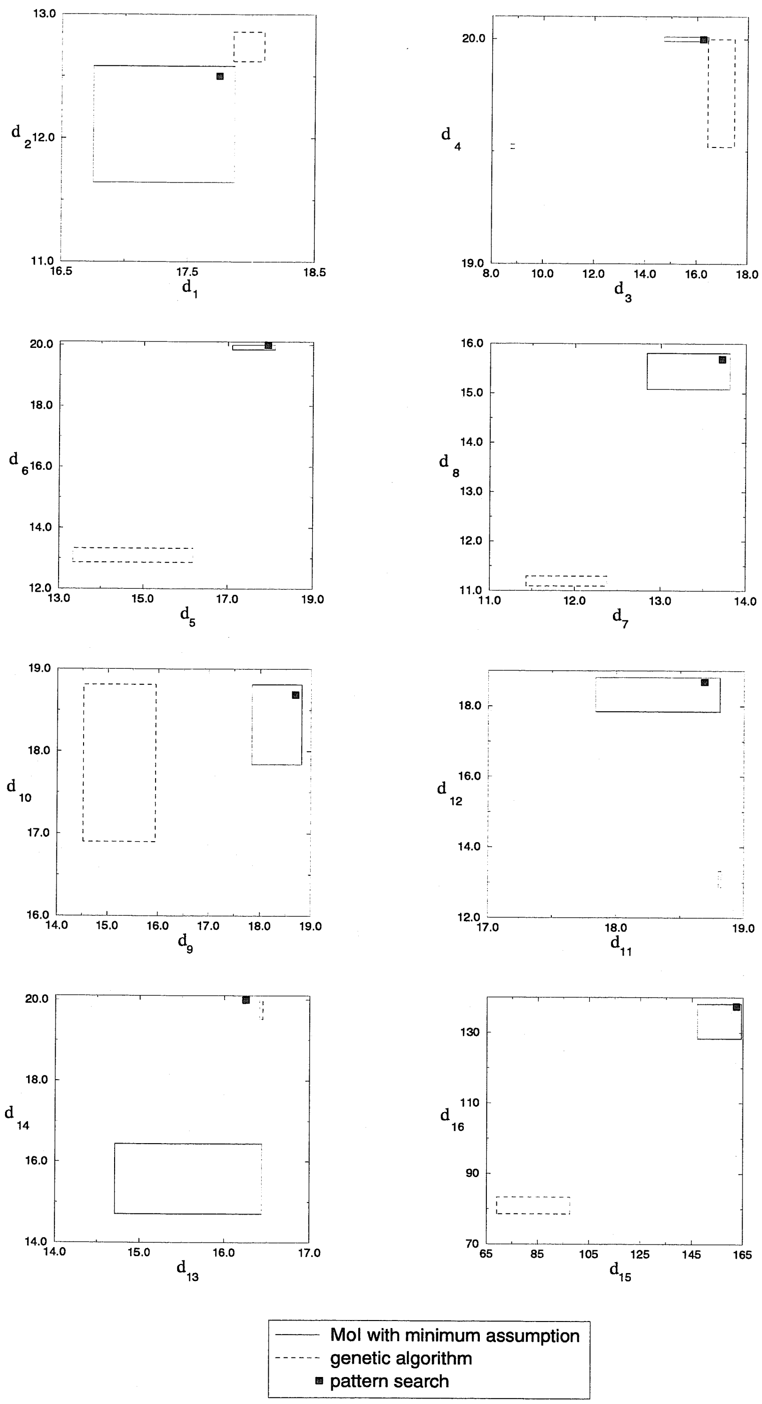

Figure 7.5: Ranges of design variables, three optimization schemes

model that forms the bulk of the thesis is a tool for decision makers, not a computational method that turns the power of decision over to an algorithm.

Design trade-offs with large numbers of variables are a compute-intensive problem, but they are often well-behaved. The test problem used in this chapter fits that description, but was not meant to be archetypal of all design problems. Thus the conclusions that can be drawn from this particular set of tests are preliminary at best. Nevertheless, it was found that blind application of pattern search made quick early progress, but then converged slowly if at all. Similarly, blind application of a genetic algorithm was slow, but did make progress. Exhaustive search can be shown *a priori* to take a prohibitively long time. Approximations were found to be useful, particularly if one pays attention to the structure of the problem. The approximation methods discussed in this chapter are covered in greater detail in William Law's doctoral thesis [46]. Linearization can work well for some design problems, either because they are quite well-behaved, or because the answers available in preliminary design are only useful to a linear approximation. Methods that return sets or ranges in general provide more information to designers.

Chapter 8

Example: Passenger Vehicle Structure Design

In the summer of 1997, an application of the M_oI to preliminary vehicle structure design was demonstrated for Volkswagen Wolfsburg. The application serves as a tutorial introduction to the M_oI and its underlying concepts.

8.1 Preliminary vehicle structure design

The general vehicle structure design problem is the engineering of a *body-in-white*, which consists of the (usually metal) frame to which components and exterior panels are fastened. While there are interesting alternative solutions such as space frames and monocoques, this chapter is concerned with the welded sheet-metal structure typical of passenger automobiles of the present day (see Figure 8.4). The vehicle structure engineers must design a body-in-white that meets certain measurable engineering targets such as stiffnesses, stress levels under load, and weight. In addition, they must satisfy many performance targets associated with less easily measured concepts such as style, manufacturability, and space and mounting requirements of other engineering groups involved in the design process. These unmeasured performances are handled informally, often by negotiation between groups working on the same vehicle. The M_oI was developed to allow for a formal approach to the incorporation of this imprecise information.

In order to avoid any difficulties involving confidential information, it was decided that an older model vehicle would provide an effective demonstration of the method. To this end a 1980 VW Rabbit (see Figure 8.1) was acquired. The vehicle was stripped to the

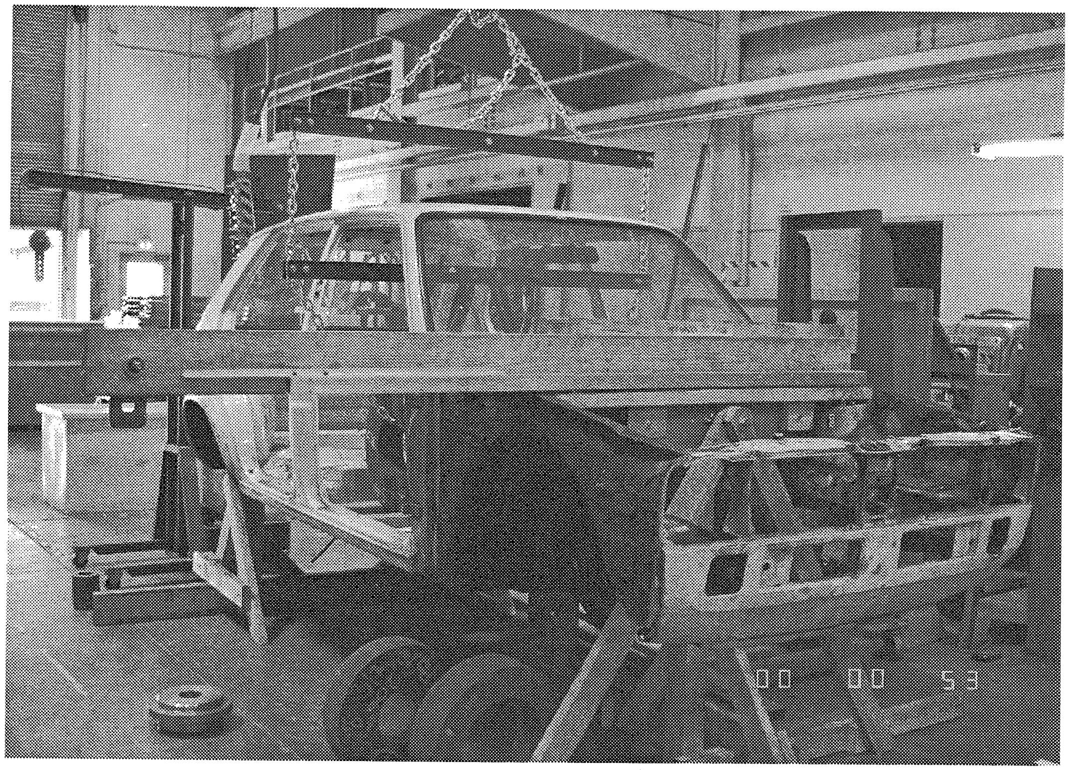

Figure 8.1: 1980 VW Rabbit in stiffness testing

structural body-in-white, and torsional and bending stiffnesses were measured. The intact body-in-white was found to have a torsional stiffness of approximately 4900 N-m/degree and a bending stiffness of approximately 2500 N/mm. Tables of data from some of the load tests are shown in Figures 8.2 and 8.3. In addition, geometric data were gathered and used to create a solid model (Figure 8.4). The solid model and the structural stiffness information together were used to create a finite element model (Figure 8.5).

Load (N)	Deflection (mm)
0.00	0.00000
284.67	0.101060
551.55	0.17780
836.22	0.33020
2001.60	0.78740
2286.27	0.91440
2837.82	1.16840

Fitting to $y = mx + c$:

$$m = 2426.16 \text{ N/mm}$$
$$c = 51.79 \text{ N}$$

Fitting to $y = mx + 0$:

$$m = 2484.80 \text{ N/mm}$$

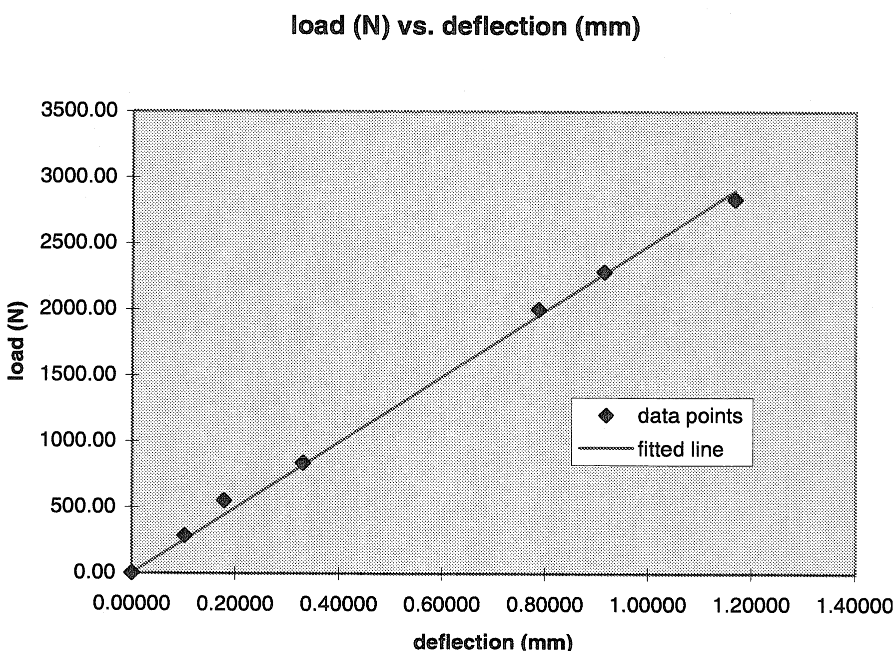

Figure 8.2: Load test, bending stiffness

Load (N)	Moment (N-m)	Deflection (mm)	Twist (deg)
0.00	0.00	0.00	0.00000
126.99	212.84	1.09	0.04407
275.78	455.03	2.41	0.09736
404.77	667.87	3.48	0.14041
551.55	910.06	4.47	0.18038
680.54	1122.90	5.72	0.23060
845.12	1394.45	6.99	0.28184
974.11	1607.28	8.08	0.32591

Fitting to $y = mx + c$:

$$m = 4960.74 \text{ N-m/deg}$$
$$c = -10.16 \text{ N-m}$$

Fitting to $y = mx + 0$:

$$m = 4917.04 \text{ N-m/deg}$$

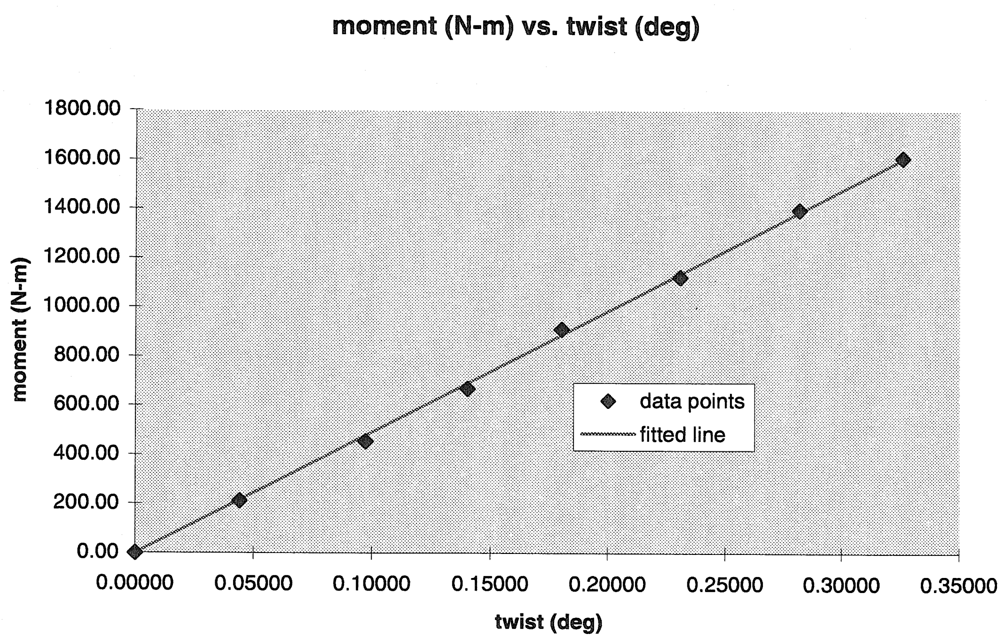

Figure 8.3: Load test, torsional stiffness

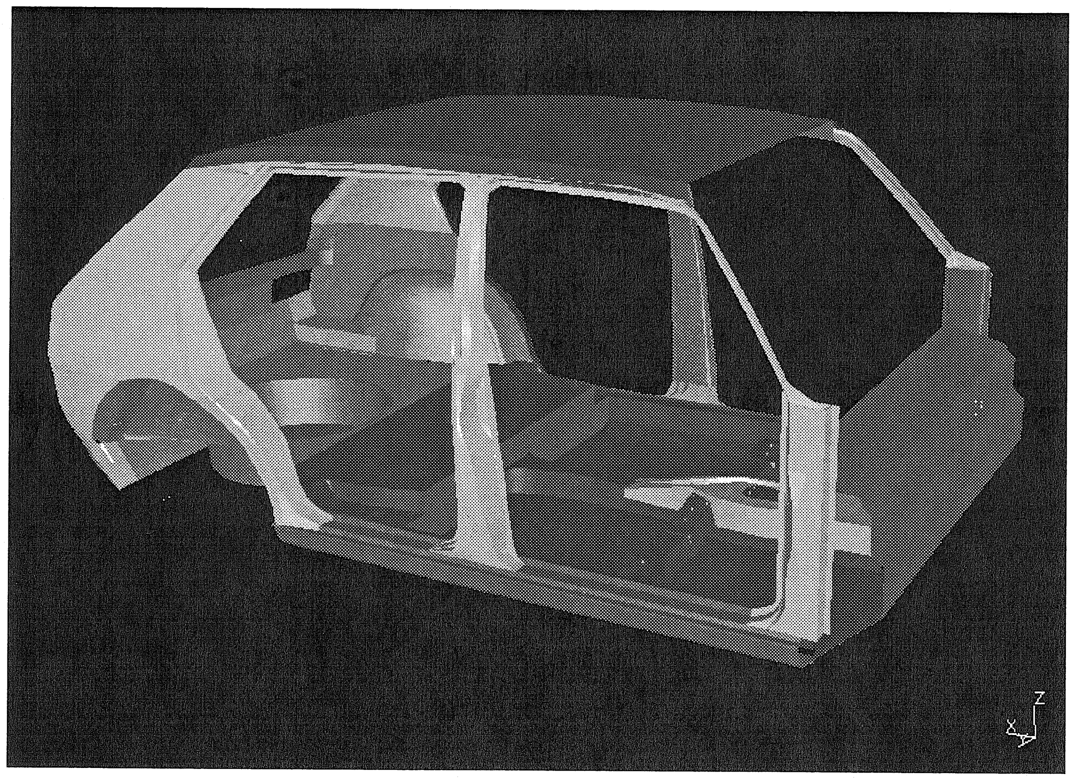

Figure 8.4: Geometric model of body-in-white in SDRC I-DEAS

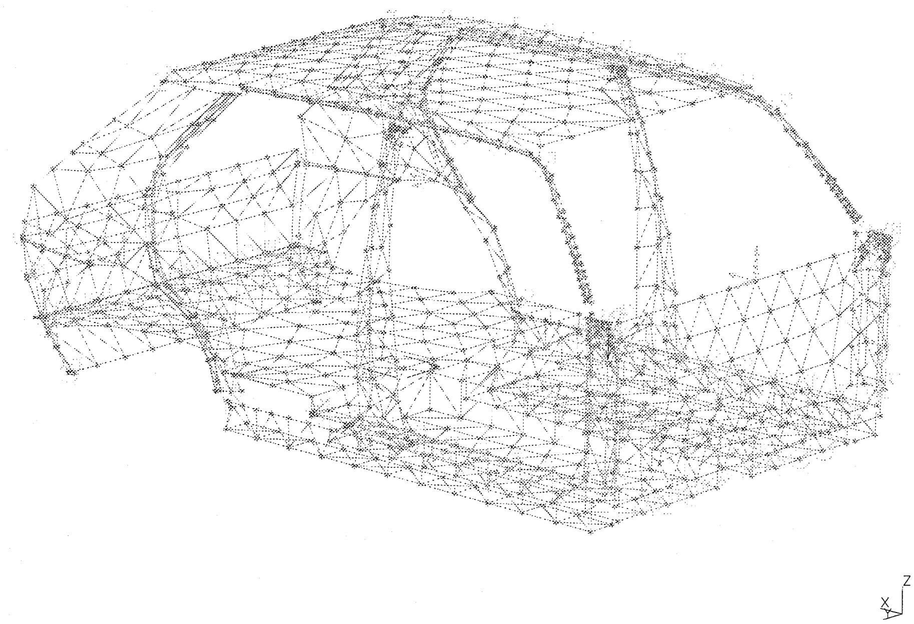

Figure 8.5: Finite element model of body-in-white

The finite element model was parameterized with five[1] design variables:

d_1: A-pillar thickness (mm)

d_2: B-pillar thickness (mm)

d_3: floor sill thickness (mm)

d_4: floor pan thickness (mm)

d_5: B-pillar location (mm aft of a nominal point chosen by stylists)

and the performance was assessed with three measures:

p_1: Bending stiffness (N/mm)

p_2: Torsional stiffness (N-m/deg)

p_3: Weight (kg)

The stated design problem was to achieve 10% improvements over the reference model in the three measured performances. In addition, it was understood that the design must not be difficult to manufacture, and that this year's model should have a "somewhat longer and sleeker look."

8.2 Applying the M_0I to include imprecise information

While standard optimization methods could be used to attempt to determine the highest achievable bending stiffness, the highest achievable torsional stiffness, or the lowest achievable weight for this analysis model, such an optimization would not tell the designer which designs are the most promising when other relevant considerations are taken into account. On the one hand, there is a necessary trade-off between the stiffnesses and the weight; it is impossible to optimize both simultaneously. Additionally, there is other (imprecise) information to consider when making the decision, such as manufacturing and styling concerns.

[1]The demonstration here was conducted using a subset of the design variables; the method can be applied directly to a larger set of variables.

The application of the M$_o$I to this problem involves constructing a different "optimization" problem that includes the imprecise information that would be left to the negotiation stage in traditional design.

8.2.1 Performance preferences

The calculated performance requirements on bending stiffness, torsional stiffness, and weight were originally expressed as targets of 10% improvements over the reference model. As was discussed above, this is unrealistically, and indeed unproductively, precise. In place of these hard targets, imprecise performance requirements were specified with a linear interpolation between two points. In the implementation of the M$_o$I, it is common to name the customer as the source of the performance preferences; in fact, it is more likely to be a manager, perhaps informed by market research, serving as the customer's proxy. To specify these imprecise requirements, the manager must answer two simple questions: "What is the lowest performance you can live with (where is $\mu = 0$)? What performance would satisfy you completely (where is $\mu = 1$)?" These bounds are clearly dependent on a number of factors, including the target market and the performance of competitors' products; discussions with employees of Ford Motor Company and VW indicated that engineering managers can answer these two questions with little more effort than is needed to settle on the initial crisp target. Figure 8.6 shows the imprecise requirements on stiffnesses and weight.

8.2.2 Design preferences

To include requirements on manufacturing, availability, style, and other things which are not calculated in the finite element analysis, designer preferences are specified on the design variables. As with the imprecise performance requirements, they range from $\mu = 0$ at the unacceptable limit to $\mu = 1$ at the most preferred. A preference is defined on each of the five design variables, as shown in Figure 8.7 and listed in Table 8.1 (as shown in Figure 8.7, preferences are interpolated between specified values in considering d_2 and d_5). Each preference is representative of imprecise information that can be incorporated using the M$_o$I:

1. The sheet steel for stamping the A-pillar is only available in certain increments, so

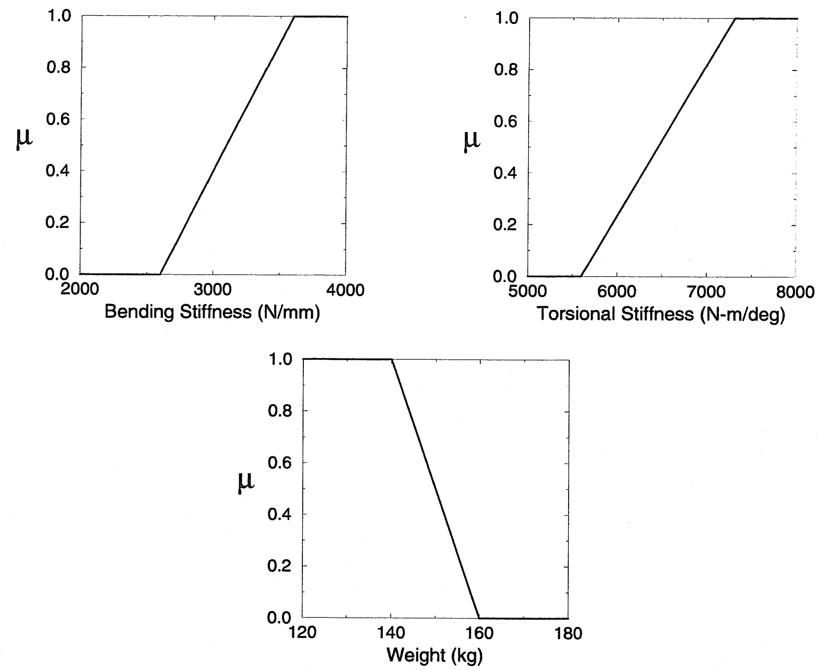

Figure 8.6: Imprecise performance requirements

this plot is discrete rather than continuous. The manufacturing engineer has a higher preference for thinner sheets, since they are easier to form; this is a design preference for manufacturability.

2. The B-pillar thickness is continuous and more complicated than the linear performance preferences. This preference does *not* indicate that the physical B-pillar might be 1.113 or 1.114 mm thick; rather it means that the designer knows that the finite element model is simplified, and that a high number for B-pillar thickness means that more reinforcing features will need to be added to the B-pillar. The designer would like to keep the B-pillar as simple as possible.

3. The floor sill thickness preference is an example of a sourcing, or availability preference; it states that some thicknesses are more easily obtained than others.

4. The floor pan thickness is preferred thicker by the designer for ease of attachments and for durability.

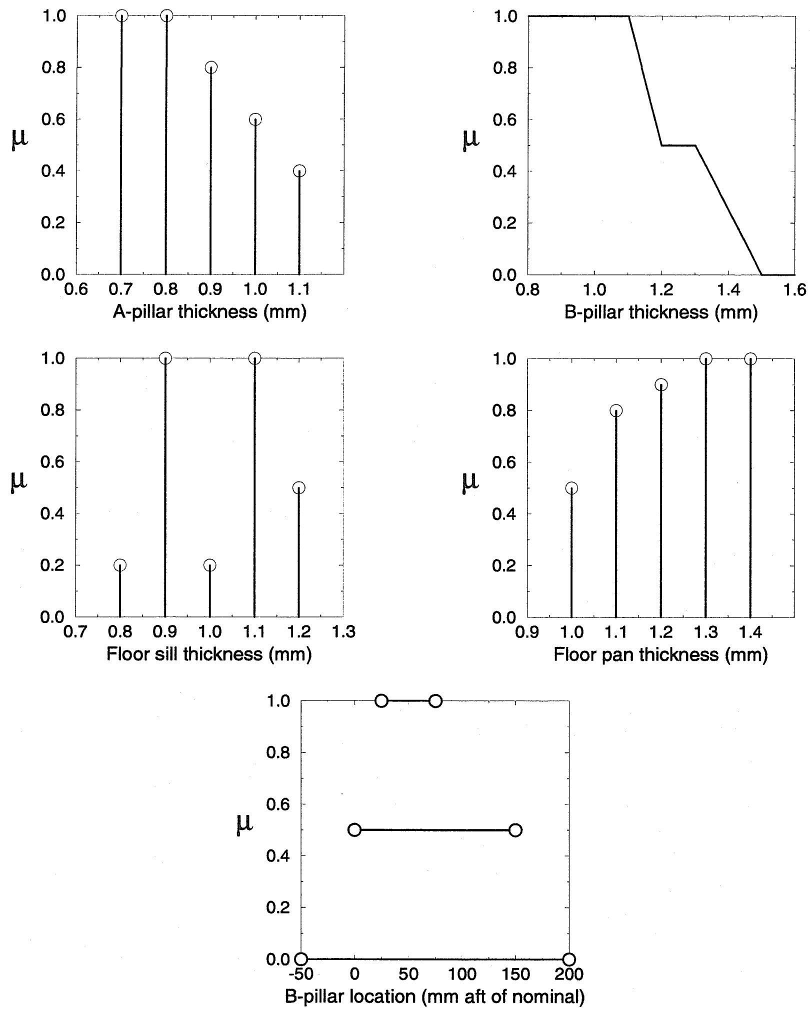

Figure 8.7: Designer preferences

5. The design preference for B-pillar location comes from the stylists, and captures the directive for a longer, sleeker look for this year's model. It has been specified differently from the other design preferences, using α-*cuts* [4], so that the stylists have given a range of perfectly acceptable values, a range of barely acceptable values, and a range of values that fall in the middle. This method of specifying preferences can have computational advantages.

d_1	0.7	0.8	0.9	1.0	1.1
μ_D	1.0	1.0	0.8	0.6	0.4
d_2	1.1	1.2	1.3	1.4	1.5
μ_D	1.0	0.5	0.5	0.25	0
d_3	0.8	0.9	1.0	1.1	1.2
μ_D	0.2	1.0	0.2	1.0	0.5
d_4	1.0	1.1	1.2	1.3	1.4
μ_D	0.5	0.8	0.9	1.0	1.0
d_5	-50	0	25–75	150	200
μ_D	0	0.5	1.0	0.5	0

Table 8.1: Designer preferences

8.2.3 Weights and strategies

In addition to these preferences, the relative importance, or weight, of each attribute must be determined, and the way in which attributes trade-off against each other must be specified. It was supposed that the interactions among preferences would be under the ultimate control of a manager. The hierarchy of preferences employed is shown in Figure 8.8. The procedure for managed negotiation described in Section 5.7.2 was applied:

1. In comparing bending stiffness p_1 and torsional stiffness p_2, it is determined that increasing either one to the best possible value would be of little benefit if the other stayed at $\mu = 0.5$:

$$\mathcal{P}_s(0.5, 1) \approx \mathcal{P}_s(0.5, 1) \approx \mathcal{P}_s(0.5, 0.5)$$

   The conclusion is that for (p_1, p_2), weights are equal and the strategy is $s = -\infty$ (*min*).

2. Next, the overall stiffness is compared to the structure weight. Here, indifference

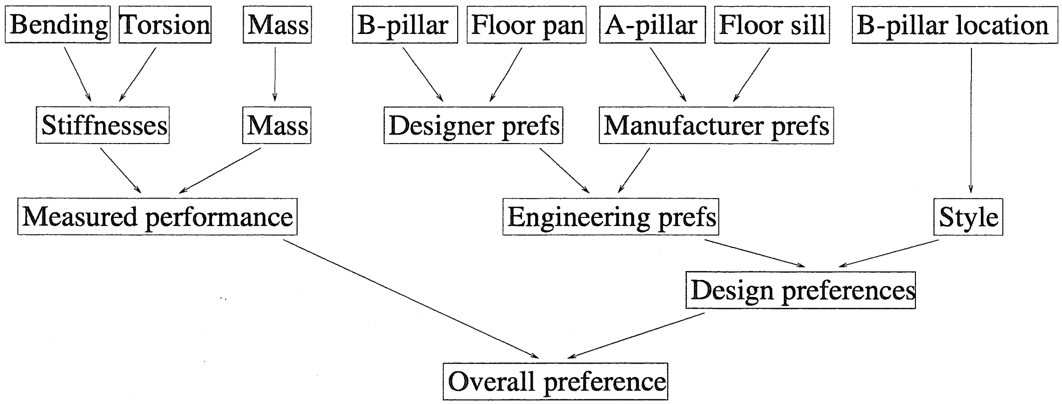

Figure 8.8: Hierarchy of preferences for VW example

points for $\mathcal{P}_s\left((p_1, p_2), p_3\right)$ are:

$$\mathcal{P}_s(0.3, 1) = \mathcal{P}_s(1, 0.2) = \mathcal{P}_s(0.5, 0.5)$$

and the calculations give $s = -0.02, \omega = 0.7$, which, rounding to one decimal place, is equivalent to the geometric mean ($s = 0$). Note that $\omega = \frac{\omega_2}{\omega_1}$ is the ratio of weights.

3. The five design variables are likewise aggregated hierarchically. First, the two that the designer specified directly, B-pillar thickness (d_2) and floor pan thickness (d_4):

$$\mathcal{P}_s(0.4, 1) = \mathcal{P}_s(1, 0.3) = \mathcal{P}_s(0.5, 0.5)$$

This indicates that the designer would be willing to see preference for floor pan thickness decrease to 0.3 in order to achieve a perfectly acceptable value for B-pillar thickness, while a perfectly acceptable value of floor pan thickness merits a decrease in B-pillar thickness only to 0.4. This translates to a strategy of $s = -1.4$, between the *min* and the geometric mean, and a weight ratio of $\omega = 0.6$.

4. Next, the two manufacturing preferences, A-pillar thickness (d_1) and floor sill thickness (d_3), are compared, and the sourcing preference on d_3 is seen to be relatively insignificant:

$$\mathcal{P}_s(0.4, 1) = \mathcal{P}_s(1, 0.1) = \mathcal{P}_s(0.5, 0.5)$$

Here, $s = -0.2$, and $\omega = 0.3$.

5. As the design preferences are combined hierarchically, the next step is to arrive at the appropriate aggregation of the two pairs of design preferences just considered: How does the combination of B-pillar thickness and floor pan thickness compare to the combination of the two manufacturing preferences A-pillar thickness and floor sill thickness? If the manager is comfortable discussing preference levels, this calculation is relatively simple:

$$\mathcal{P}_s(0.25, 1) = \mathcal{P}_s(1, 0.25) = \mathcal{P}_s(0.5, 0.5)$$

This is the case of equal weights, and compensation with the geometric mean ($s = 0$). The same result could be achieved by specifying indifference points as combinations of all four design variables.

6. The styling preference for B-pillar location (d_5) is next traded-off against the aggregate of the four other design variables. Convenience of design and manufacture are seen to be relatively less important than styling, and the two do compensate to some degree:

$$\mathcal{P}_s(0.4, 1) = \mathcal{P}_s(1, 0.3) = \mathcal{P}_s(0.5, 0.5)$$

Once again, this yields $s = -1.4$, $\omega = 0.6$.

7. The last aggregation to be considered is one between the entire set of uncalculated design preferences on the one hand, and the set of calculated performances on the other. This aggregation is also strongly compensating (close to $s = 0$), and the performances are weighted slightly more than the design considerations:

$$\mathcal{P}_s(0.2, 1) = \mathcal{P}_s(1, 0.3) = \mathcal{P}_s(0.5, 0.5)$$

This gives a compensation parameter of $s = -0.02$, and a weight ratio of $\omega = 1.3$.

Through the application of the M_oI, the design problem has been reformulated to be the maximization of the overall preference:

$$\bar{\mu}(\mathbf{d}) = \mathcal{P}_0\left(\mu_{\mathbf{D}}, \mu_{\mathbf{P}}; 1, 1.3\right)$$

where

$$\begin{aligned} \mu_{\mathbf{D}} &= \mathcal{P}_{-1.4}\Big(\mathcal{P}_0\Big(\mathcal{P}_{-1.4}\left(d_2, d_4; 1, 0.6\right), \mathcal{P}_{-0.2}\left(d_1, d_3; 1, 0.3\right); 1, 1\Big), d_5; 0.6, 1\Big) \\ \mu_{\mathbf{P}} &= \mathcal{P}_0\Big(\mathcal{P}_{-\infty}\left(p_1, p_2; 1, 1\right), p_3; 1, 0.7\Big) \end{aligned}$$

To compare the weights of the eight individual attributes, the lowest weight (floor sill thickness) is normalized to one. The normalized weights are:

ω_{d_1}	3.3
ω_{d_2}	2
ω_{d_3}	1
ω_{d_4}	3.3
ω_{d_5}	16.1
ω_{p_1}	9.9
ω_{p_2}	9.9
ω_{p_3}	13.8

The computation of $\bar{\mu}(\mathbf{d})$ for a single design point $\mathbf{d}$ is limited by the finite element stiffness calculation, which takes about a minute on a Sun Ultra1-170MHz workstation; the calculations of weight and preference aggregation are of negligible cost. Even in this relatively modest problem, where there are only five design dimensions, an exhaustive calculation of preferences over the design space is prohibitively expensive: to capture even 10 points on each dimension would require approximately 70 days. The M_oI exploits the structure of the problem to speed the search for preferred solutions.

8.3 Results

The design problem, including all imprecise information, was solved in two different ways. First, in order to demonstrate the method, the finite element analysis was run 3125 times to provide a coarse but complete check of the entire design space. The point of peak overall preference of $\bar{\mu} = 0.440$ was found by this approach to be at $\mathbf{d} = (1.0, 0.9, 0.9, 1.0, 50)$, where the design preferences μ_D are $(0.6, 1.0, 1.0, 0.5, 1.0)$; the stiffnesses and weight at this point were $\mathbf{p} = (2832, 5836, 147)$, with preferences $(0.23, 0.14, 0.62)$. The maximum achievable stiffnesses are 3365 N/mm ($\mu_\mathrm{P} = 0.77$) in bending and 6029 N-m/degree ($\mu_\mathrm{P} = 0.25$) in torsion, but the corresponding weight of 170 kg is unacceptable. Similarly, a weight of 144 kg ($\mu_\mathrm{P} = 0.78$) is achievable, but stiffnesses drop to 2803 N/mm ($\mu_\mathrm{P} = 0.20$) and 5730 N-m/degree ($\mu_\mathrm{P} = 0.08$). The combined overall preference $\bar{\mu}$ also takes into account the design preferences on style, manufacturability, and the like.

The preference numbers, as well as the ordering of alternatives, vary if the strategies and weights are varied. If it is assumed that all preferences are combined with the *min* operator, as with the comparison in Section 5.6, then the maximum achieved preference is 0.2 and occurs at $\mathbf{d} = (1.1, 0.9, 0.8, 1.2, 150)$. If all design preferences are first combined with a *min*, and the performance preferences are also combined with a *min*, and the aggregated results are combined with an equally weighted geometric mean, the two highest overall preferences are 0.361 at $\mathbf{d} = (0.9, 1, 0.9, 1.2, 50)$, and 0.361 at $\mathbf{d} = (0.7, 0.9, 0.9, 1.3, 50)$. Finally, if the above calculated weights are applied with the original compensating function, the geometric mean, the maximum preference is achieved by $\mathbf{d} = (0.9, 1, 1.1, 1, 50)$. These three different aggregation models are summarized in Table 8.2. The overall preference values are of course higher for more compensating strategies, and it is pointless to compare them. That the order of alternatives is different, however, is significant. Different negotiations lead to different results.

The power of the method lies not in an ability to find a single overall "best" point, but in the information it contains about how the total combined preference $\bar{\mu}$ varies with each of the design variables. Although it is impossible to display all five dimensions varying at once, a tool was written that uses a commercial package (Matlab) to display results interactively. Using the tool, the designer can see the change in preference that would

Aggregation strategy	d_1	d_2	d_3	d_4	d_5
calculated	1.0	0.9	0.9	1.0	50
all *min*	1.1	0.9	0.8	1.2	150
min and $\mathcal{P}_0$	0.9	1.0	0.9	1.2	50
(two maxima)	0.7	0.9	0.9	1.3	50
all $\mathcal{P}_0$	0.9	1.0	1.1	1.0	50

Table 8.2: Peak performance points, various strategies

occur by varying each design variable independently from a chosen beginning point. Results can be seen on five simultaneous plots in two dimensions (see Figure 8.9), or on a three-dimensional surface plot (see Figure 8.10) with the remaining design variables set to nominal values.

8.3.1 Approximations

Naturally, the exhaustive evaluation of points in the design space would not be performed on a real design problem. It was performed here *only* for comparison purposes. An approach that utilizes Design of Experiments (DOE) [10] to approximate the finite element calculations for bending and torsional stiffnesses reached substantially similar results in only 21 runs (or approximately 20 minutes). The average difference (from the exhaustive evaluation) in bending stiffness was approximately 1%, with a maximum difference of less than 4%, while the average and maximum differences for torsional stiffness were both less than 1%.

In some cases, the nonlinearities of the analysis function f will defeat a linear or even polynomial approximation, but in many cases, such as the example presented here, these simple approximations can drastically reduce the required computation. Since precise answers are not required for preliminary design, it is sensible to exploit approximation tools when possible. If more computation can be justified, a more thorough calculation can always be made.

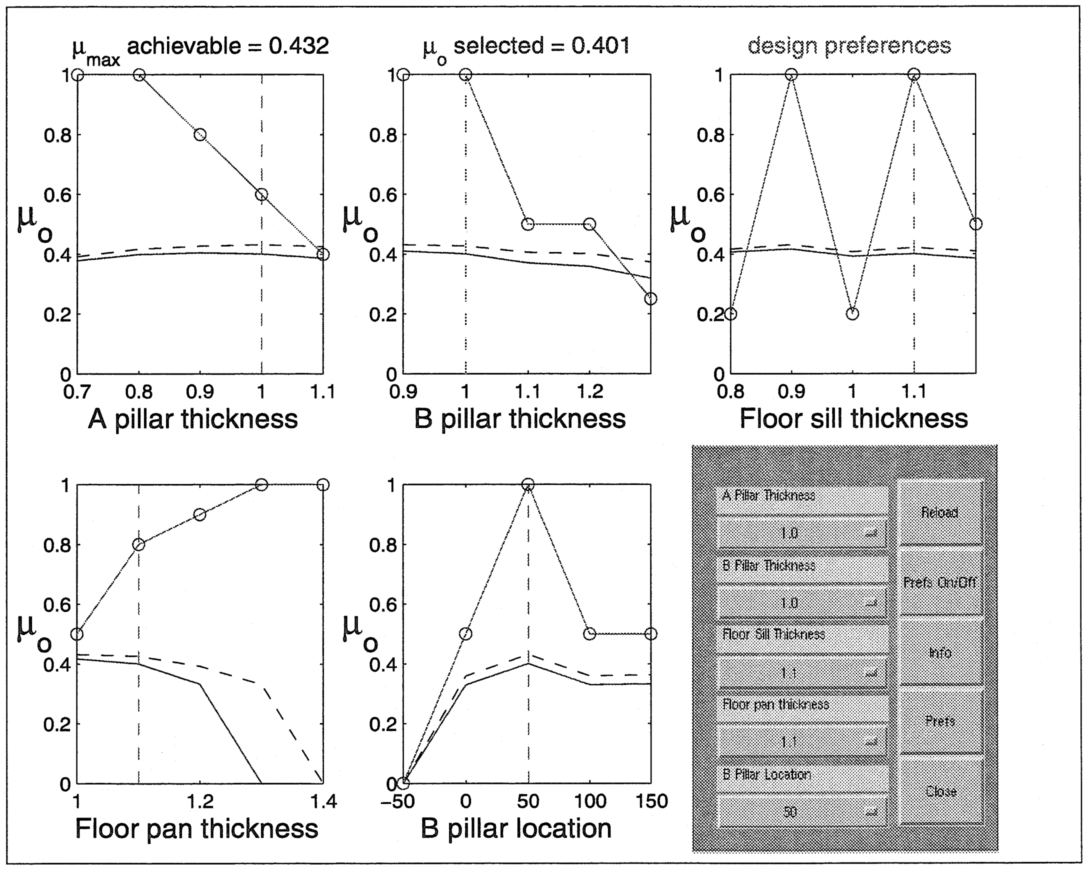

(Note that in the interface, μ_o is used in place of $\bar{\mu}$)

Figure 8.9: Graphical user interface for preference display

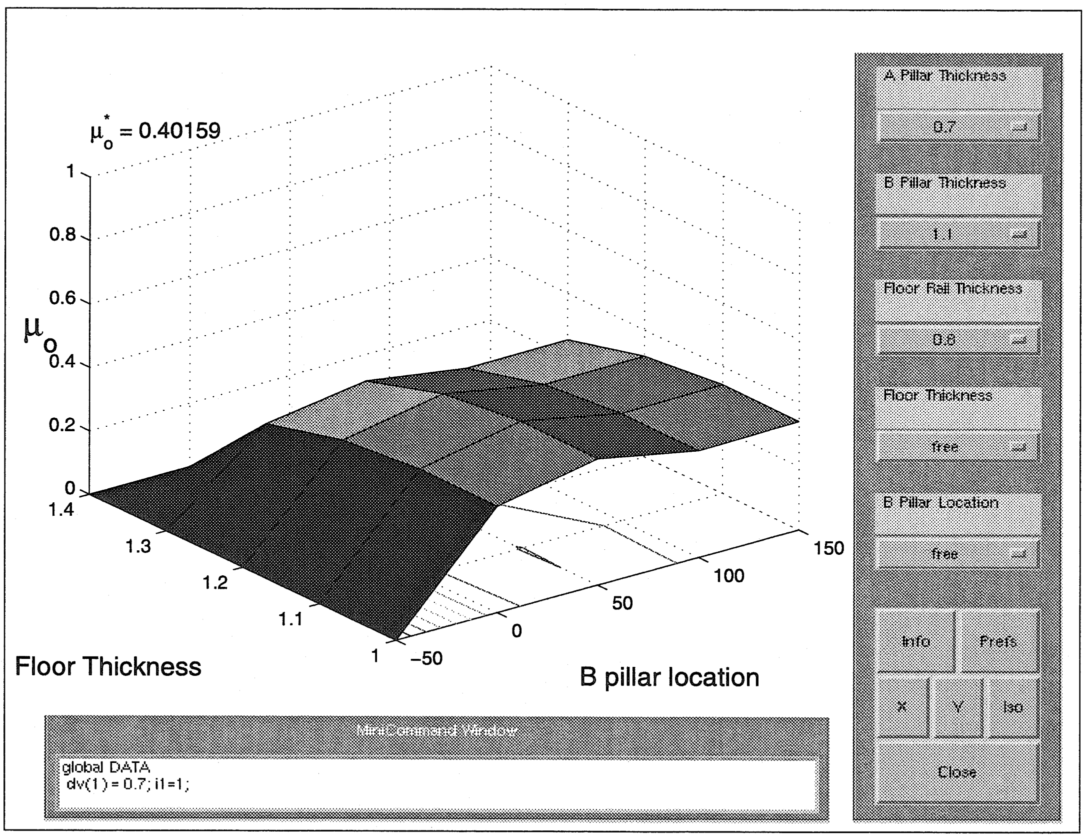

(Note that in the interface, μ_o is used in place of $\bar{\mu}$)

Figure 8.10: 3-D graphical user interface for preference display

8.4 Discussion

Engineering analysis usually requires some judgement on the part of the designer. Unless a full-scale exact prototype is to be built and tested, the accuracy of any calculated performance measure depends on the quality level of the model employed. Even when exact data are available for some attributes, final decisions about a design incorporate other, unmodelled concerns, such as manufacturing and styling.

The M_oI constructs a model of the entire decision process, expressing the calculated overall performance $\bar{\mu}(\mathbf{d})$ as a function of the design variables. It depends on many factors: the function f for calculating measurable performances, the specification of design preferences $\mu_{\mathbf{D}}$ and performance preferences $\mu_{\mathbf{P}}$, the weighting of these preferences, and the specification of trade-off aggregations between attributes. A change in any of these will affect the shape of the function $\bar{\mu}$ in design space, and thus affect the decision. The analysis f is here relatively expensive to compute, and changes in the finite element model are costly to propagate to overall preference. Changes in the other factors, on the other hand, are easily incorporated, as finite element results $f(\mathbf{d})$ are stored so that the same design point need never be analyzed more than once. This allows the M_oI to support an iterative decision process, when the information from the first round of calculations inspires a change in the preference structure.

In the example shown in this chapter, the shape of the function $\bar{\mu}(\mathbf{d})$ is sensitive to changes in the styling preference $\mu_{\mathbf{D}}(d_5)$, which is not surprising, since this preference is accorded a large weight. This and other features of the design problem can be seen in the advanced interface shown in Figure 8.9. The vertical dashed (red) lines indicate the selected values of the design variables. In this figure, $\mathbf{d} = (0.8, 1, 0.9, 1.2, 50)$. The solid (blue) lines indicate how the overall preference $\bar{\mu}$ would change by varying that design variable while holding the other four fixed at their current values. The dashed (black) lines show the maximum achievable $\bar{\mu}$ for each value of each design variable. Finally, the solid (red) lines joining the circles are the specified preferences on design variables. Thus Figure 8.9 shows that the overall preference $\bar{\mu}$ varies qualitatively, though not quantitatively, as four of the five design variables; the exception is d_4, floor pan thickness, where the overall preference $\bar{\mu}$ tends to decrease as the designer preference $\mu_{\mathbf{D}}(d_4)$ increases. The designer can interact with the preference display to examine trends in the structure of the overall preference.

8.5 Summary

This chapter presented an example application that was undertaken on behalf of VW Wolfsburg; it serves as a tutorial and an application of the M_oI and the ideas presented in this thesis. In preliminary vehicle structure design, as in all preliminary engineering design, many important decisions are made informally on the basis of imprecise information. Concerns of styling and manufacturability, for instance, can carry great weight in the design process although they are not modelled by any formal analysis. The M_oI is a tool to formally incorporate such imprecise information into the design process, and thus to make decisions on a sound basis. In this demonstration of the M_oI prepared for VW Wolfsburg, concerns of manufacturing, styling, parts availability, and design were incorporated with the engineering analysis of the structural stiffness of a VW Rabbit. The negotiation model presented earlier was applied directly to eight design and performance preferences. The results show the usefulness of the method in trading-off these conflicting attributes.

In addition to an application to the negotiation model presented in this thesis, this example shows the usefulness of incorporating approximations into the calculations of the M_oI. Also, an interactive graphical tool for preference display was developed and applied here to address the problem of displaying results graphically in several dimensions.

Chapter 9

Conclusions

The seeds of the formal study of negotiation in engineering design discussed in this thesis were planted during joint work with vehicle structure engineers from industry on the inclusion of imprecise information in preliminary vehicle structure design. These engineers noted that their most significant sources of uncertainty were the immeasurable requirements that came down to them from their managers. They had tools, techniques, skills, and intuition enough to approach the given engineering targets of stiffness, weight, acoustic response, and the like. Difficulties arose, however, when they delivered designs that seemed to meet the targets, but would be asked to modify them to accommodate other, often previously unspecified (and in general unquantified), requirements, such as styling concerns. Also, they commented that it was common for targets to be set unrealistically high at first, and that after a first design iteration, managers would adjust the targets based on the preliminary results. Both of these situations were resolved in meetings that the engineers described as "negotiations," where different engineering groups, managers, stylists, and others would attempt to divide limited resources and design freedom. They reported that it was at least as important in these negotiations to seek as many resources and as much design freedom as possible, as it was to attempt to optimize any notion of overall performance. While the negotiations could be either sufficiently open or sufficiently controlled that the results seemed productive, at times politics seemed to play a larger role than engineering productivity.

The task to formalize negotiations was then undertaken in the context of the ongoing research effort to incorporate imprecision in engineering design, which had already acquired a name, the *Method of Imprecision*, or M_oI. The M_oI used the specification of preference to formalize engineering decisions. Throughout the course of the research, more and more engineering design decision problems came to look like negotiations, at least in their mathematical models.

The essence of the model of a negotiated decision is the aggregation of preferences. Indeed, negotiation can be equated with aggregation: each group or individual is responsible for their preferences, but when those individuals or groups meet to negotiate how those preferences shall be combined at the level of the entire design, that is an aggregation problem. Even when there is no recognizable negotiation in the ordinary sense of the word, that cannot be seen in the mathematics of aggregation. An important contribution of this thesis is the recognition of the aggregation problem of negotiation.

The M_oI had previously used two different aggregation methods to combine preferences. However, a complete justification of the legitimacy of preference aggregation had not been previously given. Chapter 4 discussed some objections to any aggregation of preference raised by social choice theory; it was shown there that the particular formalism of preference used here is not subject to those general objections.

Chapter 5 answered the question of which operators are suitable for the aggregation of preference in the engineering design problem. It was seen that there are an infinite number of such operators, and that a parameterized family of operators that spans all rationally allowable levels of aggregation compensation can be constructed. Furthermore, for every aggregation problem there is only one correct aggregation operator. While small deviations from the correct operator will typically lead to the same decision, in many cases different operators lead to different decisions. This implies that results from any decision-making method that relies on a more limited set of aggregation operators are suspect. Perhaps the most significant contribution of the thesis is a mathematically complete description of such preference aggregation, and the development of techniques to apply the method and determine the correct aggregations.

A further contribution of the thesis was the exploration of conditions under which the formal aggregation model is not only a rational expression of the decision, but can be used to support set-based design. Set-based reasoning can have computational advantages, and set-based information may be of much greater use to the designer in the iterative search for better solutions. A set of sufficient conditions based on the notion of *convexity of preferences* was derived.

The thesis includes the presentation of some work that is tangentially related to the core subject of negotiation. A preliminary investigation of various optimization methods for the scalar optimization problem posed by the formal calculation of overall preference was presented in Chapter 7. Appendix A describes a slight diversion in the research to explore the extension of some aggregation methods to the realm of aeroshell analysis; it is included here to give a complete picture of the research avenues that were explored in the course of this project.

I was fortunate to be asked near the end of this research to demonstrate some of the imprecision techniques of the M_oI on an example problem for research engineers at Volkswagen Wolfsburg. Chapter 8 described the example problem, methods, and results. It shows the basic application of the M_oI and the negotiation model, and has some nice pictures as well.

The work presented in this thesis is foundational. The results will be used to develop and implement tools that can help engineers working on large projects to negotiate design solutions in a rational manner. Using the modelling techniques described here, decisions and their rationales can be scrutinized and recorded. If a dispassionate representation of a decision can uncover political motivations, it may lead to more sensible decisions. A faithful record of the negotiated design decisions made at an early stage of the design process may help reduce time spent repeating those negotiations at later stages, and may help guarantee consistency of decisions across the entire process. By providing engineers with a formalism to include uncalculated requirements into their engineering models and thus anticipate problems that commonly arise late in the design process, the need for large changes to incorporate those requirements may be reduced. The applications presented here strongly suggest that the exploitation of these ideas will lead to better designs.

Appendix A

Aeroshell Design and Analysis

The late 1980's and early 1990's have seen a change in the direction of space exploration conducted by NASA, with an emphasis on smaller, lighter spacecraft and missions with budgets and time frames an order of magnitude (or two) less than the missions of the 1960's and 1970's, allowing for less expensive, more frequent launches (see Figure A.1). This shift in emphasis highlights the need for design tools to support preliminary design, since mission designers no longer have the luxury of long times or large budgets for prototyping, testing, and redesign.

The need for simulation-based analysis tools is particularly clear in the design of spacecraft, as operating environments (*e.g.*, microgravity) are often difficult or impossible to reproduce to test prototypes. This appendix examines some of the issues of the management of information in preliminary design using simulation-based analysis and including uncertainty and data of varying reliability.

One of the unifying themes of research on the M_oI has been the aggregation of preferences. The research presented here continues this earlier work by considering preliminary design when the design space is not well modelled by a single (even a single expensive) design tool, but when significant irregularities or discontinuities in the mapping between points in the design space and their corresponding performances call for the application of different analysis tools in different regimes. The formal mathematics developed for preference and for aggregation are extended to the problem of combining output from several analyses for different (often overlapping) regimes, and the resolution of conflicts where

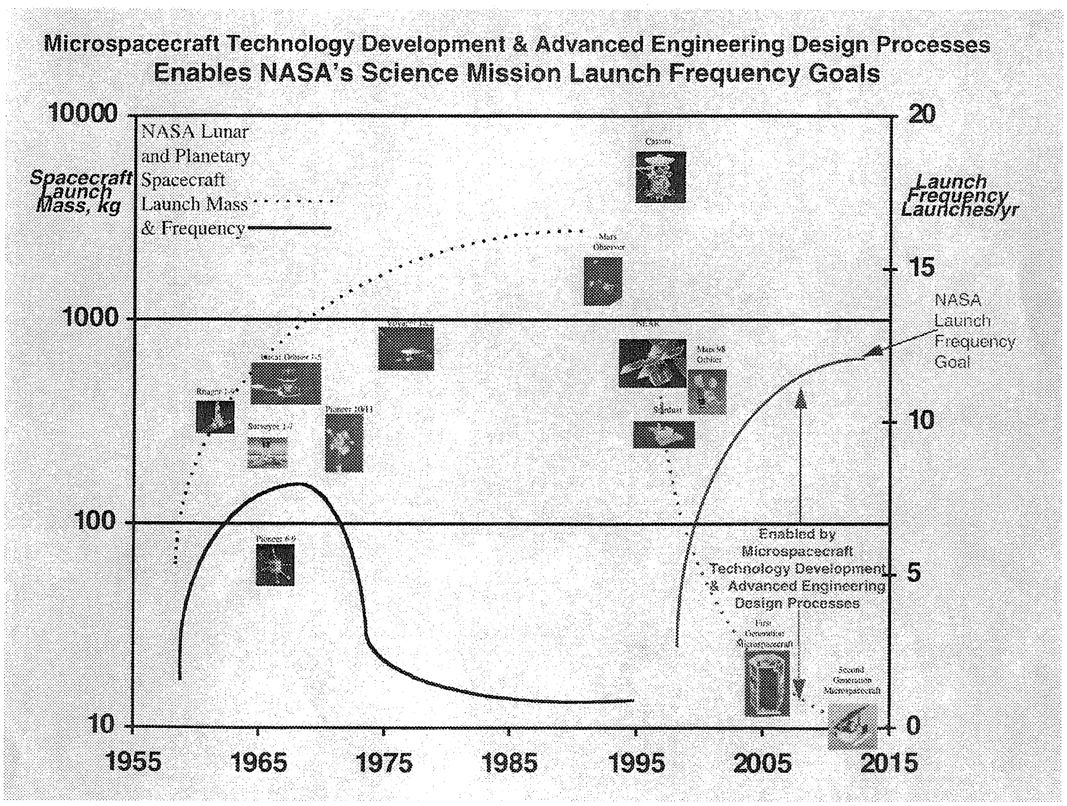

Figure A.1: Launch frequency and payload size of NASA space missions, 1955–2015 (Used by permission, John Peterson, JPL, May 15, 1997)

data from multiple sources disagree.

A.1 Aeroshell design

The example presented here is the design of a reentry aeroshell that is to be released from a spacecraft as it enters the Martian atmosphere. Two devices of this type are expected to be launched with the DS-2 Mars probe in January of 1999. After descent to the Martian surface, the aeroshell's payload, a penetrator with some instrumentation, will puncture the Martian surface so that soil experiments can be conducted and the results transmitted by radio. We shall not treat the soil penetration problem other than considering the orientation of the aeroshell at impact; the general problem here is the design of an aeroshell that will reach the surface near a desired velocity and angle of attack.

There are several sources of significant uncertainty in the operating conditions of the aeroshell during its descent:

- The shell is released from a tumbling spacecraft, so the angle of attack upon entering the atmosphere is unknown.
- The shell flies through the Martian atmosphere, the properties of which are not well characterized.
- The shell encounters unknown winds.
- The shell hits the ground, which is at an uncertain orientation and has uncertain soil properties.

These uncertainties are interesting, difficult and worthy of study, yet there are other issues of uncertainty in the analysis of the problem that take precedence. An automated or semi-automated analysis would allow for a computer search of the design space to (at least) guide preliminary design. Several levels of simulation are available, and their reliability increases with the computation time involved. If, for example, the assumption that the aeroshell reaches its terminal velocity could be made, computation time would be milliseconds, but the assumption is not correct in general. A single run of a full CFD model takes on the order of a day to set up and run on a supercomputer, which is far too computationally intensive for the project time-frames and costs envisioned, except for confirmation of a selected design. A compromise analysis program is a numerical integration of forces over the flight path, with aerodynamic coefficients determined at each time step as functions of atmospheric conditions and the attitude, velocity, and geometry of the aeroshell. The computation time required makes classical optimization, genetic algorithms, and simulated annealing procedures realistic. However, the integration routine is not simply an accurate black box: to successfully integrate over the flight path through the Martian atmosphere requires considerable engineering judgement in the calculation of the aerodynamic coefficients used at each time step of the integration. Furthermore, the output from the integration program gives no indication of how accurately the coefficients were determined.

The problem encountered here by the aeroshell designer is a common one in design analysis, that of how to guarantee good results when the problem may cover one or more of several "analysis regimes." These regimes may be inherent in the physics of the problem, as in the transition between transsonic and supersonic flow, or they may be determined by the availability of information, as in the case when experimental results are available for some (but not all) points in a design space.

A.2 Problem scope

The problem of aeroshell design involves a number of fields (*e.g.*, aerodynamics, thermodynamics, material science, structural mechanics). The aerodynamic analysis, even if considered apart from all other fields, is greatly complicated by the need to treat multiple flow regimes (hypersonic, supersonic, transsonic, subsonic, Newtonian, detached shock, free-molecular), and even if the analysis can be made tractable, the aerodynamic design problem has such a huge set of potential solutions so as to make a search for a globally "optimal" solution to the problem impractical. The present approach to aeroshell analysis is to construct an aerodynamic database for a single candidate design; the analysis is thus useful to validate a design that has already been selected, but is not seen as a tool to explore the design space (see, for example, [51]).

A long-term goal in the field of aeroshell design would be an analysis program that addressed all possible candidate configurations in all flow regimes. In order to make the problem more tractable and to address the issues of analysis in the presence of uncertainty, we shall restrict ourselves to the aerodynamic analysis of one configuration of the aeroshell. This configuration is a spherical-nosed cone with a spherical aft section, as shown schematically in Figure A.2. The distance between the centers of the two spheres is expected to be quite small. Three nondimensional parameters completely describe the idealized aeroshell: the Bluntness Ratio B (the ratio of the nose radius to the aft section radius), the Fineness Ratio F (the ratio of overall length L to maximum diameter D), and θ, the cone semiangle. The extra information gained by allowing B, F, and θ to vary is useful not only for exploration of alternative designs, but also for analysis of a single fixed design, as the geometry of the aeroshell may change: the heat shield ablates during reentry, for example.

The general design problem is to determine values of B, F, and θ that will (robustly)

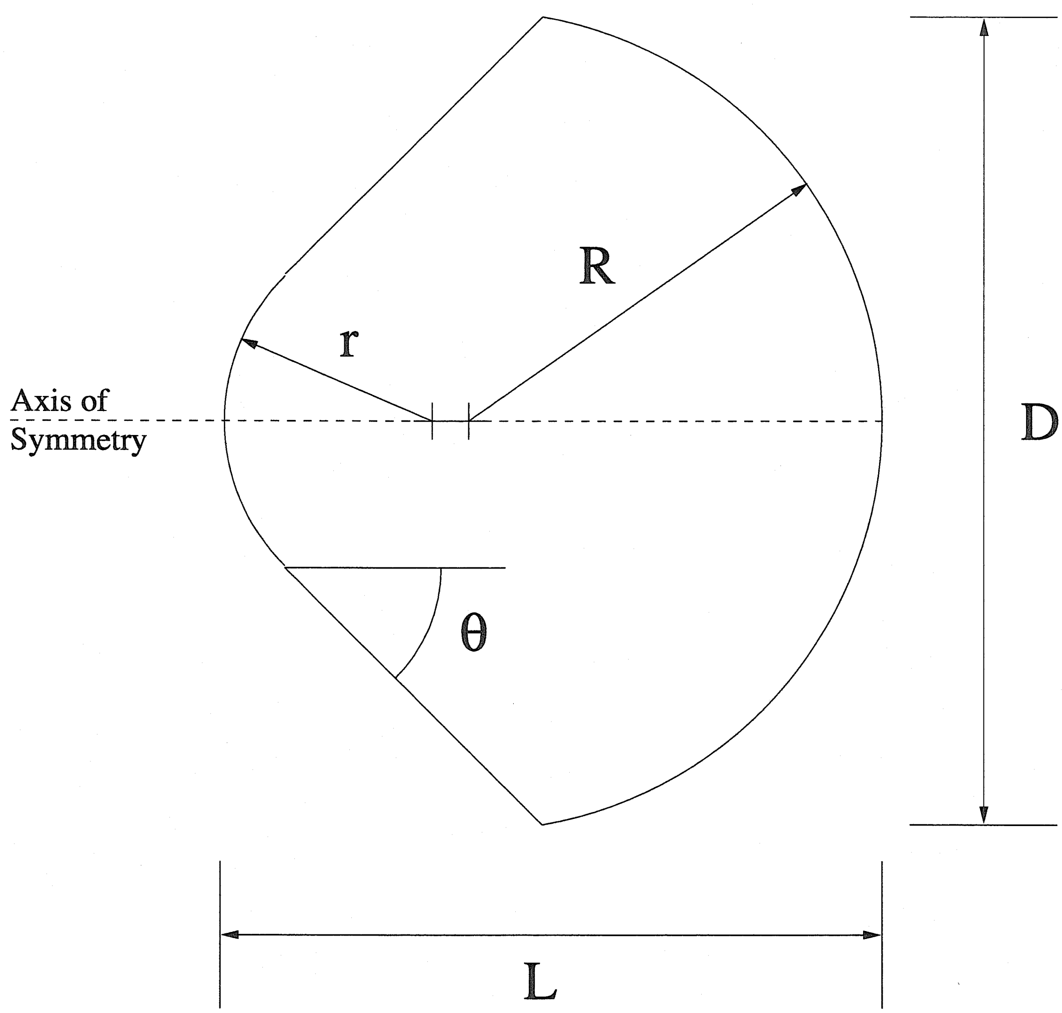

Figure A.2: Schematic of aeroshell

deliver the aeroshell to the surface at a given velocity and angle of attack, in the presence of the operating uncertainty. The more immediate problem is to deliver a reliable integration routine for computer implementation in the presence of uncertainty in the determination of aerodynamic coefficients. The designer, also referred to here as the analyst or engineer, who is interested in using an integration routine to test the performance of a design can draw upon several sources to determine the aerodynamic constants required at each time step of the integration:

- Experimental results from the literature. These might be of varying reliability. Also, the experimental results do not cover the entire design space, so interpolation between experimental points and extrapolation to unexplored areas of the design space is necessary. The reliability of an interpolated or extrapolated answer will decrease with distance from experimental points.

- Simulation data from computational fluid dynamics (CFD) computations (executed point by point — one computation for a particular configuration, angle of attack, and Mach number). These are also of imperfect reliability.

- Analytical computational models, of which there are at least three in this particular problem:

  - Newtonian flow
  - free-molecular flow
  - detached shock flow

  The aeroshell travels through several flow regimes (and is often in an indeterminate state between regimes) during the integration. In addition, most analytic models have been developed for a particular aeroshell geometry, and must be modified to provide useful information for other geometries.

- Rules of thumb: experience- and intuition-based knowledge.

The distillation of information from these sources (each of which is imperfect) is a matter of engineering judgement. As designers determine the aerodynamic constants, at the same time they refine their understanding of each of the sources (a technical reference giving experimental data that deviates significantly from a number of other experiments may be depreciated or discarded, the analytical models may be updated to better fit experimental data, *etc.*).

Fuzzy aggregation has been applied here to interpolate, extrapolate, and combine data from different analysis programs that hold in different regimes. Simultaneously, the level of quality [65] of the analysis has been explicitly represented and propagated using a mathematics of fuzzy sets [107] similar to that used to combine preferences in the MoI.

A.3 Formal treatment of the problem

The statement of the general problem is as follows: Find f such that

$$f(\mathbf{d}, \mathbf{x}) = (\mathbf{p}, \mu)$$

where $\mathbf{d}$ is a vector of design variables describing a point in design space, $\mathbf{x}$ is a vector of operating conditions, $\mathbf{p}$ is a vector of performances, and μ is some measure of the reliability or quality of the answer $\mathbf{p}$ (μ is a vector since the reliability of the components p_i of $\mathbf{p}$ need not be the same for all i, though in the example presented here the μ_i will always agree). We shall also use the more compact notation $\mathbf{y}$ to represent $(\mathbf{d}, \mathbf{x})$. In the example under consideration, $\mathbf{d} = (B, F, \theta)$ describes the geometry of the aeroshell. The operating conditions can be described by the attitude of the aeroshell and the atmospheric conditions, which for the example here can be described by the angle of attack α, the Mach number M, and the atmospheric density ρ. The performances desired are aerodynamic constants: the normal coefficient C_n, the axial coefficient C_a, the moment coefficient C_m, and the center of pressure C_p. Thus the example problem is to find f such that:

$$f(B, F, \theta, \alpha, M, \rho) = (C_n, C_a, C_m, C_p, \mu)$$

The function f is a combination of various subfunctions f_i, so that

$$f(\mathbf{d},\mathbf{x}) = \mathcal{P}\left(f_1(\mathbf{d},\mathbf{x}),\ldots,f_n(\mathbf{d},\mathbf{x})\right)$$

where $\mathcal{P}$ represents the combination. The subfunctions $f_1,\ldots,f_n$ are the sources of information available, and the set of f_i is subject to change. Adding a new source f_{n+1} to the list may make other sources unnecessary. Since a level of quality is one of the outputs of the function f, evaluation is possible with any set of f_i; indeed, it is because the analysis is uncertain and imprecise that the combination is necessary. The human designer making such a judgement in analysis will arrive at an answer and will also have an idea of how valid that answer is. This appendix presents a formal representation for both.

For the aeroshell reentry problem, data for f are available from experimental sources, CFD computations, and analysis models. At the highest level,

$$f(\mathbf{d},\mathbf{x}) = \mathcal{P}\left(f_{\text{exp}}(\mathbf{d},\mathbf{x}), f_{\text{CFD}}(\mathbf{d},\mathbf{x}), f_{\text{analytic}}(\mathbf{d},\mathbf{x})\right)$$

and subfunctions can be further refined, as for example,

$$f_{\text{analytic}}(\mathbf{d},\mathbf{x}) = \\ \mathcal{P}\left(f_{\text{Newtonian}}(\mathbf{d},\mathbf{x}), f_{\text{free-molecular}}(\mathbf{d},\mathbf{x}), f_{\text{detached\_shock}}(\mathbf{d},\mathbf{x})\right)$$

The calculations for Newtonian, free-molecular, and detached shock flow regimes are rapidly computable, and the interpolation of experimental data is a well-understood problem. Each subfunction lends itself to simple automation.

A formal solution to the problem of combining the subfunctions f_i must fulfill several purposes. Some important features of a formal solution are as follows:

- Calculation of each subfunction comes from expertise from the particular discipline. The formal solution must allow for calculation modules to be added, removed, updated, and exchanged.

- Propagation of quality is separate from combination of results, but quality information is necessary for combination of results. The formal solution should propagate and combine quality levels in a justifiable manner.

- Combination of results is easy to do, but easy to get wrong. An arithmetic mean, while computationally tractable, is often *not* the right choice, as was discussed in Chapter 5. The formal solution should use combination methods that are based in domain expertise; the methods of combination, like the calculation modules, should permit easy modification. Combination functions such as other weighted means (those between *min* and *max*) have been used in the M_OI for preference aggregation, and the fuzzy sets literature [109] has an extensive treatment of t-norms (less than *min*) and t-conorms (greater than *max*).

- One feature of this sort of analysis, when it is handled informally by a human designer, is that the subfunctions (or calculation modules) are updated when analysis by other means indicates shortcomings. The formal solution should allow for such back-propagation of information; while the ultimate goal is to provide a proposed change in a particular submodule, an acceptable intermediate step is to provide feedback to the designer, who can modify subfunctions as necessary.

The need to incorporate and modify rules points to fuzzy set theory as a candidate for the combination model. In addition, the combination of designer preference has been represented as the aggregation of fuzzy sets [4], and this problem exhibits many similarities to the designer preference problem. The aggregation will be best illustrated through the presentation of the example.

A.4 Application to the example

The aeroshell analysis problem presented here has been restricted so that the analysis space is spanned by six variables: three design variables and three variables to describe the operating conditions. While experimental data are available for some regions of the analysis space, there are no experimental data points for many regions of interest. In addition, analytic (and thus easily computed) analyses have been previously constructed to cover some of the regions in which the aeroshell will operate, usually with reference to a particular fixed geometry; in this example, the authors had access to a Newtonian analysis code and a free-molecular analysis code. These analytic codes can be applied to other aeroshell geometries if they are suitably modified. The analyst who combines these sources of information has tasks of two varieties: to interpolate and extrapolate experimental and analytic results to "new" areas of the design space, and to determine the level of quality of the interpolated results. The interpolation and extrapolation of data is well understood, and the analysis tool presented here uses polynomial and spline fitting in its implementation. The extension of analytic results to new aeroshell geometries is treated as an interpolation problem in the error.

If we consider the experimental data alone, then the analysis space can be separated into two regimes: one where the experimental data holds, and one where it is not adequate. This distinction is fuzzy; only at the actual experimental points, and only then if the experiment was reliable, can one be certain that the experimental data holds. Anywhere else in the analysis space, the quality that the designer believes the data to have will depend (at least) on how close it is to actual data points.

The designer's belief in the quality of the data is uncertain but not probabilistic; it is not the case that the analysis has a 70% chance of being right and a 30% chance of being wrong. The designer's uncertainty about the reliability of the data is naturally modelled as the degree of membership in a fuzzy set [109]. The quality (or preference, in M_oI terminology) $\mu_{\text{exp}}(\mathbf{y})$ of the applicability of interpolated experimental results to a particular point $\mathbf{y}$ is a function of the point's distance from existing experimental points, taking a value of 1 (perfect quality) at experimental points, and tailing off to 0 at some distance. The quality of interpolated results will also depend in general on the particular point; the transsonic flow regime, for example, is notoriously ill-suited to interpolation.

The calculation of μ_{exp} (or any other μ) is a matter of engineering judgement. Sometimes it may be possible to express μ_{exp} simply and analytically. For example, one might define the quality $\mu_{\text{exp}}(\mathbf{y})$ of the interpolated answer as a function of the (Euclidean) distance of the point $\mathbf{y}$ from the nearest experimental point $\mathbf{y}_{\text{nearest}}$, and let the quality tail off

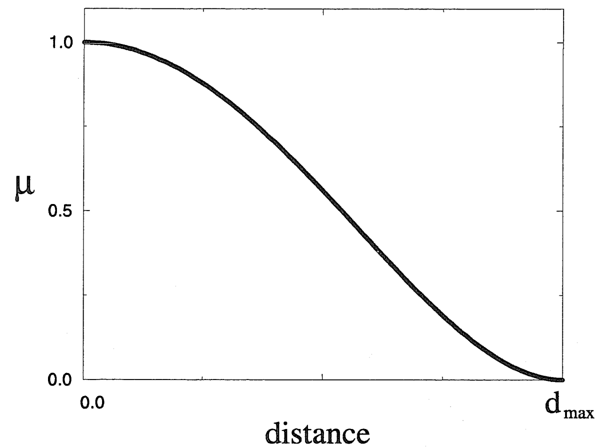

Figure A.3: Example quality curve

as some unacceptable distance $d_{\max}$ is approached, for example with:

$$\mu_{\exp}(\mathbf{y}) = \frac{1}{d_{\max}^4}(|\mathbf{y} - \mathbf{y}_{\text{nearest}}| - d_{\max})^2(|\mathbf{y} - \mathbf{y}_{\text{nearest}}| + d_{\max})^2,$$

a plot of which is shown in Figure A.3. This quartic quality curve has zero slope at μ-values of 0 and 1, and falls off fastest at the midway point. However, there is no proof that the nuances of the curve are an accurate model of the engineer's thinking. The only truly "fixed" points on the curve are those that reflect the engineer's highest belief in the quality of the data (μ=1, if the experiment is completely trusted) at those points where experiments were performed, and those that show that quality degrades to zero at some distance $d_{\max}$ which is specified by the engineer. Other researchers have argued that the human capacity to distinguish many points on a preference curve is limited [50], and so the detailed shape of the curve is unimportant. A linear interpolation is then *a priori* no worse, and has the advantage of simpler calculation:

$$\mu_{\exp}(\mathbf{y}) = \frac{d_{\max} - |\mathbf{y} - \mathbf{y}_{\text{nearest}}|}{d_{\max}}$$

Other subtleties in the determination of a level of quality of the interpolated data, some of which have been alluded to before, contribute to the difficulty of representing quality with a single curve:

1. Only a few points on a quality curve such as the one shown in Figure A.3 will be meaningful to the engineer.

2. The specification of $d_{\max}$ as a single number assumes that Euclidean distance in different dimensions of the design space are equivalent, or at least comparable. This is plainly not true in general, for even if all dimensions can be scaled so that units are comparable, it cannot be assumed that the analyst is equally concerned with "distance" in all directions.

3. The quality of interpolated data will depend on the operating point; in other words, $d_{\max}$, even if well-defined, is a function of $\mathbf{y}$. For example, the designer is likely to believe interpolated data for Mach numbers M between 3 and 8 has a higher quality than data in the transsonic range where $M \approx 1$.

4. The quality of interpolated data will depend on several nearest points, not just on the single nearest experimental point $\mathbf{y}_{\text{nearest}}$.

5. The quality of extrapolated data may be quite different from the quality of interpolated data.

6. When sufficient experimental data is available, the quality of interpolation can be checked against other experimental points; the analyst fitting a curve by hand typically uses such a check in the informal calculus of quality. The $d_{\max}$ approach obscures this.

7. Finally, the engineer may recognize the potential importance of all of these subtleties, and yet arrive at a level of quality without taking all possibilities into account. Especially in preliminary design, the engineer may proceed, considering only the most important quality criteria, and refine the calculation for more detailed analysis.

A more flexible approach to quality specification is required to capture these nuances. The natural model of membership in a fuzzy set for the quality level of a point $\mathbf{y}$ in design space indicates the use of a rule set to capture the designer's beliefs about the quality of data. The transformation of a set of IF-THEN rules into a fuzzy inference matrix is a well-known problem in fuzzy set theory, and commercial packages such as Matlab's Fuzzy Logic Toolbox [41] are available to perform this. A rule set is flexible with respect to the difficulties enumerated above, and is easily updated. In some cases the engineer may feel more comfortable circumventing the rule set and specifying quality functions directly. For instance, for Mach numbers M between 3 and 8, and ρ close to that of air (the value of ρ at which experiments were made), the engineer may wish to define a simple rule for each dimension of the design space describing the loss in quality as a function of the distance from the nearest experimental points in that dimension. The quality contribution with respect to Mach number M is $\mu_{\text{exp}:M}$, say:

$$\mu_{\text{exp}:M}(M) = \frac{d_{\text{max}:M} - |M - M_{\text{nearest}}|}{d_{\text{max}:M}}$$

Similar functions for B, F, θ, and α are combined, in this case with a multiplication:

$$\mu_{\text{exp}}(B, F, \theta, M, \alpha, \rho) = \mu_{\text{exp}:B}(B)\mu_{\text{exp}:F}(F)\mu_{\text{exp}:\theta}(\theta)\mu_{\text{exp}:M}(M)\mu_{\text{exp}:\alpha}(\alpha)$$

Note that since ρ has been assumed to be close to that of air, it has no contribution. However, when ρ is different another rule comes into play and the free-molecular analysis must also be considered.

Some data, taken from wind tunnel tests for cone angles of 10, 15, and 20 degrees [62], are shown in Figure A.4. The data are shown here as isolated experimental points. A standard interpolation scheme will generate a surface over the same range, but not all points on the surface will have the same level of quality. The rule set implemented here maintains high quality of interpolated data along the dimensions α and M (except across the transsonic region $M \approx 1$, where deviations in M are penalized strictly), but enforces relatively high penalties on deviations from experimental points in θ, B, and F. In particular, with the present data there is sufficient granularity in M and α to check curve fits; as more data becomes available in the other dimensions, the quality calculations can be updated.

Of the many physical models for fluid flow to handle different regimes, two have been

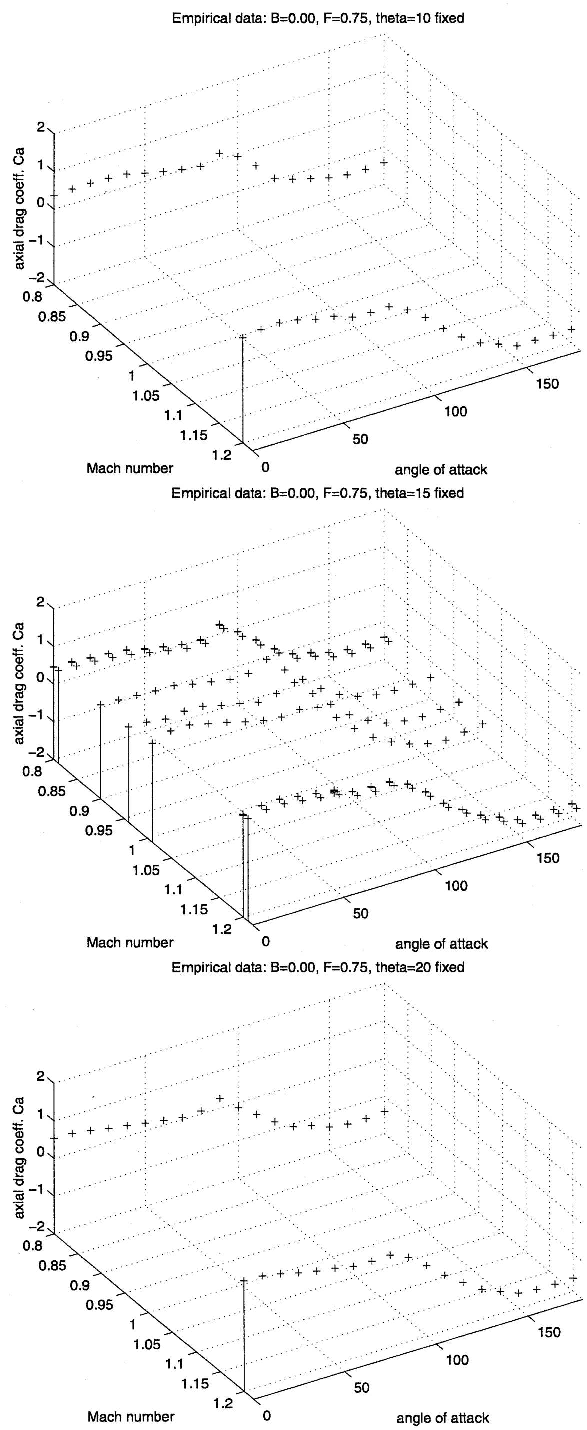

Figure A.4: Empirical data: C_a for three cone angles θ

implemented to date: a model for Newtonian flow, and a model for free-molecular flow. Each of these models covers an analysis regime likely to be encountered by the aeroshell in its descent to the surface, and as was mentioned above, each model was developed for a particular aeroshell geometry. The analyst has some confidence in the output of these models as long as two conditions are satisfied:

1. The operating point $\mathbf{x}$ is in the appropriate flow regime. For free-molecular flow, a high Knudsen number is required, which translates roughly to a low ρ, and an assumption of Newtonian flow is used for some supersonic flows when the entrained boundary layer can be assumed to stay within the shock cone.

2. The aeroshell geometry $\mathbf{d}$ must be close to one for which there are experimental data. The analysis is achieved by calculating the analytic model at a set of experimental points, and then curve fitting the error between the analytic model (which was originally developed for a different geometry) and experiment, and interpolating or extrapolating to the operational point of interest. Thus the machinery of the interpolation scheme and its attendant quality calculation are both relevant here.

Surface plots of each of the two models are shown in Figures A.5 and A.6, with continuous variation in θ, since Mach number is irrelevant for these two particular flow models. For comparison with Figure A.4, a slice of each surface at $\theta = 15°$ is also shown. These analyses are also not accurate over the entire domain; as with the experimental data, there will be varying degrees of quality.

A.4.1 Aggregation of data from disparate sources

To determine the output parameters and their quality levels for a point in analysis space, data from three calculation modules (experimental, Newtonian, and free-molecular) are considered. The varying quality of each calculation module over the space is represented by membership in a fuzzy set. This membership is determined by the application of a number of fuzzy rules. Quality of interpolated experimental data is higher near explicitly calculated data points, with greater penalties for deviations in geometry and lesser penalties for deviations in the operating parameters M and α. If the density is low then free-molecular

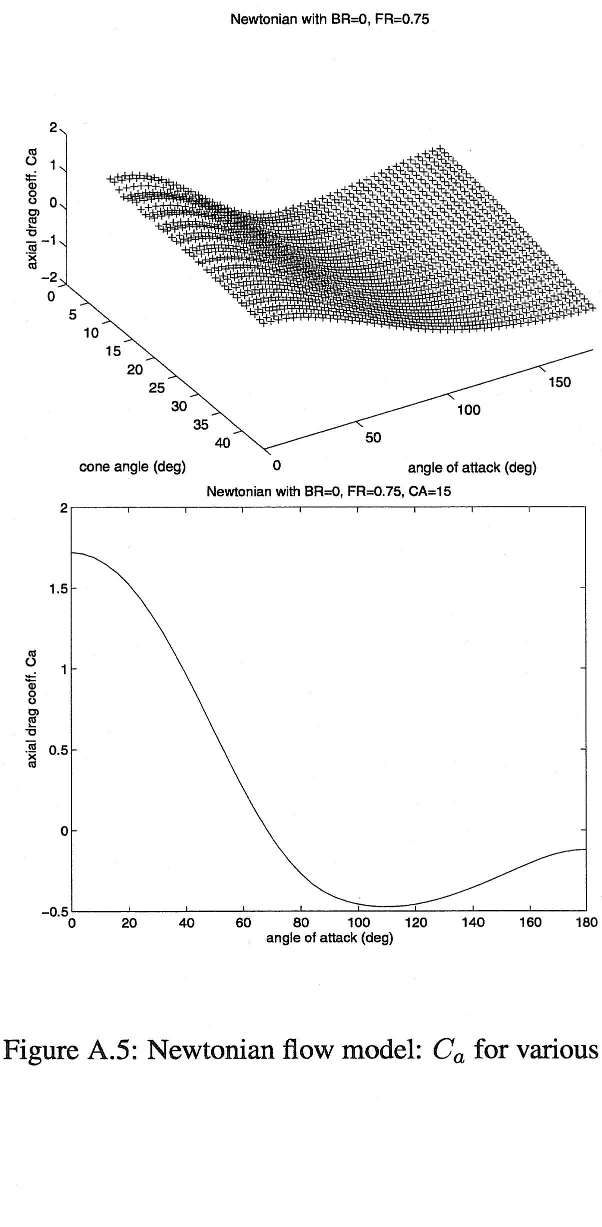

Figure A.5: Newtonian flow model: C_a for various θ

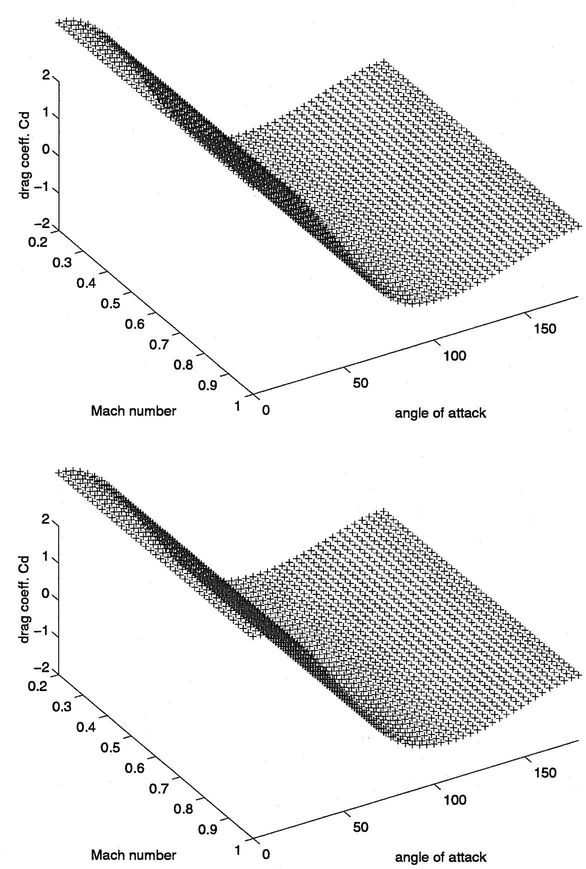

Figure A.6: Free-molecular flow model: C_a for two θ

analysis is useful: Figure A.7 embodies this fuzzy rule in a fuzzy set on density ρ expressing the applicability of the free-molecular flow analysis. The applicability of Newtonian analysis depends on the entrained boundary layer staying within the calculated shock cone. The rules determining the applicability of each regime are specified by the designer and encoded with fuzzy sets, either directly or through the construction of a fuzzy inference matrix; these sets can be updated as the designer refines the rules. Such modification is inexpensive, as it entails only a change in the aggregation problem, and does not require any expensive analysis calculations to be repeated.

The results from all analysis modules are combined, with their participation determined by the quality level of the answer. This analysis will divide the space into regions in which the different analysis modules predominate. Some regions will have high levels of quality for more than one module, as is the case when experimental data are taken in a regime for which there is an analytical model. In this case, the overall analysis includes feedback as to the legitimacy of the modules. Disagreement between modules may lead to changes in quality levels or updated models.

Where a single analysis module has a much higher quality than the others, the result from that module is used, and the quality level is returned with the result. (The analyst using the tool in an integration scheme can log the quality levels, or flag points where quality falls below a given level.) If the quality for the other modules is low but not zero, the answers from those modules can be compared with the result, and the comparison can be logged for the analyst's later use.

When two or more analysis modules return high quality levels, the results must be combined. The most straightforward way to do this is with a weighted sum (the quality levels can provide the weights). A more useful scheme is to compare the results before combining; when they agree closely, a weighted sum is acceptable, and the overall quality will be greater than either of the single quality levels (so that the computation of quality can be effected with a t-conorm [109]). If the results are not in close agreement, it is perhaps better to use the result with the higher quality, but return a lower quality level.

Especially since only three flow regimes have been implemented, there are regions of the analysis space in which all results have low quality. This shortcoming may be cor-

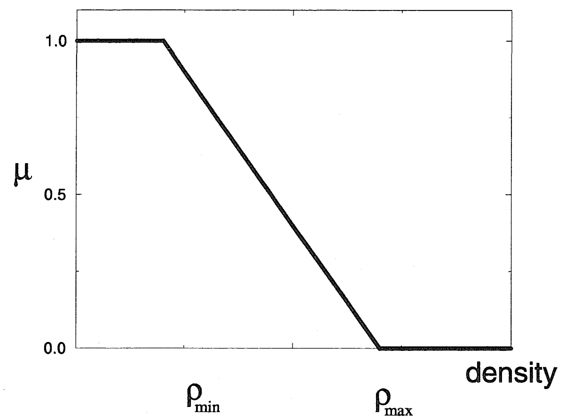

Figure A.7: Fuzzy set expressing applicability of regime

rected by incorporating other analysis models (such as the detached shock analysis or a CFD module). Nevertheless, just as the designer working informally must work with the tools available, the formal combination of results here recognizes the shortcomings in the available analysis tools and signals that with low quality levels. The analysis program presented here is designed for inclusion in an integration routine; the routine can track quality levels to determine the overall reliability of the integration and to determine which regions of the analysis space need further refinement. It would be a poor use of engineering resources to develop a detailed model of the analysis space in a region that the aeroshell never encounters.

Some aggregated results from the analysis program are shown in Figures A.8 and A.9. Each plot shows a response surface for the axial drag coefficient C_a over a region of the analysis space; Figure A.8 contains results where $\mu_{\text{free-molecular}} = 0$ (density is high, and the free molecular model does not hold) and Figure A.9 results where $\mu_{\text{free-molecular}} = 0.7$ (density is low enough that the free-molecular model has a relatively high quality level). In both cases, the empirical data dominate where experiments were performed at points that closely approximate operating conditions, and the physical flow approximation models are more important where there is less experimental data. The experimental data are less significant when $\mu_{\text{free-molecular}} = 0.7$, since these experiments were made with a density equal to that of air. The discontinuities in the plots indicate incomplete knowledge, and

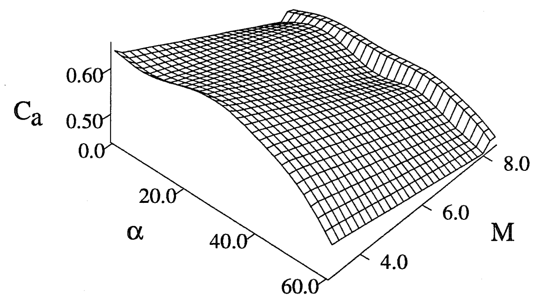

Figure A.8: Interpolated experimental results with $B = 0$, $F = 0$, $\theta = 15°$

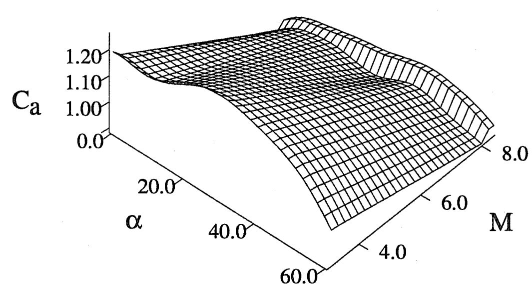

Figure A.9: Results from free molecular code with $B = 0$, $F = 0$, $\theta = 15°$

may indicate that a model needs to be refined.

A.5 Summary

This appendix described an exploration into the possible extension of some of the aggregation formalisms presented in the body of the thesis to the realm of analysis. A method to resolve disagreeing data from multiple sources was applied to the design analysis of an aeroshell reentry design problem furnished by the Jet Propulsion Laboratory. Resolution, interpolation, and extrapolation are accomplished by the use of fuzzy set aggregation functions. This approach does not add significantly to the computation cost, and thus can be refined without repeating analyses, and is therefore particularly useful when computations are expensive and must be minimized.

At the outset of the project, it was conjectured that the aggregation of data from multiple sources would bear closer resemblance to the negotiation formalisms described in the thesis than was ultimately determined to be the case. The results of this project are included here in the interests of completeness.

Bibliography

[1] J. Aczél. *Lectures on Functional Equations and their Applications*. Academic Press, New York, 1966.

[2] P. Adby and M. Dempster. *Introduction to Optimization Methods*. Chapman and Hall, London, 1974.

[3] Manish Agarwal and Jonathan Cagan. A blend of different tastes: The language of coffee makers. *Environment and Planning B: Planning and Design*, 1997.

[4] Erik K. Antonsson and Kevin N. Otto. Imprecision in Engineering Design. *ASME Journal of Mechanical Design*, 117(B) (Special Combined Issue of the Transactions of the ASME commemorating the 50^{th} anniversary of the Design Engineering Division of the ASME.):25–32, June 1995. Invited paper.

[5] Erik K. Antonsson and Kevin N. Otto. Improving Engineering Design with Fuzzy Sets. In Didier Dubois, Henri Prade, and Ronald R. Yager, editors, *Fuzzy Information Engineering: A Guided Tour of Applications*, chapter 38, pages 633–654. John Wiley & Sons, 1997.

[6] Erik K. Antonsson and Hans.-Jürgen Sebastian. Fuzzy Sets in Engineering Design. In Hans-Jürgen Zimmermann, editor, *Handbook of Fuzzy Logic*, chapter 2. Kluwer Academic Publishers, Boston, London, 1998.

[7] Kenneth J. Arrow. *Social Choice and Individual Values*. J. Wiley, 1st edition, 1951.

[8] Kenneth J. Arrow and Hervé Raynaud. *Social Choice and Multicriterion Decision-Making*. The MIT Press, 1986.

[9] Dennis Bahler, Catherine Dupont, and James Bowen. Mixed quantitative/qualitative method for evaluating compromise solutions to conflicts in collaborative design. *Artificial Intelligence for Engineering Design, Analysis and Manufacturing*, 9(4):325–336, September 1995.

[10] T. Barker. *Quality by Experimental Design*. Marcel Dekker, Inc., New York, 1985.

[11] T. Bayes. Essay towards solving a problem in the doctrine of chances. *Philosophical Transactions*, 53:370–418, 1763.

[12] Richard Ernest Bellman. *Adaptive control processes: a guided tour.* Princeton University Press, 1961.

[13] P. Biegel and M. Pecht. Design trade-offs made easy. *Concurrent Engineering*, 1(3):29–40, May/June 1991.

[14] S. Bradley and A. Agogino. An intelligent real time design methodology for catalog selection. In Stauffer L.A, editor, *Design Theory and Methods '91*, volume 31, pages 201–208. ASME, 1991.

[15] Frances M. T. Brazier, Peter H. G. van Langen, and Jan Treur. Modelling conflict management in design: An explicit approach. *Artificial Intelligence for Engineering Design, Analysis and Manufacturing*, 9(4):353–366, September 1995.

[16] D. M. Byrne and S. Taguchi. The Taguchi approach to parameter design. In *Quality Congress Transaction – Anaheim*, pages 168–177. ASQC, May 1986.

[17] Joseph E. Chen and Kevin N. Otto. A tool for imprecise calculations in engineering design. In *Proceedings of the Fourth IEEE International Conference on Fuzzy Systems (FUZZ-IEEE/IFES'95)*, volume 1, pages 389–394. IEEE, March 1995.

[18] S.-J. Chen and C.-L. Hwang. *Fuzzy Multiple Attribute Decision Making: Methods and Applications*, volume 375 of *Lecture Notes in Economics and Mathematical Systems*. Springer-Verlag, New York, 1992. In collaboration with Frank P. Hwang.

[19] Kim B. Clark and Takahiro Fujimoto. *Product Development Performance: Strategy, Organization and Management in the World Auto Industry*. Harvard Business School Press, Boston, MA, 1991.

[20] Anoop K. Dhingra. Multiple objective design optimization. In H.-J. Sebastian and E. K. Antonsson, editors, *Fuzzy Sets in Engineering Design and Configuration*, pages 53–88. Kluwer, 1996.

[21] A. R. Diaz. Fuzzy set based models in design optimization. In S. S. Rao, editor, *Advances in Design Automation - 1988*, volume DE-14, pages 477–485, New York, September 1988. ASME.

[22] A. R. Diaz. A strategy for optimal design of hierarchical systems using fuzzy sets. In J. R. Dixon, editor, *The 1989 NSF Engineering Design Research Conference*, pages 537–547, College of Engineering, University of Massachusetts, Amherst, June 1989. NSF.

[23] S. Djouad, P. Floquet, S. Domenech, and L. Pibouleau. Fuzzy set methods in separation process synthesis. In Didier Dubois, Henri Prade, and Ronald R. Yager, editors, *Fuzzy Information Engineering: A Guided Tour of Applications*. John Wiley & Sons, 1996.

[24] W. M. Dong and F. S. Wong. Fuzzy weighted averages and implementation of the extension principle. *Fuzzy Sets and Systems*, 21(2):183–199, February 1987.

[25] D. Dubois. Tuning the cutting edge conditions on a machining transfer line using fuzzy constraint propagation. In Didier Dubois, Henri Prade, and Ronald R. Yager, editors, *Fuzzy Information Engineering: A Guided Tour of Applications*. John Wiley & Sons, 1996.

[26] D. Dubois and H. Prade. A review of fuzzy set aggregation connectives. *Information Sciences*, 36:85–121, 1985.

[27] Harald Dyckhoff. Basic concepts for a theory of evaluation: Hierarchical aggregation via autodistributive connectives in fuzzy set theory. *European Journal of Operational Research*, 20:221–233, 1985.

[28] Bernard Etkin. *Dynamics of Atmospheric Flight*. Wiley, New York, 1972.

[29] H. Fargier. Fuzzy scheduling: Principles and experiments. In Didier Dubois, Henri Prade, and Ronald R. Yager, editors, *Fuzzy Information Engineering: A Guided Tour of Applications*. John Wiley & Sons, 1996.

[30] S. Finger and J. R. Dixon. A review of research in mechanical engineerign design. part i: Descriptive, prescriptive, and computer-based models of design processes. *Research in Engineering Design*, 1(1):51–67, 1989.

[31] S. Finger and J. R. Dixon. A review of research in mechanical engineering design. part ii: Representations, analysis and design for the life cycle. *Research in Engineering Design*, 1(2):121–137, 1989.

[32] S. French. *Decision Theory: An Introduction to the Mathematics of Rationality*. Ellis Horwood, Chichester, 1988.

[33] V. Gerhart. Uncertainty within the design of a wind turbine. Working paper or, RWTH Aachen, Templergraben 64, 52062 Aachen, Germany, 1997.

[34] D. E. Goldberg. *Genetic Algorithms in Search, Optimization, and Machine Learning*. Addison-Wesley, Reading, Massachusetts, 1989.

[35] B. Grabot, L. Geneste, and A. Dupeux. Tuning of fuzzy rules for multiobjective scheduling. In Didier Dubois, Henri Prade, and Ronald R. Yager, editors, *Fuzzy Information Engineering: A Guided Tour of Applications*. John Wiley & Sons, 1996.

[36] Ileana Hamburg and Peter Hamburg. Fuzzy logic, A user friendly technology for the management of uncertain knowledge in engineering design process. In H. Bandemer, editor, *Modelling Uncertain Data*, pages 153–157, Berlin, March 1993. Akademie Verlag. Mathematical Research series, Volume 68, Lectures from the GAMM-Workshop: Modelling Uncertain Data, held in Freiberg from March 21-24, 1992.

[37] Djamila Haroud, Sylvie Boulanger, Esther Gelle, and Ian Smith. Management of conflicts for preliminary engineering design tasks. *Artificial Intelligence for Engineering Design, Analysis and Manufacturing*, 9(4):313–323, September 1995.

[38] J. Hauser and D. Clausing. The house of quality. *Harvard Business Review*, 66(3):63–73, May/June 1988.

[39] Yeh-Liang Hsu, Yu-Fa Lin, and Tzyh-Li Sun. Engineering design optimization as a fuzzy control process. In *Proceedings of the Fourth IEEE International Conference on Fuzzy Systems (FUZZ-IEEE/IFES'95)*, volume 4, pages 2001–2008. IEEE, March 1995.

[40] V. Hubka and W. E. Eder. *Theory of Technical Systems*. Springer-Verlag, Berlin, 1988.

[41] J.-S. Roger Jand and Ned Gulley. *Fuzzy Logic Toolbox*. The MathWorks, Inc., 24 Prime Park Way, Natick, Mass. 01760-1500, 1995.

[42] R. Keeney and H. Raiffa. *Decisions with Multiple Objectives: Preferences and Value Tradeoffs*. Wiley, New York, 1976.

[43] R. Keeney and H. Raiffa. *Decisions with Multiple Objectives: Preferences and Value Tradeoffs*. Cambridge University Press, Cambridge, U.K., 1993.

[44] Ryszard Knosala and Witold Pedrycz. Evaluation of design alternatives in mechanical engineering. *Fuzzy Sets and Systems*, 47(3):269–280, 1992.

[45] D. Krantz, R. Luce, P. Suppes, and A. Tversky. *Foundations of Measurement*, volume I. Academic Press, New York, 1971.

[46] William S. Law. *Evaluating Imprecision in Engineering Design*. Ph.D. thesis, California Institute of Technology, Pasadena, CA, June 1996.

[47] William S. Law and Erik K. Antonsson. Including Imprecision in Engineering Design Calculations. In *Design Theory and Methodology – DTM '94*, volume DE-68, pages 109–114. ASME, September 1994.

[48] William S. Law and Erik K. Antonsson. Multi-dimensional Mapping of Design Imprecision. In *8th International Conference on Design Theory and Methodology*. ASME, August 1996.

[49] R. Luce and H. Raiffa. *Games and Decisions*. J. Wiley and Sons, New York, 1957. Dover Edition, 1985.

[50] G. A. Miller. The magic number seven, plus or minus two: some limits on our capacity for processing information. *Psychological Review*, 63:81–97, 1965.

[51] R. A. Mitcheltree, J. N. Moss, F. M. Cheatwood, F. A. Greene, and R. D. Braun. Aerodynamics of the mars microprobe entry vehicles. In *15th Applied Aerodynamics Conference*, 1997. Extended abstract submitted for the 15th Applied Aerodynamics Conference, Atlanta, Georgia, June, 1997.

[52] D. C. Montgomery. *Design and Analysis of Experiments*. Wiley, New York, 1991.

[53] R. E. Moore. *Interval Analysis*. Prentice-Hall, Englewood Cliffs, NJ, 1966.

[54] K. Müller and M. Thärigen. Applications of fuzzy hierarchies and fuzzy MADM methods to innovative system design. In *Proceedings of the Third IEEE International Conference on Fuzzy Systems*, volume 1, pages 364–367. IEEE, June 1994.

[55] Vincent Oh and John Sharpe. Conflict management in an interdisciplinary design environment. *Artificial Intelligence for Engineering Design, Analysis and Manufacturing*, 9(4):247–258, September 1995.

[56] Kevin N. Otto. *A Formal Representational Theory for Engineering Design*. Ph.D. thesis, California Institute of Technology, Pasadena, CA, June 1992.

[57] Kevin N. Otto and Erik K. Antonsson. Trade-Off Strategies in Engineering Design. *Research in Engineering Design*, 3(2):87–104, 1991.

[58] Kevin N. Otto and Erik K. Antonsson. The Method of Imprecision Compared to Utility Theory for Design Selection Problems. In *Design Theory and Methodology – DTM '93*, pages 167–173. ASME, September 1993.

[59] Kevin N. Otto and Erik K. Antonsson. Tuning Parameters in Engineering Design. *ASME Journal of Mechanical Design*, 115(1):14–19, March 1993.

[60] Kevin N. Otto and Erik K. Antonsson. Design Parameter Selection in the Presence of Noise. *Research in Engineering Design*, 6(4):234–246, 1994.

[61] G. Pahl and W. Beitz. *Engineering Design*. The Design Council, Springer-Verlag, New York, 1984.

[62] Victor L. Peterson. Motions of a short 10° blunted cone entering a martian atmosphere at arbitrary angles of attack and arbitrary pitching angles. Technical Report D-1326, National Aeronautics and Space Administration, 1962.

[63] Charles J. Petrie, Teresa A. Webster, and Mark R. Cutkosky. Using pareto optimality to coordinate distributed agents. *Artificial Intelligence for Engineering Design, Analysis and Manufacturing*, 9(4):313–323, September 1995.

[64] M. Phadke. *Quality Engineering Using Robust Design*. Prentice Hall, Englewood Cliffs, NJ, 1989.

[65] Robert M. Pirsig. *Zen and the art of motorcycle maintenance: an inquiry into values*. Morrow, 1974.

[66] Christian Posthoff and Michael Schlosser. Learning from examples for the construction of fuzzy evaluations. In *Proceedings of the Third IEEE International Conference on Fuzzy Systems*, volume 1, pages 368–371. IEEE, June 1994.

[67] S. Pugh. *Total Design*. Addison-Wesley, New York, 1990.

[68] S. S. Rao. Description and optimum design of fuzzy mechanical systems. *ASME Journal of Mechanisms, Transmissions, and Automation in Design*, 109:126–132, March 1987.

[69] S. S. Rao and A. K. Dhingra. Integrated optimal design of planar mechanisms using fuzzy theories. In *Advances in Design Automation - 1989*, volume DE-15-2, pages 161–168, New York, September 1989. ASME.

[70] S. S. Rao and A. K. Dhingra. Applications of fuzzy theories to multi-objective system optimization. NASA Contractor Report 177573, NASA Ames Research Center, 1991.

[71] S. S. Rao, A. K. Dhingra, and V. Kumar. Nonlinear membership functions in the fuzzy optimization of mechanical and structural systems. Technical Report 90-1175-CP, AIAA, New York, 1990.

[72] S. S. Rao, K. Sundararaju, B. Prakash, and C. Balakrishna. Multiobjective fuzzy optimization techniques for engineering design. *Computers and Structures*, 42(1):37–44, 1992.

[73] Yoram Reich. A critical review of general design theory. *Research in Engineering Design*, 7(1):1–18, 1995.

[74] W. M. Kim Roddis. Structural failures and engineering ethics. *ASCE Journal of Structural Engineering*, 119(5):1539–1555, May 1993.

[75] Thomas L. Saaty. *The Analytic Hierarchy Process: Planning, Priority Setting, Resource Allocation*. McGraw Hill, New York, 1980.

[76] Masatoshi Sakawa and Kosuke Kato. An interactive fuzzy satisficing methods for large-scale multiobjective linear programs with fuzzy numbers. In *Proceedings of the Fourth IEEE International Conference on Fuzzy Systems (FUZZ-IEEE/IFES'95)*, volume 3, pages 1155–1162. IEEE, March 1995.

[77] Ralf Schleiffer. Fuzzy engineering design. RWTH Aachen, 1998. Diplomarbeit.

[78] Linda C. Schmidt and Jonathan Cagan. Recursive annealing: A computational model for machine design. *Research in Engineering Design*, 7(2):102–125, 1995.

[79] Michael J. Scott and Erik K. Antonsson. Aggregation Functions for Engineering Design Trade-offs. *Fuzzy Sets and Systems*, 99(3):253–264, 1998.

[80] Michael J. Scott, William S. Law, and Erik K. Antonsson. A Fuzzy Sets Application to Preliminary Passenger Vehicle Structure Design. In Enrique Ruspini, Piero Bonissone, and Witold Pedrycz, editors, *Handbook of Fuzzy Computation*. Oxford University Press, May 1998.

[81] C. J. Shih and R. A. S. Wangsawidjaja. Multiobjective fuzzy optimization with random variables in a mix of fuzzy and probabilistic environment. In *Proceedings of the Fourth IEEE International Conference on Fuzzy Systems (FUZZ-IEEE/IFES'95)*, volume 3, pages 1163–1170. IEEE, March 1995.

[82] J. N. Siddall. *Analytical Decision Making in Engineering Design*. Prentice-Hall, 1972.

[83] G. Stiny. Computing with form and meaning in architecture. *Journal of Architectural Education*, 39(1):7–19, 1988.

[84] U. Thole, H.-J. Zimmerman, and P. Zysno. On the suitability of minimum and product operators for the intersection of fuzzy sets. *Fuzzy Sets and Systems*, 2(2):167–180, 1979.

[85] D. L. Thurston and J. Carnahan. Fuzzy ratings and utility analysis in preliminary design: evaluation of multiple attributes. *ASME Journal of Mechanical Design*, September 1990. Submitted for review.

[86] D. L. Thurston and J. Carnahan. Fuzzy ratings and utility analysis in preliminary design evaluation of multiple attributes. *ASME Journal of Mechanical Design*, 114(4):648–658, December 1992.

[87] Deborah L. Thurston. A formal method for subjective design evaluation with multiple attributes. *Research in Engineering Design*, 3(2):105–122, 1991.

[88] T. Tomiyama and H. Yoshikawa. Extended general design theory. In H. Yoshikawa and E. Warman, editors, *Design Theory for CAD*, pages 95–130. IFIP, Elsevier Science Publishers, 1987.

[89] Myron Tribus. *Rational Descriptions, Decisions, and Designs*. Pergamon Press, New York, 1969.

[90] I. B. Türkşen. Scheduling system design: Three fuzzy theory approaches. In Didier Dubois, Henri Prade, and Ronald R. Yager, editors, *Fuzzy Information Engineering: A Guided Tour of Applications*. John Wiley & Sons, 1996.

[91] Design Verein Deutscher Ingenieure (VDI) Society for Product Development and Marketing: Committee for Systematic Design, editors. *Systematic Approach to the Design of Technical Systems and Products*, volume VDI 2221 of *VDI Guidelines*. Beuth Verlag, D-1000 Berlin, Germany, translation of the german 11/1986 edition, August 1987. Copyright: VDI-Verlag GmbH, D-4000 Dusseldorf.

[92] T. Vincent. Game theory as a design tool. *ASME Journal of Mechanisms, Transmissions, and Automation in Design*, 105:165–170, June 1983.

[93] John von Neumann and Oskar Morgenstern. *Theory of Games and Economic Behavior*. Princeton University Press, second edition, 1953.

[94] Peter Wakker and Daniel Deneffe. Eliciting von Neumann-Morgenstern utilities when probabilities are distorted or unknown. *Management Science*, 42(8):1131–1150, August 1996.

[95] Allen C. Ward, Jeffrey K. Liker, Durward K. Sobek, and John J. Cristiano. The second Toyota paradox: How delaying decisions can make better cars faster. *Sloan Management Review*, 36(3):43–61, Spring 1995.

[96] Allen C. Ward, T. Lozano-Pérez, and Warren P. Seering. Extending the constraint propagation of intervals. *Artificial Intelligence in Engineering Design and Manufacturing*, 4(1):47–54, 1990.

[97] Allen C. Ward and Warren P. Seering. The performance of a mechanical design compiler. memo 1084, MIT Artificial Intelligence Laboratory, 1989.

[98] Allen C. Ward and Warren P. Seering. Quantitative inference in a mechanical design compiler. memo 1062, MIT Artificial Intelligence Laboratory, 1989.

[99] D. E. Whitney. Manufacturing by design. *Harvard Business Review*, 66(4):83–91, July 1988.

[100] Kristin L. Wood. *A Method for Representing and Manipulating Uncertainties in Preliminary Engineering Design*. Ph.D. thesis, California Institute of Technology, Pasadena, CA, 1989.

[101] Kristin L. Wood and Erik K. Antonsson. A Fuzzy Sets Approach to Computational Tools for Preliminary Engineering Design. In S. S. Rao, editor, *Advances in Design Automation - 1987, Volume One: Design Methods, Computer Graphics, and Expert Systems – DE-Vol. 10-1*, pages 263–271, New York, September 1987. ASME.

[102] Kristin L. Wood and Erik K. Antonsson. Computations with Imprecise Parameters in Engineering Design: Background and Theory. *ASME Journal of Mechanisms, Transmissions, and Automation in Design*, 111(4):616–625, December 1989.

[103] R. Yager. Decision making with fuzzy sets. *Decision Sciences*, 6(3):590–600, 1975.

[104] R. Yager. Fuzzy decision making including unequal objectives. *Fuzzy Sets and Systems*, 1:87–95, 1978.

[105] H. Yoshikawa. General design theory and a CAD system. In T. Sata and E. Warman, editors, *Man-Machine Communication in CAD/CAM*, pages 35–58. IFIP, North-Holland, 1981.

[106] H. Yoshikawa. Automation of thinking in design. In *Computer Applications in Production and Engineering*, pages 405–417, Amsterdam, 1983. North-Holland.

[107] L. A. Zadeh. Fuzzy sets. *Information and Control*, 8:338–353, 1965.

[108] L. A. Zadeh. Fuzzy logic and approximate reasoning. *Synthese*, 30:407–428, 1975.

[109] Hans-Jürgen Zimmermann. *Fuzzy Set Theory - and Its Applications*. Management Science/Operations Research. Kluwer-Nijhoff Publishing, Boston, MA, 1985.

[110] Hans-Jürgen Zimmermann. Media selection and other applications of fuzzy linear programming. In Didier Dubois, Henri Prade, and Ronald R. Yager, editors, *Fuzzy Information Engineering: A Guided Tour of Applications*. John Wiley & Sons, 1996.

[111] Hans-Jürgen Zimmermann and Hans-Jürgen Sebastian. Optimization and fuzziness in problems of design and configuration. In *Proceedings of the Second IEEE International Conference on Fuzzy Systems*, pages 1237–1240, San Francisco, CA, June 1993. IEEE, Pergamon Press.

[112] Hans-Jürgen Zimmermann and Hans-Jürgen Sebastian. Fuzzy design - Integration of fuzzy theory with knowledge-based system-design. In *Proceedings of the Third IEEE International Conference on Fuzzy Systems*, volume 1, pages 352–357. IEEE, June 1994.

[113] Hans-Jürgen Zimmermann and Hans-Jürgen Sebastian. Intelligent system design support by fuzzy-multi-criteria decision making and/or evolutionary algorithms. In *Proceedings of the Fourth IEEE International Conference on Fuzzy Systems (FUZZ-IEEE/IFES'95)*, volume 1, pages 367–374. IEEE, March 1995.

Index

A-pillar, 120, 121

acoustic mode, 31

aeroshell, 138, 140–142, 145

affine transformation, 34

aggregation, 3, 4, 16, 20, 39

- design-appropriate, 6, 39, 46, 47, 49, 56, 80
- design-appropriate, maximal, 56
- function, 17, 39, 41, 43, 44, 58, 62, 63, 66, 80
- hierarchical, 76
- hierarchy, 75, 124
- of fuzzy sets, 10
- of more than two preferences, 75
- of preference, 5, 39, 40
- operator, 6, 40, 42, 47, 58

AHP, *see* Analytic Hierarchy Process

α-cut, 58, 85, 91

Analytic Hierarchy Process, 8, 46, 72

annihilation, 45, 47, 50, 56–58, 62, 71

Arrow's Theorem, *see* Impossibility Theorem

Arrow, Kenneth J., 6, 20, 21, 23, 24, 82–85

artificial intelligence, 4, 11

- compared to M_oI, 12

averaging operators, 10

axioms, 5, 22

- annihilation, 42, 44, 45, 47
- bisymmetry, 48
- continuity, 42–44
- for aggregation functions, 44
- for aggregation operators, 42
- for MADM problem, 33
- for social choice problem, 23
- for the M_oI, 39, 41
- for weighted means, 48
- homogeneity of weights, 48
- idempotency, 42, 44, 45, 48
- increasing in variables, 48
- increasing in weights, 48
- independence of irrelevant alternatives, 23
- internality, 48
- limited imposed orders, 32
- monotonicity, 42–45
- no dictatorial orders, 23
- no imposed orders, 23, 32, 84
- positive response, 23
- self-scaling weights, 42, 44
- symmetry, 42–44
- unrestricted domain, 23, 31, 84
- zero weights, 42, 44

B-pillar, 120, 122, 123

Bahler, 11
balance of trade, 59, 60
Bayes, 9
bending stiffness, 78, 115, 116, 120, 124
bluntness ratio, 140
body-in-white, 114, 115, 119
bolt, 2
Bradley and Agogino, 9

Caltech, 5, 6, 9, 13
CFD, 139
chassis, 78
Chen and Hwang, 46, 47
compensation, 53, 54, 56–58, 63, 67, 80
    parameter, 50, 56, 57, 67, 70, 72, 81
cone semiangle, 140
connected set of designs, 85
constraint, 31, 60, 79
convexity, 4, 25, 82, 85, 86
crisp specification, 14, 78
curse of dimensionality, 99

decision
    group decision making, 9, 20
    probabilistic, 9
    theory, 9
    under uncertainty, 9, 20, 35
    with multiple criteria, 9, 20, 29
decision making, 4, 9
design
    analysis, 7
    as pre-theory field, 8
    computer-based models, 7
    decision problem, 17
    descriptive models, 7
    environments, 7
    for manufacture, 8, 79
    languages, 7
    life-cycle, 8
    preference, 15, 36, 121, 123
    prescriptive models, 7
    representations, 7
    set-based, 5, 6, 8, 82–85, 89, 100, 101, 110, 135
    theory and methodology, 7
    variable, 13
Design of Experiments, 17, 111, 129
Design Variable Space, *see* *DVS*
Dhingra, 11
Diaz, 11
DOE, *see* Design of Experiments
dominated point, 62
drawing board, 2
DS-2, 138
Dubois, 11
Duell, Charles, 1
DVS, 13, 16, 36, 85, 98, 111

Ecclesiastes, 1
economics, 9
exhaustive search, 99
expectation, mathematics of, 34
experimental design, 9
extension principle, 16, 36, 108

Fargier, 11
fineness ratio, 140
Finger and Dixon, 7
finite element, 115, 119, 120
floor pan, 120, 122
floor sill, 120, 122
functional equations, 47
fuzzy
    constraint propagation, 11
    design with evolutionary algorithms, 11
    evaluations, 11
    linear programming, 59, 62
    MADM methods, 11, 46

  multi-criteria decision making, 11
  optimization, 11
  scheduling, 11
  set, 62
fuzzy sets, 4, 6, 9, 85, 108, 145
  applications to decision making, 11
  as generalization of classical sets, 10
  membership function, 10, 46
  theory, 10

game theory, 9, 37
GDT, *see* General Design Theory
General Design Theory, 8
General Possibility Theorem, *see* Impossibility Theorem
generating function, 49
genetic algorithm, 100, 106, 139
geometric mean, 49, 56, 66
Gerhart, 11
Grabot, 11
grammars, *see* shape grammars

Hamburg and Hamburg, 11
heavy machinery, 79
hierarchy, *see* aggregation hierarchy
homeomorphism, 49
Hsu, 11
Hubka and Eder, 8
hypercube in $\mathbb{R}^n$, 85

idempotency, 56, 70
Impossibility Theorem, 6, 20, 21, 24, 25, 29, 33, 82–85
  proof of, 24
  statement of, 24
imprecision, 4, 14
index, 76
indifference, 34, 70, 71, 92
interval analysis, 84, 85

Jet Propulsion Laboratory, 156

Kant, 24
Keeney and Raiffa, 9
Knosala and Pedrycz, 11

lottery, 36
lottery question, 34
Luce and Raiffa, 9

MADM, *see* multi-attribute decision making
majority method
  intransitivity of, 24
majority method of decision making, 21
manipulability, 37, 75
Mars, 138
Matlab, 128, 149
maximum stress, 31
Method of Imprecision, *see* $\mathrm{M_oI}$
microgravity, 137
military aircraft, 79
Miller, 58
MODM, *see* multi-objective decision making
$\mathrm{M_oI}$, 9, 12, 13, 16, 33, 39, 45–47, 56, 58, 80, 82, 97, 101, 114
$\mathrm{M_oI}$ preference, *see* preference
monocoques, 114
MSC Nastran, 102
μ-convex, 86, 87, 89–92
Müller and Thärigen, 11
multi-attribute decision making, 29, 39, 45, 47, 59, 80, 82
multi-criteria decision making, 4, 30
multi-objective decision making, 29, 39, 46
multi-objective optimization, 4
multiobjective scheduling, 11

NASA, 137
natural selection, 100
negotiation, 3, 5, 20, 35, 60, 66, 67, 70, 72, 75
  hierarchical, 75
noise, 17

objective function, 60
Oh and Sharpe, 11
optimization, 97, 105, 139
  classical, 98
ordinal ranking, 20
Otto, 11

Pahl and Beitz, 8
paradox of voting, 22
Pareto optimality, 12
pattern search, 99
performance
  preference, 15, 36
  unmodelled, 15
  variable, 13
Performance Variable Space, *see DVS*
Petrie, 12
political spectrum, 25
Posthoff, 11
Powell's method, 99
preference, 3, 4, 15, 62, 66, 70
  aggregation, 39, 72
  function, 40, 85
  individual, 19, 34, 36, 66
  induced, 16
  inter-attribute
    comparison of, 36
  $M_O I$, 34, 36, 46
  numerical scale for, 19, 34, 35
  overall, 17, 19, 36, 40
  single-peaked, 31, 85
preliminary design, economic importance of, 4
press fit, 2
probability, 34
process planning, 9
product of powers, 49, 66
profit, 59, 60
proximity of designs, 86
Pugh charts, 8, 47
PVS, 13, 16, 36

QFD, *see* Quality Function Deployment
quality, 143
quality control, 9
Quality Function Deployment, 8, 33, 47

Rao, 11
rationality, 42, 45, 82, 84, 85, 92
rivet, 2

Saaty, *see* Analytic Hierarchy Process, 72
Sakawa and Kato, 11
satisfaction, 62
Schemebuilder, 11
Schleiffer, 11
SDRC I-DEAS, 118
set-based reasoning, *see* design, set-based
7 ± 2 categories, 36, 58
shape grammars, 8
Siddall, 9
simulated annealing, 99, 139
single-peaked preference, *see* preference
single-peakedness, 24
snap fit, 2
social choice, 21, 33
social choice, differences between MADM and, 29
solid model, 115
Solomon, 1
space exploration, 137
space frames, 114
stiffness, 114, 115, 125
Stiny, 8
strategy, *see* trade-off strategy
stress, 114
strictly increasing homeomorphisms on $\mathbb{R}$, 49

Structural Dynamics
Research
Corporation, 118
style, 15, 114, 121
Sun workstation, 102, 127

t-conorm, 10, 47, 145
t-norm, 10, 47, 145
Taguchi methods, 8
Thurston, 9
and Carnahan, 11
torsional stiffness, 78, 115, 117, 120, 124
trade-off, 39, 66, 98, 120
trade-off strategy, 39, 67, 69, 70, 72, 74–76, 80
Tribus, 42
tuning parameters, 17
tunneling algorithms, 99
Turksen, 11

undominated
point, 60, 66, 67, 69, 70, 72, 74
set, 60, 68, 69, 74
solution, 60, 67
utility, 22, 34
function, 34
interpersonal
comparison of, 22, 30, 35
lottery, 34
von Neumann–Morgenstern, 22, 34
utility theory, 9, 11, 32, 47

value function, 20, 22, 34
VDI, *see* Verein Deutscher Ingenieure
Velcro®, 2
Verein Deutscher Ingenieure, 8
veto, 33
Volkswagen, 114, 133
von Neumann and Morgenstern, 9, 34
VW Rabbit, 114, 115, 133

Ward, 8
weak order, 19, 34
weight, 120, 125
weighted means, 47, 49, 56, 57, 80, 145
weighted sum, 27, 66
weights, 40, 62, 66, 67, 69, 70, 72, 75
in hierarchies, 77
predetermined, 70, 72
weld, 2, 114
windturbine, 11
Wolfsburg, 114, 133

Yoshikawa, 8

Zimmermann, 46, 59, 60
Zimmermann and Sebastian, 11