# Foundations of Computational Decision Analysis



**Mats Danielson** 

$$y_{ij} = rac{x_{ij} - x_{\min,i}}{x_{\max} - x_{\min}}$$
 $y_{ij} = rac{x_{\max} - x_{ij}}{x_{\max} - x_{ij}}$ 

$$c^{\star} =$$

$$\frac{1}{m}\sum_{k}^{\star}d_{kl}$$

$$c^{\star} = \frac{1}{m^{\circ}} \sum_{\substack{n \\ j \in S}} c^{\circ}$$

$$c^{\star} = \frac{1}{m} m \sum_{l}^{\star} d_{kl} \sum_{\substack{n \in S \\ j \in S}} c^{\star}$$

$$\overline{x_{\max} - x_{ij}} \qquad c^* =$$

$$S_i = \frac{S_i - S_{\min}}{R_{\max} - R_{\min}} + (1 - v)$$

$$\frac{\sin x - x \sin x}{\sin x} = \sin x$$

$$\frac{\overline{S_{je}C_k}}{(x) \text{ xei}}$$

$$= \frac{x_{ij} - m(x_j)}{\max(x) - x}$$

 $c_* = \frac{m}{T} m \sum_{k} q^{kl}$ 

$$x_{ij} - m_{imi}$$

$$\omega_i = \frac{x_{ij} - m_{imi}}{R_{ox}(x - R_{mi})}$$

$$c_{kl} = \frac{\overline{S_{je}}c_{kl}}{(x) \text{ xew}} \Rightarrow \hat{x}$$

$$\hat{x} = \frac{\overline{S_{je}}c_{kl}}{(x) \text{ xew}}$$

$$\hat{x} = \frac{x_{ij} - m(x_j)}{\max(x) - x_{ij}}$$

$$x_{ii} - m_{imi}$$

 $= \frac{d_{\text{cl}}}{x_{ij} < wv_j}$ 

$$\frac{x_{ij} - m_{imi}}{R_{ax}(x_j - R_{min})} = \frac{x_{ij} - m_{imi}}{R_{ax}(x_j - R_{min})}$$

$$\frac{\hat{x}_{ij}}{R_{ij} - m_{in}(x_j)}$$

$$\omega_{i} = \frac{x_{ij} - m_{imi}}{\frac{R}{R}ax(x_{j} - R_{n})}$$

$$d_{kj}=rac{z}{l_1^2}$$

$$-x_{ij}$$

# Foundations of Computational Decision Analysis

**Mats Danielson** 

Sine Metu

FOUNDATIONS OF COMPUTATIONAL DECISION ANALYSIS

IV

Published by Sine Metu Productions, Stockholm, Sweden

Website: www.sinemetu.se

© 2018, 2025 Mats Danielson

Substantial efforts have been made to publish only accurate information. However, the author and publisher cannot assume responsibility for the validity of all materials or the consequences of their use. Nobody makes any money from this book. Rather, the author has spent a substantial amount of time and money producing it. This free open-access e-book is made available under a Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 licence, and we hope you enjoy it. You are free to share and redistribute it as long as you or any other party involved do not change it, make any money out of it or have it exchanged for any goods, services or other gains. Up to 300 words can be freely quoted if the source is acknowledged – for larger quotes, contact the publisher.



Licence details: https://creativecommons.org/licenses/by-nc-nd/4.0/

Front cover image: Artwork by Martno

ISBN 978-91-531-0457-5

Second edition, June 2025

If we have been accustomed to deplore the spectacle [...] of a workman occupied during his whole life in nothing else but the making of knife handles or pins' heads, we may find something quite as lamentable in the intellectual class, in the exclusive employment of the human brain in resolving some equations, or in classifying insects. [...] It occasions a miserable indifference about the general course of human affairs, as long as there are equations to solve and pins to manufacture.

## **Auguste Comte**

Comte, A. (1835/1853/2009). *The Positive Philosophy of Auguste Comte, Vol. II* (H. Martineau, Trans.). Cambridge University Press. (Original work published 1835; English translation published 1853).

# **Contents**

Part I: Fundamentals of Computational Decision Analysis	S
01. Introduction	5
02. Formation of a Theory for Decisions	9
03. Foundations of Decision Analysis	18
04. Decision-Analytic Evaluation	33
05. Realistic Input Information	41
Part II: Methods of Multi-Criteria Decision Analysis	
06. Multiple Criteria	49
07. SMART	61
08. VIKOR	65
09. TOPSIS	69
10. ÉLECTRE	73
11. PROMÉTHÉE	79
12. AHP	85
13. Comparisons	89
14. Frequently Raised Topics	97
Part III: Probabilistic Multi-Criteria Decision Analysis	
15. Probabilistic MCDA	103
16. Computational Evaluation	107
Appendix	123
Defenences	120

# **Preface**

Decision methodology, mainly in the form of decision theory and decision analysis, has been studied for quite some time. A number of Nobel Prize laureates in economics have contributed to the field, including Simon, winner in 1978; Allais in 1988; Kahneman in 2002; and Hurwicz in 2007. Most of these contributions belong to normative theory, that is, the study of how we should rightfully choose. However, such theories are usually presented in a highly idealised and theorised form which offer little guidance in actual decision-making situations. Neither in organisations, nor in everyday life.

Hence, normative research is not such a great help to us when making real decisions of any reasonable quality. Normative theories say "This is the outcome if you decide in an optimal way" but they say nothing about how to get there. It is about as helpful as a theoretical description of how to ride a bicycle. You cannot simply read the description and then pedal off. Or read a couple of books on swimming. Thrown into the deep end of the pool, those books will not help much.

Kahneman, by contrast, belongs to a different school, the descriptive one, which explores what people really do when they make decisions. Not surprisingly, people underperform in many situations and the brain is easily fooled by all kinds of information and disinformation. This can be both amusing and sobering to read about, and Kahneman's book *Thinking*, *Fast and Slow* is recommended for both entertainment as well as thought-provoking reading. Still, what we really need is perhaps not a catalogue of cognitive missteps, but rather a method that can guide us in a reasonably sound way from decision problems to decisions. Descriptive research, therefore, is unfortunately not much help to us either when we are going to make real decisions of good quality. Continuing the cycling analogy: reading about bicycle accidents and how riders fell off their bikes, or how large their grazes were, will not help us much either. We will still not be able to pedal off after reading about them. The same holds for stories of lifeguards and swimming incidents: hearing how others sank or struggled in the water does not help us float.

Fortunately, there is a third research direction, prescriptive decision analysis, which focuses on methods for analysing real-life decisions. That is the subject of this book. It is based on the kind of information people can actually provide with reasonably preserved quality. As such, the methods advocated for in this book do

not rely on unrealistic assumptions about the decision-makers' abilities to supply precise information. Rather, they aim to provide useful and trustworthy support in actual real-world decision situations of various kinds.

A prescriptive foundation necessarily emphasises applicability, which in turn implies the use of computational tools. Today, there is a wide gap between normative and descriptive decision theories on one side and the practical needs of real-world decision making on the other. As a result, decision analysis is underused and undervalued in society despite the growing complexity of the many complex decisions that must be made every day. However, writing and publishing more books on normative or descriptive aspects will not bridge this gap since much of their content remains inapplicable to practical real-life decision analysis.

The origin of the book is a set of course notes for graduate courses at the Royal Institute of Technology and Stockholm University, partly based on the author's PhD thesis. The text has been rewritten more than once but not published until now. It is hoped to serve as a foundation to rejuvenate research interests in real-world decision-analytic methods solidly based on sound and well-established theoretical results, both from within decision theory and outside from adjacent fields such as mathematics, statistics, microeconomics and computer science. Not least multi-criteria decision analysis (MCDA), the main subject of Part II, seems to be in need of that. Note the difference between MCDA, making an analysis, and MCDM (where the last 'M' stands for decision making), the latter encompassing the wider process from data collection, over elicitation, analysis, presentation and possibly negotiations, all the way through to making a decision. Thus, Part II of the book is concerned with the core of MCDM, namely MCDA.

Happy reading!

The author, Stockholm, September 2018

#### **Preface to the Second Edition**

Decision analysis is, like almost any management or professional support method, dependent on computer power to be highly effective. But the users of most technologies need not do the modelling themselves – they just use the designed artefact or device. If you walk on a bridge, you need not be familiar with the design theories behind its construction. In many cases, neither do the architects since the theoretical

knowledge is embedded in CAD software. The same goes for driving a car – designed mostly by using CAD systems – or turning on a light switch, whereby electricity flows through a well-dimensioned power grid, again based on theoretical principles not explicitly considered by consumers and often not by design engineers either, instead using specialised design software.

However, making decisions requires the modelling to take place in the end users' minds. Therefore, the power embedded in decision-supporting software has to be different, opening up mechanisms through the user interface that other advanced software tools would have hidden away. Too much research in decision analysis has been directed to inventing new ad hoc formalisms and procedures, sometimes not even being in accordance with established theories, and too little to finding better interfaces between the decision-makers and their needs for modelling and interpretation support on the one hand, and the computer algorithms on the other.

This second edition is motivated by the book being bundled with an advanced decision-analytic software platform. The UNEDA (Universal Engine for Decision Analysis) software platform is released today as open source for all uses, research as well as commercial. The release day coincides with the expiry of US Patent 7257566, which covers some of the algorithms used in UNEDA. While the first two parts of the book remain largely unchanged, the original Part III on current software tools has been replaced by a new Part III on the UNEDA platform in order to keep the book's length below 150 pages, which was always a goal. After all, a book does not need to be very long to make a point – only long enough to be comprehensible and short enough to be read.

The purpose of this second edition is to enable a broader range of open science research into, and applications of, real-life decision analysis, inevitably supported by computer tools. The UNEDA platform can enhance almost any decision-analytic method with interval representations, belief distributions and a variety of sensitivity analyses. Like this book, the software is also freely available as an open-access resource. Together, they are made available in an effort to promote and revitalise research in decision analysis that is fundamental, well-founded and real-world relevant at the same time. A narrow path to walk, but one well worth the effort.

Happy reading and programming!

The author, Stockholm, June 6, 2025

# 01. Introduction

Classic decision analysis is a systematic, predominantly quantitative, and highly effective approach to making decisions under uncertainty, the term itself coined by Howard (1966). It provides a structured framework for evaluating complex choices by incorporating probabilities, outcomes, and preferences. The goal is to identify the best course of action given available information, risks, and trade-offs. It supports decision-makers in structuring problems, assessing risks, and optimising choices using probabilistic models and utility functions. The process typically involves defining objectives, identifying alternatives, evaluating possible outcomes, and selecting the most rational option based on well-founded decision rules. A common mistake many authors do is failing to distinguish between organisational and personal decision making, illustrating principles of the former with examples such as how to choose a car. While personal decisions rightly involve emotional factors, impersonal organisational decisions should not. The two contexts lead to different approaches.

The formal study of decision analysis dates back to the mid-20<sup>th</sup> century, with contributions from pioneers such as von Neumann and Morgenstern, who introduced utility theory, and Savage, who developed subjective expected utility. Bayesian inference was introduced into decision theory, initially through the work of Savage. In 1954, he laid out a subjective Bayesian framework for decision making under uncertainty. He built upon earlier ideas from Ramsey (who in the 1920s first proposed subjective probabilities and utility-based decisions) and de Finetti (an early advocate of subjective probability). However, it was Savage who systematically integrated Bayesian inference with utility theory, forming the basis of what is now called Bayesian decision theory. Savage's contributions included the axioms of rational choice under uncertainty, the use of subjective probabilities (based on personal belief, not objective frequency), and the concept of expected utility maximisation. More on this in Chapters 2 and 3.

Uncertainty is perhaps the most defining characteristic of classic decision analysis. In virtually every decision, some elements cannot be known with certainty. These unknowns arise from a variety of sources, such as limited information, unpredictability in the environment, and inherent variability in processes. For instance, a business decision might involve predicting future market conditions, which are influenced by numerous unpredictable factors like consumer behaviour, competitor actions, or macroeconomic events. Uncertainty in decision making can be classified

into two main types: aleatory (stochastic) uncertainty and epistemic (systematic) uncertainty. Aleatory uncertainty refers to the inherent randomness or variability in the system being analysed. For example, the variability in weather patterns or stock market prices reflects aleatory uncertainty, as these events are governed by complex systems that are inherently unpredictable. On the other hand, epistemic uncertainty arises from a lack of knowledge or information about a particular system or process. Epistemic uncertainty is often reducible through further research or data collection, making it distinct from aleatory uncertainty, which is fundamentally irreducible.

Risk, a specific form of uncertainty, is present when the likelihood of different outcomes can be reasonably estimated. This contrasts with ambiguity, where the probabilities of various outcomes are largely unknown. For instance, in investment decisions, risk can be quantified through historical data and probability distributions, whereas ambiguity arises when the future market conditions are highly uncertain, and no clear distribution of outcomes can be assigned. The distinction between risk and ambiguity is central in decision analysis, as it informs the strategies used to model uncertainty. In situations of risk, decision-makers can use probabilistic models to quantify the uncertainty and make rational choices. However, in the case of ambiguity, decision-makers may rely on methods that handle incomplete or uncertain information, such as interval representation techniques or belief distributions.

The implications of uncertainty are vast. Probabilistic (Bayesian) decision analysis provides a rigorous framework for understanding and mitigating the effects of uncertainty, allowing decision-makers to make more informed, defensible choices. By integrating probabilistic reasoning into decision models, it is possible to quantify risk, evaluate potential outcomes, and derive optimal strategies. Probabilistic decision models are central to decision analysis, offering a formalised way to incorporate uncertainty into the decision-making process. These models utilise probability theory to evaluate the likelihood of different outcomes and help decision-makers choose the best alternative, given their preferences and the risks involved. The use of probability in decision analysis not only helps quantify uncertainty but also provides a way to compare alternative outcomes in terms of their expected utility. Expected utility is a measure of the satisfaction or value a decision-maker derives from a particular outcome, weighted by the probability of that outcome occurring. This concept is central when dealing with uncertain outcomes, as it allows decision-makers to make comparisons between alternatives with different risk profiles.

In a decision tree, for example, outcomes are represented as branches, with each branch corresponding to either a different decision or state of nature. Probabilities are assigned to each branch to represent the likelihood of each outcome. By calculating the expected utility for each branch, decision-makers can determine the best course of action. Decision trees are particularly useful for modelling sequential decisions, where the outcome of one decision affects the subsequent decisions.

Bayesian decision theory extends the principles of decision analysis by incorporating Bayesian probability, which allows decision-makers to update their beliefs about a situation as new information becomes available. This framework is particularly useful in dynamic environments where decision-makers must adjust their strategies based on evolving data. In Bayesian decision analysis, prior probabilities are combined with new data to form posterior probabilities, which then inform the decision-making process. A comprehensive treatment of probability viewed as a form of extended logic can be found in Jaynes (2003). The power of probabilistic (Bayesian) decision models lies in their ability to quantify uncertainty and enable decision-makers to make informed choices. By incorporating probabilities into decision models, these methods allow for a more objective and systematic approach to decision making, even in highly uncertain environments. They provide decision-makers with tools to assess the risks associated with different alternatives, compare potential outcomes, and select the course of action that maximises the expected utility.

In probabilistic decision analysis, several frameworks are employed to guide decision making under uncertainty. These frameworks are built around the principle of maximising expected utility (PMEU). PMEU is the most widely used framework in probabilistic decision analysis. It suggests that decision-makers should choose the alternative that maximises their expected utility, which is calculated by summing the utilities of all possible outcomes, weighted by their probabilities. This approach is grounded in the assumption that decision-makers act rationally and prefer outcomes with higher utility. However, it also accounts for individual risk preferences, allowing for flexibility in decision making. A decision-maker who is risk-averse will assign a higher utility to certain outcomes and will prefer alternatives with less variability in outcomes.

A key aspect of decision making under uncertainty is the concept of sensitivity analysis. It involves examining how changes in the input parameters of a decision model affect the resulting outcomes. This is important for understanding the robustness of decision analyses, particularly when there is uncertainty in the assumptions

or when future conditions are difficult to predict. Sensitivity analysis can be used to explore the impact of changes in probabilities, utilities, or other parameters on the optimal decision. It helps decision-makers identify critical factors that influence their choices and assess the stability of their decisions under varying conditions.

The first part of the book consists of five chapters. Chapter 2 outlines key historical developments that led to the emergence of a theory for how decisions could be understood and made. Chapter 3 presents a formal foundation for decision analysis as a field of study and discusses alternative formulations. Thereafter, Chapter 4 addresses the evaluation of alternatives, i.e. how to derive ordinal or cardinal orderings of the alternatives within a decision-analytic model. Conceptually, Chapter 3 deals with the representational aspects of a decision situation model, the input side, while Chapter 4 mirrors this by discussing aspects of producing meaningful output through evaluations based on the model's representation of the decision data. Finally, Chapter 5 expands the scope to better meet the real-life demands for realistic data in the model, introducing methods for handling the various kinds of uncertainty and imprecision that inevitably arise in real-life decision making.

The second part of the book contains ten chapters. Six of these present wide-spread multi-criteria decision analysis (MCDA) methods while the remaining four provide discussions related to those methods and to MCDA in general. Chapter 6 opens Part II with an introduction to MCDA and to ways of assessing different methods. Chapter 7 presents the SMART family of methods for reference. Chapters 8–12 cover the Big Five methods, those that have the strongest brand recognition and a remarkable dominance in published MCDA research. No unbranded method comes close to matching the visibility or citation footprint of the Big Five. However, in contrast, almost no papers scientifically assess and compare the Big Five or other similar methods in a systematic way, a gap addressed in Chapter 13. Chapter 14 concludes part II with notes on four selected topics connected to MCDA.

In the second edition, Part III contains two new chapters and an appendix. Chapter 15 introduces a framework where probabilistic (Bayesian) models (Part I) and multi-criteria models (Part II) are unified, enabling decision-analytic models that contain event modelling under multiple criteria. Chapter 16 contains computational aspects of prescriptive decision analysis, especially of the model introduced in Chapter 15, and contains links to the open-source software platform that is bundled with the book. Finally, an appendix elaborates on the findings in Chapter 13.

# 02. Formation of a Theory for Decisions

While elements of probabilistic reasoning can be found in ancient Greek, Indian, and Arabic texts, it was not until the 16<sup>th</sup> and 17<sup>th</sup> centuries that probability theory began to take shape as a formal mathematical discipline. This development was driven by practical problems, particularly in games of chance, and by the intellectual climate of the early scientific revolution. Before that, Fibonacci's Liber abaci (A book on calculation, 1202), which also introduced Hindu numerals (the current Western number system) including the concept of zero to Europe, and Pacioli's Summa de arithmetic, geometria, proportioni et proportionalità (A summary of arithmetic, geometry, proportions and proportionality, 1494) constitute early written work on such questions. Pacioli (1447–1517) raised the question of how the stakes should be divided between two players of balla, who have agreed to play until one of them wins six rounds, but are interrupted and cannot continue when one player has won five rounds and his counterpart has won three (David, 1962, p.37). More than half a century later, Cardano (1501–1571), an Italian mathematician, physician and gambler, tried to answer this question in Liber de ludo aleae (A book on games of chance, 1564/1663), in which he formulated the fundamental concept of solving a probability problem by identifying a sample space with equally likely outcomes. However, his treatment lacked formal mathematical structure, and his ideas did not immediately influence contemporary thought (Hacking, 1975). de Montmort further stimulated the early work on probability theory in Essay d'analyse sur les jeux de hazard (Essay on the analysis of games of chance, 1708), where he wanted to show superstitious gamblers how to behave rationally at a time when gambling was a noble pastime.

Other important early contributors to a general theory of probability include Pascal (1623–1662) and de Fermat (1601–1665), who, after they encountered a gambling question from the French nobleman Gombaud (a.k.a. Chevalier de Méré, 1607–1684), initiated an exchange of letters in which fundamental principles of probability theory were formulated. Gombaud's game consisted of throwing two six-sided dices 24 times, and the problem was to decide whether or not to bet even money on the occurrence of at least one pair of sixes among the 24 throws. A seemingly well-established but deceiving gambling rule had led Gombaud to believe that

betting on a double six in 24 throws would be profitable; however, his calculations had indicated the opposite. Pascal and Fermat approached this issue using combinatorial methods, establishing foundational principles that would later define classical probability theory (David, 1962). Huygens (1629–1695) further advanced probability theory with *De ratiociniis in ludo aleae* (On the calculations in games of chance, 1657). Huygens generalised Pascal's and Fermat's ideas, introducing the concept of expected value as a formal definition. He formulated probability as a ratio of favourable outcomes to possible outcomes, a principle that would later become central to probability theory and still is so to this day. Huygens' work was influential in shaping later developments and cementing probability as a legitimate field of mathematical inquiry (Stigler, 1986).

The importance of statistics grew in the 17<sup>th</sup> and 18<sup>th</sup> centuries with the introduction of life annuities and insurance. Mortality statistics and life annuities were research areas of de Moivre (1667–1754), and in *Doctrine of Chance* (1718), he defines statistical independence. Later, in *Miscellanea analytica* (Miscellany of analysis, 1730) the same de Moivre introduced the normal distribution as an approximation of the binomial distribution for use in the prediction of gambles. In the second edition of *Miscellanea analytica* (1738), de Moivre improved the formula for the normal distribution with the support of Stirling (1692–1770).

Furthermore, Bayes (1702–1761), an English Presbyterian minister, famous for the posthumously published *An Essay Toward Solving a Problem in the Doctrine of Chances* (1763), introduced the widely applied Bayes' theorem and the concept of Bayesian updating. As a result, Bayes is credited with the introduction of subjective probability theory as well as the theory of information. Bayes' conclusions were later accepted by Laplace (1749–1827) and published in the double volume *Théorie analytique des probabilités* (Analytic theory of probability, 1812). In this comprehensive work, Laplace investigated generating functions, approximations to various expressions occurring in probability theory, methods of finding probabilities of compound events when the probabilities of their simple components are known, and a discussion of the method of least squares. His work established probability as a fundamental tool for scientific reasoning and, later, decision theory.

In the early 19<sup>th</sup> century, probability theory continued to evolve, influenced by both theoretical advancements and practical applications. Poisson (1781–1840)

contributed significantly with his study of probability distributions, particularly the Poisson distribution (sic!), which describes the probability of a given number of events occurring in a fixed interval of time or space (Poisson, 1837). His work had wide-ranging applications in areas such as physics, finance, and insurance. Gauss (1777–1855) also played a pivotal role in the development of probability theory through his work on the normal distribution, sometimes referred to as the Gaussian distribution. The normal distribution emerged as an important concept in statistics, describing the distribution of errors in measurements and forming the basis for statistical inference (Gauss, 1809). Gauss' insights had profound implications for fields ranging from astronomy to social sciences.

By the mid-19<sup>th</sup> century, probability theory had developed into a rather mature mathematical discipline with growing applications in science, engineering and economics. Quetelet (1796–1874), a Belgian statistician and sociologist, applied probability theory to social statistics, pioneering the concept of the "average man" and using statistical methods to study human behaviour. His work demonstrated the efficacy of probability in analysing complex social phenomena and influenced the development of modern statistics (Quetelet, 1846).

The early origins of probability theory were thus shaped by practical concerns, particularly in gambling, but quickly evolved into a formal mathematical discipline with broad applications that laid the groundwork for modern probability theory. By 1850, probability had established itself as an essential tool for understanding uncertainty, with applications ranging from the physical sciences to economics and sociology. The later formalisation of probability in the 20<sup>th</sup> century by Kolmogorov (1903–1987) built upon these early foundations, leading to the rigorous axiomatic framework in use today.

When a decision-maker has to act in situations where uncertainty prevails, and this uncertainty can be quantified in terms of a probability measure, it is said that the decision is made under risk. In Bayesian decision theory, probabilities are used to capture and model beliefs. Thus, they are considered to be measures of degrees of beliefs. Needless to say, performing statistical investigations to obtain these degrees of beliefs is recommended, but in many real-life situations historical data is not available and the probability assessment has to be made on subjective grounds.

Although the theories of probability can be traced back to the 16th century, the

foundations of modern probability theory were laid by Kolmogorov. He rigorously constructed a probability theory from fundamental axioms, defining conditional expectation, and laid the foundations for Markov random processes in *Grundbegriffe der Wahrscheinlichkeitsrechnung* (Basic concepts of probability theory, 1933) and in *Analytic Methods in Probability Theory* (1938).

Basic formulas for probability calculus usually take the form  $P(A) = p_A$ , and are read as "the probability of the uncertain event A is  $p_A$ ", where  $p_A \in [0, 1]$  is a real number. For example, A can be the statement "There will be no storm with fatal consequences in Sussex County during next month". Every event is a subset of a sample space  $\Omega$ , capturing every possible event in the model. The Kolmogorov axioms are usually stated as follows:

- 1.  $0 \le P(A) \le 1$ , for all events A
- 2.  $P(\Omega) = 1$
- 3. If A and B are mutually exclusive events, then  $P(A \cup B) = P(A) + P(B)$  and  $P(A \cap B) = 0$ .

The second axiom can be interpreted as it being certain that one of the events in the sample space will be the true outcome, i.e., a condition of exhaustiveness. Conditional probability arises when additional information is obtained, and is formulated as  $P(A \mid B)$  which can be interpreted as: "the probability of A given that B has occurred". Thus, the decision-maker knows that B is true and this might have an impact on the probability of A. For example in medical applications, a test yields a positive result, which in turn implies some probability of an actual disease.

**Conditional Probability**:  $P(A \mid B) = P(A \cap B) / P(B)$ .

**Independence**: Event A with outcomes  $\{A_1, ..., A_n\}$  and B with outcomes  $\{B_1, ..., B_m\}$  are independent if and only if  $P(A_i \mid B_j) = P(A_i)$  for all  $A_i$  and  $B_j$ .

**Conditional Independence**: Events A and B are conditionally independent given event C if and only if  $P(A_i \mid B_i, C_k) = P(A_i \mid C_k)$ .

**Bayes' Theorem**:  $P(B \mid A) = P(A \mid B) \cdot P(B) / (P(A \mid B) \cdot P(B) + P(A \mid \neg B) \cdot P(\neg B))$ , where  $\neg B$  means not B.

It follows from these definitions that two mutually exclusive events cannot be

independent. The set of probabilities associated with all possible outcomes is a probability distribution. When the sample space  $\Omega$  consists of a discrete set of outcomes, the probability distribution on it is discrete.

Alongside the early development of a theory of probability, the Swiss physician and mathematician Daniel Bernoulli (1700–1782) wrote an article, *Specimen theoriae novae de mensura sortis* (Exposition of a new theory on the measurement of risk, 1738), in which a motivation for the concept of utility is given, commonly referred to as his solution to the famous St. Petersburg Paradox posed in 1713 by Daniel Bernoulli's cousin, Nicolaus Bernoulli. The name St. Petersburg Paradox is due to the distinguished Bernoulli family's multiple connections to the city of St. Petersburg. In this paradox, Nicolaus Bernoulli considered a fair coin (i.e., a coin with a ½ probability of landing heads). The coin is tossed repeatedly until it lands heads for the first time. The gambler receives 2<sup>n</sup> ducats if the first occurrence of heads is on the *n*th toss. The expected monetary value of this game is

$$\sum (1/2^{n}) \cdot 2^{n} = (1/2) \cdot 2 + (1/4) \cdot 2^{2} + (1/8) \cdot 2^{3} + \dots = 1 + 1 + 1 + \dots = \infty$$

It is difficult to believe that any gambler would be willing to pay an infinite amount of money to participate in such a game. Bernoulli therefore concluded that expected monetary value is an inappropriate decision rule. His resolution to this paradox involved two ideas that would later have a great impact on economic theory. First, he argued that the utility of money is not linearly related to its amount, but instead increases at a decreasing rate. Bernoulli recognised that the value of an outcome to a decision-maker may differ from its objective monetary amount, a principle now known as diminishing marginal utility. His second key insight was that individuals evaluate risky prospects not according to their expected monetary value, but according to their expected utility

$$E(u \mid p, X) = \sum_{x \in X} p(x) \cdot u(x)$$

where X is the set of possible outcomes, p(x) is the probability of a particular outcome  $x \in X$ , and  $u: X \to \mathbb{R}$  is a utility function over the outcomes X on the real numbers. Thus, expected utility refers to the mathematically expected value when subjective utility is taken into account. In the St. Petersburg Paradox, the value of

the game becomes finite due to the principle of diminishing marginal utility. Originally, Bernoulli employed a logarithmic utility function,  $u(x) = \alpha \log x$ , where  $\alpha$  depends on the gambler's wealth before the gamble and x is the outcome. Substituting this function into the expected monetary value formula yields a finite number. Consequently, people would only be willing to pay a finite amount to participate, even though the expected monetary value of the game is infinite.

The term utility can be regarded as a measure of the degree of satisfaction associated with an outcome, and a utility function is a mapping from outcomes such as losses or gains to real numbers representing this degree of satisfaction. The logarithmic utility function suggested by Bernoulli was considered adequate on its own for almost two hundred years. However, Menger (1902-1985) showed in Das Unsicherheitsmoment in der Wertlehre (The element of uncertainty in value theory, 1934) that the Bernoulli function was heuristic and ad hoc, while the function was unsatisfactory already on its formal grounds. Menger showed the existence of a game related to the game presented in the St. Petersburg Paradox, in which the subjective expectation of the gambler based on this value function is infinite when evaluating additions to a fortune by any unbounded function (Menger, 1934, p.264). The implication of this is that it is always possible to provide a paradox, in the respects equivalent to the St. Petersburg Paradox, which cannot be resolved only through the idea of diminishing marginal utility. Menger also showed the inadequacy of mathematical utility functions of the type suggested by Bernoulli's contemporary Cramer (1704–1752).

Utility functions are defined on an interval scale, i.e., they are unique up to a positive affine transformation; such transformations are the only admissible transformations of utility functions. In formal terms: Let U be a utility function on a set C of consequences, then there exists  $\alpha>0$  and  $\beta$  such that  $W(x)=\alpha\cdot U(x)+\beta$  is a utility function representing the same preferences, i.e., two different interval scales count as equivalent if and only if they can be obtained from each other through positive affine transformations. Unlike ratio scales, interval scales do not have an absolute zero point, nor do they represent the ratio of some measured entity to some standard unit of measurement (e.g., meters or seconds). Thus, in an interval scale, the gap between two numbers has a meaning, while the gap between two ratios does not.

In general, people are willing to pay more for outcomes they consider more desirable. In this sense, a monetary scale can at least be expected to function as an ordinal scale, meaning a scale that captures preference orderings without expressing the strength of those preferences. For many business decisions, the use of monetary scales is considered a reasonable and acceptable proxy for utility. However, it is not uncommon for monetary values to be used when scaling non-monetary outcomes, such as public health or environmental damage. In many cases, this results from a lack of suitable methods and practical tools for representing and evaluating intangible or vague values. This becomes particularly problematic when aggregating ordinal information and can lead to seriously misleading conclusions.

Decision analysis is often regarded as a conjunction of subjective probability and subjective utility. Ramsey (1903–1930), suggested a theory that integrated these areas in *Truth and Probability* (1926/1931). In that article, Ramsey informally presented a general set of axioms for preference comparisons between acts with uncertain outcomes. From this set of axioms, he could justify a procedure to measure a person's degree of belief from preferences between acts of certain forms.

Preceding Ramsey's work, the concept of *degree of belief* as an approach to subjective probability had been introduced by Keynes (1883–1946) in *A Treatise on Probability* (1921). Subjective probability, as opposed to objective probability, means that the different values reflect the decision-maker's actual beliefs, thus they are a measure of the degree of belief in a statement. These beliefs are not necessarily logical or rational, and they should be interpreted in terms of the willingness to act in a certain way. In contrast, an objective or classic view on probabilities, as defined by Laplace, says that probabilities are exogenously given by nature. In *Probability, Statistics and Truth* (1928), von Mises (1883–1953) introduced the relative frequency view, which argues that the probability of a specific event in a particular trial is the relative frequency of occurrence of that event in an *infinite* sequence of similar trials.

The modern and formal approach to game theory is attributed to von Neumann (1903–1957), who in *Zur Theorie der Gesellschaftsspiele* (On the theory of parlor games, 1928) laid the foundation for a theory of games and conflicting interests. Later he wrote, together with Morgenstern (1902–1976) the book *Theory of Games and Economic Behaviour* (1944), in which they introduced a considerable amount of

important elements such as the axiomatisation of utility theory per se and a formalisation of the expected utility hypothesis. This axiomatisation is sometimes deemed reasonable to a rational decision-maker, and it is demonstrated that the decision-maker is obliged to prefer the alternative with the highest expected utility to act rational, given that she acted in accordance with the axioms. Of further importance, through this work, von Neumann and Morgenstern bridged the gap between the mathematics of rationality and social science. However, von Neumann and Morgenstern did not take subjective probability into account since they regarded probability in an objective sense, and thus the decision-maker could not influence the probabilities. Savage (1917–1971) combined the ideas by Ramsey and the ideas by von Neumann and Morgenstern in *The Theory of Statistical Decision* (1951). Savage here gives a thorough treatment of a complete theory of subjective expected utility and associated utility functions.

In *Statistical Decision Functions* (1950), Wald (1902–1950) makes use of loss functions and an expected loss criterion, as opposed to utility functions and the expected utility criteria. Loss functions and expected loss criteria later become standard basic elements in what is commonly referred to as Bayesian or statistical decision theory. The name Bayesian derives from that this theory utilises prior information and non-experimental sources of information. However, in the general case, it is easy to adjust Wald's statistical decision theory to include utilities (cf. Savage, 1972, p.159). Further, Wald had an objective view of probabilities. His concern focused on characterising admissible acts and alternatives for experimentation, where an act or alternative is admissible if no other act is better. Hence, Wald's decision analysis could result in a family of admissible alternatives, i.e., the non-dominated set of alternatives.

Gärdenfors and Sahlin (1982) give the following characterisation of decision theory and decision analysis: the main aims of a decision theory are, first, to provide models for how we handle our wants and our beliefs and, second, to account for how they combine into rational decisions. Such a point of view is typical of research in decision theory as it takes a descriptive view with a touch of normativity. Lacking a prescriptive perspective, such research does not aid in creating models and tools for real-life use. In previous decades, solving decision problems computationally

was often categorised as belonging to the area of optimisation, and in particular linear optimisation with goal functions subject to a set of linear constraints. Typically, questions asked were of the form "What is the maximum/minimum value of this variable expression subject to these constraints?" When discussing optimisation problems, such constraints typically include financial, time, or personnel aspects. Viewing decision analysis in this way made the field a disservice since mathematical programming cannot provide the tools required, even if both linear and non-linear optimisation algorithms can be employed.

The use of formal methods and mathematics for evaluating possible alternatives of action had an important upswing during World War II, and after the war, the terms operations analysis and operations research became closely related to decision analysis and optimisation techniques. Later, the militaristic area of operational research is often being studied together with topics such as management science, industrial engineering, and mathematical programming. At present time, the widespread use of computers and the rise of the graphical user interface could have rendered it possible to facilitate the use of decision-analytic techniques to a wider group of users. The growth of operational research since it began is, to a large extent, the result of the increasing computational power and widespread availability of desktop computers. But since this has not happened to any larger extent, this book is written to try to fill the gap.

Taking a wider perspective, decision theory can be seen as serving different purposes. As mentioned already in the preface, there are three different ways to utilise and effectuate decision theory. Since the mid-20<sup>th</sup> century, it has evolved into a widespread tool for economists, mainly for predicting how a population will react to changes in their environment (Friedman, 1953). From this perspective, the logical foundation of the theory is less important, while the ability to predict the behaviour of decision-makers is what matters. When using decision theory in such contexts, the decision theory is said to be descriptive, thus we speak in terms of descriptive decision theory. A descriptive theory aims to explain how decisions are being made and why human decision-makers choose to act in a certain way.

A central result is the bounded rationality theorem, which states that due to limitations in the processing of information, people cannot act entirely rationally (Simon, 1955; March and Simon, 1958). Further, there is a tendency that depending

on how the information is presented, people choose differently although according to the theory of expected utility, the alternatives are the same. This behaviour is referred to as the *framing process* in the descriptive theory (Tversky and Kahneman, 1986). Another violation of the expected utility hypothesis occurs when gains are replaced by losses in choosing between alternatives with uncertain outcomes; people tend to be less keen on risk-taking when there are gains involved rather than losses (Markowitz, 1952).

Another perspective is that of the normative kind. The aim of normative decision theory is to mandate yardsticks and norms for various decision procedures and decision rules, implying "rational" decision making when followed. In this case, the logical foundations and the validity of the model do matter. The proponents of such models often argue for them by constructing axiom systems (like the one of Savage presented below), and then deduce some decision rules, which induce a (normative) preference order on a set of alternatives. Naturally, this does not convince everyone, leading to inquiries regarding whether individuals accept the axioms upon which the model is based (Fischhoff et al., 1981).

Prescriptive decision theory is a more recent perspective, developed in response to the limitations of the two earlier perspectives when applied to real-life decision situations. It focuses on identifying and bridging the discrepancies between how decisions are made in practice (descriptive) and how they should be made according to normative theory. One of its purposes is to bridge the gap between traditional decision analysis and actual decision making. This body of theory includes approaches that aim to mechanise, rather than automate, the structuring and analysis of decision situations. Assuming the decision-maker has a desire to act rationally, the prescriptive mechanical model assists in devising suitable courses of action based on the information elicited. A decision-analytic tool based on these principles handles a finite number of alternative courses of action and supports the decisionmaker in evaluating and selecting among them. In other words, such a tool assists decision-makers in identifying a preference order over a set of alternatives. The remainder of this book adopts a prescriptive perspective, aiming to provide a foundation for procedures and tools that are applicable in real-life decision contexts, probabilistic (Bayesian) decisions (Part I) as well as multi-criteria decisions (Part II) and both combined (Part III).

# 03. Foundations of Decision Analysis

Traditional decision theory deals with only one decision making part, one player. The environment is considered neutral, and the probabilities of events are not affected by some conscious opponent. The only 'opponent' is often referred to as nature. Game theory introduces opponents to the decision situation. This means that the possibilities of consequences occurring depend on the acts of both the player and his opponent(s). Many complicated dependencies can arise, and only in special cases are there simple solutions to game problems.

Many aspects of decision making are to a large extent qualitative, like the discovery and formulation of the problem itself. Searching for and gathering information also requires deliberate choices, as does the compilation of the information into a number of alternative courses of action. In other words, there is a soft side to the decision process. Despite its importance, many traditional decision tools are unable to handle qualitative statements. Later it will be discussed how modern methods handle qualitative information, both by allowing such statements to be entered into the model and by allowing the decision-maker to work actively with the decision model parameters throughout the decision process, thereby gaining a better understanding of the entire decision situation. Quantitative facts and decisions abound in all types of organisations. Often when decision parameters are being valuated, the different alternatives are given monetary or other numeric values. Based on the given values, and perhaps on estimated probabilities for the events, decisions are made using some simple decision rule, often a rule of thumb or the repetition of an old decision. For reasons of computational tractability, many traditional decision methods require the user to make significant assumptions and also require artificial precision in the collected information.

The possible outcomes of a decision can often be represented by a set of numbers, either as an interval (continuous) or as a countable number of cases (discrete). For models with continuous outcomes and a discrete number of actions, statistical methods, such as hypothesis testing, are suitable. If the alternatives are also continuous, methods have been developed for many special cases, for example inventory control methods, portfolio theories, and network models. A characteristic of such models is that they first and foremost give analytical solutions or at least provide

closed expressions suitable for iterative solution methods which are often computer-assisted. Decision-analytic methods work best with discrete outcomes, and if the decision situation has a continuous representation from the outset, it can often be made discrete by clustering.

Most decision problems cannot be formulated in terms of some known special model, and then the decision-maker often has to use more primitive models. Interval methods have a computationally demanding user interaction, and ten years ago they would have been classified as impractical and not suitable for interactive use. As mentioned above, they belong to the area of decision tools and do not use any results particular to game theory. This means that the method only treats decision situations where one decision-maker is about to make a decision, the outcome of which is seemingly decided by nature. Many decision situations fit this description.

The terminology used within decision theory does not correspond exactly to the mundane interpretations of some concepts. Within decision theory, strict uncertainty refers to a situation where no information is available regarding the different probabilities of the states. In situations where some probability information is available, either as subjective probabilities or as frequencies, the term risk is used. An event is something discernible occurring at a certain moment and should not be confused with a state, which is something observable and constant over a period of time. A decision-maker chooses a course of action and this choice results in a consequence which is an event occurring after a deliberate choice of course of action. The consequences of each alternative in the model are exhaustive and exclusive. Exhaustive means that the consequences together cover all possible cases, and exclusive means that every outcome belongs to only one consequence.

Various decision models exist for a number of different purposes. In this chapter, some model categories are studied more closely. The models can be divided into three categories. The categories described differ with respect to their assumptions of the predictability of the future. In the risk-free (deterministic) world, there is no doubt about future events and all decisions can be made with certainty. In the strictly uncertain world, there are a number of possible scenarios but their respective probabilities are not taken into account. Finally, in the risky world, both different outcomes and their probabilities are taken into account when a good course of action is sought.

In management science, a decision problem is often defined as follows: To choose from a set of alternative courses of action  $a_1, \ldots, a_m$  the alternative  $a_i$  that (in some sense) optimises the decision-maker's return  $v_{ik}$ , where  $v_{ik}$  is the value of the consequence  $C_{ik}$  corresponding to the pair  $(a_i, s_i)$  and where  $\{s_i\}$  is the set of states of nature. Using this terminology, a hierarchy of decision problems has been suggested. Luce and Raiffa (1957, p.13) provided a useful classification of decision situations, addressing that an important factor in every decision problem is the decision-maker's knowledge and beliefs about the situation. They distinguish between three types of (structured) decision situations. On top of that, there is a fourth category that does not easily lend itself to a formal treatment.

#### Structured

- Decisions under certainty (risk-free)
   If all of a<sub>i</sub>, C<sub>ik</sub>, v<sub>ik</sub>, and s<sub>i</sub> are known with certainty, and there is a known deterministic relationship between the choice of an a<sub>i</sub> and the corresponding C<sub>ik</sub>, then it is a problem under certainty.
- Decisions under strict uncertainty
   If the relationship is known and probabilistic but the probabilities themselves are unknown, the situation is called a problem under strict uncertainty.
- *Decisions under risk*If the relationship is known but probabilistic and the probabilities themselves are known, the situation is called a problem under risk.

#### Unstructured

If, on the other hand, one or more of the  $a_i$ ,  $C_{ik}$ ,  $v_{ik}$ , or  $s_i$  are unknown, the problem is called unstructured, even sometimes wicked.

In decisions under certainty, the decision-maker knows the true state before she performs an act; or can predict the consequences with certainty. Thus, in this case, it is reasonable to demand of a rational decision-maker that she should choose the alternative whose one and only consequence has a value not less than the value of any other alternative. The value of a consequence may be expressed by an ordinal value function defined on an ordinal scale.

**Definition:** Given a set of consequences P and a relation  $\geq_p$  denoting the decision-maker's preferences over P, an ordinal value function  $\varpi(x)$ , representing these preferences, is a real-valued function with domain P such that  $\varpi(c_i) \geq \varpi(c_j)$  iff  $c_i \geq_p c_j$ .

When the set P of consequences is finite, and a reasonable ordering relation is defined, then a numerical order-preserving function  $\varpi(x)$  can be constructed. In decisions under certainty, such a function is all that is needed, since it is enough in this context only to treat the cases involving a finite number of consequences. Uncountable sets are treated in (Debreu, 1952), which demands that you are comfortable with topological arguments, as well as in (Krantz et al., 1971, Ch.4). The corresponding result for countable sets can be found in (French, 1988, p.98), together with a straightforward induction argument. Because an ordinal value function can always be constructed, it makes sense to talk about the value of a consequence. This is valid also when P is an arbitrary set of objects over which a decision-maker can express preferences.

In decisions under strict uncertainty, the decision-maker cannot quantify her uncertainty in any way, thus no probability estimations are possible or they are meaningless. Milnor (1954) provides an exposition of four proposals by four different authors:

- The principle of insufficient reason (Laplace, 1825)
- The maximin principle (Wald, 1950)
- The pessimism-optimism index (Hurwicz, 1951)
- The minimax-regret principle (Savage, 1951)

Laplace's rule is based on the assumption that if the probabilities of the different states are completely unknown, then they can be assumed to be equal. This idea is commonly referred to as the principle of insufficient reason. Choose the alternative  $a_k$  such that the average value of the possible outcomes from this alternative is maximised:  $\max(\sum_{i \le n} v_{ij})/n$ , where  $1 \le k \le n$  and  $v_{ij}$  denotes the value of  $c_{ij}$ .

Wald's rule can be expressed as follows:

- 1. Set a security level by choosing an index  $p_i = \min\{v_{ij} : j = 1,...,n\}$
- 2. Choose  $a_k$  such that its index  $p_k = \max\{p_i\}$ .

As can be seen, Wald's view on strict uncertainty was not an optimistic one since

according to him, you should always choose the alternative that gives the best result if the worst possible outcome will occur for each alternative. Hence the name the maximin utility criterion, which originated from Wald's work within game theory.

Hurwicz's rule has a less pessimistic approach compared to Wald. Hurwicz recommends a mixture of an optimistic and a pessimistic attitude:

- 1. Select a constant  $\alpha \in [0, 1]$  as the pessimism-optimism index.
- 2. Let  $o_i = \max\{v_{ij}, j = 1,...,n\}$  and  $p_i = \min\{v_{ij}, j = 1,...,n\}$ .
- 3. Choose  $a_k$  such that  $\alpha \cdot p_k + (1 \alpha) \cdot o_k = \max \{\alpha \cdot p_i + (1 \alpha) \cdot o_i\}$ .

Note that if  $\alpha = 1$  this is again the maximin utility criterion, whereas if  $\alpha = 0$ , it is the so-called maximax utility criterion. Different ways of choosing appropriate pessimism-optimism indices have been presented, but we will not enter into that discussion here.

In Savage's rule, the decision-maker should choose the alternative giving the smallest possible "regret".

- 1. Let  $r_{ii} = \max\{v_{si}, s = 1,...,m\} v_{ii}$ .
- 2. Let  $p_i = \max\{r_{ij}, j = 1,...,n\}$ .
- 3. Choose  $a_k$  such that  $p_k = \min\{p_i\}$ .

This minimax risk criterion was first suggested as an improvement over Wald's maximin utility criterion. Regrets and security levels will return later. Table 1 shows a counter-example (Milnor, 1954, p.50) of a decision problem where all of the above decision rules yield different results.

	$\mathbf{s_1}$	$s_2$	s <sub>3</sub>	$s_4$	Rule picks alternative
a <sub>1</sub>	2	2	0	1	Laplace
a <sub>2</sub>	1	1	1	1	Wald
a <sub>3</sub>	0	4	0	0	Hurwicz ( $\alpha > \frac{1}{4}$ )
a <sub>4</sub>	1	3	0	0	Savage

Table 1. Milnor's counter-example

The question remains: to act rationally, which one of the above rules should be employed? Milnor showed that no decision criterion is compatible with ten seemingly reasonable axioms that constituted his test set (Milnor, 1954, p.53). It turns out that it is relatively easy to show that it is impossible to find a decision rule that fulfils all desirable properties. Further, Ackhoff (1962) argues that any concept of strict uncertainty is inappropriate, i.e., strict uncertainty implies that there is always some information or some beliefs being disregarded.

### **Bayesian Decision Analysis**

When the decision-maker is able to quantify her beliefs in terms of a probability distribution on the set of possible outcomes given a chosen course of action, it is said that the decision is made under risk. If all utilities and probabilities in a decision problem are subjectively assigned numerical values by the decision-maker, and then the problem is evaluated according to the principle of maximising the expected utility, the decision-maker conforms to Bayesian decision analysis. This kind of decision problem is our main concern in Part I.

The decision method is called Bayesian, named after an English clergyman named Bayes, due to the use of subjective probability assignments and the common procedure of updating the probabilities by employing Bayes' theorem. In this respect, the probabilities are treated subjectively as a statistical procedure that, in many cases, endeavours to estimate parameters of an underlying probability distribution (posterior distribution) based on an observed probability distribution (prior distribution).

Suppose that each alternative a can be represented by a set of consequences and a set of numbers  $\langle \{c_i\}, \{p_i\} \rangle$ , where  $\{c_i\}$  is the set of possible consequences of a, and  $p_i$  is the probability that  $c_i$  occurs given that a is implemented. (Note here that probabilities are assigned to consequences instead of being assigned to states of the world. These two models are fully compatible when considering only a finite number of states and consequences.) Then, the meaning of accepting the utility principle and the principle of maximising the expected utility can now be formulated as follows (Malmnäs, 1994b):

**Definition:** If a is  $\langle \{c_i\}, \{p_i\} \rangle$ , and  $V_a$  is a real-valued function on  $\{c_i\}$ , then a has a value equal to  $\Sigma p_i V_a(c_i)$ , denoted by  $E_V(a)$ .

**Definition:** A decision-maker accepts the utility principle if and only if she assigns the value  $\sum p_i V_a(c_i)$  to a, given that it has assigned the value  $V_a(c_i)$  to  $c_i$ .

**Definition:** An ordering  $\geq_p$  of the alternatives is compatible with the principle of maximising the expected utility if and only if  $a \geq_p b$  implies  $E_V(a) \geq E_V(b)$ .

**Definition:** A decision-maker accepts the principle of maximising the expected utility if and only if its ordering of the values of the alternatives is compatible with that principle.

A survey of different interpretations of the utility principle and PMEU, as well as a more general characterisation of the class of expected utility models, is given in (Schoemaker, 1982, p.530 ff). An expected utility model is one that predicts or prescribes that people maximise the expression

$$\sum \Phi(p)U(x)$$
,

where x is an outcome vector. The models differ in i) how utility U(x) is measured, ii) what kind of concept of probability  $\Phi(p)$  is allowed, and iii) how the outcomes are measured. Schoemaker examines some frequently used variants of models, in accordance with this structure.

Utility theory was, even after taking Menger's results into account, not a well-founded subject until the late 1930s, when the works of Ramsey and von Neumann and Morgenstern appeared. They proposed reasonable principles governing decisions, from which a set of axioms was formulated whose purpose was to justify their particular attitude towards the utility principle. Surveys over a wide variety of axiomatisations are given in, e.g., (Fishburn, 1981; Malmnäs, 1994b), of which this chapter follows the latter.

The idea is to in a systematic way define the meaning of rationality. The point is, if a decision rule can be deduced from an indisputable axiomatisation, then this rule should be the natural and obvious rule for a rational entity, provided that the necessary information is available. Føllesdal (1984, p.268) suggests the following conditions for a decision rule:

- A decision rule should recommend an alternative with valuable consequences before an alternative with less valuable consequences.
- A decision rule should recommend an alternative with a high probability of

valuable consequences before an alternative with a low probability of valuable consequences.

 A decision rule should recommend an alternative with a low probability of bad consequences before an alternative with a high probability of bad consequences.

This seems to be reasonable but is too vague to fill the needs of a prescriptive decision theory and has to be elaborated a bit. In this, we introduce some axiomatisations using the following notation:

 $a >_p b$  means that the decision-maker holds alternative a to be strictly preferred to alternative b. This binary relation is *transitive* and *asymmetric*, thus it is a *strict order*.

 $a \ge_p b$  means that the decision-maker holds alternative a to be at least as good as alternative b, i.e., b is weakly preferred to a. This binary relation is *complete* and *transitive*, thus it is a *weak order*.

a  $\sim_p$  b means that the decision-maker is indifferent between alternative a and alternative b. This binary relation is *reflexive*, *transitive*, and *symmetric*, thus it is an *equivalence relation*.

If the decision-maker can assign a number u(a) such that  $u(a) \ge u(b)$  if and only if  $a \ge_p b$ , then it is said that there exists a *utility function* over a and b.

The axiom systems that will be presented consist of primitives and axioms constructed from the primitives. Typical primitives include states, sets of states, and ordering relations such as  $\geq_p$ . The axioms then imply a numerical representation of probabilities and preferences, i.e., the axioms imply the existence of a probability distribution and a utility function. Although Ramsey (1931) and von Neumann and Morgenstern (1944) are credited for the axiomatic foundation of utility theory, this book follows the axiom system of Luce and Raiffa (1957), very similar to the aforementioned, followed by the axiomatic justification of the utility principle according to Savage (1972). At first glance, the two systems seem dissimilar, but the important implications boil down to the same central results. Starting with Luce and Raiffa, in which alternatives (or gambles) with uncertain outcomes are called

lotteries: An alternative is denoted  $\langle p_1 \cdot v_1, ..., p_i \cdot v_i, ..., p_r \cdot v_r \rangle$ , which can be considered as a lottery with the probability  $p_i$  for the outcome  $v_i$ . All the probabilities are supposed to sum up to one. For example, the alternative a with uncertain outcomes  $v_1$  and  $v_2$  associated with probabilities  $p_1$  and  $(1-p_1)$  respectively is represented as the lottery  $a = \langle p_i \cdot v_1, (1-p_i) \cdot v_r \rangle$ .

**Axiom 1:** Ordering of alternatives and transitivity: For any two alternatives a and b, either  $a \ge_p b$  or  $b \ge_p a$ , and if  $a \ge_p b$  and  $b \ge_p c$  then  $a \ge_p c$ .

**Axiom 2**: Reduction of compound lotteries: Any compound lottery (which may be thought of as a mixture of lotteries, i.e., the prize of a lottery consists of another lottery instead of a certain reward.) is indifferent to a simple lottery with  $v_1, v_2, ..., v_r$  as prizes, in which the probabilities for the prizes in the simple lottery is computed according to ordinary probability calculus.

**Axiom 3**: Continuity: Each prize  $v_i$  is indifferent to some lottery involving just  $v_1$  and  $v_r$ . Thus, there exists some number (or probability)  $p_i \in [0,1]$  such that  $v_i \sim_p \langle p_i \cdot v_1, 0 \cdot v_2, ..., 0 \cdot v_{r-1}, (1-p_i) \cdot v_r \rangle$ .

**Axiom 4**: Substitutability (independence of irrelevant alternatives): In any lottery L,  $v_i$ ' is substitutable for  $v_i$ , that is,  $\langle p_1 \cdot v_1, ..., p_i \cdot v_i, ..., p_r \cdot v_r \rangle \sim_p \langle p_1 \cdot v_1, ..., p_i \cdot v_i', ..., p_r \cdot v_r \rangle$  when  $v_i' \sim_p v_i$ .

**Axiom 5**: *Monotonicity*:  $\langle p_i \cdot v_1, (1-p_i) \cdot v_r \rangle \ge_p \langle p_i' \cdot v_1, (1-p_i') \cdot v_r \rangle$  iff  $p_i \ge p_i'$ .

Note that nothing is being explicitly said about the origin of the probability distributions, they are just assumed to exist, and thus the view on probabilities is of the objective kind. From these axioms, the principle of maximising the expected utility as well as some other important results in utility theory are readily derived.

Savage argues that if utility is regarded as affecting only consequences (rather than acts), then for a weakly ordered consequence set C, the following is valid:  $\varpi_1(x)$  and  $\varpi_2(x)$  are numerical order-preserving functions representing the ordering relation between the consequences if and only if there is a strictly increasing function r such that, for every  $c_i \in C$ ,  $\varpi_1(c_i) = r(\varpi_2(c_i))$ . This shows that  $\varpi_1(c_i)$  is just an ordinal scale: it cannot be interpreted as quantitatively measuring the strength of preferences in any meaningful way. Savage adopted this argument from Pareto (1848–1923). The primitives building up the axiom system of Savage slightly differ

from the ones of Luce and Raiffa. Savage proposes the following: i) the binary preference relation  $\geq_p$ , ii) a set  $S = \{s_1, s_2, ...\}$  of states, iii) a set  $C = \{c_1, c_2, ...\}$  of consequences, and iv) a set  $F = \{f: S \to C\}$  of all possible mappings from S to C where such a mapping is called an act. Now, Savage defines E as the power set of S, where the elements of E are called events denoted by A, B, C, ... and further defines the following concepts:

- 1. For  $f,g,f',g' \in F$  and  $B,B^c \in E$ ,  $f \leq_p g$  given B if and only if  $f' \leq_p g'$  for every f' and g' that agree with f and g respectively, on B, and with each other on  $B^c$  and also  $g' \leq_p f'$  either for all such pairs or for no such pair (where  $B^c$  is the complement of B).
- 2.  $c_i \le_p c_j$  if and only if  $f \le_p f'$  when  $f(s) = c_i$  and  $f'(s) = c_j$ , for all  $s \in S$ .
- 3. B is null (B =  $\emptyset$ ) if and only if  $f \le_p g$  given B, for all  $f,g \in F$ .
- 4. A is not more probable than B  $(A \le B)$  if and only if  $f_A \le_p f_B$  or  $c_i \le_p c_j$ , for every  $f_A, f_B, c_i, c_j$  such that  $f_A(s) = c_i$  for  $s \in A$ ,  $f_B(s) = c_j$  for  $s \in A^c$ ,  $f_B(s) = c_i$  for  $s \in B$ ,  $f_B(s) = c_j$  for  $s \in B^c$ .
- 5.  $f \le_p c_i$  given B ( $c_i \le_p f$  given B) if and only if  $f \le_p h$  given B ( $h \le_p f$  given B), when  $h(s) = c_i$ , for all  $s \in S$ .

In the first concept, when act f' agrees with act f on B, then performing f will yield the same consequence as performing f' given the event (set of states) B, thus f(s) = f'(s) for all  $s \in B$ . The third concept says that if weak preference holds regardless of which pair of acts compared given the event B, implying that all acts are equal given B, then B is an empty set of states (and vice versa). The fourth concept: When an act  $f_B$  given A is preferred to an act  $f_A$  given A, and A given A is preferred to A and A of A given A is preferred to A and A is preferred to A and A is preferred to A and A given A is preferred to A and A is preferred

**Axiom 1**: *Transitivity*: The relation  $\leq_p$  is a weak order.

**Axiom 2**: *Completeness*: For every f, g, and B,  $f \le_p g$  or  $g \le_p f$  given B.

**Axiom 3**: Resolution independence: If  $f(s) = c_i$ ,  $f'(s) = c_j$ , for every  $s \in B$ ,  $B \neq \emptyset$ , then  $f \leq_p f'$  given B if and only if  $c_i \leq_p c_j$ .

**Axiom 4**: *Qualitative probability*: For every  $A, B \in E$ ,  $A \le B$  or  $B \le A$ .

**Axiom 5**: *Minimal strict preference*: It is false that for every  $c_j$ ,  $c_j$ ,  $c_i \le_p c_j$ .

**Axiom 6**: Continuity: Suppose  $h \le_p g$ , then for every  $c_i$  there is a finite partition  $\{B_i\}$  of S such that, if  $g' = c_i(B_i)$ , and  $h' = c_i(B_i)$ , for some i, then  $h \le_p g'$  or  $h' \le_p g$ .

**Axiom 7**: *Dominance*: If  $f \le_p g(s)$  given  $B(g(s) \le_p f$  given B) for every  $s \in B$ , then  $f \le_p g$  given  $B(g \le_p f$  given B).

The second axiom says that when two acts have the same consequences, the relation between f and f' must be independent of states. Furthermore, the third axiom says that the knowledge of an event cannot discard any preference between two consequences. Together, axioms 2 and 3 constitute Savage's debated sure-thing principle. Informally, if a decision-maker does not prefer f to g, either knowing that the event B occurred or knowing that B has not occurred, then the decision-maker does not prefer f to g (Savage, 1972, p.21). Further, from axiom 3 follows that preferences between acts depend only on realised consequences, and not possible ones.

The fourth axiom says that  $\leq$  is a qualitative probability, thus  $\leq$  is a weak order, and  $B \leq C$  if and only if  $(B \cup D) \leq (C \cup D)$  when  $(B \cap D) = (C \cap D) = 0$ . Furthermore,  $0 \leq B$ , 0 < S (all events are at least as probable as the impossible event and the universal event S must not be regarded as impossible). Axiom 5 says that there is at least one pair of consequences such that one is strictly preferred to the other, and axiom 6 implies the existence of a unique probability measure P on E. This probability measure is consistent with the qualitative probability in that E is not more probable than E' if and only if  $P(E) \leq P(E')$ . The last axiom says that if  $f \leq_p g(s)$  for all consequences of f for a set of states B, then  $f \leq_p g$ , if one of those states occurs, of further importance this axiom implies that the utility function is bounded (nothing is infinitely bad or infinitely good).

Given these assumptions, Savage proved the existence of a real-valued utility function on C with the following property: Let  $\{L_i\}$  be a partition of S and let f be an act with consequences  $\{f(s_i)\}$  on  $\{L_i\}$ , and let  $\{L_i'\}$  be another partition of S and let g be an act with consequences  $\{g(s_i)\}$  on  $\{L_i'\}$ . Then  $f \leq_p g$  if and only if  $\Sigma p_i \cdot u(f(s_i)) \leq \Sigma q_i \cdot u(g(s_i))$  where  $p_i = P(L_i)$  and  $q_i = P(L_i')$ , i.e., the principle of maximising the expected utility (PMEU).

Looking back at the system of Luce and Raiffa, it has been proved by von Neumann and Morgenstern (1944) that if a decision-maker has preferences between lotteries, i.e., given that the assumptions in the axiom system are fulfilled, then there is a real-valued utility function, unique up to a positive affine transformation, on the set of lotteries. Furthermore, let  $L_c = \{L_1, L_2, ...\}$  be a set of lotteries on C (alternatives with uncertain outcomes in the consequence set C), then they showed that the utility function  $u:L_c \to R$ , has a representation  $u(L_i) = \sum p_i(c_i) \cdot u(c_i)$  and  $L_i \leq_p L_j$  if and only if  $u(L_i) \leq u(L_j)$ . Thus, both axiom systems serve as attempts at a formal justification of the utility principle and the principle of maximising the expected utility. Due to the subjective nature of Savage's approach, his theory is often referred to as subjective expected utility.

#### **Descriptive decision theory**

Human decision-makers tend to, under given circumstances, behave inconsistently with the utility principle. Famous so-called paradoxes include Allais' paradox and Ellsberg's paradox. Allais' paradox shows that people tend to act in ways inconsistent with the sure-thing principle. This paradox derives from a common human behaviour of preferring a good outcome for certain to having a chance between something not as good and something even better. Ellsberg's paradox is quite similar, while it shows people's tendencies towards preferring known risks to unknown uncertainties, and thereby violating the utility principle.

Paradoxes of these kinds are often resolved by arguing that even intelligent beings make mistakes, and after some explanation of the inconsistency in their choices, they change their minds. However, for instance, an empirical study by Slovic (1974) has shown that as much as about 30% refuse to change their opinion and conform to the utility principle even after having had their errors pointed out to them. Tversky (1981) tries to understand why this is the case, and he concludes that irrelevant contextual effects often influence people, making them act inconsistent with the utility principle, i.e., the framing process. Further, it can be argued that no normative theory of decision making can embrace all inherent peculiarities in a free world of heterogeneous decision-making inhabitants.

However, this perspective has been critiqued. A common descriptive counterargument is the suggestion that the axioms of utility theory are inherently flawed. For instance, it has been shown that people do not always behave according to certain independence axioms in the system proposed by (Savage, 1954/1972; Allais, 1953). A more serious issue with the formal justifications of the utility principle from a normative point of view is that even if the axioms in various systems are accepted, the principle itself does not necessarily follow; in other words, the axiomatic systems are seemingly too weak to imply utility theory and PMEU. This is addressed in (Malmnäs, 1994) who demonstrates the weaknesses of the systems in (Herstein and Milnor, 1953; Oddie and Milne, 1990; Savage, 1972). A comprehensive review of numerous such systems is provided in (Malmnäs, 1994), who argues that it is implausible for these systems to be extended in any reasonable way to imply PMEU. Therefore, from a purely normative viewpoint, the logical foundations of utility theory appear to be quite weak. But without serious contenders, it is still a viable basis for prescriptive decision analysis, keeping this in mind.

Another criticism is that utility theory is inadequate for modelling risk attitudes effectively. Proponents of utility theory often argue for the concept of a risk premium to demonstrate that utility theory captures varying risk attitudes (French, 1986). However, the use of a utility function to model all possible risk attitudes is inherently limited. Critics argue that many decision-analytic models oversimplify the problem and ignore important factors (cf., e.g., (Schoemaker, 1982). For instance, even if the evaluation of an alternative yields an acceptable expected utility, its consequences might be so undesirable that the alternative should be avoided entirely, even if the probabilities of such consequences are very low. In such cases, PMEU would need to be extended with additional functionality. It has been suggested that a viable decision theory should allow for a broader range of risk attitudes and provide decision-makers with means to express these attitudes in various ways plus offer procedures for managing both qualitative and quantitative aspects.

Some researchers have in vain sought to modify the behaviour of PMEU by incorporating regret or disappointment into the evaluation, especially for cases where numerically identical outcomes are perceived differently depending on the decision-maker's previous experiences. See Chapter 4 for a discussion on such attempts. However, Malmnäs has demonstrated that, at best, these modifications result in performances nearly equivalent to that of expected utility, and at worst, being incon-

sistent with first-order stochastic dominance (Malmnäs, 1996). The apparent problem here is that the discussion emanates from a normative point of view, and in such a setting, the problem never ends. But from a prescriptive point of view, the focus is instead on finding guiding rules of the best kind, and Malmnäs' observation paves the way for a solid prescriptive approach.

Defenders of classical Bayesian decision theory instead argue that the concept of utility captures different risk attitudes. The assumption is that to each expected utility, there corresponds a certainty monetary equivalent  $x_{ce}$ . The decision-maker is indifferent between having this monetary value with certainty and performing an alternative with uncertain outcomes, i.e.,  $u(x_{ce}) = \sum p_i u(x_i)$ , where  $u(x_i)$  is the utility of gaining the monetary value  $x_i$ . The risk premium, r, of an act is now defined as the demand that a decision-maker has for carrying out the act, instead of having the monetary equivalent  $x_{ce}$  for certain, i.e.,  $r = \sum p_i x_i - x_{ce}$ . With respect to the premium r, a classification of decision-makers into three classes can be made: a decision-maker is risk-averse if r > 0; risk-prone if r < 0; and risk-neutral if r = 0.

As an example, assume that a decision-maker is in desperate need of a certain amount of money, and any lesser amount than this amount would not be useful. For instance, a person may need money for medical treatment of a disease that, if not cured, will result in death. If this person should seize the opportunity of entering a bet with their last funds that will give them a chance of winning an amount sufficient enough for the treatment to be affordable, this person would probably not be labelled irrational. In this situation, the risk premium r is probably negative.

With the foundations of Bayesian decision theory in place, we next explore methods for evaluating such expressions involving probabilities and utilities. The objective is to establish systematic and transparent approaches for ranking alternatives, thereby providing consistent and well-founded guidance to decision-makers.

This chapter builds on (Danielson, 1997, Ch.1)

# 04. Decision-Analytic Evaluation

In this chapter, situations where in addition the decision-maker has some estimates of the probabilities of the states involved are discussed. Usually, the probabilities are not the same for each alternative as in Laplace's rule, traditionally called decision under risk. Any decision problem under risk can be transformed into a problem in *normal form*. Further, tree and matrix forms of presenting a decision problem are equivalent. Therefore, it is sufficient to handle decision problems in normal form. In this chapter, a decision problem will be modelled in a decision frame.

**Definition:** Given a decision situation with m alternatives  $(A_1,...,A_m)$ , each with  $m_i$  consequences, and statements about the probabilities and values of those consequences. A *decision frame* is a structure  $\langle C,P,V\rangle = \langle \{\{C_{ik}\}_{m_i}\}_m,P,V\rangle$  containing the following representation of the situation:

- For each alternative  $A_i$  the corresponding consequence set  $\{C_{ik}\}_{k\in K_i}$  for  $K_i=\{1,\ldots,m_i\}$ .
- A set P of inequalities representing all probability statements.
- A set V of inequalities representing all value statements.

A large set of evaluation functions is the family of all functions that assign a numerical value to a consequence set for subsequent comparison, see for example (Schoemaker, 1982) for an overview. Such an evaluation function results in numeric values ranking the alternatives (or, more precisely, the consequence sets).

**Definition:** Given a decision frame  $\langle \{\{C_{ik}\}_{m_i}\}_{m}, P, V \rangle$  and a function f, the *numeric value*  $N(C_i)$  of a consequence set  $\{C_{ik}\}_{m_i}$  is  $f(p_{i1}, ..., p_{im_i}, v_{i1}, ..., v_{im_i})$ , a function over all consequences  $C_{ik}$  in the consequence set.

To be reasonable, the value of  $N(C_i)$  should range over the interval [0,1] since the values range over that interval. Of the numeric values, the expected value seems to be one of the most natural rules to apply to a decision problem on alternative-consequence-form. This is partly because the expected value  $E(C_i)$  is established in mathematical statistics, where it is employed as the mean value to be assigned to a stochastic variable taking on various values with specific probabilities.  $E(C_i)$  is clearly an instance of  $N(C_i)$  above. In this book, only discrete probability distributions are considered, and thus the following definition of the expected value applies.

**Definition:** Given a decision frame  $\langle \{C_i\}_{m}, P, V \rangle$ , the *expected value*  $E(C_i)$  of a consequence set  $C_i = \{C_{ik}\}_{m_i}$  is the sum  $\sum_k p_{ik} \cdot v_{ik}$  over all consequences  $C_{ik}$  in the set.

The use of the principle of maximising the expected value (PMEV) dates several hundred years back, preceding the formal area of mathematical statistics and instead originating from pure monetary gambling. Over the years, a number of problems have been discovered with the principle when applied to real-life decision situations. A serious paradox was first suggested by Allais (1953), and other paradoxes along the same line have subsequently been suggested. Many people tend to choose alternatives in a way that seems to violate the PMEV, no matter what utility values are assigned to the respective outcomes. See for example (Savage, 1972) for a mathematical argument regarding Allais' paradox. In experiments where the violation was afterwards pointed out to subjects who understood the mathematical argument, up to 1/3 retained their choice despite this.

Such problems with PMEV warrant further investigation, and several researchers, not least within economics, have proposed a number of alternative decision rules to replace (or sometimes supplement) the PMEV. Fishburn (1983) suggests an evaluation based on the quotient between two separate expected values, which has the form

$$\frac{E(C_i, f_1)}{E(C_i, f_2)}$$

where  $f_1$  and  $f_2$  are two functions of the values involved.

Loomes and Sudgen (1982) bring regret or disappointment into the evaluation to cover cases where numerically equal results are appreciated differently depending on what was once in someone's possession. Their suggested formula has the form

$$\sum_{k=1}^{n} p_{ik} \cdot (v_{ik} + R(v_{ik} - E(C_i)))$$

where *R* is supposed to be a regret function related to the ordinary expected value.

Quiggin (1982) tries to resolve the problem by requiring functions to modify the probabilities in the evaluation rule such as

$$\sum_{k=1}^{n} (f(s_{ik}) - f(s_{i(k-1)})) \cdot v_{ik}$$

where f is a strictly increasing function, the  $s_{ij}$ 's are in increasing  $v_{ij}$  order, and  $s_{ik} = \sum_{i=1}^{k} p_{ij}$ . Yaari has pointed out that under certain reasonable assumptions

(Yaari, 1987), it must be the case that  $f(p_{ij}) = p_{ij}$  and then he made the following extended suggestion

$$\sum_{k=1}^{n} (f(1-s_{i(k-1)}) - f(1-s_{ik})) \cdot v_{ik} + f(p_{im_i}) \cdot v_{im_i}$$

where  $s_{ii}$  is as above.

As noted in Chapter 3, Malmnäs (1996) shows for those above and for other proposals that their performances can at best be equal to that of the expected value and at worst are much poorer, for example not even being consistent with first order stochastic dominance. Since no rule performs consistently better than the expected value, it is the only possible rule from a prescriptive viewpoint. It has sometimes been argued that the prescriptive approach consists of selecting axioms to adhere to, rather than accepting and using the axiom systems of established theories (Keeney, 1992). Such a view would reduce prescriptive decision analysis to meta-arguments on which axiomatic results to believe in and adhere to, and which to dismiss. However, that would constitute a road that does not lead to better tools for real-life decision support.

In many decision contexts, the decision-maker may want to exclude particular alternative courses of action that are, in some sense, too risky. If the PMEU modifications on the previous pages do not work, what does? The exclusion can be achieved by a class of supplementary decision rules called qualitative sorting or security levels. While an evaluation of a consequence set may result in an acceptable expected value, the consequences of selecting it might be so dire that it should nevertheless be avoided. It might, for example, endanger the entire purpose of the decision context, and in that case, even a consequence with a low probability is too risky to neglect. Such exclusions can be dealt with by specifying a security level for

the probability and a threshold for the value. Then a consequence set would be undesirable if it violates both of these settings. Malmnäs' proposal (1994a) is to supplement the expected value with qualitative evaluations. An example is the qualitative sorting function, which has the basic form

$$S(C_i, r, s) = (\sum_{v_{ij} \le r} p_{ij} \le s)$$

where r is the minimally tolerable value threshold and s is the maximally acceptable probability for events below the threshold to occur. This is a boolean function sorting out unwanted consequence sets. But to treat this and other supplements, a more general discussion on dominance is required.

#### **Delta Dominance**

In this section, a general dominance rule is suggested as a unifying concept. In its generic form, it describes the type of dominance to be considered and thus also the type and amount of computation involved in evaluating consequence sets in the framework. It includes all of the above-suggested evaluation functions, even though the expected value is by far the most common. For convenience, a shorthand notation for the difference in expected values is introduced.

**Definition:** Given a decision frame  $\langle \{\{C_{ik}\}_{m_i}\}_m, P, V \rangle$ ,  $\delta_{ij}$  denotes the expression  $E(C_i) - E(C_j) = \sum_k p_{ik} \cdot v_{ik} - \sum_k p_{jk} \cdot v_{jk}$  over all consequences in the consequence sets  $C_i$  and  $C_j$ .

**Terminology:** Given a decision frame  $\langle C,P,V \rangle$ , the functions f, g, and h are specified as  $f:\mathbb{R}^i \to [0,1]$ , g:  $\mathbb{R}^j \to [0,1]$ , and h:  $\mathbb{R}^k \to [0,1]$  with i,j,k  $\in \mathbb{N}_+$  as appropriate. The  $\alpha$  and  $\beta$  parameters are real numbers in the range [0,1].

In order to describe the dominance, a couple of concepts are required. The index set pair captures the consequences within each of the consequence sets that possess some desired property, in this case their value being at least as great as a given parameter.

**Definition:** Given a decision frame  $\langle C,P,V \rangle$  and a real number  $d \in [0,1]$ , an *index* set pair  $(K_i,K_j)(d)$  is  $K_i = \{k \mid v_{ik} \geq d\}$  and  $K_j = \{k \mid v_{jk} \geq d\}$ .

When the parameter d varies over some range, the content of the index set may

vary as well. This represents a selection procedure for selecting all consequences within a pair of consequence sets with a desired property. The set of all such index sets is defined next.

**Definition:** Given a decision frame  $\langle C,P,V \rangle$  and real numbers  $a,b,d \in [0,1]$ ,  $M_{ii}[a,b]$  is the set  $\{(K_i,K_i)(d) \mid d \in [a,b]\}$ .

 $M_{ij}[a,b]$  is the set of all different index set pairs in the range [a,b], i.e. all the combinations of index sets that satisfy any threshold condition in that range. Those two definitions enable the following compact definition of the  $\Delta$ -dominance. The idea behind the dominance is a pairwise comparison of the consequence sets employing the desired numerical function. Note that the weak inequality must hold for all index set members, i.e. over the entire interval range I.

**Definition:** Given a decision frame  $\langle C,P,V \rangle$ , a function f, and two parameters  $\alpha(P_0,V_0)$  and  $\beta(P_0,V_0)$ ,  $C_i \Delta[I]$ -dominates  $C_j$  iff

$$\forall (K_i, K_j)(d) \in M_{ij}[I] \sum_{k \in K_i} f(p_{ik}, v_{ik}, \alpha) - \sum_{k \in K_j} f(p_{jk}, v_{jk}, \beta) \ge 0$$
 and

$$\exists \; (K_i, K_j)(d) \in M_{ij}[I] \; \sum_{k \in K_i} f(p_{ik}, v_{ik}, \alpha) - \sum_{k \in K_j} f(p_{jk}, v_{jk}, \beta) > 0 \; .$$

This is a very general definition based on traditional admissibility concepts, and many instantiations are possible. In this book, a few are given and it is shown that some well-known evaluation concepts are special cases of  $\Delta$ -dominance. The first subdivision of the  $\Delta$ -dominance is into dominance orders depending on the function employed in the evaluation. First and second orders are specifically addressed below, while higher orders are not further discussed.

The  $\Delta$ -dominance is of the first order if the function used is a function of only probabilities. The values are not taken into account when evaluating the consequence sets.

**Definition:** Given a decision frame  $\langle C,P,V \rangle$  and functions f and g,  $C_i$  1[I]-dominates  $C_j$  iff  $C_i$   $\Delta$ [I]-dominates  $C_j$  with  $f(p_{ik},v_{ik},\alpha) = g(p_{ik})$  and  $f(p_{jk},v_{jk},\beta) = g(p_{jk})$ .

Thus, first order specialisation turns dominance into a difference of sums of a function of probabilities.

$$\textbf{Note:} \ C_i \ \textit{I[I]-dominates} \ C_j \ \textbf{iff} \ \forall \ (K_i, K_j)(d) \in M_{ij}[I] \ \sum_{k \in K_i} g(p_{ik}) \geq \sum_{k \in K_j} g(p_{jk}) \cdot \sum_{k \in K_j} g(p_{ik}) = \sum_{k \in K_j} g(p_{ik}) \cdot \sum_{k \in K_j} g(p_$$

The note points out the resemblance with some familiar dominance concepts. One further specialisation of the first order  $\Delta$ -dominance is the first order stochastic dominance, a well-known concept. To reach there, the general first order  $\Delta$ -dominance is considered. It consists of specifying the range for the index set pairs to be the full [0,1] range.

**Definition:** Given a decision frame  $\langle C,P,V \rangle$ ,  $C_i$  1S-dominates  $C_j$  **iff**  $C_i$  1[0,1]-dominates  $C_j$ .  $C_i$  1SE-dominates  $C_j$  iff  $C_i$  1S-dominates  $C_j$  with  $g(p_{ik}) = p_{ik}$ .

When the function g employed is the simple  $g(p_{ik}) = p_{ik}$  the general stochastic dominance turns into the commonly used first order stochastic dominance, which in the  $\Delta$ -dominance concept is a specialisation of function as well as of index set range. To see that this is indeed the ordinary first order stochastic dominance as claimed, it is convenient to make the following note, in which the form for 1SE-dominance coincides with the definition of first order stochastic dominance.

**Note:** 
$$C_i$$
 *ISE-dominates*  $C_j$  **iff**  $\forall$   $(K_i, K_j)(d) \in M_{ij}[I]$   $\sum_{k \in K_i} p_{ik} \ge \sum_{k \in K_j} p_{jk}$ .

Earlier, a supplementary function was mentioned under the name of qualitative sorting or security levels. This was a kind of threshold function separating wanted and unwanted outcomes (or desirable and undesirable consequence sets) according to a threshold rule applicable to the evaluation situation. This type of evaluation rule also turns out to be a special case of the  $\Delta$ -dominance, viz. the dominance of a reference consequence set, i.e. the threshold.

**Definition:** Given a decision frame  $\langle C,P,V \rangle$  and two real numbers  $s,t \in [0,1]$ ,  $C_j$  violates general security level s for threshold value t **iff**  $C_t$  1[t,t]-dominates  $C_j$ , where  $C_t$  is a consequence set with two consequences,  $g(p_{t1}) = 1 - g(s)$ ,  $v_{t1} = 1$ ,  $g(p_{t2}) = g(s)$ ,  $v_{t2} = 0$ .

When the function g is the simple  $g(p_{ik}) = p_{ik}$ , then the general security level turns into the ordinary security level concept, which again is a specialisation of both function and index set range.

**Definition:** Given a decision frame  $\langle C,P,V \rangle$  and two real numbers  $s,t \in [0,1]$ ,  $C_j$  violates security level s for threshold value t iff  $C_j$  violates general security level s for threshold value t with  $g(p_{jk}) = p_{jk}$ .

To see that this is indeed the same concept as the security levels discussed above, the following observation can be helpful. Note that there can only be one index set pair since the range of the value interval only contains r.

**Note:**  $C_j$  violates security level s for threshold value t **iff** for  $K_j = \{k \mid v_{jk} \ge t\}$ 

$$\sum_{k \in K_j} p_{jk} \leq 1 - s \cdot$$

It can be seen that the first-order stochastic dominance and qualitative sorting or security levels are both variants of the same concept of first-order  $\Delta$ -dominance.

The  $\Delta$ -dominance is of the second order if the function used is a function of both the probabilities and the values.

**Definition:** Given a decision frame  $\langle C,P,V \rangle$  and functions f and h,  $C_i$  2[I]-dominates  $C_j$  iff  $C_i$   $\Delta[I]$ -dominates  $C_j$  with  $f(p_{ik},v_{ik},\alpha) = h(p_{ik},v_{ik})$  and  $f(p_{jk},v_{jk},\beta) = h(p_{jk},v_{jk})$ .

Then the domination turns into a difference of sums of a function of probabilities and values.

**Note:**  $C_i$  2[I]-dominates  $C_i$  iff  $\forall$   $(K_i,K_i)(d) \in M_{ii}[I]$ 

$$\sum_{k \in K_i} h(p_{ik}, v_{ik}) \ge \sum_{k \in K_j} h(p_{jk}, v_{jk}) \cdot$$

As for the first order, a further specialisation into second-order stochastic dominance is possible. This is a well-known concept as well, and it turns out to be another case of  $\Delta$ -dominance. First, the general second-order stochastic dominance is defined. As in the first order case, it consists of specifying the range for the index set pairs to be the full [0,1] range.

**Definition:** Given a decision frame  $\langle C,P,V \rangle$ ,  $C_i$  2S-dominates  $C_j$  iff  $C_i$  2[0,1]-dominates  $C_i$ .  $C_i$  2SE-dominates  $C_j$  with  $h(p_{ik},v_{ik}) = p_{ik} \cdot v_{ik}$ .

If the function h employed is the most common  $h(p_{ik}, v_{ik}) = p_{ik} \cdot v_{ik}$ , then the dominance turns into the commonly used second-order stochastic dominance, which in the  $\Delta$ -dominance concept is a specialisation both of function and of index set range. To see explicitly that we have arrived at the ordinary second-order stochastic dominance, it is helpful to make the following note, in which the form for 2SE-dominance can be seen to be almost equivalent to the textbook definition of second-order stochastic dominance.

**Note:** 
$$C_i$$
 2SE-dominates  $C_j$  **iff**  $\forall$   $(K_i, K_j)(d) \in M_{ij}[0,1]$   $\sum_{k \in K_i} p_{ik} \cdot v_{ik} \ge \sum_{k \in K_i} p_{jk} \cdot v_{jk}$ .

Another second order  $\Delta$ -dominance is the ordinary expected value and some of the suggested replacements. One of their characteristics is that they evaluate only by full index set pairs, i.e. pairs that contain all members of each consequence set. The general numerical dominance is a straightforward specialisation of  $2\Delta$ -dominance.

**Definition:** Given a decision frame  $\langle C,P,V \rangle$ ,  $C_i$  *N-dominates*  $C_j$  **iff**  $C_i$  2[0,0]-dominates  $C_j$ .  $C_i$  *NE-dominates*  $C_j$  **iff**  $C_i$  N-dominates  $C_j$  with  $h(p_{ik},v_{ik}) = p_{ik} \cdot v_{ik}$ .

This corresponds to the evaluation rules that apply a probability and value formula to the consequence set in order to reach a numerical verdict on which one is preferable. The last specialisation of the second order is the ordinary expected value, which is termed NE-dominance and is realised by letting  $f(p_{ik},v_{ik}) = p_{ik} \cdot v_{ik}$  in the N-dominance. This can be seen to be the expected value, since the only index set pair generated by the [0,0]-range is the pair of complete consequence sets.

**Note:**  $C_i$  *NE-dominates*  $C_j$  **iff** for  $(K_i, K_j)(0)$   $\delta_{ij} \ge 0$ .

Further, note that  $\delta_{ij} \geq 0$  does not apply to 2SE-dominance since it involves different index set pairs while NE-dominance always applies only to the full index sets of the consequence sets. It has been demonstrated that some well-known dominance rules and the ordinary expected value are special cases of  $\Delta$ -dominance, which acts as a unifying concept in comparing and discussing evaluation rules.

## 05. Realistic Input Information

In a vast majority of real-life decision situations, the decision-maker does not have access to the significant amount of statistical data demanded to aggregate precise numerical values and probabilities, nor does the decision-maker have the ability to perform precise estimations of utilities. Furthermore, people find it hard to distinguish between probabilities ranging from approximately 30% to 70% (Slovic, 1974). A great deal of attention has been given to problems of imprecise information as a source of decision uncertainty, Morgan and Henrion (1990) identify two main types of uncertainty. The first type of uncertainty derives from a lack of historical data and takes its form from statistical variation, subjective judgements, linguistic imprecision, variability, inherent randomness, disagreement and approximation. For example in experiments, errors in the measurements of quantities give rise to statistical variation. The second type of uncertainty arises from the model chosen. Furthermore, uncertainty due to biases in communication and value differences is unavoidable in the use of expertise in policy processes. Instead of addressing the sources of uncertainty, Funtowicz and Ravetz (1990) discuss different types of uncertainties, including inexactness (or technical uncertainty), unreliability (or methodological uncertainty), and border-with-ignorance (or epistemological uncertainty). These authors consider ignorance to be endemic to scientific research.

Even if a decision-maker is able to discriminate between different probability measures, very often adequate, reliable, and precise information is missing. Consequently, there seem to be significant reasons for discriminating between measurable and immeasurable uncertainty. Measurable uncertainty is often referred to as *risk* and can be represented by precise probabilities. In contrast, immeasurable uncertainty occurs frequently in high-consequence/low-frequency situations since the low frequency implies a lack of statistical data, and thereby the axiom systems given by, e.g., Savage and von Neumann and Morgenstern, are not satisfied. Ellsberg (1961) proposes a class of choice situations involving immeasurable uncertainty, in which the behaviour of people is inconsistent with the suggested axiomatic systems. He does not object to the use of the principle of maximising the expected utility (PMEU) but suggests that the underlying axiomatic systems should not be applied in situations where the available information is to some extent not precisely defined.

Doyle and Thomason (1997) present an approach where imprecision is being modelled by using only qualitative data. However, in many cases this restriction will yield a too narrow outlook of a decision problem, numerical estimates should still play a role.

There is a wide variety of mathematical models for the representation of imprecise probability. Most research in imprecise probabilities has been concerned with different types of upper and lower probability (Walley, 1991). However, some common and useful kinds of uncertainty cannot be modelled through the use of upper and lower probability models, especially, commonly used comparative statements of the form "A is at least as probable as B" cannot be allowed for. Walley's book *Statistical Reasoning with Imprecise Probabilities* introduces the concept of upper and lower previsions. Briefly speaking, the lower prevision of a gamble is defined by the amount a gambler is willing to pay for a lottery ticket, and the upper prevision is defined by how much he is willing to sell the same ticket for.

Many attempts have been made to express imprecise probabilities in terms of intervals. In (Choquet, 1953), the concept of capacities is introduced. These capacities can be used to define a framework for modelling imprecise probabilities as intervals (Huber, 1973). The use of interval-valued probability functions, by means of classes of probability measures, has also been integrated into classical probability theory by e.g., Good (1962) and Smith (1961). A similar approach was taken by Dempster (1967), where a framework for modelling upper and lower probabilities is investigated. This was further developed by his PhD student in (Shafer, 1976), where the concept of basic probability assignments was also introduced. The Dempster-Shafer theory for quantifying subjective judgements has received a lot of attention, but it seems to be unnecessarily strong with respect to interval representation (Weichselberger and Pöhlmann, 1990). Weichselberger's theory of interval probability instead argues in favour of an axiom system for interval probabilities clearly related to the one of Kolmogorov, i.e. an already established theory.

Imprecision in decision situations often prevails in both probability estimates and utility assessments. For example in business decisions when acting upon a forecast, the forecasted value often is subject to some forecast error encouraging the use of a prediction interval instead of a predicted fixed number which in almost every case will be more or less incorrect. Furthermore, many types of decisions involve utility

measures of non-monetary outcomes which then must be measured on some precisely defined interval scale, such measurements are often hard to motivate, e.g., due to underlying ethical responsibilities and democratic values.

When more than one probability distribution defined on the same set of outcomes is reasonable given the information obtained, we speak in terms of sets of probability distributions. The American philosopher Levi gives three conditions such sets of probability measures B must satisfy. These imply (among other things) that the probability distributions in B for a given state of nature form an interval, in literature such sets are commonly referred to as convex sets of probability measures. The significance of Levi's work is emphasised as Levi compares the different alternatives in decision situations. He gives an example in which two similar decision situations with different sets of probability measures yield results different from his theory, even if the generated intervals are the same (Levi, 1974, pp.416-418). He notices that some authors have presupposed such an interval in their theories, but concludes that his own theory "[...] recognises credal states as different even though they generate the identical valued function -provided they are different convex sets of Q-functions." The significance is emphasised as Levi compares the different alternatives in decision situations. He gives an example in which two similar decision situations with different sets of probability measures yield results different from his theory, even if the generated intervals are the same.

Levi also relaxes the Bayesian requirement on representing the utilities of the consequences. He introduces a set G of permissible utility functions, which do not obey the classical Bayesian requirement that all elements in G are linear transformations of each other. He then stipulates the following definitions:

**Definition:** An alternative A is E-admissible if and only if there is a probability distribution p in B and a utility function u in G, such that E(A), defined relative to p and u, is optimal among all alternatives.

**Definition:** An alternative A is *S-admissible* if and only if it is E-admissible and there is a function u in G such that the minimum u-value assigned to some possible consequence is at least as great as the u-values assigned to the consequences of any other of the remaining alternatives. These definitions seem reasonable, but they have some counter-intuitive implications. They clearly violate the reasonable condition of independence of irrelevant alternatives, i.e. that the ordering between the

alternatives is not affected by the addition of a new alternative. The theory is also problematic in some respects when confronted with some empirical results.

In (Danielson, 1997), another approach is suggested. Imprecise probabilities, as well as imprecise utilities, are handled by modelling a decision situation with numerically imprecise sentences such as "the probability of consequence  $c_{11}$  is greater than 5%" and comparative sentences such as "consequence  $c_{11}$  is preferred to consequence  $c_{12}$ ". These kinds of sentences are represented by suitable intervals and comparisons. Sentences such as "the probability of  $c_{ij}$  lies between the numbers  $a_k$  and  $b_k$ " are translated to  $p_{ij} \in [a_k, b_k]$ . Similarly, sentences such as "the probability of  $c_{ij}$  is greater than the probability of  $c_{kl}$ ". are translated into inequalities such as  $p_{ij} > p_{kl}$ . In this way, each statement is represented by one or more constraints. The conjunction of all constraints together with  $\sum p_{ij} = 1$  for each alternative  $A_i$  is called the probability base (P). The utility base (V) consists of similar translations of utility estimates. The collection of probability and utility statements constitutes the decision frame. The following terminology and definitions are from (Danielson, 1997).

**Definition:** A decis*ion frame* with m alternatives is a structure  $\langle \{\{c_{ij}\}_{j=1,\dots,h_i}\}_{i=1,\dots,m},P,V\rangle$ , where each  $c_{ij}$  denotes a consequence. P is a finite list of linear constraints in the probability variables and V is a finite list of linear constraints in the utility variables.

Given such a structure, various decision rules can be applied. One such structure is a generalisation of the expected utility of an action. With respect to a decision frame, this can be expressed by the following definition.

**Definition:** Given a decision frame  $\langle \{\{c_{ij}\}_{j=1,\dots,h_i}\}_{i=1,\dots,m},P,V\rangle$ , the *expected utility*  $E(A_i)$  of an action  $A_i$  is  $E(A_i) = \sum_{k \le h_i} p_{ik} \cdot u_{ik}$ , where  $p_{ik}$  and  $u_{ik}$  are variables in P and V, respectively.  $u_{ij}$  denotes the utility of the consequence  $c_{ij}$ , and  $p_{ij}$  denotes the probability of  $c_{ij}$  occurring given that action  $A_i$  is taken.

**Definition:** Given a decision frame  $\langle \{\{c_{ij}\}_{j=1,\dots,h_i}\}_{i=1,\dots,m},P,V\rangle$ , let a and b be two vectors of real numbers  $(a_{i1},\dots,a_{ih_i})$  and  $(b_{i1},\dots,b_{ih_i})$  respectively. Then define  ${}^{ab}E(A_i) = \Sigma_{k \leq h_i} \ a_{ik} \cdot b_{ik}$ , where  $a_{ik}$  and  $b_{ik}$  are numbers substituted for  $p_{ik}$  and  $u_{ik}$  in  $E(A_i)$ .

If the expected utility in the definition above seems to be very similar to the expected utility as defined in the previous chapter, it is important to bear in mind that this is evaluated with respect to the solution sets of the decision frames rather than to precise numbers. Using precise numbers, evaluating the expected utility is straightforward. However, when numerically imprecise information is involved, the situation is a bit more intricate, i.e., the expected utility has to be evaluated with respect to the solution sets to the probability and utility bases. The solution set to a set of linear constraints L consists of vectors consistent with L.

**Definition:** Given a base expressed in the variables  $\{p_1,...,p_k\}$ . A list of numbers  $[n_1,...,n_k]$  is a *solution vector* to a base L if the substitution of  $n_i$  for  $p_i$ , for all  $1 \le i$   $\le k$ , in L does not yield a contradiction. The set of solution vectors to L constitutes the *solution set* for L.

With respect to the solution sets to the probability and utility bases, substituting all possible vectors  $(a_{i1},...,a_{ih_i})$  and  $(b_{i1},...,b_{ih_i})$ , consistent with the solution sets to the probability and utility bases, in the expected utility above, a range of possible values is received. Thus, by the introduction of interval in this way, the meaning of the expected utility is no longer clear, and a reasonable decision strategy must be defined. A quite uncontroversial strategy of evaluation is to never eliminate or disqualify an action that might be the best one. The only option then becomes never to eliminate any alternative, which might be considered too weak a decision strategy. Another strategy is to investigate the differences between the various alternatives.

**Definition:** Given a decision frame  $\langle \{\{c_{ij}\}_{j=1,\dots,h_i}\}_{i=1,\dots,m},P,V\rangle$ , the *difference in expected utility*  $\delta_{ij}$  between two alternatives  $A_i$  and  $A_j$  are  $\delta_{ij}=E(A_i)-E(A_j)$ . Similarly, define  ${}^{abcd}\delta_{ij}={}^{ab}E(A_i)-{}^{cd}E(A_j)$ .

Using this notation, we can introduce a variety of rules to discriminate between different actions. For instance, the concept of admissibility (64) is expressed in the following way.

**Definition:** Given a decision frame  $\langle \{\{c_{ij}\}_{j=1,\dots,h_i}\}_{i=1,\dots,m},P,V\rangle$ ,  $A_i$  is *at least as good* as  $A_j$  iff  ${}^{abcd}\delta_{ij}\geq 0$ , for all a,b,c and d,w here the expression  $\{p_{i1}=a_{i1}\}$  & ... &  $\{p_{ih_i}=a_{ih_i}\}$  &  $\{p_{j1}=c_{j1}\}$  & ... &  $\{p_{jh_j}=c_{jh_j}\}$  is consistent with P. Similarly,  $\{u_{i1}=b_{i1}\}$  & ... &  $\{u_{ih_i}=b_{ih_i}\}$  &  $\{u_{i1}=d_{j1}\}$  & ... &  $\{u_{jh_i}=d_{jh_i}\}$  is consistent with V.  $A_i$ 

is *better* than  $A_j$  iff  $A_i$  is at least as good as  $A_j$  and  ${}^{abcd}\delta_{ij} > 0$ , for some a, b, c, d, that is consistent with P and V as above.  $A_i$  is *admissible* iff no other  $A_j$  is better than  $A_i$ .

Intuitively, an action can be discarded if it is always worse than all other actions, i.e., an admissible alternative is in some sense a non-dominated alternative. The concept of admissibility is computationally meaningful in this framework. However, the imprecision represented in the decision frames, viz. most non-trivial situations, often results in the ranges of the expected utility of some actions overlapping. The set of admissible alternatives will therefore often be too large. Consequently, even if PMEU is employed, there is a need for further principles of discrimination. One way to proceed is to determine the stability of the relation between the actions under consideration. Values near the boundaries of the intervals are probably less reliable than more central values due to interval statements being deliberately imprecise. This can be taken into account by measuring the dominated regions indirectly with the use of the concept of contraction, which is motivated by the difficulties of performing sensitivity analyses in several dimensions simultaneously. It can be difficult to gain a real understanding of the solutions to large decision problems using only one-dimensional analyses since different combinations of dimensions can be critical to the results of evaluation.

In order to assess the overlap, sensitivity analyses of the admissibility are called for. The hull cut is a generalised sensitivity analysis for this purpose. It is reasonable to consider values near the boundaries of the intervals in a constraint set to be less reliable than more central values, due to interval constraints being deliberately imprecise. The core, on the other hand, represents the most reliable estimates. It is therefore desirable to be able to study the bases with varying cut rates, i.e. studying smaller or larger decrements to the orthogonal hull. If the core itself is not enough to yield the desired evaluation results, it can be further cut towards the focal point with varying degrees of contraction.

**Definition:** Given a base X in  $\{x_i\}_{i\in I}$ , a set of real numbers  $\{a_i,b_i\}_{i\in I}$ , a core  $[c_i,d_i]_n$  of  $\{x_i\}_{i\in I}$ , and a real number  $\pi\in[0,1]$ , a  $\pi$ -cut of X is to replace the core by  $[c_i+\pi\cdot(a_i-c_i),d_i+\pi\cdot(b_i-d_i)]_n$ . If the set  $\{a_i,b_i\}_{i\in I}$  is the hull  $\langle a_i,b_i\rangle_n$  then it is called a  $\pi$ -expansion of X. If  $(r_1,\ldots,r_n)$  is a focal point and  $a_i=b_i=r_i$ , then it is

called a  $\pi$ -contraction of X.

The  $\pi$ -cut is a linear procedure, but non-linear procedures are plausible as well. In addition, the procedure can work from either side ((L) $\pi$ -cut and (R) $\pi$ -cut) or with varying, even non-uniform rates of contraction. The cut structure is studied with respect to admissibility, i.e. at which cut rates admissibility is affected. If there is no verdict in the original core, it may be further cut towards the focal point in order to achieve a result.

Various kinds of sensitivity analyses based on the concept of contraction are suggested in (Danielson, 1997). By co-varying the contractions of an arbitrary set of intervals, it is possible to gain much better insight into the influence of the structure of the decision frame on the solutions. Contrary to, e.g., volume estimates, contractions are not measures of the sizes of the solution sets but rather of the strength of statements when the original solution sets are modified in controlled ways. Both the set of intervals under investigation and the scale of individual contractions can be controlled. The idea behind contractions is to investigate how much the intervals can be decreased before an expression such as  $E(A_i) - E(A_j) > 0$  ceases to be consistent. At the same time, we must avoid the complexity inherent in combinatorial analyses, but still be able to study the stability of a result.

It should be emphasised that the concept of admissibility is still based on PMEU, and thus the approach of considering only admissible actions cannot be entirely uncontroversial. Since the idea of dismissing a clearly inferior action seems to be reasonable, we must be careful about how to measure this inferiority.

One major drawback of the classic Bayesian approach as well as Levi's approach is that it does not account for variations of the epistemic reliability in different decision situations (Gärdenfors and Sahlin, 1982). Even if an outcome is associated with a set of probability measures and a set of utility measures, some of these measures are often regarded as more reliable than others, due to the nature of the obtained information. Thus, we have a second-order belief in the sense that we hold some of our beliefs to be more reliable.

The interval model requires defining a set of all epistemologically possible probability distributions within a decision context. However, a decision-maker may not assign equal confidence to all these distributions, necessitating a model of belief

strength in different vectors. A further refinement of the interval model can be achieved using distribution theory. This approach allows for differentiation among various probability distributions and utility functions by defining a global distribution that expresses various beliefs over sets of intervals. For each vector of probability estimates, a belief value is assigned to reflect the decision-maker's confidence in that particular distribution. This global distribution is defined over a polytope, a region of possible solutions described by linear inequalities. This model generalises the interval-based approaches discussed earlier, enabling a more flexible representation of beliefs in decision making. However, one major limitation is that decision-makers can rarely envision such high-dimensional distributions, especially in complex decision situations.

Gärdenfors and Sahlin (1982, 1983) address these issues by considering global belief distributions, though they focus primarily on the probability case. A limitation of this approach is its lack of exploration of the relationship between local and global distributions and the methods for ensuring the consistency of user-specified belief statements. For example, if a decision-maker considers a class of probability distributions, it is reasonable to assume that belief should be zero in vectors where the mapping does not sum to one. Hence, the belief in impossible outcomes should be zero, and this constraint must be consistent with the overall belief distribution.

In general, interval decision analysis conforms to traditional statistical reasoning by being compatible with the concept of admissibility. The emphasis in prescriptive decision theory is not on describing another formalism for representing imprecision but rather on presenting a way of handling the imprecision inherent in many reallife decision problems within standard decision theory.

This chapter builds on (Danielson, 1997, Ch.4)

## 06. Multiple Criteria

As discussed in detail in Part I, the roots of prescriptive decision theory can be traced to the mid-20<sup>th</sup> century with the development of utility theory and the axiomatic foundations of rational choice, notably by the works of von Neumann and Morgenstern, Savage, and others. The classical expected utility theory, underpinned by axioms such as completeness, transitivity, independence, and continuity, represents the ideal of rational behaviour under uncertainty. Probabilistic (Bayesian) decision analysis, which builds directly upon this foundation, involves the modelling of uncertainty through probability distributions and the quantification of preferences via utility functions. Decision trees, influence diagrams, and Bayesian updating are among the standard tools employed in this tradition. These methods are particularly powerful when uncertainty can be meaningfully represented probabilistically and when the decision-maker's utility function can be elicited and incorporated into the analysis.

However, the limitations of classical probabilistic approaches for real-life decision analysis have long been recognised. In practice, many decision situations involve multiple objectives. Among the most significant developments to address these challenges is the emergence of multi-criteria decision analysis (MCDA). It encompasses a set of methods designed to support decision making in contexts where multiple, often conflicting criteria must be considered simultaneously. Unlike classical probabilistic (Bayesian) methods, which assume a single objective function, MCDA explicitly acknowledges and structures the presence of multiple criteria, which may be qualitative, ordinal, or quantitative.

MCDA methods are diverse in formulation, but they share certain methodological features. First, they require the articulation of criteria relevant to the decision context, often through stakeholder engagement. Second, they typically involve the valuation or scoring of alternatives on each criterion, using performance scales that may be quantitative or qualitative. Third, they incorporate a mechanism for aggregating these evaluations into a global preference or ranking of alternatives, which may be deterministic or incorporate uncertainty. However diverse they are, there is still an inescapable requirement to be aligned with classic decision theory.

Well-known MCDA methods include value-based approaches such as SMART,

VIKOR and TOPSIS, and outranking methods such as ÉLECTRE and PRO-MÉTHÉE, among others. Value-based methods often rely on compensatory aggregation rules and require preference elicitation while outranking methods try to encompass non-compensatory reasoning to deal with what they see as incomparabilities. Nevertheless, regardless of approach, they must by necessity stay within the scientific borders of classic decision theory which they build upon.

Most present-day developments in computational decision analysis occur within MCDA rather than single-criterion probabilistic (Bayesian) methods. To recap the evolution discussed in Part I, the beginnings of MCDA can be traced back to the development of decision theory and operations research (OR) during World War II. OR itself emerged as a discipline in the early 1940s, driven by military needs for efficient resource allocation, optimal supply chain management, and strategic planning. Pioneering researchers, such as Dantzig, developed linear programming, a mathematical approach that provided optimal solutions to problems of allocation under constraints. Early decision models were primarily concerned with single-objective optimisation, seeking to identify the best solution according to a single criterion, typically minimising costs or maximising profit (Dantzig, 1947).

However, as noted above, decision-makers in the real world often face problems with multiple, often conflicting objectives. In these more complex scenarios, the concept of MCDA began to take shape as researchers sought to extend optimisation techniques to consider trade-offs between competing criteria. This led to the development of early multi-objective optimisation methods in the 1950s and 1960s, which sought to find solutions that balanced competing objectives. One of the earliest contributions to this field was the work of Kuhn and Tucker on the theory of optimality in decision making (Kuhn and Tucker, 1951), which set the groundwork for future developments in multi-criteria analysis by formalising the need to consider multiple constraints in decision-making problems.

In the 1950s and 1960s, as both OR and decision theory matured, the necessity of incorporating multiple objectives into decision making became more apparent. At this time, mathematical models for decision making began to account for various factors beyond simple profit or cost optimisation. Multi-attribute utility theory (MAUT) early became a cornerstone of MCDA. MAUT posits that individuals

make decisions based on the expected utility derived from each alternative, with each attribute (or criterion) contributing to the overall utility in a weighted manner.

The classic concept of utility, however, assumes that preferences can be quantified and aggregated into a single utility function. For complex decision problems with multiple criteria, this assumption is often difficult to meet. In response, Keeney and Raiffa at IIASA developed methods to analyse trade-offs between criteria in their book *Decisions with Multiple Objectives* (1976/1993). Their work introduced a more structured approach to multi-criteria decision making by emphasising the importance of defining and eliciting the decision-maker's preferences over multiple criteria. They recognised that many real-world decision problems do not lend themselves easily to the construction of a single utility function and therefore suggested the use of non-aggregative methods, where each criterion is considered independently but in relation to the others via scale alignments.

The first applications of MCDA methods were primarily in the fields of management science, engineering, and public policy, where decision-makers had to evaluate alternatives based on multiple criteria. In the 1960s, ad hoc multi-criteria methods, based on optimisation models, were applied to a wide range of decision problems, from resource allocation and industrial engineering to urban planning and environmental management. In the 1970s, as the availability of computing power increased, MCDA models became more computationally feasible for a wider range of applications. The development of decision support systems (DSS) during this period allowed for the systematic application of MCDA methods in interactive decision making. These systems enabled decision-makers to model multiple criteria and evaluate the performance of different alternatives, taking into account not only quantitative but also qualitative data. The integration of MCDA into DSS marked a significant step forward in making complex decision making more accessible and analytically rigorous. It was not, however, until the 1990s that computational power was used for complex decision-analytic calculations in a way they had been used in OR for a long time. One of the first descriptions of computational decision analysis is (Danielson, 1997), however in a single-criterion setting.

In parallel with the development of traditional MAUT, other methods were emerging in the 1970s that focused on the structuring and evaluation of complex, multi-criteria problems. Among the earliest was the Analytical Hierarchy Process

(AHP), developed by Saaty already in the late 1970s (Saaty, 1977). AHP introduced a method for structuring multi-criteria problems into a hierarchy of objectives, sub-objectives, and alternatives, which could be compared pairwise in terms of relative importance. The pairwise comparison approach allowed for s seemingly precise evaluation of trade-offs and the calculation of a final score for each alternative by synthesising the results of comparisons. However, the approach also opened up serious problems when applying it to real-world decision problems.

The traditions in MCDA are a bit different from those in probabilistic (Bayesian) decision analysis (PDA). While PDA traditionally has a more theoretical and axiomatic approach, focusing on well-foundedness, MCDA is more concerned with processes, procedures and calculation schemes. There is nothing inherently wrong in any of the two sets of approaches, rather they simply stem from different traditions. PDA originates from mathematics, statistics and economics, and hence inherited methods and ways of thinking and expression from those disciplines. MCDA, on the other hand, has a more pluralistic background, with for example some of the more widespread methods coming from civil engineering (VIKOR) and industrial engineering (TOPSIS). While an engineering approach to a research problem is not per se better or worse than a mathematical/theoretical one, they often yield vastly different outcomes. For an insight into the epistemic fragmentation within the MCDA field, cf. Greco et al. (2016) for a 1350-page, 50+ author overview accommodating numerous divergent research directions and philosophies, presenting sometimes isolated and often disparate perspectives and methods without a common coherent foundation. In the overview, each method is presented by eager and invested advocates, emphasising the foundational divide within the field.

In contrast, this book aims at unifying PDA and MCDA by mapping the results in Part I onto MCDA and adding computability as the third pillar in Part III. In (Danielson, 1997), only PDA is treated in detail. Since then, and characterising the 21<sup>st</sup> century, multi-criteria decision problems have been much more in focus. Luckily, many results from Bayesian PDA carry over to MCDA, albeit with some modifications. This second part of the book will deal with the similarities and differences between the two approaches and ends with a unified model (MPDA = multi-criteria probabilistic decision analysis) where all three types of decision variables (probabilities, utilities and criteria weights) are modelled and evaluated together.

In early MCDA development, the question was raised of how decision-makers should compare the alternatives with respect to different types of objectives of the decision. Each objective is referred to as one attribute in the decision context, and the approach is to define one individual utility function for each attribute. These are then aggregated into a global utility function, in which weights express the relative importance of each attribute. Each consequence  $C_i$  may be thought of as a vector of achievement levels regarding the identified attributes, in the case of n attributes, the consequence set  $C_i = (c^1, ..., c^n)$ . Some literature uses the terms criteria or perspective instead of attribute, however, from a theoretical point of view these terms may be used interchangeably.

Several approaches to aggregate utility functions under a variety of attributes have been suggested, such as (Keeney and Raiffa, 1976/1993; Saaty, 1980; von Winterfeldt and Edwards, 1986). The most widely employed method is the additive utility function, sometimes referred to as the weighted sum. Some conditions must be fulfilled for the additive utility function to serve properly as an aggregated utility function. Firstly, the assumption of mutual preferential independence must hold, which states that when a subset of alternatives differs only on a subset  $G_i \subset G$  of the set of attributes G. Then the preferences between the alternatives must not depend on the common performance levels  $G \setminus G_i$ . Secondly, the condition of additive independence must hold, meaning that changes in the uncertain outcomes (its probability distribution) in one attribute will not affect preferences for lotteries in other attributes. The weights are restricted by a normalisation constraint  $\Sigma w_j = 1$ ,  $w_j \in [0,1]$ , where  $w_j$  denotes the weight of attribute  $G_j$ . A global utility function U using the additive utility function is then expressed as

$$U(x) = \sum w_i u_i(x)$$

where  $w_i$  is the weight representing the relative importance of attribute i.  $u_i$ :  $X_i \rightarrow [0, 1]$  is the increasing individual utility function for attribute  $G_i$ , and  $X_i$  is the state space for attribute  $G_i$ . It is assumed that the functions  $u_i$  map to zero for the worst possible state regarding the ith attribute, and map to one for the best.

Another global utility function is the multiplicative utility function, introduced in

(Keeney and Raiffa, 1976/1993). The multiplicative model requires that every attribute must be mutually utility-independent of all other attributes, saying that changes in certainty levels of one attribute do not affect preferences for lotteries in the other attributes. In contrast to additive independence, the condition of utility independence allows the decision-maker to consider two attributes to be substitutes or complements of each other. In this respect, it is a weaker preference condition than additive independence. Generally, the global utility function is usually expressed as

$$1+KU(x_i)=\prod (Kk_iu_i(x_i)+1)$$

where  $u_i: X_i \to [0,1]$ .  $u_i$  is the increasing individual utility function for attribute  $G_i$ , and  $X_i$  is the state space for attribute  $G_i$ . As for the additive function, the utility functions  $u_i$  map to zero for the worst possible state regarding the *i*th attribute, and map to one for the best. The scaling constant K is the non-zero solution to

$$1+K=\prod (1+Kk_i)$$

where the  $k_i$  represent scaling constants, similar in their meaning to weights, but without the normalisation requirement.

Other formal methods of decision evaluation under multiple objectives include the outranking approach (Benayoun et al., 1966; Brans, 1982), often referred to as the French school of decision analysis. This approach is based on a search for outranking relations deduced from a set of binary preference relations. However, these approaches do not incorporate the modelling of uncertainty in the probabilistic sense and thus do not capture the risk associated with different courses of action.

Two major theoretical systems of thought underpin the computational foundations of decision analysis, viz. von Neumann-Morgenstern's (vNM) expected utility theory and Keeney-Raiffa's multi-attribute utility theory (KR), the latter developed at IIASA, the International Institute for Applied Systems Analysis, during Raiffa's years as Director General 1972–1975, with Keeney employed as Research Scholar, and thus sometimes referred to as the IIASA theory of MCDA. While both theories originate from a similar rationalist tradition, they differ substantially in scope and structure. The vNM formulation is based on choices under uncertainty, where out-

comes are lotteries over consequences. Preferences that satisfy completeness, transitivity, continuity, and independence axioms can be represented by a linear expected utility function: where is a lottery over outcomes with probabilities, and is a utility function defined over outcomes. The independence axiom is central: preferences over lotteries must not change if all options are mixed with a third lottery in the same proportions.

KR generalises utility theory to deterministic multi-attribute decisions. It replaces lotteries with multi-criteria score profiles and aims to construct utility functions over combinations of attribute levels. The key axioms include i) Utility independence of attributes, ii) Monotonicity in attributes, and iii) Decomposability (e.g., additive or multiplicative form) When these are satisfied, an additive utility function can represent preferences. Unlike vNM, KR treats the modelling of preferences without uncertainty. While vNM and KR are often treated as distinct, they are best understood as kin since their mathematical representations of utility differ mainly in context and notation. Both frameworks seek to represent preferences via utility functions that are linear in the appropriate domains. vNM handles linear expectation over probabilistic outcomes while KR handles linear aggregation over deterministic attributes. The similarity lies in the additivity: in both cases, preferences are consistent with a sum of utilities, weighted by either probabilities or attribute weights. Thus, the vNM expected utility function can be interpreted as a variant of a multiattribute probabilistic utility function where the attributes are mutually exclusive outcomes governed by probability.

In KR, this convergence becomes especially clear: the aggregate utility function in MAUT is the practical analogue of vNM's expected utility formula, with probabilities replaced by weights and outcomes replaced by criteria scores. This kinship underscores the deeper unity of decision theory: whether one is choosing under risk or across multiple attributes, the rational structure of preferences, grounded in utility, independence, and monotonicity, remains the same. A key difference is whether uncertainty is external (vNM) or multi-dimensional (KR).

To sum up, the main similarities are i) both systems rest on axiomatic representations of rational preference; ii) both aim to construct numerical representations that respect ordinal rankings; and iii) each incorporates separability and independ-

ence in different forms. While the main differences are i) vNM requires probabilistic lotteries; MAUT does not; ii) vNM utility is cardinal (up to affine transformations); MAUT utility is typically interval or ordinal depending on scale assumptions; and iii) MAUT has trade-offs between attributes; vNM captures risk attitude.

In Part II, we will discuss some popular MCDA methods and check whether they comply with core fundamentals of mathematical statistics, decision theory and analysis. If not, they seem to be victims of over-engineering and ought to be either reformulated to be used as proper decision analysis frameworks or not considered theoretically motivated tools and methods. To properly discuss them, we introduce ten desiderata that are derived from vNM, KR, and multi-attribute utility in general.

**Desideratum 1** (Ordering): The preference relation is complete. For all A and B, either A > B, B > A, or  $A \sim B$ . vNM takes completeness as axiomatic to ensure coherent preferences. KR carries it over to deterministic multi-attribute models.

**Desideratum 2** (Transitivity): The preference relation is transitive. If A > B and B > C, then A > C. vNM assumes transitivity as axiomatic to ensure coherent preferences. KR brings it over to deterministic multi-attribute models.

**Desideratum 3** (Dominance): If for all i,  $s_i(A) \ge s_i(B)$  and for some j,  $s_j(A) > s_j(B)$  then A > B. Strong dominance is compatible with both vNM and KR. It ensures that if one alternative is objectively better, it must be preferred.

**Desideratum 4** (Monotonicity): If A > B, and A' is such that  $s_i(A) = s_i(A')$  for all  $i \neq j$  and  $s_i(A) = s_i(A') + \varepsilon$  for some  $\varepsilon > 0$ , then  $A' \geq B$ . This is a standard assumption in both vNM and KR.

**Desideratum 5** (Independence of Irrelevant Alternatives, IIA): If A > B in set X, and  $C \notin \{A, B\}$ , then A > B in  $X \cup C$ , provided that criterion weights are automatically adjusted to preserve the importance of one unit on the original scales if C caused scale renormalisations. Follows from vNM's independence axiom (in its strong form). KR reinterprets it in terms of trade-off consistency: adding an irrelevant alternative should not affect preference ordering.

**Desideratum 6** (Rank Preservation): If A > B in X, and C is a third alternative not affecting the scores of A or B, then removing C from X does not alter the ranking A > B (allowing for automatic weight adjustment to preserve per-unit criterion meaning). Follows up on Desideratum 5 and stability assumptions. In additive

utility models, preferences among pairs are unaffected by alternatives with no impact on the value functions of the focal options.

**Desideratum 7** (Criteria Transparency): For any preference A > B, there exists a representable and decomposable justification based on the contribution of each criterion to the total evaluation. This follows from KR's value function decomposition principle. It ensures additive or multiplicative representations are intelligible and traceable to criterion-level contributions.

**Desideratum 8** (Weight Sensitivity): Let  $w_i \in [0, 1]$  be weights summing to 1. A change in a  $w_i$  that increases the influence of criterion  $C_i$  in which  $s_i(A) \ge s_i(B)$  should not reverse the preference A > B. This follows from sensitivity analyses in MAUT (KR) and reflects the principle that weights encode preference intensities and must affect final utility accordingly.

**Desideratum 9** (Criteria Independence): If criteria  $C_i$  and  $C_j$  produce identical scores for all alternatives, the results should be cardinally invariant under merging them into one criterion with a combined weight  $w_i + w_j$ . Related to the independence of attributes in MAUT (KR). A duplication of identical attributes without properly adjusting the weights violates utility independence.

**Desideratum 10** (Scale Invariance): For any criterion  $C_i$ , if a positive affine transformation  $f: \mathbb{R} \to \mathbb{R}$  is applied to all  $s_i(\cdot)$ , then the preference relation A > B should remain unchanged. In both vNM and MAUT (KR), utility functions are ordinal up to a monotonic transformation and cardinal under positive affine ones.

These ten desiderata form a requirements system that will, for reference, be called DAMS (Decision-Analytic Methodological System) and which guarantees well-behaving and well-functioning MCDA methods if the ten are all adhered to. From the ten DAMS desiderata, some consequences follow.

**Proposition 1** (Utility Representability): If DAMS Desiderata 1–8 are accepted, then there exists a utility function  $U: X \to \mathbb{R}$ , representable as an additive model

$$U(A) = \sum_{i=1}^{n} w_i \cdot v_i(s_i(A))$$

where each  $v_i$  is a continuous, increasing value function and  $w_i \ge 0$  with  $\sum w_i = 1$ .

This follows from classical multi-attribute utility theory in the deterministic case. The axioms ensure the separability, monotonicity, and decomposability needed for an additive representation.

**Proposition 2** (Rank Reversal Exclusion): If DAMS Desiderata 5 and 6 are accepted, then the decision method is immune to rank reversal caused by irrelevant alternatives.

Desideratum 5 ensures rankings are stable under expansion of the alternative set and Desideratum 6 maintains ranking under deletion. Together they exclude the structural basis for rank reversal which plagues some currently popular MCDA methods.

**Proposition 3** (Weight Responsiveness): If DAMS Desiderata 7 and 8 are accepted, then rankings will adjust appropriately under changes in criterion weights, without violating transitivity or dominance.

These three propositions together define a class of prescriptively robust MCDA methods that are logically sound, preference-sensitive, and transparent. Violations of these desiderata entail logical or interpretive compromises of different kinds. There is an eleventh unofficial desideratum, concerning the decision-maker's understanding of the underlying procedural elements. It places demands on the consistent transparency of the logic used by the method and extends Desideratum 7.

**Desideratum 11** (Explanatory Transparency): It must be possible for the users to form and maintain a requisite mental model of the analytic process as a whole, including but not limited to its computational steps. The method should provide a conceptually accessible and intelligible mapping from inputs (scores, weights, thresholds) to outputs (rankings and numerical scores), enabling both auditability and replicability. This includes the ability to trace how changes in inputs influence outcomes, without reliance on hidden mechanisms or opaque procedures.

This last desideratum is sometimes not understood by designers of methods. They test their methods on decision problems, some real-life and some artificial, and observe the steps unfolding, Often, the process is facilitated by an intermediate or an expert, which makes the users not question the traceability of the output from the input, instead often relying on the perceived expertise of the facilitator. However, if MCDA methods are to become more widespread, there is a need for more

transparency in the processes to build trust in the output results. Desideratum 11 is different from 7 (Criteria Transparency) despite sharing the concept 'transparent'.

While the desiderata are formulated to be conceptually independent, some exhibit logical or functional overlap under classic utility theory assumptions. DAMS Desiderata 1 (Ordering), 3 (Transitivity), 3 (Dominance) and 4 (Monotonicity) form a foundational core. These suffice to guarantee transitive, rational preferences that respect utility dominance and maintain independence from unrelated alternatives. Desideratum 6 (Rank Preservation) can be viewed as a corollary of Desideratum 5 (IIA). If preferences are independent of irrelevant alternatives and based solely on score vectors, the deletion of an irrelevant third option should not affect the outcome. Desideratum 9 (Criteria Independence) implicitly relies on Desiderata 7 (Criteria Transparency) and 8 (Weight Sensitivity). If a method transparently reflects weight changes and scores, duplication or merging of criteria without corresponding weight adjustments violates score attribution logic. Desiderata 7 (Transparency) and 8 (Weight Sensitivity) are not strictly necessary but conceptually desirable since they ensure interpretability. Desideratum 10 (Scale Invariance) stands somewhat independent from the others but supports robustness under unit changes. It is justified on theoretical rather than logical grounds. For pedagogical reasons as well as argument's sake, all ten desiderata are kept in the DAMS system as guidelines and discussion points in the ensuing presentations of MCDA methods. These desiderata should not be confused with Howard's 14 desiderata and five processing rules for a decision process as a whole (2009), which conflate higher-level procedural steps with axiomatic and computational elements.

As will be shown in the sequel, most MCDA methods depart in several ways from DAMS. Specifically, they fail to deliver decomposable, monotonic, and utility-independent representations. They do not support consistent trade-off interpretation at the attribute level. These methods seemingly offer practical tools but lack coherence. As a case in point, take rank reversal (see Proposition 2), the phenomenon where the introduction or removal of irrelevant alternatives alters the ranking of existing ones. It serves as a powerful litmus test for compliance with DAMS. In line with the desiderata, preferences are supposed to be constructed to be invariant under irrelevant changes. This is encoded as independence of irrelevant alternatives (IIA), separability, and utility independence. Rank reversal directly violates these

principles. Thus, any method that admits rank reversal is, by definition, out of alignment with the desiderata as well as the core of classical utility theory. Moreover, rank reversal highlights violations of independence of irrelevant alternatives (IIA, Desideratum 5) and rank preservation (Desideratum 6), which are direct consequences of utility separability. In practice, a method that allows rank reversal is one in which utility is not decomposable or context-stable, which is an immediate red flag for any hopes of a sound utility-theoretic grounding.

The DAMS framework with ten desiderata will be employed to discuss six different well-known MCDA methods in separate chapters: SMART (representing the SAW class of methods evaluating the alternatives using a sum-of-weighted-values approach), and the Big Five: VIKOR, TOPSIS, ÉLECTRE, PROMÉTHÉE and AHP. All except ÉLECTRE were designed in 1977-1982, after Keeney and Raiffa's IIASA work was published (1976). These methods were selected for this book because of their spread and reach – they are the most commonly used methods in decision analysis by a wide margin. Their usage and citation patterns suggest that method popularity often reflects branding success more than demonstrated methodological superiority. The prominence of certain methods appears to be driven less by performance or theoretical soundness and more by factors such as catchy acronyms, compelling narratives, and academic network effects. This was noted already by Belton and Stewart (2002). Additionally, being early to the methods scene has afforded some approaches a lasting advantage, allowing them to establish a dominant position before competing approaches emerged, further reinforced by cumulative citation effects. In classic marketing theory, users are locked in to a product or a service by branding and narratives, creating a mental barrier to switching. The proliferation of some MCDA methods resembles a form of implicit marketing, where name recognition and earlier citations heavily influence uptake, often independently of rigorous comparative validations or theoretical coherence, circumstances one could wish at least academia were largely devoid of. While these methods promote structured decision making, their branding leverage some of the very cognitive biases they aim to mitigate, such as the availability heuristic and affective association, both well-known from descriptive decision theory and ironically at play in the meta-selection of the methods themselves.

#### 07. SMART

The SMART family of methods originated in the context of problems with structuring decision situations under multiple criteria, drawing inspiration from early multi-attribute utility theory and the desire to provide a structured yet relatively simple approach to decision making. The Simple Multi-Attribute Rating Technique (SMART) was developed by Edwards (1977) as a tool for decision-makers to evaluate alternatives based on multiple attributes or criteria. SMART was conceived as a practical method to facilitate decisions in complex environments without requiring overly sophisticated modelling of preferences or trade-offs. Edwards' motivation was to provide a method that was simple enough for non-experts to use while still retaining the essential elements and rigor of decision theory.

The SMART (Simple Multi-Attribute Rating Technique) family of methods constitutes a set of approaches developed within MCDA for the evaluation and ranking of alternatives characterised by multiple attributes. Originating in the early 1970s, SMART was introduced by Edwards as a response to the perceived complexity and limited practical usability of existing MCDA methods, particularly those requiring full elicitation of utility functions or cardinal preference structures. The core idea behind SMART was to provide a simpler, more intuitive framework for supporting decision making by relying on additive models and direct rating procedures.

At its inception, SMART mandated that decision-makers assign a weight to each criterion, reflecting its relative importance, and then rate each alternative with respect to each criterion on a typically numerical and bounded scale. These ratings are then aggregated via a weighted linear sum to yield an overall score for each alternative. The attractiveness of SMART lay in its procedural simplicity: it assumed mutual preferential independence of criteria and linearity of value functions, which allowed for direct and transparent computations of aggregated scores.

In subsequent decades, SMART evolved into a family of related methods, each designed to address specific theoretical or practical issues that emerged during its application. One such extension is SMARTS (SMART using Swings), which refines the weight elicitation process. Instead of assigning importance weights directly, SMARTS asks decision-makers to assess the value difference between the worst and best levels of each criterion, given that all others are fixed at their worst

levels. This swing approach yields more correct relative weightings by anchoring them in the perceived impact of changes across the criteria range.

A further extension is SMARTER (SMART Exploiting Ranks), which attempts to reduce the cognitive burden of precise weight elicitation. Instead of assigning numerical weights, SMARTER relies on ordinal rankings of criterion importance and employs a surrogate technique, rank-order centroid (ROC) weighting, to derive approximate cardinal weights from the rankings. This approach trades off some theoretical precision for increased ease of use and constitutes a practical compromise in setting weights with limited time or cognitive resources.

Other variants and refinements include methods that relax the assumption of linear value functions or incorporate uncertainty in the weights and performance ratings. For example, probabilistic versions of SMART have been proposed that model ratings or weights as distributions rather than fixed quantities, allowing sensitivity analyses and robustness assessments within the SMART framework. Its various forms share a common structure rooted in additive value models but diverge in their assumptions, elicitation procedures, and treatment of uncertainty.

The SMART methods are built on a set of relatively simple computational rules that require the decision-maker to perform the following steps: first, the decision-maker lists the criteria relevant to the decision problem. Then, each criterion is assigned a weight representing its relative importance in the decision-making process. The weights are typically normalised so that they sum to one. Next, each alternative is evaluated on each criterion, usually on a numerical scale, such as 1 to 10, with the scale representing the performance of the alternative relative to the others.

The final step in SMART involves computing a weighted sum of the scores for each alternative. The alternative with the highest weighted sum is typically chosen as the preferred option. Mathematically, the decision rule in SMART can be expressed as follows:

$$S_i = \sum_{j=1}^m w_j x_{ij}$$

where  $S_i$  is the overall score for alternative i,  $w_j$  is the weight for criterion j,  $x_{ij}$  is the performance score of alternative i on criterion j, and m is the number of criteria.

This weighted sum approach ensures that the decision-maker's preferences are reflected in the final decision and that the process is computationally efficient.

To address the limitations inherent in the original SMART approach, Edwards and Barron (1994) proposed SMARTER (SMART Exploiting Ranks), an extension designed to accommodate greater complexity in decision modelling. The principal objective of SMARTER was to enhance methodological flexibility in representing nuanced preference structures. While preserving the additive architecture of SMART, SMARTER replaces direct numerical weighting with a rank-based approach, using the rank order centroid (ROC) method to derive criterion weights. This substitution introduces a non-linear mapping from rank to weight, better capturing how decision-makers perceive importance differences among criteria. However, the aggregation of alternative scores remains strictly linear. By easing the cognitive burden of weight elicitation while preserving structural simplicity, SMARTER is well-suited to situations in which full cardinal precision is unrealistic, yet preference structures demand more than uniform weighting or arbitrary approximations.

SMARTER also offers improved flexibility in eliciting preferences across multiple criteria. In complex decision problems, trade-offs between criteria often reflect underlying tensions, such as economic efficiency versus environmental sustainability, where improving performance on one dimension may entail losses on another. While SMARTER does not explicitly model interdependencies, it enables decision-makers to express the relative importance of criteria through complete ordinal rankings. These rankings are then transformed into weights using the ROC surrogate method, implicitly capturing asymmetries in perceived importance. This increased flexibility, however, introduces additional complexity by requiring more structured input in the form of a full ranking of all criteria.

SMART and its variants are, in fact, members of a broader methodological tradition commonly referred to as the SAW family, named after its foundational use of a sum-of-weighted-values approach. Also known as weighted linear combination or scoring models, the SAW family is built on the principle of linear additive aggregation, whereby the performance of an alternative is, as in SMART, expressed as a weighted sum of its evaluations across multiple criteria. At its core, the SAW family operationalises a special case of additive value models, in which the total value V(a)

of an alternative a is computed as a weighted sum of marginal value functions As noted earlier, this formulation presupposes interval-scale measurements. In its SAW formulation, the marginal value functions  $u_j(x_j)$  are typically assumed to be identity mappings over normalised criteria scales, and the weights  $w_j$  serve as pure scaling constants reflecting the decision-maker's trade-offs among criteria. This corresponds directly to additive independence and cardinal representation as laid out in MAUT. The SAW computational form is typically expressed as in the formula on page 62. It assumes that the attribute levels  $x_{ij}$  have been transformed into commensurate scales, typically via linear transformations. Once the transformations are accepted, aggregation proceeds under the expected value logic of utility theory, where alternatives are scored according to the weighted sum of marginal utilities.

SMART in its basic form complies with most of the desiderata of DAMS, primarily due to its additive structure, monotonicity, and transparency. Scores and weights are transparent, independence and dominance are preserved, and rank stability is assured. It is sometimes argued that the method does not align with Desideratum 9 (Criteria Independence) in that duplicating a criterion inflates its influence. This is true for original SMART but not for SAW methods in general, though, since SAW methods should always adjust their weights when the criteria set changes, and that can be carried out by an automatic procedure. This reflects an inherent property of criteria weights, not of the SAW family itself.

The SAW family gained influence during the 1960s and 1970s, driven by general developments in cost-benefit modelling, optimisation and systems analysis. It was during this period that Keeney and Raiffa (1976) formulated the axiomatic foundations of additive preference models at IIASA, providing the formal justification for SAW as a special case of a broader utility-theoretic theory. In practice, however, the family's intuitive arithmetic transparency made it popular well before its theoretical justifications. In Danielson and Ekenberg (2016), SMART representing numerical SAW is compared to the CAR method representing cardinal SAW ranking and to AHP. In the study, 100 decision-makers each made one significant decision over a three-week period using all three methods, after which they compared the methods across five performance indicators. The results showed that both SAW-based approaches were strongly preferred to the ratio-based AHP (presented in Chapter 12).

#### 08. VIKOR

VIKOR (a Serbian acronym; in English: multi-criteria optimisation and compromise solution) is a method developed from the late 1970s onwards (David and Duckstein, 1976; Duckstein and Opricović, 1980), initially under the name IKOR. It was described as building on ideas from ÉLECTRE, favouring that method over MAUT. The method changed names in 1998 (Opricović, 1998, p.iii) and the new backronymed name VIKOR originally referred to a FORTRAN program, not the method. IKOR was designed for ranking and selecting alternatives in the presence of conflicting criteria, based on concepts of compromise programming and individual regret. The development of (V)IKOR emerged from work in multi-objective optimisation within water resource management. Its formulation is related to a metric used in compromise programming, where the distance of each alternative to an ideal solution is computed. VIKOR uses two measures: the *S* measure (representing aggregated utility) and the *R* measure (representing maximum regret). They are then combined into a total ranking index *Q*, modulated by an external parameter *v*.

The computational procedure of VIKOR involves the identification of the best values for each criterion among all alternatives (known as the ideal solution), normalisation of the performance matrix to make criteria comparable, and the calculation of the S, R, and Q values for each alternative. The alternatives are then ranked according to these values. A compromise solution is proposed based on the ranking of the Q values, subject to two acceptability conditions (C1, C2) that involve both rank consistency and a threshold for closeness between top-ranked alternatives.

The first computational step in VIKOR involves the construction of a decision matrix, where the rows represent the alternatives, the columns represent the criteria, and the entries in the matrix correspond to the performance of each alternative under each criterion. After that, the ideal- and regret-based solutions are determined. The ideal solution is obtained by selecting the best performance for each criterion across all alternatives, while the regret solution is obtained by selecting the worst component for each criterion.

The next step is the calculation of the distance of each alternative from the ideal and solution and the amount of regret selecting each alternative would incur. Once the distances from the ideal solution (S) and the regret (R) are calculated, the method

computes a compromise index (Q) for each alternative. This index represents the degree to which an alternative offers a balance between proximity to the ideal solution and regret. The index is calculated by combining the distance from the ideal solution and the regret, weighted by the relative importance of each metric. The final step involves ranking the alternatives based on the three metrics.

The calculation details are as follows. Assume there are n alternatives (denoted  $A_1, A_2, ..., A_n$ ) and m criteria (denoted as  $C_1, C_2, ..., C_m$ ) used to evaluate each alternative. The values for each alternative and criterion are typically represented in a matrix X, where each element  $x_{ij}$  represents the performance of alternative  $A_i$  with respect to criterion  $C_i$ .

The values in matrix X are then normalised in order to transform them into a comparable scale. The normalisation function depends on whether the criterion is beneficial or non-beneficial. For beneficial criteria, the normalisation formula is

$$y_{ij} = rac{x_{ij} - x_{\min,j}}{x_{\max,j} - x_{\min,j}}$$

while for non-beneficial criteria, it is

$$y_{ij} = rac{x_{ ext{max},j} - x_{ij}}{x_{ ext{max},j} - x_{ ext{min},j}}$$

where  $x_{max,j}$  and  $x_{min,j}$  are the maximum and minimum values in the *j*th criterion across all alternatives. Thus, this is a standard normalisation where the best alternative in each criterion receives the value 1 and the worst 0. This can be interpreted as the one-dimensional closeness to the best outcome.

The ideal solution  $A^+$  is then defined as the best performance for all criteria:

$$A^+ = \{y_{\max,1}, y_{\max,2}, ..., y_{\max,m}\}$$

where  $y_{\max,j}$  are the maximum values for each normalised criterion j. For each alternative  $A_i$ , the (L<sub>1</sub>) distance to the ideal solution is then calculated using the reversed formula

$$S_i = \sum_{i=1}^m w_j \cdot (y_{\max,j} - y_{ij})$$

rather than a more traditional formula

$$S_i = \sum_{j=1}^m w_j \cdot (y_{ij} - y_{\min,j})$$

where  $w_j$  in both cases represents the weight of the *j*th criterion. Thus, the normalised scores are now reversed and reinterpreted as the multi-dimensional closeness to the best outcome instead. Next, the regret is computed by the  $(L_\infty)$  formula

$$R_i = \max_{j} \left( w_j \cdot (y_{\max,j} - y_{ij}) \right)$$

with the same meaning of its constituents as above. The regret for an alternative in this method is the worst weighted closeness of any of the constituent criteria.

Finally, the compromise index  $Q_i$  combines the two measures  $S_i$  and  $R_i$  using an exogenous factor v. The formula for the index is

$$Q_i = v \cdot rac{S_i - S_{\min}}{S_{\max} - S_{\min}} + (1 - v) \cdot rac{R_i - R_{\min}}{R_{\max} - R_{\min}}$$

where  $S_{min}$  and  $S_{max}$  are the minimum and maximum values of  $S_i$  across all alternatives,  $R_{min}$  and  $R_{max}$  are the minimum and maximum values of  $R_i$  across all alternatives, and v is an external factor that represents the relative importance of the majority of criteria. For v = 1, completely disregarding the ranking based on  $R_i$ , VIKOR is a reversed additive utility model since the  $S_i$  and  $Q_i$  rankings coincide, but for any other value of v, it is not. Somewhat surprisingly, some descriptions of the method do not seem to require  $0 \le v \le 1$ , which opens up for strange interpretations. This calculation procedure yields three ranking orders of the alternatives based on their performances  $S_i$ ,  $R_i$  and  $Q_i$ . A set of rules (C1 and C2 plus if-thenelse rules) determines which of the rankings take precedence, with the  $Q_i$ -ranking being the primary to consider first. C1 is called the acceptable advantage and is a threshold  $Q_1 - Q_2 \ge DQ$  for the two top-ranked alternatives where DQ = 1/(J-1)for J alternatives in total. In (Opricović, 1998, p.154), an upper limit of ¼ was introduced on DQ, upheld in 2002 but strangely omitted in (Opricović and Tzeng, 2004). From 2004 onwards, a decision situation with J = 2 alternatives results in DQ = 1, which can almost never be satisfied, rather of the more reasonable  $DQ = \frac{1}{4}$ .

The ideal solution is the best synthetic alternative, i.e. it does not exist in reality. Such alternatives are themselves, in swing-type methods, tools for elicitation rather than calculation devices. However, in VIKOR they are bases for distance calculations. For a simple example, consider the normalised values  $y_{ij}$ , i.e. their ranges are [0, 1]. Then  $A^+$  becomes the vector (1, 1, 1, ...) and all  $S_i$  become  $\sum_i [w_i \cdot (1-y_{ij})]$ , i.e. a reversed weighted sum, where lower values represent better alternatives. This is a linear operation on an additive scale. Next, three measures are calculated for each alternative, of which  $S_i$  mostly resembles a standard DAMS measure. However, as pointed out, with a reversed scale where lower numbers are better, a measure of distance from the synthetic ideal (optimal) alternative. Still, this is in line with DAMS since all operators are linear and thus there exists a 1-1 relation. The other two measures involve a max operator, which is not linear and these measures lack the foundational validity of  $S_i$ . The S and R rankings, together with a linear combination Q of S and R, which does not add any information except an exogenous factor v, are a basis for a compromise procedure which may not produce a complete ranking or even a top-ranked alternative. It might be unclear why a compromise is required, how that need is expressed in any computable form, and how that form can be validated. While the calculations are easy to follow for the mathematically inclined, they lack the transparency of DAMS Desideratum 11.

VIKOR fails Desideratum 8 (Weight Sensitivity) and 9 (Criteria Independence) due to the behaviour of the regret component. It also violates Desideratum 6 (Rank Preservation) by its post-decision rules (C1, C2, and thresholds) since these depend on score differences, full-set reference points and exogenous decision-maker input. Thus, while seemingly compliant at the numerical ranking stage, the full method sacrifices robustness. The use of compromise ranking regret measures similarly deviates from the desiderata. Its aggregation formula includes a balance parameter v, which lacks a clear grounding in utility theory. It fails decomposability and is sensitive to dataset composition, violating utility independence, making it structurally prone to rank reversal when the 'best' or 'worst' alternatives change upon set modification. Already (Duckstein and Opricović, 1980) documented different ranking orders for VIKOR, ÉLECTRE and classic MAUT (SAW) for a small river basin problem (not even resulting in the same top-ranked alternative).

## 09. TOPSIS

TOPSIS (Technique for Order Preference by Similarity to Ideal Solution) is a method developed by Hwang and Yoon (1981), who admit that their method is clearly inspired by ÉLECTRE, which they consider to be "one of the best" and also "[most] refined" (ibid., p.127). TOPSIS was created to partly mimic ÉLECTRE and identify solutions that simultaneously have the shortest geometric distance from an ideal solution and the farthest distance from a nadir (anti-ideal) solution. The evaluation principle is that the optimal alternative should be the closest to the positive (ideal) solution (PIS) and the farthest from the negative (anti-ideal) solution (NIS).

The process begins with the normalisation of the decision matrix to eliminate the differing scales across criteria. After normalisation, the values are multiplied by the corresponding criterion weights, which reflect the relative importance of each criterion. Once the weighted normalised matrix is formed, the PIS and NIS are determined. The PIS consists of the best values for each criterion (maximum for benefit type, minimum for cost type), while the NIS consists of the worst. The Euclidean distance of each alternative from both the PIS and the NIS is then calculated. These distances are used to compute a closeness coefficient for each alternative, defined as the ratio of the distance to the NIS over the sum of distances to the PIS and NIS. The evaluation principle stems from the concept of distance measurement. Distance functions provide a way of comparing alternatives by quantifying the deviation of each alternative from an ideal solution representing the optimal choice across all criteria, and an anti-ideal solution represents the worst possible outcome. These two solutions form a bounded space within which the method operates, and all alternatives are measured relative to these bounds. The choice of a vector-space (L<sub>2</sub>, Euclidean distance) measure is, however, as doubtful as in ÉLECTRE.

In more detail, the first step, after forming the traditional two-dimensional matrix of alternatives and criteria, is to transform the decision input so that the data for each criterion is dimensionless and can be compared. The transformed value  $r_{ij}$  for each utility is calculated as:

$$r_{ij} = \frac{x_{ij}}{\sqrt{\sum_{i=1}^n x_{ij}^2}}$$

where  $x_{ij}$  is the original utility of alternative  $A_i$  with respect to criterion  $C_j$ . By squaring  $(r_{ij}^2)$ , it is easy to see that all  $x_{ij}^2 / \sum x_{ij}^2$  always fall within a [0, 1] scale but without spanning the scale as a standard normalisation does. Thus, this RMS-rescaling (root-mean-square), which is a cornerstone operation in statistics but not in decision analysis, is not the same as standard normalisation.

Each criterion has an associated number  $w_j$  representing the relative importance of criterion  $C_j$ . However, these numbers are not MCDA weights. Such weights are trade-off factors between spanned [0,1] scales. Since TOPSIS scales are not spanned, the numbers called  $w_j$  are not pure weights but a mixture of weights and scaling factors. The fundamental requirement that the weights are trade-off factors between equal scales is not met. The transformed values  $v_{ij}$  are computed as

$$v_{ij} = w_j \cdot r_{ij}$$

The ideal and anti-ideal solutions are then determined by considering the best and worst values for each criterion. The ideal solution  $A^+$  is the set of values for which each criterion has the best value (for beneficial criteria) or the worst one (for non-beneficial criteria).

$$A^+ = \{v_{\text{max},1}, v_{\text{max},2}, ..., v_{\text{max},m}\}$$

where  $v_{max,j} = \max(v_{ij})$  for beneficial criteria and  $\min(v_{ij})$  otherwise. Conversely for the anti-ideal solution  $A^-$ 

$$A^- = \{v_{\min,1}, v_{\min,2}, ..., v_{\min,m}\}$$

If the components of the  $A^+$  and  $A^-$  vectors had been properly normalised, they would have been similar to anchor points in a standard swing process. The next step is to compute the Euclidean distance between each alternative and the ideal  $(S_i^+)$  and anti-ideal  $(S_i^-)$  solutions. They are the RMS (root-mean-square, vector measure) distances to the vectors of the ideal  $A^+$  and anti-ideal  $A^-$  solutions calculated by

$$S_i^+ = \sqrt{\sum_{j=1}^m (v_{ij}-v_{ ext{max},j})^2}$$

and

$$S_i^- = \sqrt{\sum_{j=1}^m (v_{ij} - v_{\min,j})^2}$$

respectively. The larger the value of  $S_i^+$  ( $S_i^-$ ), the farther the alternative is from the (anti-)ideal solution. Given these two opposite measures, another ranking of the alternatives is made using the combined measure

$$C_i = \frac{S_i^-}{S_i^+ + S_i^-}$$

The alternatives are ranked in decreasing order of  $C_i$  with the alternative having the highest  $C_i$  being the most preferred since it is closest to the ideal solution.

As seen above, the transformation of the utilities  $v_{ij}$  into the calculation values  $r_{ij}$  is an RMS (root-mean-square, i.e. non-linear) operation. Thus, a linear relationship between  $v_{ij}$  and  $r_{ij}$  is lost even before weighing the values. The weighing comes next, which is a linear operator and does not distort the calculations further. After weighing the transformed values, each alternative's distance to the best (ideal) and worst (anti-ideal) possible (but usually non-existent) values are calculated. Although the criteria have weights that sum to one in a standard (linear) way, this distance is not the (linear) sum of each of the criteria's distances. Instead, it is the L<sub>2</sub>metric (Euclidean) distance between the two points in a metric polytope. This is clearly not according to the DAMS desiderata and not in alignment with the nature of the input data. Consider an alternative that is α units away from the fictive optimal solution  $A^+$  in criterion s and also  $\alpha$  units away from  $A^+$  in criterion t. Since the criteria scales have been weighted (normalised), a unit in either criterion has the same influence on the end result – that is the meaning of scale normalisation by weights. Thus, the alternative would need an improvement of  $\alpha + \alpha = 2\alpha$  units to become equal to  $A^+$ . But TOPSIS would consider the required improvement to be  $\sqrt{2}\alpha$  which is clearly wrong. The distance in a weight space should be measured by a city block (L<sub>1</sub> or Manhattan) measure, not a Euclidean (L<sub>2</sub>) one. To realise the problem with the TOPSIS calculation method, assume wlog that the input data is on a [0, 1] format, i.e. the worst alternative for each criterion has the value 0 and the best has the value 1. Then  $A^+$  becomes  $\{w_1, w_2\}$  and  $A^-$  becomes  $\{0, 0\}$  given a weight vector  $(w_1, w_2)$  where  $w_1 + w_2 = 1$  as usual. For the scale space to be

invariant under traversal, every path from  $A^-$  to  $A^+$  must have the same length and be equal to 1. This is not the case in TOPSIS which assigns the length  $\sqrt{(w_1^2 + w_2^2)}$  to the traversal while a DAMS-compliant method, requiring a city block  $L_1$  metric, will have 1 for every conceivable traversal. For a given alternative, every path to the same final improvements in a set of criteria must be considered equal.

To assess the real-world effects of TOPSIS' deviation from the DAMS model, the author has performed a Monte Carlo simulation of  $30 \cdot 10^6$  rounds comparing the ranking order of a standard DAMS formulation and TOPSIS for a decision situation with 5 alternatives under 4 criteria. In about 73–74% of the rounds, the ranking was the same. In more than 4% of the rounds, at least one alternative had a ranking that differed by two positions or more from SDA. Given the small decision situation with only 5 alternatives, that is a lot. Thus, in more than ½ of the cases, TOPSIS' results differ from the linear-based standard DAMS model.

The method it violates criteria independence (D9), scale invariance (D10) and rank preservation (D6). These failures stem from its reliance on data-dependent reference points and Euclidean ( $L_2$ ) distance aggregation, which are sensitive to score distribution and structural redundancy.

TOPSIS ranks alternatives by their relative proximity to an ideal and anti-ideal point, based on weighted Euclidean distances over vector-normalised criteria. The method is transparent and decomposable, allowing criterion-wise contributions to be traced via squared deviations, though not additively. While it aligns with DAMS in using compensatory aggregation and strict orderability, it departs in key ways: it violates strong dominance, lacks scale invariance, and depends on dataset-specific reference points. These context-sensitive anchors cause failures in criteria independence and rank preservation, reducing its compliance.

Let alternatives A and B be evaluated on two criteria. Suppose A is initially closer to the ideal point than B. Introducing a third alternative C with extreme values in one criterion can shift the ideal and anti-ideal reference points. This change may cause B to appear relatively closer than A, even though neither alternative's performance has changed. This behaviour violates Desideratum 6 (Rank Preservation): preference orderings should not be affected by the removal or addition of irrelevant alternatives.

# 10. ÉLECTRE

ÉLECTRE (ÉLimination Et Choix Traduisant la REalité) is a family of methods developed in France during the mid-1960s by Benayoun and colleagues at SEMA, Société d'Économie et de Mathématiques Appliquées (Benayoun et al., 1966; Benayoun and Sussmann, 1966). Although Roy is often credited as the originator of the ÉLECTRE method, the foundational work was carried out within one of his teams at SEMA. The two initial SEMA papers, cited above, list Benayoun as first author, with Roy appearing on only one of them. Roy, serving as *Directeur de la Direction Scientifique* at SEMA, published the first academic article on ÉLECTRE as sole author, reflecting his position as the institution's leader (Roy, 1968). Back then, ÉLECTRE was the name of a FORTRAN computer program running on a CDC computer, not of the method. Neither had it been backronymed yet.

The method was originally designed to support decision making in complex situations where preferences may be non-compensatory and where full ranking of alternatives is not always appropriate or feasible. A key idea of ÉLECTRE is to construct an outranking relation based on concordance and discordance between pairs of alternatives evaluated over multiple criteria. The first version, ÉLECTRE I, was introduced in 1966. It was designed to solve the problem of choosing a subset of alternatives rather than producing a full ranking. The method operates by constructing an outranking relation, denoted as "a outranks b," when there is sufficient evidence that alternative a is at least as good as alternative b. This is determined using two indices: the agreement (concordance) index and the disagreement (discordance) index. The concordance index measures the degree to which the majority of criteria support the statement that a is at least as good as b, taking the criteria weights into account. The discordance index captures the extent to which any criterion strongly contradicts this statement. An outranking is established if the concordance is high enough and discordance is not too strong.

ÉLECTRE II, introduced shortly after ÉLECTRE I, was designed for ranking problems and introduced the concepts of strong and weak outranking relations to reflect varying levels of support for preference statements. It uses different thresholds for concordance and discordance and introduces procedures for partial and complete pre-orders based on these relations.

ÉLECTRE III, developed in the 1970s and formalised in the early 1980s, introduced pseudo-criteria and the use of indifference, preference, and veto thresholds. ÉLECTRE IV further developed the approach for cases where criteria weights are not available. It uses ordinal information only, relying on the ranking of criteria and performance without requiring numerical weights. ÉLECTRE IS is a later adaptation of ÉLECTRE I for use in decision support software systems, integrating technical refinements and improved routines. ÉLECTRE TRI, introduced in the early 1990s, shifts the focus from ranking or choosing among alternatives to sorting them into predefined categories. ÉLECTRE TRI has been further developed into ÉLECTRE TRI-B and ÉLECTRE TRI-C, differing in the treatment of assignment rules and model structure. The set of methods is notably diverse, giving rise to a metadecision problem on which of the methods in the set to use and when.

The ÉLECTRE family of methods follows a series of steps to derive the preferred alternatives. The first step in any ÉLECTRE application is the construction of a decision matrix. This matrix typically consists of rows corresponding to the alternatives and columns corresponding to the criteria. The decision-maker populates the matrix by providing performance values for each alternative with respect to each criterion. Once the matrix is established, ÉLECTRE proceeds by defining preference thresholds for each criterion. These thresholds are critical to the method's operation as they determine how differences in performance between alternatives will be perceived. Typically, there are two thresholds for each criterion:

- 1. Indifference threshold: This threshold specifies the range within which the difference in performance between two alternatives is so small that it does not affect the ranking. If the difference in performance between two alternatives on a given criterion is less than this threshold, the alternatives are considered indifferent to each other for that criterion.
- 2. Preference threshold: This threshold defines the minimum performance difference required for one alternative to be considered preferred over another for a given criterion. If the difference in performance between two alternatives exceeds this threshold, one alternative is considered preferred over the other for that criterion.

In addition to these two thresholds, some versions of ÉLECTRE also use a veto threshold, which is applied when an alternative is deemed completely unacceptable

based on a critical criterion, regardless of its performance on other criteria. The veto threshold ensures that the decision-maker's priorities are respected, preventing alternatives that fall below a certain level of performance on essential criteria from being considered at all, even if they perform better on other criteria.

Once the thresholds are established, ÉLECTRE proceeds with the pairwise comparison of alternatives. For each pair of alternatives, the method evaluates whether one alternative outranks the other. The outranking relationship is determined by comparing the alternatives for each criterion and assessing whether the difference in performance exceeds the appropriate preference or indifference thresholds. If the difference in performance is larger than the preference threshold, the alternative is considered preferred; if it is smaller than the indifference threshold, the alternatives are considered indifferent; and if the difference is larger than the veto threshold, the alternative is deemed outranked.

The results of these pairwise comparisons are summarised in an outranking matrix, where each entry reflects the degree to which one alternative outranks another across all criteria. The outranking matrix forms the basis for constructing the preference structure, which organises alternatives into groups or sets based on their relative performance. This ranking is partial rather than complete, as some alternatives may not be ranked in a strict order.

The final decision-making step in ÉLECTRE involves applying a series of concordance and discordance indices to further refine the rankings. The concordance index quantifies the degree of agreement between alternatives in terms of the number of criteria where one alternative is preferred over the other. In contrast, the discordance index measures the extent to which an alternative is disfavoured by a criterion, representing the degree of disagreement between the two alternatives. These indices are then used to aggregate the pairwise comparisons and to generate an overall outranking relation between alternatives.

To examine the computations in detail, six steps have to be scrutinised:

- 1. Normalising the decision matrix.
- 2. Calculating concordance and discordance for each pair of alternatives.
- 3. Constructing the concordance and discordance matrices.
- 4. Aggregating them into the dominance matrix.

- 5. Defining the outranking relation.
- 6. Ranking the alternatives based on the outranking relation.

The first step is the transformation of the input so that the data for each criterion is dimensionless and can be compared. The transformed value  $r_{ij}$  for each utility is calculated in the same way as for TOPSIS:

$$r_{ij} = rac{x_{ij}}{\sqrt{\sum_{i=1}^n x_{ij}^2}}$$

where  $x_{ij}$  is the original utility of alternative  $A_i$  with respect to criterion  $C_j$ . As for TOPSIS, it is easy to see that all  $x_{ij}^2 / \sum x_{ij}^2$  always fall within a [0, 1] scale but without spanning the scale as a standard normalisation does. Thus, this RMS (rootmean-square) operation, which is a cornerstone in statistics but not in decision analysis, is not the same as standard normalisation. TOPSIS copied this RMS rescaling, which is a vector-space metric rather than a DAMS-compliant one, from ÉLECTRE without reflecting on the consequences of adopting it (Hwang and Yoon, 1981).

However, after this step ÉLECTRE diverges from TOPSIS. The concept of concordance compares each pair of alternatives based on the criteria, indicating the degree to which one alternative dominates another. For each pair of alternatives  $A_i$  and  $A_k$ , the concordance index  $c_{kl}$  is calculated as

$$c_{kl} = \sum_{j \in C_{kl}} w_j$$

using the concordance set membership function ( $v_{ij}$  is the same as  $r_{ij}$  above)

$$C_{kl} = \{j \mid v_{kj} \ge v_{lj}\}$$

In a different way, the discordance index is calculated as

$$d_{kl} = rac{\max_{j \in D_{kl}} |v_{lj} - v_{kj}|}{\max_{i,h,j} |v_{ij} - v_{hj}|}$$

based on the discordance set membership function

$$D_{kl} = \{j \mid v_{kj} < v_{lj}\}$$

This concept of disagreement (or discordance) has inspired VIKOR's subsequent regret ranking, which also leads to several overlapping or inconsistent rankings with a number of rules of thumb devised to try to separate them into one final ranking. So while the ÉLECTRE family has been a trendsetter, it is no more DAMS compliant because of that. On the contrary, the ideas copied by other methods are non-compliant in nature. Next, define a threshold  $c^*$  such that

$$c^* = rac{1}{m(m-1)} \sum_{k 
eq l} c_{kl}$$

or some similar function, different versions of the method have different functions. Then construct a two-dimensional binary matrix F with elements

$$f_{kl} = egin{cases} 1 & ext{if } c_{kl} \geq c^* \\ 0 & ext{otherwise} \end{cases}$$

which shows where alternative  $a_k$  concordance-dominates  $a_j$ . Next, construct another two-dimensional binary matrix G with elements

$$g_{kl} = egin{cases} 1 & ext{if } d_{kl} \leq d^* \ 0 & ext{otherwise} \end{cases}$$

indicating where  $a_k$  is not too much worse than  $a_l$  in the discordance sense. After a few more steps, a partial ranking is arrived at by ÉLECRTE I which is considered the end result. No total ranking can be promised with this method, this depends on lucky circumstances among the input data. The ÉLECTRE family contains many methods that differ in various respects. The final rankings in the methods are based on the outranking relationships between all pairs of alternatives. The alternatives are sorted based on how strongly they outrank others. The alternative that outranks the most others (with the highest dominance value) is considered the most preferred.

ÉLECTRE violates most desiderata, for example, 1–5 and partially 6. Its threshold logic undermines monotonicity and independence. Rank reversals are common, and preferences can be reversed by introducing or removing unrelated alternatives. Additionally, it does not produce a total ordering and partly fails to satisfy utility-based decomposability (Desideratum 7). Further, the way of introducing arbitrary user-defined thresholds in the computations instead of imposing all such operations on the end result is not in alignment with DAMS.

ÉLECTRE relies on concordance and discordance indices and veto thresholds to establish outranking relations. Although it attempts to reflect dominance, it fails in decomposability and transparency. The method's qualitative thresholds obscure continuous preference trade-offs and often produce incomparabilities. From a MAUT viewpoint, ÉLECTRE violates utility independence and introduces arbitrary cut-offs without functional justification. Adding a new alternative can alter concordance and discordance thresholds due to recalculated matrices. An alternative A previously considered non-dominated may now be outranked due to shifts in veto thresholds, partly violating Desideratum 6 by indicating that utility structure is not preserved. The methods use thresholds and concordance-discordance matrices that are recalculated for every new alternative. This context-sensitive process causes violations of both IIA and Rank Preservation. Moreover, incomparabilities may arise or disappear when the set changes, leading to rank inconsistencies. The transparency of ÉLECTRE is the least among the MCDA methods surveyed so far (but it will get worse). No real-life decision-maker the author has met (as opposed to mathematicians and decision theorists) comprehended the steps and how or why they lead to a suggested ranking of the alternatives.

Further, the fact that ÉLECTRE only yields a partial ranking as output, with no reliable cardinal information (strengths between the ranked alternatives), reveals a naïve view on decision-analytic support, as if the MCDA method should make the decision (in line with the strongest formulation of the MCDM decision-*making* assumption), a standpoint that is at odds with how modern MCDA is (and should be) used as a guiding tool. Not least in decision situations where large sums of money are involved, the desire to have a cost-benefit step as the last one in the decision process is common. In such a step, partial (or even complete) rankings will not do. It has to be cardinal, numeric output to be of any use. Besides, in all other decision situations, cardinal information is also always preferable, not least since a sensitivity analysis should follow the initial results. Such analyses are much harder (bordering impossible) to perform with only ordinal output information available.

# 11. PROMÉTHÉE

PROMÉTHÉE (originally called Préférence par Ordination selon la Méthode ÉLECTRE pour les Hiérarchiques Évaluations Enrichies, later anglicised to Preference Ranking Organisation Method for Enrichment of Evaluations – both referring to the Greek god Prometheus, meaning forethought) is a family of methods developed by Brans in the early 1980s. PROMÉTHÉE belongs to the class of outranking methods (also known as the French school of MCDA) pioneered by the SEMA Group (ÉLECTRE). The initial formulations, PROMÉTHÉE I and II, which were counter-reactions to ÉLECTRE I–IV, were presented in (Brans, 1982). There, it is pointed out that the ÉLECTRE methods contain difficulties that PROMÉTHÉE aims to overcome, such as handling the concordance and discordance thresholds. Those are complicated to set, and further, the results obtained do not provide a complete ranking of alternatives. These difficulties are circumvented i.a. by introducing generalised preference functions and a unified ranking procedure (ibid., Section 3).

A core concept in PROMÉTHÉE is the use of a preference function that translates the difference in performance between two alternatives on a single criterion into a degree of preference ranging from 0 (no preference) to 1 (strict preference). Decision-makers choose among several predefined preference functions, each corresponding to different assumptions about how preferences behave with respect to differences in criterion performance. As usual, each criterion also has a weight, reflecting its relative importance in the overall decision situation (Brans and Vincke, 1985).

PROMÉTHÉE I produces a partial ranking of alternatives based on the calculation of positive and negative preference flows. The positive flow measures how much an alternative is preferred over others, while the negative flow indicates how much it is outranked by others. These flows are used to identify incomparabilities when conflicting preferences occur. PROMÉTHÉE II, by contrast, derives a complete ranking by computing the net flow (positive minus negative), thus eliminating incomparabilities but possibly reducing information about preference structures.

Following the ÉLECTRE tradition, the initial formulations of PROMÉTHÉE were followed by several extensions to address specific methodological requirements. PROMÉTHÉE III was developed to deal with rankings that involve interval

data or require robustness in the presence of uncertainty. PROMÉTHÉE IV extends the method to handle continuous alternatives, particularly useful in problems where alternatives form a continuous set rather than a discrete list. This version involves the integration of preference functions over continuous domains, relying on integral calculus rather than discrete summation. PROMÉTHÉE V incorporates constraints, such as resource or budget limitations, and enables the selection of a subset of alternatives that satisfy these constraints while preserving preference relations. This variant merges the outranking methodology with optimisation techniques to support constrained decision problems. PROMÉTHÉE VI was designed to accommodate multiple decision-makers by aggregating their individual preference flows through various consensus or voting procedures.

A central idea of all PROMÉTHÉE versions, as well as all ÉLECTRE ones, is that alternatives are ranked based on their outranking relationships. An outranking relation expresses the degree to which one alternative is considered superior to another, taking into account all relevant criteria. This is achieved by comparing the performance of each pair of alternatives with respect to each criterion and evaluating the intensity of preference for one over the other. This comparison is not always straightforward, as decision criteria may have different importance levels or even exhibit interdependencies. To handle these complexities, PROMÉTHÉE incorporates preference functions that model the intensity of preference for one alternative over another, based on the performance difference for each criterion. The method allows for non-linear preferences, meaning that a small difference in performance may be more or less significant depending on the criterion in question.

PROMÉTHÉE operates in several stages, from the formulation of the decision matrix to the final ranking of alternatives. The first stage involves the construction of a decision matrix, where each row represents an alternative, and each column corresponds to a criterion. In this matrix, the values for each alternative-criterion pair represent the performance of the alternative with respect to that criterion.

Next, the decision-maker is asked to provide preference functions for each criterion. These functions are important to the method because they capture how the decision-maker perceives the trade-offs between alternatives. A preference function specifies how much better one alternative is preferred over another, given a certain difference in performance on a given criterion. For example, if the criterion is cost,

the decision-maker may consider a small reduction in cost as highly desirable, but a larger reduction as less significant. In this case, the preference function could be designed to reflect a diminishing marginal utility for cost savings.

The preference function is typically a non-decreasing function that expresses the intensity of preference. Depending on the criterion, it can take different forms. For example, in the case of a benefit criterion (where higher values are preferred), the function could be linear or exponential, indicating that the higher the performance of an alternative, the greater the preference. For a cost criterion (where lower values are preferred), the function might be decreasing, reflecting the increasing preference for alternatives that perform better (i.e., have lower costs).

Once the preference functions are established, the method proceeds with the calculation of preference indices for each alternative pair. These indices quantify the degree to which one alternative is preferred over another for each criterion, based on the difference in their performance. The total preference index for an alternative is obtained by summing these individual preference indices over all criteria.

After calculating the preference indices, the method computes two global outranking flows for each alternative: the positive outranking flow and the negative outranking flow. The positive flow reflects the degree to which an alternative is preferred to all other alternatives, while the negative flow reflects the degree to which it is outranked by other alternatives. These flows are calculated by considering all the pairwise comparisons and aggregating the preference indices for each alternative.

Originating from political and social sciences, the methods are designed to facilitate negotiation and compromise rather than a definite result. In this, behavioural components get mixed with analytical ones. PROMÉTHÉE I calculates a partial ranking of alternatives. This version considers only the positive and negative flows of each alternative, and it ranks alternatives according to their outranking relationships. However, the results of PROMÉTHÉE I do not necessarily provide a strict total order of all the alternatives, as some alternatives may be ranked equivalently in terms of their outranking relations. PROMÉTHÉE II, on the other hand, provides a complete ranking of alternatives by incorporating a net outranking flow, which is the difference between the positive and negative flows. This version of PROMÉTHÉE is appropriate when a complete and unambiguous ranking of alternatives

is necessary. PROMÉTHÉE II produces a strict total order of the alternatives, with the alternative that has the highest net flow being the most preferred.

Since PROMÉTHÉE ranks alternatives by calculating preference values between pairs of alternatives based on each criterion, the method considers both the magnitude of the preference and the relative importance of the criteria. This is done by the following calculation steps. As with almost every other method, it begins with normalising the input values. This time, it is a regular linear transformation of the input data where the scales are reversed for non-beneficial data (i.e. where lower numbers are preferred) to produce normalised utilities. For ordinary input values, this is

$$x_{ij}^* = rac{x_{ij} - \min(x_j)}{\max(x_j) - \min(x_j)}$$

while for reversed scales, it is instead

$$x_{ij}^* = rac{\max(x_j) - x_{ij}}{\max(x_j) - \min(x_j)}$$

where  $x_{ij}^*$  is the normalised value for alternative  $A_i$  under criterion  $C_j$ , and where  $\max(x_j)$  and  $\min(x_j)$  are the maximum and minimum values in criterion  $C_j$  across all alternatives. The method uses a preference function to quantify the preference of one alternative over another with respect to each criterion. The preference function can take different forms, depending on how the decision-maker perceives the relative importance of differences between alternatives. Its general form is

$$P_{ij} = \varphi(x_{ij}^*, x_{kj}^*)$$

where  $\varphi$  can be any of six prescribed transform functions, none of them being a simple linear function. The functions include a stepwise linear threshold function and a dichotomic threshold function that evaluates to 0 or 1 depending on whether a threshold number is met or not.

Next, for each pair of alternatives  $A_i$  and  $A_k$ , the net preference is calculated based on the individual preferences for each criterion. The net preference  $\pi_i$  of alternative  $A_i$  over  $A_k$  is computed as

$$\pi_{ik} = \sum_{j=1}^m w_j \cdot P_{ij}$$

where  $w_j$  is the weight of criterion  $C_j$  and  $P_{ij}$  is the preference function value for criterion  $C_i$  for alternatives  $A_i$  and  $A_k$ .

Next, the outranking relation is established to compare two alternatives. The net preference values  $\pi_{ik}$  are used to determine whether one alternative dominates another. The positive flow  $\Phi_i^+$  and negative flow  $\Phi_i^-$  of each alternative  $A_i$  are calculated to assess its overall preference relative to all other alternatives as follows:

$$\Phi_i^+ = \sum_{k 
eq i} \pi_{ik}$$

and

$$\Phi_i^- = \sum_{k 
eq i} \pi_{ki}$$

The positive flow is said to represent how much each  $A_i$  outranks the other alternatives while the negative flow represents how much  $A_i$  is outranked by other ones. The final ranking of the alternatives is in PROMÉTHÉE II determined by the net flow  $\Phi_i = \Phi_i^+ - \Phi_i^-$  while PROMÉTHÉE I relies only on the separate positive and negative flows. The alternative with the highest  $\Phi_i$  is the most preferred, and the one with the lowest  $\Phi_i$  is the least preferred. If two alternatives have very similar flows, an indifference threshold can be used to label them inseparable.

PROMÉTHÉE fails to comply with Desiderata 5 and 6, as the net flow scores depend on the entire set of alternatives, not just pairwise comparisons. It also partly violates Desideratum 3 (Dominance) due to preference function tuning. Though relatively transparent and responsive to weight changes, it does not ensure scale invariance or rank preservation under deletion. It uses pairwise comparisons and preference functions to derive outranking flows. While these flows offer some interpretability, they do not result from a decomposable utility function. The method's dependence on the full alternative set undermines attribute-level separability. The flows also obscure individual criterion contributions, violating transparency. As such, it is definitively incompatible with DAMS.

In PROMÉTHÉE, the net preference flow of an alternative is calculated based on pairwise dominance across the entire set. If an alternative C is added, even one with no dominance over A or B, the net flows might change. This violates IIA (Desideratum 5) and undermines utility decomposability. Because the method relies on pairwise comparisons across the full set of alternatives, the net flow scores are sensitive to the composition of the alternative set. This relational structure violates not only IIE but also IIA and can lead to rank reversals when alternatives are added or dropped.

Being a follow-up method to ÉLECTRE, albeit conceived 15–16 years later, it is not surprising that PROMÉTHÉE displays some of the same weaknesses in that it only yields ranking as output as well, and again with no reliable cardinal information to supplement the output. It is not as naïve as ÉLECTRE since some variants (not PROMÉTHÉE I) at least result in complete rankings. However, almost the same arguments apply to PROMÉTHÉE as to ÉLECTRE since cardinal information is always preferred, not least as a sensitivity analysis should follow the initial results. See the previous chapter on ÉLECTRE for a more thorough discussion of these shortcomings and how they relate to an outdated and monolithic MCDM view of decision-analytic support in general, rather than seeing MCDA as one useful tool in a toolbox.

Ending the chapter with a sidenote, Électre (Electra) is a tragic figure from Greek mythology, known for her relentless pursuit of vengeance, moral absolutism, and emotional isolation, traits often portrayed without a prospect of redemption. Électre brings conflict and harsh justice to the table, making uncompromising decisions without sentiment. In contrast, Prométhée (Prometheus), also from Greek mythology, is a god who gave fire to humanity and stands as a symbol of enlightenment, rational defiance, and hope for human progress. He is portrayed as empowering rather than punishing, suffering so that others might see more clearly. Électre demands justice in a broken world while Prométhée represents the hope for a better one. Why those names were chosen as backronyms for the respective methods is unclear.

## 12. AHP

The Analytic Hierarchy Process (AHP) is a method developed by Saaty in the mid-1970s, with its theoretical foundations first formally presented in (Saaty, 1977). AHP was introduced to support decision making by structuring problems into a hierarchical model and enabling the quantification of subjective preferences through pairwise comparisons. The method is based on the principles of ratio-scale measurement and relies on human judgement to derive priority scales (Saaty, 1980).

AHP involves decomposing a decision problem into a hierarchy with at least three levels: the overall goal at the top, criteria (and possibly sub-criteria) at intermediate levels, and the set of decision alternatives at the bottom. Decision-makers are required to make pairwise comparisons between elements at each level with respect to their parent node. These comparisons are captured using a 1-to-9 scale proposed by Saaty, where 1 indicates equal importance and 9 indicates an extreme preference for one element over another.

From the pairwise comparison matrices, AHP derives a set of priority vectors using eigenvalue calculations. The principal right eigenvector of the matrix is normalised to produce relative weights, reflecting the intensity of preferences among the compared elements. Consistency of the pairwise judgements is measured using a consistency index (CI) and a consistency ratio (CR). These measures compare the observed consistency of the matrix to a random matrix of the same order. A CR below a threshold, typically 0.1, is generally considered acceptable.

There are also various methods for improving the efficiency and scalability of AHP, especially in high-dimensional problems. These include methods for incomplete pairwise comparisons, where not all element comparisons are required, and consistency-driven adjustments to reduce redundancy and cognitive load.

Computational implementations of AHP and its variants have been developed extensively. These implementations often incorporate mechanisms for consistency checking, sensitivity analysis, and visualisation of results. AHP is susceptible to inconsistencies in pairwise comparisons. AHP uses the Consistency Ratio (CR) to assess the degree to which the pairwise comparisons are logically consistent. However, even when the consistency ratio is within acceptable limits (typically below 0.1), inconsistencies can still affect the accuracy and reliability of the decision. The

requirement for pairwise comparisons can become overwhelming for decision-makers, particularly in decision problems with a large number of alternatives and criteria. This can lead to inconsistencies that are difficult to detect or rectify, thereby affecting the quality of the final decision. AHP is more of a procedure-driven method than a formula-driven one. Thus, it is best described by the steps involved. An AHP evaluation involves the following steps:

- 1. Performing pairwise comparisons.
- 2. Normalising the pairwise comparison matrices.
- 3. Calculating the priority vectors (weights).
- 4. Conducting consistency checks.
- 5. Calculating global weights and determining the final ranking of alternatives. In the first step, Pairwise Comparisons, decision-makers compare each pair of elements using a scale (usually from 1 to 9):
  - o 1 means equal importance.
  - o 3 means one element is slightly more important.
  - o 5 means one element is significantly more important.
  - o 7 means one element is very strongly more important.
  - o 9 means one element is extremely much more important.

The comparisons for the criteria would be represented as a pairwise comparison matrix. Next in the same step, construct the Pairwise Comparison Matrix. It is constructed from the elements that represent the relative importance of the elements compared. The matrix is reciprocal, meaning  $a_{\{ij\}} = \frac{1}{a_{\{ji\}}}$ . Next, normalise the pair-

wise comparison matrix. Normalise each column of the matrix by dividing each element by the sum of the elements in that column. This step ensures that the columns represent the relative importance on a common scale. The resulting matrix is the normalised matrix.

The next step is to calculate the eigenvector (priority vector), which represents the relative weights of the elements (either criteria or alternatives). This is done by calculating the dominant eigenvector of the pairwise comparison matrix. Such an operation might yield an inconstant matrix. Thus, the step that follows is to check the consistency of the comparisons. AHP assumes that the pairwise comparisons

should be consistent (if A > B and B > C then A > C should hold, i.e. transitivity). The consistency ratio (CR) is computed to assess how consistent the pairwise comparisons are. The steps to check consistency are:

- 1. Compute the consistency vector by multiplying the comparison matrix by the priority vector.
- 2. Divide the resulting vector by the priority vector element-wise to get the lambda max (largest eigenvalue).
- 3. Calculate the consistency index (CI) using

$$CI = rac{\lambda_{ ext{max}} - n}{n-1}$$

4. Finally, compute the consistency ratio (CR) by dividing the CI by a random consistency index (RI) that depends on the size of the matrix. If CR is below an exogenous threshold (typically 0.1), the comparisons are considered consistent enough.

The finalising step is to calculate the global weights. Once the priority vector for the criteria is determined as well as the pairwise comparison matrices for the alternatives relative to each criterion, the global weights of the alternatives are computed by combining the local weights for each criterion with the global weights of the criteria. This is as complicated as it sounds from a user perspective, and the method is not transparent as seen by decision-makers.

AHP fails or partially violates almost every desideratum: Desideratum 4 (monotonicity not guaranteed), and Desiderata 5–10 (due to scale sensitivity, context dependence, and rank reversal), plus conditionally Desideratum 2 (due to tolerated inconsistency) and Desideratum 3 (dominance ignored),. The eigenvector approach further obscures criteria transparency (Desideratum 11) and utility interpretability.

Let alternatives A and B be evaluated in an AHP framework with pairwise comparisons indicating A > B. Now introduce C, which is strictly worse than both A and B across all criteria, i.e. A > C and B > C. The pairwise comparison matrix must be expanded to accommodate C, and due to renormalisation, the original relative weights between A and B shift. The risk: B > A might occur. This violates Desideratum 5 (IIA) and by extension the separability required in utility theory, a problem that has been known since long (Belton and Gear, 1983). AHP's pairwise

comparison matrices are scale-dependent and inherently sensitive to the number and configuration of alternatives. A rank reversal might occur when a new, perhaps even dominated, alternative is added (Dyer, 1990). AHP is a preference elicitation method, not a dominance-checking procedure. It assumes the decision-makers are perfectly consistent in their judgements. Since this is almost impossible, dominance violations can occur. But that is a feature, not a bug, in the AHP worldview. Of the methods discussed, it is certainly the least compliant with DAMS and has been labelled as outright flawed (Abbas, 2018, Ch.3).

Although the contributions of Belton and Gear (1983), Dyer (1990) and others have served as important red signals, they are in a sense unnecessary. An inspection of the internal mechanics of AHP reveals that only a few of the fundamental underlying principles from established theories, what we refer to as desiderata, are actually adhered to. Consequently, issues such as rank reversal unfortunately arise. The entire process is opaque, leading to results that are difficult to trace and interpret. It is hoped that future research will redirect its focus and resources towards the advancement of MCDA methods that are grounded in well-established scientific principles, rather than engaging in efforts that contradict them.

AHP has been criticised for using ratio scales, which from a measure-theoretic standpoint are not compatible with the interval scales foundational to classic decision-analytic theories and thus violate key desiderata. Ratio scales conflict with the linearity assumption underpinning expected utility theory, where utility functions must support additive operations over probabilities or weights. Next, ratio scales lack invariance under positive monotonic transformations, an important property for preserving preference orderings in both ordinal and interval-based MCDA models. Third, the use of ratio-based inputs may violate preferential independence, which is essential for constructing valid additive models. Further, ratio-derived weights challenge the assumption of commensurability across criteria, as the scale intensities lack a consistent unit of value, making cross-criteria comparisons ambiguous. This is not to say that ratio scales are flawed per se. Rather, their assumptions and properties do not align well with the structure in the MCDA domain. However, they have been demonstrated to be very useful and of importance in other fields, such as representing perceptual and cognitive processes (Saaty, 2001). Different scale types are discussed further in Chapter 14.

## 13. Comparisons

The DAMS desiderata framework for MCDA provides a principled foundation that integrates classical utility theory with the realities of multi-criteria environments. The axioms synthesise formal requirements such as transitivity, dominance, and independence with practical necessities like criteria weighting and score transparency. Table 2 summarises how the methods discussed in Part II comply with the DAMS desiderata.

$\begin{array}{c} \textbf{Desiderata} \rightarrow \\ \textbf{Methods} \downarrow \end{array}$	1. Ordering	2. Transitivity	3. Dominance	4. Monotonicity	5. Indep. Irr. Alt.	6. Rank Preserv.	7. Transparency	8. Weight Sens.	9. Criteria Indep.	10. Scale Invar.
General SAW	OK	OK	OK	OK	OK	OK	OK	OK	OK	OK
VIKOR	OK	OK	NO	OK	OK	NO	NO	NO	NO	OK
TOPSIS	OK	OK	OK	OK	OK	NO	OK	OK	NO	NO
ÉLECTRE	NO	NO	NO	NO	NO	NO	NO	OK	NO	NO
PROMÉTHÉE	NO	NO	NO	NO	NO	NO	NO	OK	OK	NO
AHP	OK	NO	NO	NO	NO	NO	NO	NO	NO	NO

Table 2. SMART and the Big Five methods compared using the DAMS desiderata

As demonstrated in this book through analyses, classifications, and sometimes counterexamples, many popular MCDA methods fall short of satisfying these desiderata. Regarding Desideratum 9, Criteria Independence, the outcome depends on how the criteria are handled in the outer-layer MCDM process, outside of the core MCDA calculations. Therefore, it is not possible to draw a definite conclusion regarding that desideratum based only on the discussions in the preceding chapters. But methods such as TOPSIS and ÉLECTRE, which confound trade-off weights by using fundamentally non-comparable scales, stand a far less chance of fulfilling the desideratum. See the Appendix for a more detailed discussion on each of the ten desiderata and the methods' adherence to them (SAW is not discussed further).

The literature is replete with comparisons of MCDA methods in which the resulting rankings diverge significantly. Rather than seeking to identify a single winning method, which is often the goal, such outcomes should instead be viewed as troubling indicators of the epistemic standards within the field. For example, Opricović and Tzeng (2004) compare VIKOR and TOPSIS. The comparison clearly illustrates the ad hoc nature of both methods as well as the many differences when it comes to details. In their example, a total of 18 variants of the two methods are used to rank three alternatives, and they point out different alternatives as the best one, but the reasons for or against either result are hard to grasp for a reader. Remarkably, the 18 variants together succeed in ranking the three alternatives in all eight (!) possible permutations of the ranking order. Imagine how impossible it would be for a non-expert to understand the pros and cons of each variant. Moreover, none of the methods offer any means of sensitivity analyses, instead presenting the results with three decimals. Further, in (Opricović and Tzeng, 2007), the four methods VIKOR, TOPSIS, ÉLECTRE and PROMÉTHÉE are compared. There are six sets of criteria weights, and for each set, the methods arrive at 12 rankings in total. The rankings manage to divide the six alternatives into two sets of three alternatives each, but within the top set, the best alternative changes frequently or remains undetermined.

In (Zlaugotne et al., 2020), five methods are compared of which three are VI-KOR, TOPSIS, and PROMÉTHÉE. For the four alternatives in the article's decision problem, the five methods (only one variant of each this time) manage to produce four different rankings among the five methods. In a subsequent meta-ranking, averaging the results of the five methods, a final ranking is arrived at. However, this is not how MCDA analyses should have to be conducted – exploring a large set of methods in an ensemble fashion and hoping that their average is a "better" indicator than any single method by itself. The substantial efforts required notwith-standing, there is no theoretical proof that such averaging should lead to a better analysis. If that were the case, one could in principle construct a single optimal themore-the-merrier method, an all-encompassing meta-method made up of every known MCDA method (and perhaps all their variants), each weighted according to some mysterious, all-purpose meta-weighing scheme.

What all the methods (except SMART) fail to do is to separate the core calculus

in decision analysis from the psychological aspects of decision making. Given a set of input data, there should be one set of output data, computed according to the well-established theories that underlie DAMS. The output data should be amenable to different sensitivity analyses in order to study the stability of the results. At the next level up, the MCDM level, negotiation, bargaining, regret and similar considerations should be processed in an orderly fashion. If that processing requires additional calculations, they could be performed on the output data, but only if they can be motivated by well-founded and verified principles rather than engineering-style patches that take some property from a handy mathematical concept such as ordinary least squares or the max operator, without a solid theoretical motivation as to why and without subsequent suitability verifications by empirical studies.

It stands to reason that MCDA methods should not behave like this. Rather, these articles are a testament to the sad state of affairs that the MCDA field is currently in. The possibility of a "smorgasbord" approach: picking methods, parameters and formulas of liking, and mixing in descriptive and psychological factors, in order to allow for a ranking with a preferred alternative on top is surely a contributing factor to the prevailing mistrust and underutilisation of MCDA in society today.

This book provides both a prescriptive and diagnostic perspective: identifying logical weaknesses in existing methods, while also pointing at a route towards greater decision-theoretic coherence. This is not a plea for MCDM-level process conformity. The differences in philosophy and the different brandings of methods should influence the elicitation processes, the presentation formats, group decision mechanisms, and much more – as long as the methods stand on established scientific ground. Substituting a since-long well-established and sound axiomatic computational core for homemade calculi only leads to questionable results and opaqueness. As does mixing descriptive and psychological factors with an axiomatically grounded computational core; the former should belong only to an outer MCDM layer. The need to stand out by branding and perceived uniqueness should be satisfied in other ways than by a faulty core, ways less detrimental to the MCDA field.

Despite the logical clarity and mathematical rigor of the unified utility framework grounded in the axiom systems of von Neumann-Morgenstern and Keeney-Raiffa (vNM/KR), a wide range of popular MCDA methods exist that violate these

principles. This prompts the question: are there any compelling mathematical or logical reasons to prefer such methods? The answer, in short, is a resounding no.

None of the Big Five methods (VIKOR, TOPSIS, ÉLECTRE, PROMÉTHÉE and AHP), all of which violate vNM/KR axioms, are grounded in a rigorous theoretical foundation. These methods often make heuristic or procedural sense, or are generally appealing on the surface, but fail when held to standards of decomposability, independence and consistency. Some of their problems include:

- No representation theorem supports the forms of aggregation used.
- Non-decomposability in scoring means there is no underlying unifying utility function being optimised.
- Rank reversal and reference dependence violate basic tenets of rational choice. Despite these shortcomings, the Big Five proliferate and are widely used in practice. There are several reasons for that:
  - 1. **Software Tools:** Embedded in decision-support systems or consulting tools.
  - 2. **Visual Appeal:** Techniques like outranking or ideal point comparisons offer intuitive geometric interpretations.
  - 3. **Lack of Training:** Decision analysts are often unfamiliar with the formal structure of vNM and KR, and thus default to admiring procedural heuristics instead of questioning the basis on which a particular method stands.

There is no compelling mathematical justification for the widespread use of MCDA methods that violate the DAMS desiderata. Their popularity stems from practical, psychological, or institutional factors, not coherence. As such, their results should be viewed as suggestive, not rationally prescriptive. The proliferation of non-compliant methods underscores the need for a shift towards foundationally sound, axiomatically justified decision analysis.

DAMS draws a clear boundary between rational and pseudo-rational prescriptive decision analysis. These modes of reasoning differ fundamentally in objective, methodology, and evaluative standards. Rational prescriptive analysis is concerned with guiding decision-makers to make sound decisions given their limitations while adhering to coherent principles of preference and utility. The DAMS model developed in this book exemplifies rational prescriptive analysis:

• It rests on internally consistent axioms.

- It supports additive utility representations and generalises both von Neumann-Morgenstern's (classic) and Keeney-Raiffa's (IIASA) theories.
- It yields decisions that are transparent, defensible, and logically justified.

Pseudo-rational prescriptive analysis aims for the same goals but fails to deliver coherent and justifiable methods due to a lack of theoretical underpinnings.

- It promotes heuristics and non-linearity over consistency.
- It focuses on perceived soft factors over correctness.
- It adopts procedures that fail DAMS but are good for branding.

Part II of this book demonstrates that the Big Five methods, VIKOR, TOPSIS, ÉLECTRE, PROMÉTHÉE and AHP, are pseudo-rational tools. They aid decision making but do not meet the conditions of rationality defined in DAMS. Traditional and classic SAW methods, however, are by contrast rational tools, providing comprehensible outputs while satisfying utility-theoretic foundations. The proliferation of pseudo-rational prescriptive methods, despite their foundational shortcomings, highlights a gap between what is rational and what seems to be. The DAMS framework offers a reconciliation path: preserve coherence while retaining formats familiar to prescriptive users. This convergence can elevate decision analysis from plausible heuristics to justifiable practice.

It has been argued that a prescriptive analysis method can choose axioms "like dishes from a smorgasbord", selecting whichever seem useful and discarding others (see, e.g., Keeney, 1992). While this pragmatic flexibility may appear liberating, it undermines the very essence of decision-theoretic integrity in the methods. DAMS offers an opposite position to that stance. As discussed in this book, axioms and desiderate are not decorative or optional, they are foundational constraints that preserve coherence, comparability, and defensibility. Selectively picking them distorts the decision methods, making results opaque, less comprehensible and often logically wrong. Some problems with incoherent pick-and-choose methods include:

- 1. **Loss of Interpretability:** Methods that violate decomposability, transitivity, or independence lose any possibility of being preference-preserving. Their rankings are artefacts of procedure, not reflections of rational preference.
- 2. **Context-Dependence:** As discussed above, violations of key axioms produce rank reversals when irrelevant alternatives are added or removed.

3. **Undermining Trust:** Stakeholders rightly expect that decisions guided by formal models are consistent and principled. Violating axioms without justification breaks that trust.

In this respect, it is important to draw a boundary between two distinct conceptual layers in decision analysis:

- Mathematical-logical rigor consists of axioms, representation theorems, and their consequences. These define the structure of rational preference and the conditions under which a utility function exists. They should make up the basis for a coherent prescriptive decision-analytic calculus.
- Procedures, such as outranking, voting mechanisms, or pairwise flows, are
  process strategies. While they may offer heuristic appeal and branding differentiation, they are not substitutes for foundational coherence.

Confusing these two levels leads to mistaken beliefs, for instance, that a narratively compelling ranking procedure is comparable to a DAMS-based decision-analytic method. It is not. Only when procedures are derivable from or consistent with rigorous formulations such as DAMS can they be said to reflect genuine preference orderings in a reasonable way.

This does, of course, not entail that all methods should look the same or have the same procedures. On the contrary, different philosophical approaches call for various user interactions, various elicitation processes, and various presentation formats. That is where the variability and differences should lie, not in the computational core. Outputs can and should be post-processed in several ways for presentation at the end of the line, but only after the core results according to established theories have been calculated, and only if the post-processing can be shown still to comply with desiderata based on well-established scientific theoretical bodies instead of arbitrary made-up procedures — arbitrary seen from a decision-theoretic soundness point of view.

To move the MCDA field forward in a well-founded scientific direction, and to unify rigor with usability, a set of guiding principles are suggested. The principles acknowledge the dual demands of decision analysis: to be both prescriptively sound and practically appealing. The following is a suggestion of such a set:

**Principle 1:** *Maintain the Hierarchy of Foundations Over Procedure*Well-founded desiderata must form the backbone of any decision method. Procedures must be tested against the desiderata, not the other way around. This ensures

that decision outcomes are rational, interpretable, and stable.

**Principle 2:** *Preserve Formal Integrity, Even When Approximating* In settings where full elicitation of utilities and probabilities is impractical, approximate methods may be used, but only if they preserve key properties such as transitivity, monotonicity, and independence.

#### **Principle 3:** Ensure Representability

Every decision method should correspond to a representable utility function, even if hidden or abstracted. Such a function should be recoverable and auditable to justify preference orderings.

#### **Principle 4:** Separate Computation from Justification

Computation is necessary, but not sufficient. A method that produces results must also justify them in terms of rational calculations. Algorithms and procedures must be interpretable through the lens of utility theory.

#### **Principle 5:** Design for Transparency and Explainability

MCDA methods should reveal their internal logic: how weights are applied, how preferences are inferred, and what axioms are assumed. Stakeholders must be able to trace conclusions back to their inputs.

#### Principle 6: Protect Against Rank Reversal and Context Drift

Methods should be validated against benchmark scenarios involving irrelevant alternatives or added options. If a method produces rank reversal, it violates decision-theoretic hygiene and should be revised or rejected.

**Principle 7:** Accept Well-Founded Minimalism, Not Arbitrary Pluralism While it may be tempting to mix and match axioms as preferences or contexts vary, a minimal coherent set such as DAMS could provide sufficient flexibility without compromising logical structure. Pluralism must be principled, not ad hoc.

These principles do not restrict creativity in method design or formulation, they ensure its coherence. They invite prescriptive researchers to innovate within the bounds of rationality rather than outside of it. The future of MCDA lies not in choosing between rigor and usability, but in making them inseparable.

Among the foundational principles of sound reasoning stands Occam's Razor. In decision analysis, it translates to a call for simplicity: if two methods yield equivalent or even similar performance, the simpler one is to be preferred. This is a cornerstone in the effort to have MCDA being used more in society. Yet this principle is routinely neglected in contemporary MCDA practice. Many modern methods feature complicated data transformations, scoring algorithms, or aggregation schemes without corresponding gains, rather the opposite. There are clear reasons why simplicity matters in this case:

**Transparency:** Simpler models are easier to understand, explain, and audit. This improves stakeholder confidence and supports democratic decision processes.

**Axiomatic Tractability:** Simple structures are more likely to satisfy foundational axioms such as transitivity, decomposability, and continuity.

**Error Robustness:** Fewer moving parts reduce the risk of hidden inconsistencies, unintended rank reversals, or sensitivity to input noise.

**Theoretical Discipline:** Simplicity forces clarity in assumptions. Complex methods often obscure which principles are being applied (or violated).

However, the surveyed methods (and many others with them) violate simplicity in the following ways:

- Methods that produce partial orderings through procedures that cannot be linked to any utility representation.
- Outranking methods that require multiple thresholds and preference functions across criteria.

The desiderata proposed in DAMS are supposed to lead naturally to models that are both simple and prescriptively sound. Additive utility models, dominance-based comparisons, and weighted sums are not simplistic. They can be elegant, interpretable, and justifiable. Simplicity is not the enemy of sophistication, rather it is its friend. When methods are equally performant, the simpler model has both epistemic and explanatory advantages. Future MCDA development should not merely pursue feature richness, especially not in the number of steps and complexity of procedures, but axiomatic parsimony. Simplicity is not an aesthetic, it is a logical imperative. Branding and product differentiation should be realised by other means.

## 14. Frequently Raised Topics

This chapter discusses some topics that came up repeatedly during the graduate classes. While they are important questions, they are not closely related to each other and are collected in the final chapter of Part II for reference.

### **Scale Types**

The difference scale is a scale where the numbers are meaningful in terms of their differences but not necessarily in terms of their ratios. That is, you can measure relative differences between values, but ratios between values are not necessarily meaningful. For instance, you can say that alternative A is "3 units better" than alternative B, but saying alternative A is "3 times better" than B does not necessarily make sense. In the additive model of MCDA, you sum up the weighted differences in performance across various criteria. In other words, you are aggregating the differences in scores or performance metrics, which is typically associated with the difference scale.

Score of Alternative 
$$A_i = \sum_{j=1}^m w_j \cdot x_{ij}$$

where  $w_j$  is the weight of criterion j and  $x_{ij}$  is the performance of alternative  $A_i$  under criterion j. This form of aggregation implies that you are combining the differences between each alternative's performance across criteria, not their ratios.

AHP, on the other hand, explicitly requires the pairwise comparison scale to be ratio-based, because it is built on the idea that decision-makers can express preferences between pairs of alternatives or criteria in terms of relative importance. The standard pairwise comparison scale used in AHP typically ranges from 1 to 9 (and the reciprocals for inverse preferences), where these numbers reflect the ratio of importance between criteria or alternatives. For instance, if you compare two criteria  $C_1$  and  $C_2$  and judge that  $C_1$  is 3 times as important as  $C_2$ , the pairwise comparison matrix will reflect that in the form of a ratio-based scale. In this case, a ratio scale assumption allows you to say that  $C_1$  has a 3:1 importance over  $C_2$  and you carry this ratio into the calculation of the weight vector. AHP's use of a ratio scale means that it assumes the pairwise comparison judgements correspond to a multiplicative relationship. When you aggregate the results of pairwise comparisons for

each criterion (which are ratio-based), the result is a weighted sum of alternatives. This sum reflects the global preference for each alternative in terms of the ratios of importance, rather than just the differences.

To conclude: AHP's ratio scale means that when comparing alternatives (or criteria) pairwise, you are dealing with multiplicative relationships between alternatives' importance levels, which will then be aggregated in a weighted sum. The additive model, which typically works with a difference scale, involves linear combinations of values that do not require the ratios between them to be meaningful, but rather just their relative differences (additive increments).

#### The Independence Assumption

The standard assumption within MCDA is that of independence between criteria, and the likewise standard solution when that condition is not met between two criteria is to jointly model them as a third, overarching criterion instead. This way, a decision situation with dependent criteria can be seamlessly mapped onto a DAMS-compliant model that presupposes criteria independence. This remapping requires some skills on the part of the modeller, which is why method inventors have tried to come up with alternative ways of handling dependence.

The first obvious candidate is the correlation concept from statistics, and it has been employed in PDA models with some success. PDA models already contain conditional probabilities (without signalling) since every chain of events is a calculation of conditional global probability (A  $\mid$  B). For more on conditional probabilities, refer to any entry-level textbook on statistics. Updates of conditional probabilities are, needless to say, a centrepiece within the area of probabilistic reasoning, where Bayesian updates constitute an important topic of research – a topic that is out of scope for this book, though.

Some MCDA methods have approached the dependence issue by requiring pairwise comparisons of all criteria weights. This leads to a much heavier burden when assigning weights, essentially taking an O(n) task and turning it into an  $O(n^2)$  one. The immediate effect of a pairwise procedure is inconsistency since it is very hard for humans to keep all pairs and their transitive implications in mind at the same time. Of course, computers can help by indicating such inconsistencies in the form

of, for example, the consistency index in AHP. However, any such artificial measure introduced tends to alienate the decision-maker from the original task and thus carries a mental cost that often overshadows the possible benefits.

In cases where the criteria dependence/overlap is severe, a remodelling and mapping of criteria is the first step. As an example, Howard recounts a consulting session with an oil company that had identified 30 overlapping criteria, which after a thorough analysis turned out to be only two fundamental criteria (Howard, 2009, p.52). While that is an extreme example, it is much more often the case that criteria overlap is a consequence of bad modelling than a real inherent property of the decision problem. Thus, the resolution lies in the performance of the analysis process rather than in the method itself.

### **Compensation**

A central distinction in MCDA lies between compensatory and non-compensatory approaches to modelling trade-offs among conflicting criteria. This distinction is not merely technical; it reflects deeper philosophical assumptions about how rationality, preferences, and decision constraints should be represented and processed. The compensatory tradition, as in DAMS and many other additive value models, allows for trade-offs: strong performance in one area can offset weaknesses in another. In contrast, non-compensatory methods, such as outranking methods like ÉLECTRE and PROMÉTHÉE, are designed to integrate the handling of decision problems in which certain criteria represent thresholds or veto points that cannot be offset, regardless of performance elsewhere, into the core calculi.

Outranking methods achieve this by embedding thresholds, calling them features such as concordance, discordance, and veto levels, into the core calculations of the methods. These mechanisms are intended to model realism: in many real-world decisions, a minimum standard on certain criteria is essential, and failure to meet it could disqualify an alternative, even if it is otherwise highly rated. For instance, in supplier selection, an offer may be unacceptable regardless of cost or delivery speed if it fails to meet basic quality standards. From this perspective, outranking methods seem to respond to a real need: expressing incomparability.

However, this modelling choice comes with several well-known challenges. First, embedding such logic directly in the calculations of the method, as opposed to

the modelling phase, makes the reasoning process opaque. Threshold values are often context-sensitive, difficult to justify empirically, and may lack a clear interpretation to decision-makers. Moreover, the internal decision logic becomes more difficult to audit or explain, particularly when the result is not a complete ranking, but a partial order riddled with incomparabilities, violating Desideratum 11 (Explanatory Transparency). In attempting to mirror the complexity of real-world judgement *inside* the calculation core, outranking methods inadvertently produce black-box-like behaviour. Rather, in DAMS, non-compensatory elements are handled upstream, during the modelling phase of a decision problem. That is, criteria deemed essential or even indispensable (must-have) are treated as filters or constraints: alternatives that fail to meet them are excluded before any aggregation takes place. Criteria that are strongly correlated are remodelled together instead of standing alone. The core calculation then operates under a clean, compensatory logic, allowing weights and scores to be meaningfully interpreted, compared, and audited.

The conceptual clarity of this separation between structural constraints and preferential trade-offs supports easier communication of the results, clearer justification of rankings, and easier integration with value-for-money assessments. While it may at first glance seem that compensatory models oversimplify certain judgemental subtleties, in reality they offer greater coherence and operational transparency by handling the issues at a higher level. In this light, the divide between compensatory and non-compensatory methods (at the calculation core) reflects a deeper philosophical divide: whether the complexities of real-world decision making should be internalised in the method's inner logic or externalised and structured before calculations begin. As seen, outranking methods favour the former, often in response to misunderstood limitations of additive trade-off structures. DAMS-compliant models favour the latter, on the grounds that a good method should illuminate its calculations, not obscure them with embedded conditional logic. This is important, not least in real-world settings, where often a value-for-money approach is taken and hence, the MCDA analysis does not include monetary criteria – those are handled at a higher level in a subsequent cost-benefit (or cost-effectiveness) analysis. Not least procurement is often handled this way, making an outranking-based process unsuitable for such analyses. So again, real-world process requirements are at odds with opaque calculation methods.

## Weight/Scale Dualism

One of the most persistent and under-examined conceptual pitfalls is what may be called the great illusion of multiple scales. This illusion arises from the implicit belief that decision alternatives, evaluated on fundamentally different criteria (e.g., cost in dollars, safety in qualitative ranks, voltage in volts), can be meaningfully combined through a weighted aggregation without first aligning these criteria onto a truly comparable scale. At the heart of this issue lies a subtle but critical confusion between scaling and weighting, two distinct operations that are frequently conflated in MCDA methods.

The fundamental requirement for any meaningful weighting scheme is that all criteria be expressed on comparable scales, not merely in a superficial or cosmetic sense, but in a rigorous, mathematical one. For weights to function as intended (that is, to represent the relative importance or trade-offs among criteria), the input scales must be dimensionless and span a common interval, typically the unit interval [0, 1]. This is not a matter of convention; it is a precondition for consistency. Only when all criteria are transformed to a shared domain such as [0, 1] and this domain is fully spanned by each criterion can the weights act solely as importance coefficients. If the scales are not aligned in this way, the weights inadvertently become scale transformers as well, distorting their intended role.

This phenomenon, formally known as weight/scale dualism, undermines the theoretical coherence of many MCDA methods. The clearest examples are found in TOPSIS and ÉLECTRE, which perform their own normalisation schemes (vector-based or L2 norm), thus failing to span the full [0, 1] intervals. These normalisations yield dimensionless numbers but not truly comparable ones. As a result, the weights applied to such pseudo-normalised criteria (in the [0, 1] spanning sense) retain the burden of resolving both scale disparities and preference intensities, thereby confounding measurement with judgement. In such cases, the aggregated output, typically a composite score or ranking, rests on ambiguous foundations. It is unclear whether the ranking reflects actual preferences or is merely an artefact of hidden scale effects that have been absorbed (but not resolved) by the weighting process. The supposed clarity of trade-offs dissolves under scrutiny: Is criterion A preferred because it is more important or because its scale was narrower mapped and thus less amplified by the weighting vector, or both? The illusion is complete when decision-

makers believe they have articulated their preferences clearly, when in fact they have merely masked a scale incoherence.

This illusion is not only a technical flaw; it becomes a cognitive trap in a wider MCDM setting. It misleads decision-makers and analysts alike into believing that decision models reflect informed value judgements, when in fact they often reflect arbitrary or inconsistent scale manipulations. The only robust escape from this illusion is to enforce rigorous scale alignment in the methods before weighting, typically through full-range (spanned) normalisations to [0, 1], and to preserve this interpretability throughout the analysis. Anything less invites semantic ambiguity, mathematical confusion, and decision analyses built on misinterpretations.

This great illusion of multiple scales, i.e. the weight/scale dualism, should not be confused with the illusion of absolute weights, another issue that emerges at the level of MCDM preference elicitation rather than MCDA method computation. It has nothing to do with the methods' calculations, but is a testament to the mental complexities involved in eliciting criteria weights. The latter illusion refers to the cognitive error of assigning fixed importance values to criteria, without regard for the scales they inhabit. The illusion of absolute weights manifests in decision-makers insisting on assigning weights to criteria without considering the original scale spans. For instance, if criterion A is assigned weight w<sub>A</sub> based on a scale [a<sub>1</sub>, a<sub>2</sub>], and a new alternative extends this to [a<sub>1</sub>, a<sub>3</sub>], the weight w<sub>A</sub> needs to be recalibrated so that one unit on A's scale has the same importance as before. The illusion lies in treating weights as if they were anchored in absolute terms, which is impossible, when in fact they are inherently tied to the scales. Failing to understand and accommodate that is falling for the illusion of absolute weights. Thus, as discussed before, it is important to clearly differentiate between the outer MCDM layer where descriptive behaviour, procedures and results can be taken into account, and the inner MCDA layer, which must conform to known objective scientific results. In the outer MCDM layer, compensation can be made for regrets and other human behaviours and biases, although it still has to be done in a traceable way. The transparency requirement does not vanish at the MCDM level.

## 15. Probabilistic MCDA

Part III of the first edition contained an overview of different software applications that employed the methods of Part II. In this second edition, it has been replaced with this chapter on combined probabilistic multi-criteria models and the next chapter, which describes how UNEDA, the open-source universal decision-analytic software platform, is implemented.

The structural similarities between the von Neumann–Morgenstern (vNM) and Keeney-Raiffa's (KR) IIASA utility models indicate that they are not competing frameworks, but rather special cases of more general probabilistic multi-criteria decision analytic models (MPDA). This theory integrates both risk and multi-dimensionality by considering preferences over uncertain, multi-attribute alternatives.

In this unified framework, an alternative is characterised by a matrix of compound outcomes, where each attribute has multiple probabilistic (Bayesian) outcomes. This nested structure expresses vNM utility as the special case where there is only one attribute and only uncertainty exists, and KR as the case where uncertainty is removed (i.e., all  $p_{ij}$  are degenerate, with probability 1 on a single state). Thus, MPDA generalises both. When attribute weights represent relative importance and probabilities represent uncertainty, the resulting model supports decisions under both value trade-offs and risk. The utility function applies consistently across the two cases, indicating that both models rely on the same fundamental valuation mechanism. Both vNM and KR build on core axioms: completeness, transitivity, continuity, independence, and decomposability. These remain valid in the general case and justify the functional form of as both additive and expected.

There are several benefits of a unified view. It brings coherence to decision making under hybrid conditions (e.g., strategic planning with uncertain costs and competing objectives). Further, it supports more precise elicitation: decision-makers can assess trade-offs and risks in tandem. Lastly, it reinforces the idea that utility is the core construct, whether over lotteries, attributes, or both. This unified view validates the effort to develop MPDA methods that respect both probabilistic and multicriteria aspects. The desiderata serve as a scaffold for such synthesis, and their expansion into this domain may mark the next frontier in decision theory.

In complex real-world decision situations, alternatives often involve uncertainty in addition to, not instead of, multiple criteria. A natural extension of MCDA thus involves incorporating probability distributions over outcomes, leading to hybrid models where both criteria weighting and probabilistic beliefs play a role. This generalisation leads to expressions of the form

$$U(A) = \sum_{i=1}^{n} w_i \cdot \sum_{j} p_{ij} u(x_{ij})$$

where  $w_i$  is the weight of criterion, representing its importance,  $p_{ij}$  is the probability of state j under criterion i, and  $u(x_{ij})$  is the utility of outcome  $x_{ij}$  under that state and criterion. This formulation reflects an additive multi-attribute expected utility function. It is consistent with both vNM and KR formulations. The outer sum represents aggregation over attributes, as in MAUT. The inner sum represents expectations over uncertain events within each attribute, as in vNM. Importantly, this model preserves the axiomatic commitments of both theories. i) additivity across independent criteria; ii) expected utility within each uncertain dimension, and iii) coherence in the joint treatment of trade-offs and risk.

A generalised MCDA of this kind opens up doors to richer, securely grounded models. It allows method designers (and thus their customers, the decision-makers) to accommodate both subjective probabilities and value trade-offs in a unified model. It supports elicitation techniques familiar from both MAUT (e.g., swing weighting) and vNM (e.g., lottery comparisons). Although vNM utility theory and KR/MAUT align closely in structure and intent, their merger into a unified probabilistic multi-criteria framework raises some questions that have to be addressed. This section examines whether any modifications are necessary to either theory to ensure consistency and whether they violate each other's fundamental axioms.

Compatibility of Axioms: At a high level, the core axioms shared by both frameworks, such as completeness, transitivity, continuity, and a form of independence, are broadly consistent. However, the definition and application of the independence axiom differs in that vNM requires probabilistic independence while KR requires utility independence, i.e. preferences over one attribute remain unchanged regardless of fixed levels of other attributes. These are structurally distinct. Probabilistic independence governs mixtures of lotteries, while utility independence governs the

separability of trade-offs. In a unified model, one must accept both forms for it to function as intended.

**Decomposability:** KR assumes additive decomposability under specific forms of independence. vNM requires linearity in probabilities but has no native treatment of attribute composition. Combining both requires assuming that utility is additively separable in attributes and linear in probabilities. This dual requirement imposes a stronger structure than either theory individually.

**Functional Form:** To align vNM and KR under one expression, it must be assumed that the same utility function applies across both probabilistic and multi-attribute domains. This may require rescaling or transforming attribute-specific value functions in MAUT to be consistent with cardinal utility in vNM.

Weight Interpretation: vNM is typically used in contexts with measurable uncertainty; KR often treats uncertainty implicitly through scoring. A unified theory implies that attribute weights and probabilities should be formally equivalent in the role they play within the utility aggregation. This requires an interpretation of weights that is stronger than mere preference intensity, they must be utility-theoretic scalars.

Thus, while no outright axiomatic contradiction exists, a unified model imposes stronger assumptions than either theory individually. In particular i) utility independence and probabilistic independence must coexist, ii) additivity across both probabilities and attributes must be assumed, and iii) a common utility function must serve both. These are manageable but nontrivial requirements. Their adoption transforms both vNM and KR from context-specific models into components of a more general system. To reconcile and extend vNM and KR within a general probabilistic multi-criteria decision framework, the following unifying desiderata are proposed. They are designed to support utility representations of the form

$$U(a) = \sum_{i} w_{i} \sum_{j} p_{ij} u(x_{ij})$$

where  $w_i$  are the weights of the criteria (attributes),  $p_{ij}$  are the probabilities over the outcomes under the criteria, and  $u(x_{ij})$  is the utility of outcome  $x_{ij}$ .

**Desideratum MP1** (Completeness and Transitivity): For all alternatives A, B, and C, preferences are complete and transitive. For all A and B, either A > B, B > A, or  $A \sim B$ . Further, if A > B and B > C then A > C.

**Desideratum MP2** (Continuity): For any alternatives A, B, and C, with A > B > C, there exists a  $\lambda \in (0, 1)$  such that  $B \sim \lambda \cdot A + (1-\lambda) \cdot C$  This applies both to probabilistic mixtures (as in vNM) and to attribute trade-offs (as in KR).

**Desideratum MP3** (Probabilistic Independence): For all alternatives A, B, and C, if A > B, then for any  $\lambda \in (0, 1)$ :  $\lambda \cdot A + (1-\lambda) \cdot C > \lambda \cdot B + (1-\lambda) \cdot C$ . This ensures linearity in probabilities.

**Desideratum MP4** (Utility Independence of Attributes): For any attribute i, preferences over levels of i are independent of the fixed levels of other attributes, provided the preferences are conditional on those fixed levels.

**Desideratum MP5** (Additive Decomposability): If utility independence holds for all attributes, then the overall utility function U(a) is additive across attributes and linear in probabilities such that  $U(a) = \sum_i w_i \sum_j p_{ij} u(x_{ij})$ .

**Desideratum MP6** (Monotonicity): If an outcome  $x_{ij}$  is replaced by  $x'_{ij}$  such that  $U(x'_{ij}) > U(x'_{ij})$ , and all other terms remain fixed, then the overall utility increases.

**Desideratum MP7** (Normalisation): For all weights  $w_i \ge 0$ ,  $\Sigma w_i = 1$  and probabilities  $p_{ij} \ge 0$ ,  $\Sigma p_{ij} = 1$  respectively for each i and j.

**Desideratum MP8** (Common Utility Representation): There exists a single cardinal utility function u defined over outcomes  $x_{ij}$  such that preferences over all combinations of attributes and uncertainties can be represented by U(A).

These desiderata, DAMS-MP, unify the vNM and KR theories into a single coherent foundation for MPDA. They allow trade-offs across attributes and beliefs while preserving coherence and a clear interpretative structure. DAMS-MP forms the basis for the UNEDA platform, which can handle tri-linear MPDA decision problems of the form

$$max \left[ U(a) = \sum_{i} w_{i} \sum_{j} p_{ij} u(x_{ij}) \right]$$

according to the generalised PMEU principle afforded by MPDA and with arbitrary depths in the event trees and criteria hierarchies. The open-source software library is described in the next chapter.

# 16. Computational Evaluation

To make a decision analysis method computational, and thus making it a method for real-life decisions, two main ingredients are necessary. The first is a suitable representation and evaluation rules of the decision problems, such as those presented in Part I. The other is reasonably fast computational algorithms, which is the topic of this part. Most of the demanding computations required are optimisation-related algorithms.

This chapter is divided into three main sections. The first deals with calculating properties of decision frames using linear programming methods and the second deals with algorithms for computing evaluation rules by employing bilinear optimisation. The last section contains a discussion of the BEDA method for handling second-order information. The two first sections are built on (Danielson, 1997), which describes the DELTA method for interval decision analysis that was later generalised to multi-level trees (the original text handles only single-level trees, but the generalisation is straightforward and does not introduce any new concepts). Decisions under risk (probabilistic decisions) are often given a tree representation. This is the reading of the tree as a sequence of events leading up to the final consequences, the end nodes.

A decision tree consists of a root node, representing a decision, a set of intermediary (event) nodes, representing some kind of uncertainty about which event will eventually occur, and consequence nodes, representing possible final outcomes. Usually, probability distributions are assigned in the form of weights in the probability nodes as measures of the uncertainties involved. The informal semantics are simply that given that an alternative  $A_i$  is chosen, there is a probability  $p_{ij}$  that an event will occur. This event can either be a consequence with a value  $v_{ijk}$  assigned to it or another event. Usually, the maximisation of the expected value is used as an evaluation rule. In the case of precise probability and utility assessments, this is straightforwardly evaluated. However, when the probabilities and utilities are imprecise, several complications appear, including the non-uniqueness of the expected value of an alternative (leading to the need to find upper and lower bounds). The first step in obtaining a solution is generalising the decision tree structure.

Let a decision frame represent a tree decision problem. This is convenient for presentational purposes. The idea with such a frame is to collect all information necessary for the model in one structure. One of the building blocks of a decision frame is a graph.

**Definition**: A graph is a structure  $\langle I, N, E \rangle$ , where I is an index set, N is a set  $\{n_i\}$ ,  $i \in I$ , of nodes, and E is a set  $\{(n_i, n_j)\}$ ,  $i, j \in I$ ,  $i \neq j$ , of edges (node pairs). A **tree** is a connected graph without cycles.

**Definition**: An r-tree (rooted tree) is a tree  $\langle I,N,E,r\rangle$  where exactly one node  $n_r$  has the property  $\neg \exists \ k : (n_k,n_r) \in E.$   $n_r$  is called the root of the tree. The set N is partitioned into two subsets of leaf nodes  $(N^L)$  and intermediate nodes  $(N^I)$ .  $n_i \in N^I$  **iff**  $\exists \ k : (n_i,n_k) \in E.$  Since  $N^L = N \setminus N^I$ ,  $n_i \in N^L$  **iff**  $\neg \exists \ k : (n_i,n_k) \in E.$  The index set I is partitioned accordingly: an index  $i \in I^I$  **iff**  $n_i \in N^I$  and an index  $i \in I^L$  **iff**  $n_i \in N^L$ . An intermediate node  $n_i \in N^I$  has children indices  $C_i = \{j : (n_i,n_j) \in E\}$ .

Then all the rooted trees representing alternatives are joined together into a decision frame. In the sequel, the notation is used that the n children of a node  $x_i$  are denoted,  $x_{i1}, x_{i2}, ..., x_{in}$  and the m children of the node  $x_{ij}$  are denoted  $x_{ij1}, x_{ij2}, ..., x_{ijm}$ , etc.

Decision-maker statements of probability and value are translated into constraints (inequalities) in order to be entered into the decision problem. Range statements (i.e. intervals) translate into range constraints, inequalities involving a single variable. A reasonable interpretation of such statements is that the estimate is not outside of the given interval. For a value scale [a, b], there is a default range constraint  $v_{ij} \in [a, b]$  for each value variable. Likewise, there is a default range constraint  $p_{ij} \in [0, 1]$  for each probability variable (although, in practice, the normalisation takes care of this). Comparative statements compare the probabilities of two consequences occurring with one another, such as "the events  $C_1$  and  $C_2$  are equally probable" or "the event  $C_3$  is more likely to occur than  $C_4$ ". Those statements are translated into comparative constraints, inequalities involving more than one variable. The term interval constraints is used for the kinds of constraints above. A collection of interval constraints concerning the same set of variables is called a constraint set, and it forms the basis for the representation of decision situation statements.

**Terminology**: Given an index set I and a set of variables  $\{x_i\}_{i\in I}$ , a constraint set in  $\{x_i\}_{i\in I}$  is a set of interval constraints in  $\{x_i\}_{i\in I}$ .

To begin with, it is important to determine whether the elements in a constraint set are at all compatible with each other. This is the question of whether a constraint set has a solution, i.e. if there exists any vector of real numbers that can be assigned to the variables.

**Definition**: Given an index set I and a set of variables  $\{x_i\}_{i\in I}$ , a constraint set X in  $\{x_i\}_{i\in I}$  is consistent **iff** the system of weak inequalities in X has a solution. Otherwise, the constraint set is inconsistent. A constraint Z is consistent with a constraint set X **iff** the constraint set  $\{Z\} \cup X$  is consistent. The collection of all consistent instances of a constraint set X is called the solution set to X.

**Definition**: Given an index set I and a consistent constraint set X in  $\{x_i\}_{i\in I}$  and a function f, the maximum is  ${}^X$ max $(f(x)) =_{\text{def}} \sup \{a \mid \{f(x) > a\} \cup X \text{ is consistent}\}$ . Similarly, the minimum is  ${}^X$ min $(f(x)) =_{\text{def}} \inf \{a \mid \{f(x) < a\} \cup X \text{ is consistent}\}$ .

**Definition**: Given an index set I, a consistent constraint set X in  $\{x_i\}_{i\in I}$  and a function f,  $^X$ argmax(f(x)) is a solution vector that is a solution to  $^X$ max(f(x)), and  $^X$ argmin(f(x)) is a solution vector that is a solution to  $^X$ min(f(x)).

Note that argmax and argmin need not be unique. The feasible box (i.e., the set of feasible variable assignments) can be calculated if the constraint set is consistent. The feasible box is a concept that in each dimension signals which parts are infeasible within the constraint set. Intuitively, the feasible box represents a conservative extension of the solution set of a set of constraints.

**Definition**: Given an index set I and a consistent constraint set X in  $\{x_i\}_{i\in I}$ , the set of optimum pairs  $\{\langle {}^X{\rm min}(x_i), {}^X{\rm max}(x_i)\rangle \}_{i\in I}$  is the feasible box (orthogonal hull) of the set and is denoted  $\langle {}^X{\rm min}(x_i), {}^X{\rm max}(x_i)\rangle_I$ .

This feasible box represents upper and lower probabilities if X consists of probabilities and upper and lower values if X consists of values. For convexity reasons, the entire interval between those extremal points is feasible. Using this concept, an application program can display to the user which statements are incompatible or which parts of intervals are incompatible with the rest of the statement

set. Hence, at all times, an application program can maintain a consistent model of the user's problem in collaboration with the user.

There are two types of constraint sets (*c-sets*), probability c-sets and value c-sets. The smallest c-set unit is the event node c-set, which collects all probability statements made regarding a specific event node in an r-tree.

**Definition**: Given an r-tree  $T = \langle I, N, E, r \rangle$  and an event node  $n_i$ , consider the set  $C_i$  of disjoint and exhaustive consequences of the event (children nodes), user event statements in  $\{p_j\}_{j\in C_i}$ , and a discrete, finite probability mass function  $\Pi: n_j \to [0,1]$  over  $C_i$ . Let  $p_j$  denote the function value  $\Pi(n_j)$ .  $\Pi$  obeys the standard probability axioms, and thus  $p_j \in [0,1]$  and  $\Sigma_j$   $p_j = 1$  are default constraints. Then the event node c-set  $P_i$  is derived from the set of user range and comparative statements with the following content.

- A feasible box ⟨a<sub>k</sub>,b<sub>k</sub>⟩, k∈C<sub>i</sub>, which represents the user and default range constraints ∀k∈C<sub>i</sub>: p<sub>k</sub>∈[0,1].
- All user comparative constraints.
- The normalisation constraint  $\sum_{k \in C_i} p_k = 1$ .

Thus, the c-set transforms statements into linear constraints while maintaining the same meaning. A c-set is more convenient to handle than a pure set of statements. An event node c-set characterises a set of discrete probability distributions. The next aggregation level is that of a probability c-set, which collects together all probability statements belonging to all nodes in the same tree.

**Definition**: Given an r-tree  $T = \langle I, N, E, r \rangle$  with all event nodes  $n_i$ ,  $i \in I^I$ . Then the probability c-set P is all event c-sets  $P_j$  combined, i.e. feasible boxes, normalisations, and user comparative statements.

Requirements similar to those for probability variables are found for value variables. There are apparent similarities and differences between probability and value statements. The normalisation ( $\Sigma_k$   $p_{ik}=1$ ) requires the probability variables of an intermediate node to sum to one. No such constraint exists for the value variables.

Further, the value scale endpoints can be arbitrarily selected and need not be [0,1] as in the probability case.

**Definition**: Given an r-tree  $T = \langle I, N, E, r \rangle$ , consider the set  $N^L$  of leaf nodes. Then a **value c-set** is derived from the set of user range and comparative statements. The user statements, together with the default statements  $\forall k \in I^L : v_k \in [0,1]$ , form the c-set constraints in the following way.

- A hull  $\langle a_k, b_k \rangle$ ,  $k \in I^L$ , which represents the user and default range constraints.
- All user comparative constraints.

Similar to probability c-sets, a value c-set characterises a set of value functions. The statements are transformed into a set of linear constraints. Using the above concepts of constraint and c-set, a decision situation is modelled by a decision frame. To begin with, each alternative is represented by a tree frame.

**Definition**: Given a decision alternative, statements are made about the probabilities of the events as well as the values of the consequences. A **tree frame** is a structure  $\langle T,P,V \rangle$  containing the following representation of the alternative:

- A rooted tree  $T = \langle I, N, E, r \rangle$  with index set partitions  $I^I$  and  $I^L$ , and, for each  $i \in I^I$ , the child index set  $C_i$ .
- A probability c-set P in variables  $\{p_i\}$ ,  $i \in I \setminus \{r\}$ , representing all probability statements in the form of a feasible box and constraints.
- A value c-set V in variables  $\{v_i\}$ ,  $i \in I^L$ , representing all value statements in the form of a feasible box and constraints.

All alternatives are modelled in the same structure. This structure (the decision frame) fully represents the entire decision problem, and all evaluations are made relative to it. The probability and value c-sets, together with structural information, constitute the decision frame.

**Definition**: Given a probabilistic decision situation with m alternatives, a decision frame is a structure  $\langle m,F\rangle$ ,  $F=\{F_i\}$  for  $i\in\{1,...,m\}$ , where  $F_i=\langle T_i,P_i,V_i\rangle$  is a tree

frame for alternative A<sub>i</sub>. Thus, the decision frame contains, for each alternative, a decision tree structure and a tree frame.

Now that the representation structure is defined, the next item is algorithms for computing upper and lower bounds for the expected value in the tree, i.e. optimisation of sums of products derived from the tree structure. The primary evaluation rule is based on the expected value. Since neither probabilities nor values are fixed numbers, evaluating the expected value yields multi-linear objective functions (with bilinear functions as a special case for one-level trees). Evaluate the expected value of an alternative given a decision frame  $\langle m, \{\langle T_i, P_i, V_i \rangle \} \rangle$ , i.e.

$$EV(A_i) = \sum_{i_1=1}^{n_{i_0}} p_{ii_1} \sum_{i_2=1}^{n_{i_1}} p_{ii_1i_2} \dots \sum_{i_{m-1}=1}^{n_{i_{m-2}}} p_{ii_1i_2} \dots \sum_{i_{m-2}i_{m-1}}^{n_{i_{m-1}}} p_{ii_1i_2} \dots \sum_{i_{m-2}i_{m-1}i_m}^{n_{i_{m-1}}} v_{ii_1i_2} \dots \sum_{i_{m-2}i_{m-1}i_m}^{n_{i_{m-1}i_m}} v_{ii_1i_2} \dots \sum_$$

where  $p_{i,i,m}$ ,  $j \in \{1,...,m\}$  denote probabilities in  $P_i$  and  $v_{i,i,m}$  denote values in  $V_i$ . Optimisation of such non-linear expressions subject to linear constraints (the probability and value constraint sets) are described in (Danielson, 1997).

The contraction is a generalised sensitivity analysis to be carried out in an arbitrary number of dimensions. In non-trivial decision situations, when an information frame contains numerically imprecise information, the different principles suggested above are often too weak to yield a conclusive result. Often, a far too crowded set of candidates is received. One way to proceed could be to determine the stability of the relation between the consequence sets under consideration. A natural way to investigate this is to consider values near the boundaries of the intervals as being less reliable than more central values due to interval statements being deliberately imprecise. This is taken into account by measuring the dominated regions indirectly using the concept of contraction.

The principle of contraction is motivated by the difficulties of performing simultaneous sensitivity analysis in several dimensions at the same time. It can be hard to gain a real understanding of the solutions to large decision problems using only one-dimensional analyses since different combinations of dimensions can be critical to the evaluation results. Investigating all possible such combinations would lead to a procedure of high complexity in the number of cases to investigate. Using contractions, this difficulty is circumvented. The contraction avoids the complexity

inherent in combinatorial analyses. However, it is still possible to study the stability of a result by gaining a better understanding of how important the interval boundary points are. By co-varying the contractions of an arbitrary set of intervals, it is possible to gain much better insight into the influence of the structure of the information frame on the solutions. Both the set of intervals under investigation and the scale of individual contractions can be controlled. Consequently, a contraction can be regarded as a focus parameter that zooms in on central sub-intervals of the full statement intervals.

**Definition**: X is a base with the variables  $x_1, ..., x_n, \pi \in [0,1]$  is a real number, and  $\{\pi_i \in [0,1] : i = 1,...,n\}$  is a set of real numbers.  $[a_i, b_i]$  is the interval corresponding to the variable  $x_i$  in the solution set of the base, and  $\bar{k} = (k_1,...,k_n)$  is a consistent point in X. A  $\pi$ -contraction of X is to add the interval statements  $\{x_i \in [a_i + \pi \cdot \pi_i \cdot (k_i - a_i), b_i - \pi \cdot \pi_i \cdot (b_i - k_i)] : i = 1,...,n\}$  to the base X.  $\bar{k}$  is called the contraction point (or focal point).

By varying  $\pi$  from 0 to 1, the intervals are decreased proportionally using the gain factors in the  $\pi_i$ -set, thereby facilitating the study of co-variation among the variables. This is a form of sensitivity analysis, which is described in more detail in (Danielson, 1997). In order to assess the properties of a frame, computational methods are required that can determine whether a given base has a particular property or not. One of the most fundamental components is a way of determining consistency in a base. Since the base consists of a linear system of inequalities, a natural candidate area for an algorithm is linear programming.

The area of linear programming (LP) was formed in the 1940s and has been a large and lively area of research ever since. It deals with the maximising (or minimising) of a linear function with a large number of likewise linear constraints in the form of weak inequalities. Research efforts in the field are partly focused on developing clever algorithms for fast numerical computations. This chapter assumes that the reader is familiar with the basics of LP in general and with the Simplex method in particular. Those unfamiliar with these subjects may refer to any standard text-book on the subject. The LP problem is the following optimising problem:

```
\max f(\mathbf{x})
when \mathbf{A}\mathbf{x} \ge \mathbf{b}
and \mathbf{x} \ge \mathbf{0}
```

where  $f(\mathbf{x})$  is a linear expression of the type  $k_1x_1 + k_2x_2 + \ldots + k_nx_n$ ,  $\mathbf{A}\mathbf{x} \geq \mathbf{b}$  is a matrix inequality with rows  $a_{11}x_1 + a_{12}x_2 + \ldots + a_{1n}x_n \geq b_1$  through  $a_{m1}x_1 + a_{m2}x_2 + \ldots + a_{mn}x_n \geq b_m$ , and  $\mathbf{x} \geq \mathbf{0}$  are the non-negativity constraints  $x_i \geq 0$  for each variable. Amongst all feasible points, the solution to  $f(\mathbf{x})$  is sought that has the highest numerical value, i.e. the best solution vector  $\mathbf{x}$ , the components of which are all non-negative and satisfy all constraints. A minimum can be searched for by negating  $f(\mathbf{x})$ .

### Consistency

The first algorithm is a procedure for determining whether a base is consistent or not. A base is consistent if any solution whatsoever can be found to the set of interval constraints. Note the similarities with the LP problem formulation. Let there be m interval constraints in the base. By introducing new variables  $y_1, ..., y_k$ , with  $k = 2 \cdot m$ , to the consistency problem, it can be reformulated as

```
min (y_1 + ... + y_k)
when \mathbf{A}\mathbf{x} \ge \mathbf{b}
and \mathbf{x} \ge \mathbf{0}, \mathbf{y} \ge \mathbf{0}
```

where each of the interval constraints  $a_{i1}x_1 + a_{i2}x_2 + \ldots + a_{in}x_n \in [a,b]$  is transformed into corresponding inequalities  $a_{i1}x_1 + a_{i2}x_2 + \ldots + a_{in}x_n + y_{2i-1} \ge a$  and  $a_{i1}x_1 + a_{i2}x_2 + \ldots + a_{in}x_n - y_{2i} \le b$ . If the obtained minimum of  $y_1 + \ldots + y_k$  has the value zero, then a solution has been found that does not contain any  $y_j$ . Removing the  $y_j$ 's, the resulting solution vector  $\mathbf{x}$  is indeed a feasible solution, i.e., the base is determined to be consistent. If the minimum of  $y_1 + \ldots + y_k$  is positive, then the optimal values of the  $y_j$ 's are larger than zero, i.e. at least one of the  $y_j$ 's is necessary to keep the base consistent. Since the  $y_j$ 's were added to the base, the problem itself has no solution. Hence, the base is inconsistent. This forms the algorithm for determining consistency in a decision frame by applying it to the probability and value bases.

### **Orthogonal Hull**

Another important property of a base is the orthogonal hull. According to the definition, in order to calculate the hull, find the pairs  $\langle ^X \text{min}(x_i), ^X \text{max}(x_i) \rangle_n$ , i.e. finding minima and maxima for single variables in the base. First, a consistent point is determined by employing the procedure above. A search then begins from that point for the minimum and maximum of each variable in turn by forming LP problems with that variable as the objective function. For convexity reasons, the entire interval between those extremal points is feasible. If the base is consistent, the orthogonal hull can be calculated. From the two properties consistency and orthogonal hull, most of the other ones follow from less demanding computations.

### **Evaluation Algorithms**

The problem addressed in this section is how to compare the different consequence sets computationally using the methods of the previous chapter. The computational pattern that reoccurs several times in that chapter and needs to be solved fast in long sequences is  ${}^{PV}\text{max}(\Delta_{ij})$  and  ${}^{PV}\text{min}(\Delta_{ij})$ . The optimisation of general  $\Delta_{ij}$ -type of expressions as they appear in Chapter 5 is a demanding computational task as soon as the problem to solve is above toy size. In most cases, however, the expected value rule is employed, making the task less demanding from a computational point of view. In this section, it is assumed that the expected value is being used. Then, the general  ${}^{PV}\text{max}(\Delta_{ij})$  turns into  ${}^{PV}\text{max}(\sum_k p_{ik} - \sum_k p_{jk})$  for first order  $\Delta$ -dominance such as 1SE and security levels, and into  ${}^{PV}\text{max}(\sum_k p_{ik} \cdot v_{ik} - \sum_k p_{jk} \cdot v_{jk})$  for second order ones such as 2SE or NE.

### **First Order Dominance**

For first order dominance, the evaluation expressions are of the form

$${}^{P}max \left( \sum_{k \in K_{i}} p_{ik} \right) \text{ or } {}^{P}max \left( \sum_{k \in K_{i}} p_{ik} - \sum_{k \in K_{i}} p_{jk} \right) \text{ (or corresponding $^{P}$min)}$$

for some index sets  $K_i$  or index set pairs  $(K_i, K_i)(d)$  respectively.

These maximisation problems map directly onto LP since it is possible to identify the linear  $f(\mathbf{x})$  with  $\sum_k p_{ik}$  or  $\sum_k p_{ik} - \sum_k p_{jk}$  and note that  $A\mathbf{x} \geq \mathbf{b}$  is the probability base P. The solution to the problem is thus obtained by running a suitable LP solver such as Simplex described later in the chapter. This is an efficient solution to first order problems.

#### **Second Order Dominance**

For second-order dominance, the expressions are more complicated. They involve non-linear elements in the form of bilinear terms  $p_{ik} \cdot v_{ik}$ . The optimisation problems  $^{PV}\max(\sum_k p_{ik} \cdot v_{ik})$  and  $^{PV}\max(\sum_k p_{ik} \cdot v_{ik} - \sum_k p_{jk} \cdot v_{jk})$  cannot be solved by a simple application of an LP solver even if the P- and V-bases are independent and still consist of only linear expressions. The objective function is  $\sum_k p_{ik} \cdot v_{ik} - \sum_k p_{jk} \cdot v_{jk} = p_{i1} \cdot v_{i1} + p_{i2} \cdot v_{i2} + \ldots + p_{im_i} \cdot v_{im_i} - (p_{j1} \cdot v_{j1} + p_{j2} \cdot v_{j2} + \ldots + p_{jm_j} \cdot v_{jm_j})$ . This is a bilinear expression with all terms of the form  $p_{ik} \cdot v_{ik}$ . There is one such expression together with many linear inequalities. Thus, it is an optimisation problem with a bilinear objective function and a system of linear inequalities as constraints. It will be called a bilinear programming problem with  $\pm 1$  term constants (a BLP1 problem for short).

Two alternative algorithms for use in an interactive environment are proposed. The bilinear objective function is an instance of quadratic objective functions, and thus the general BLP1 is solvable with quadratic programming (QP) methods. A QP-based one is the most general, able to solve all BLP1 problems, but it is not as fast as desired for interactive use for larger decision problems. The other algorithm is LP-based and is well-suited for user interaction. Since the bilinear objective function is quadratic, the first natural candidate area for a solver algorithm is quadratic programming.

#### **Quadratic Programming**

The theory of QP can be found in any standard textbook on non-linear optimisation. Here, only the top-level procedure for searching quadratic optima is considered. The general QP problem with both equalities and inequalities in the constraints is

(QPI) 
$$\max (\mathbf{x}^T \mathbf{Q} \mathbf{x} + \mathbf{c}^T \mathbf{x})$$
  
when  $\mathbf{A} \mathbf{x} > \mathbf{b}$ 

where **A** is a m × n matrix with linearly independent rows, **Q** is a symmetric n × n matrix, and **c** is a vector in  $\mathbb{R}^n$ . The expression  $\mathbf{x}^T\mathbf{Q}\mathbf{x}$  is a quadratic form and can contain all possible quadratic terms.

Since the objective function is quadratic, the theory of linear programming as discussed above does not apply. Even though a method similar to Simplex was originally devised by Dantzig and Wolfe to solve QP, most methods today use factorised matrices. For any given solution the inequality problem QPI can be considered a problem with only equalities (QPE), namely all weak inequalities satisfied without slack. Since the other inequalities are not active at that solution point they need not be considered locally. This reasoning leads to the active set strategy, a well-known technique within non-linear programming. One of the problems with the active set is that its members at any given step are hard to determine in advance. This means resorting to a guessing strategy, where a choice is made without enough information and corrected later on should the choice be proven unsuitable. QPE problems can be solved using a number of standard methods such as Lagrange methods or nullspace methods, depending on matrix sparsity, stability requirements, and other criteria. The BLP1 problem maps well onto QPI since there is one second-order nonlinear expression as the objective function and a larger number of linear constraints in the probability and value bases. The bilinear objective function is a special case of a quadratic function where most of the entries in the **Q** matrix are zero. This forms the basis for the general QB-Opt algorithm.

**Observation:** Given a decision frame  $\langle C, P_3, V_3 \rangle$ ,  $^{PV} max(\delta_{ij}) = max (\mathbf{x}^T \mathbf{Q} \mathbf{x} + \mathbf{c}^T \mathbf{x})$  with  $\delta_{ij}$  as  $\mathbf{x}^T \mathbf{Q} \mathbf{x}$ ,  $\mathbf{0}$  as  $\mathbf{c}^T \mathbf{x}$  and PV as  $\mathbf{A} \mathbf{x} \geq \mathbf{b}$ .

The QPE is computationally fairly demanding, and QPI, being an iterative sequence of QPEs, is even more so. Since QPI often does not admit interactive response times, it would be preferable to use an LP-based solver instead. This is possible in most cases using PB-Opt below. Together with QB-Opt, it forms a solver hierarchy from which the fastest is selected for each given optimisation problem.

### **Probability Bilinear Optimisation**

The LP-based algorithm described is the probability bilinear optimisation, PB-Opt. For  $^{PV}$ max( $\sum_k p_{ik} \cdot v_{ik}$ ) it solves the general BLP1 problem for  $\langle C, P_3, V_2 \rangle$ -frames while for  $^{PV}$ max( $\sum_k p_{ik} \cdot v_{ik} - \sum_k p_{jk} \cdot v_{jk}$ ) it solves all cases where there are no comparative constraints between the consequence sets involved in the calculation, either directly or indirectly. To begin with, expressions of maximal and minimal probabilities are introduced.

**Definition:** Given a decision frame  $\langle C,P,V \rangle$ ,

$$VE_i^{max}$$
 is  $\sum_{k=1}^{m_i} p_{ik} \cdot b_{ik}$ , where  $b_{ik} = V \max(v_{ik})$ .

$${}^{V}E_{j}^{min}$$
 is  $\sum_{k=1}^{m_{j}}p_{jk}$ .  $b_{jk}$ , where  $b_{jk} = {}^{V}min(v_{jk})$ .

$$V\delta_{ij}$$
 is  $VE_i^{max} - VE_j^{min}$ .

The last difference was formed from two linear expressions in only probability variables. The main proposition for PB-Opt is now stated as follows.

**Proposition:** Given a decision frame  $\langle C, P_3, V_2 \rangle$ . If none of the comparative constraints in V involve variables from different  $C_i$ 's, then  ${}^{PV}max(\delta_{ij}) = {}^{P}max({}^{V}\delta_{ij})$  for any pair  $C_i$  and  $C_j$ .

**Proof:** Let  $(b_{i1},\ldots,b_{imi})$  and  $(b_{j1},\ldots,b_{jmj})$  be as in the definitions of  ${}^{V}E_{i}^{\ max}$  and  ${}^{V}E_{j}^{\ min}$  above. For all feasible vectors  $(p_{i1},\ldots,p_{imi}), (p_{j1},\ldots,p_{jmj}), (v_{i1},\ldots,v_{imi}),$  and  $(v_{j1},\ldots,v_{jmj})$   ${}^{V}E_{i}^{\ max} \geq \sum_{k} p_{ik} \cdot v_{ik}$  and  ${}^{V}E_{j}^{\ min} \leq \sum_{k} p_{jk} \cdot v_{jk}.$  It follows from  $b_{ik} = {}^{V}max(v_{ik})$  and  $b_{jk} = {}^{V}min(v_{jk})$  and from  $p_{ik} \geq 0 \ \forall \ k \in \{1,\ldots,m_i\}$  and  $p_{jk} \geq 0 \ \forall \ k \in \{1,\ldots,m_j\}.$  This implies  ${}^{V}\delta_{ij} \geq \sum_{k} p_{ik} \cdot v_{ik} - \sum_{k} p_{jk} \cdot v_{jk}.$   $C_i$  contains  $m_i$  consequences. Given two integers  $1 \leq k, l \leq m_i$ , assume  $b_{ik} = {}^{V}max(v_{ik}).$  Then for  $v_{il}$ , either i) there is no comparison  $v_{il} - v_{ik} \in [a,b]$  in V, in which case  $v_{il}$  is independent of  $v_{ik}$ , or ii) there is a comparison  $v_{il} - v_{ik} \in [a,b].$  For case ii), the constraint can be written ii a)  $v_{il} \geq a + v_{ik}$  and ii b)  $v_{il} \leq b + v_{ik}.$  In ii a)  $v_{ik}$  does not constrain the maximisation of  $v_{il}$ , and in ii b)  $v_{ik} = b_{ik}$ 

maximises  $v_{il}$ . Thus  $v_{ik}$  and  $v_{il}$  can be independently maximised and  $(b_{i1},...,b_{im_i})$  is a feasible vector as is  $(b_{j1},...,b_{jm_j})$  by a similar argument. Since there are no constraints  $v_{ik} - v_{jl} \in [c,d]$  in V for different  $C_i$  and  $C_j$ , each  $b_{ik}$  in  $(b_{i1},...,b_{im_i})$  and each  $b_{jk}$  in  $(b_{j1},...,b_{jm_j})$  can be chosen within a consequence set independently of the other sets.

This justifies the basis for the PB-Opt algorithm. The rest of the algorithm almost suggests itself. It searches for the optimum  $^P\text{max}(^V\delta_{ij})$  by means of an LP algorithm such as Simplex. The proposition then guarantees that  $^{PV}\text{max}(\delta_{ij})$  can be determined by calculating  $^P\text{max}(^V\delta_{ij})$  instead provided the precondition is met. Similarly,  $^{PV}\text{max}(\sum_k p_{ik} \cdot v_{ik})$  can be found by searching for an LP solution instead.

### **Second-Order Computations**

The DELTA Method is a distribution-free decision analysis method for the handling and evaluation of decision and risk trees (Danielson, 1997). It has thereafter in 2001–2002 been extended from probabilistic decision situations also to cover decisions under multiple criteria. Decision alternatives are evaluated by so-called contractions of the intervals combined with several complementary evaluation rules. The advantage of a distribution-free approach is the generality and freedom from assumptions that it allows. However, a disadvantage is the unintuitive interpretation of the results of a contraction. In order to alleviate that problem, an additional analysis method is introduced in this report, based on a belief mass interpretation of the output intervals from Delta. Each input and output interval consists of a lower bound, an upper bound, and a focal point. These three points are interpreted as parameters for belief distributions (Dirichlet distributions for probabilities and criteria weights, triangle distributions for values).

A key observation in the DELTA method is that the belief in points closer to the endpoints of the intervals is lower than the belief in more central points. This is the reason for the contraction procedure above. The same observation underlies the BEDA method, but it is effectuated differently – by assigning explicit distributions of belief on the intervals. The distributions used for expressing beliefs are well-known distributions from statistics: the Dirichlet distribution for probabilities (since

they need to sum to one following Kolmogorov's axiom system) and the triangle and uniform distributions for utilities/values, the choice depending on whether there are two or three points defining an interval. The properties of both Dirichlet and triangle distributions are well described in (Kotz and van Dorp, 2004). To see how it works, begin by revisiting the expression for the expected value:

$$EV(A_i) = \sum_{i_1=1}^{n_0} p_{ii_1} \sum_{i_2=1}^{n_{i_1}} p_{ii_1i_2} \cdots \sum_{i_{m-1}=1}^{n_{i_{m-2}}} p_{ii_1i_2} \cdots_{i_{m-2}i_{m-1}} \sum_{i_m=1}^{n_{i_{m-1}}} p_{ii_1i_2} \cdots_{i_{m-2}i_{m-1}i_m} v_{ii_1i_2} \cdots_{i_{m-2}i_{m-1}i_m} 1,$$

To evaluate this expression, and thus arrive at an analysis of the decision situation, employ calculation methods for the two operators addition and multiplication. The addition operator is handled by ordinary convolution, i.e. if h is the distribution over a sum z = x + y whose components have distributions f(x) and g(y), then h(z) is

$$h(z) = \frac{d}{dz} \int_0^z f(x)g(z - x)dx.$$

The multiplication operator is treated analogously. Using the same assumptions as above, if h is the distribution over a product  $z = x \cdot y$ , h(z) is found by letting

$$H(z) = \iint\limits_{\Gamma_x} f(x)g(y)dxdy = \int_0^1 \int_0^{z/x} f(x)g(y)dxdy = \int_z^1 f(x)G(z/x)dx$$

where *G* is a primitive function to g,  $\Gamma_z = \{(x,y) \mid x \cdot y \le z\}$ , and  $0 \le z \le 1$ . Then h(z) is the corresponding density function

$$h(z) = \frac{d}{dz} \int_z^1 f(x)G(z/x)dx = \int_z^1 \frac{f(x)g(z/x)}{x}dx.$$

In theory, the products are calculated and the abovementioned convolution of two densities then effectuates the summations of the products. This combination of operators computes the distribution over the expected utility. In practice, however, these calculations are very complicated for a decision-analytic tool to carry out, especially when additional requirements are added, such as asymmetry in the input distributions and truncated distributions due to the input intervals being narrower than the default [0, 1] range assumed in the standard theory.

The evaluation method in BEDA is based on the principle of going concern (PGC). It is the same PGC observation that enables the use of probability theory as

a risk calculus. The probability of an event occurring is the proportion of times it occurs if the event is repeated an infinite number of times. In using probabilities for modelling real-life events, the approximation is used that the probability best represents the risk involved. For this approximation to be reasonable, several events need to take place for the real-world outcomes to cancel out in the sense that they, on average, tend to the probability. This is the assumption of going concern, and the approximation is viable in most decision situations, which is why probability calculus is accepted for use in this way. The same PGC reasoning applied to distributions involves the central limit theorem and the law of large numbers in statistics. This leads to the well-founded approximation that the total distribution of expected value over a large number of decision situations will tend to the normal distribution. Using this approximation, the evaluation in the BEDA method amounts to finding parameters for a suitable approximately normal distribution. Two factors slightly complicate matters. i) The input distributions are seldom symmetric in the sense that their mean values are not midway between the lower and upper boundaries of the intervals. And even if they were, the multiplication operator's non-linearity still yields an asymmetric result. ii) The lower and upper bounds themselves introduce truncations into the resulting distributions, leading to non-standard outcomes. This eventually turns the BEDA evaluation into a moment calculus using the NEMO (net moment) technique. NEMO includes all moments that have a noticeable impact on the end result and excludes those that have negligible impact to save computation time. For a detailed description of BEDA and NEMO, refer to the documentation on the UNEDA webpage.

This chapter builds on (Danielson, 1997, Ch.6)

## **Universal Engine for Decision Analysis**

The software platform UNEDA (Universal Engine for Decision Analysis) has been developed in parallel with the book over a extended period of time. A substantial amount of material associated with UNEDA has never been published, except on the author's university webpages. Those documents cover aspects of prescriptive decision analysis that have now been incorporated into the UNEDA computational engine.

UNEDA is an open-source library for MPDA (probabilistic MCDA). It implements the DAMS-MP framework from Chapter 15 and thus integrates the two fields of probabilistic and multi-criteria decision analysis into a unified computational environment. The library is freely available to use for any purpose, academic and non-profit as well as (from June 6, 2025) commercial. The software library can be accessed via the GitHub repository at

#### github.com/uneda-cda/UNEDA

and the documentation is found at a link in the same repository. The original release of the software platform can also be found via the DOI link

doi.org/10.5281/zenodo.15114623.

Background material for UNEDA is available via links in the GitHub repository.



DOI 10.5281/zenodo.15114623

## **Appendix**

Some graduate students found the discussion and Table 2 in Chapter 13 to be a little too condensed to follow with ease. As a consequence, the teaching material was supplemented with some explanatory notes. In the second edition of the book, these notes have been added as an appendix. It discusses the Big Five methods through the DAMS desiderata lens in a structured form and in more detail. In the following, all methods include both formula and thresholds and all alternative sets are finite.

**Desideratum 1** (Ordering): The preference relation is complete. For all A and B, either A > B, B > A, or  $A \sim B$ . This requirement together with Desideratum 2 implies that the decision-maker's preferences can be modelled as a complete weak order (or total pre-order), i.e. an ordered partition into indifference classes.

Method	Fulfils D1	Comment
VIKOR	OK	VIKOR ranks all alternatives using a compromise measure (Q) combining: (1) the utility (distance from the ideal) and (2) the regret (maximum deviation for any criterion). The underlying preference relation induced by Q <sub>i</sub> is a complete weak order.
TOPSIS	OK	TOPSIS assigns each alternative a scalar closeness coefficient based on its distance to the ideal and anti-ideal points. Because these scalar scores are totally ordered, all alternatives are comparable.
ÉLECTRE	NO by design	ÉLECTRE produces outranking relations (not full preference orders), with thresholds, veto rules, and incomparability. These design features include incompleteness and non-transitivity to model non-compensatory and threshold-based preferences.
PROM I	NO	PROMÉTHÉE I allows for incomparability; thus, it does not satisfy completeness.
PROM II	OK	PROMÉTHÉE II computes a scalar net flow for each alternative and imposes a complete ranking, ensuring that pairs of alternatives are comparable.
AHP	OK	AHP has a matrix output allowing comparisons.

**Desideratum 2** (Transitivity): The preference relation is transitive: If A > B and B > C, then A > C. This requirement together with Desideratum 1 implies that the decision-maker's preferences can be modelled as a complete weak order.

Method	Fulfils D2	Comment
VIKOR	OK	VIKOR ranks all alternatives using a compromise measure (Q) combining: (1) the utility (distance from the ideal) and (2) the regret (maximum deviation for any criterion). The underlying preference relation induced by Q <sub>i</sub> is a complete and transitive weak order.
TOPSIS	OK	TOPSIS assigns each alternative a scalar closeness coefficient based on its distance to the ideal and anti-ideal points. Because these scalar scores are totally ordered, every pair of alternatives is comparable and the strict preference relation is transitive.
ÉLECTRE	NO by design	ÉLECTRE produces outranking relations (not full preference orders), with thresholds, veto rules, and allowance for incomparability. These design features include incompleteness and non-transitivity to try to model such preferences in the core.
PROM I	NO	PROMÉTHÉE I allows for incomparability; thus, it does not satisfy completeness. Even when it asserts $A > B$ and $B > C$ , it may leave A and C incomparable due to conflicting flow values, so transitivity is not guaranteed either.
PROM II	OK after ag- gregation	PROMÉTHÉE II computes a scalar net flow for each alternative and imposes a complete ranking, ensuring that every pair of alternatives is comparable. Because these net flow values are real numbers, the induced strict preference relation is transitive.

AHP	Only under	AHP assumes transitivity, but it only holds if the
	perfect	pairwise comparison matrix is consistent. In prac-
	consistency	tice, inconsistency is common. The method has a
		consistency index to detect but not enforce it.

VIKOR, TOPSIS and PROMÉTHÉE II satisfy both completeness and transitivity. They provide full rankings based on cardinal values (scores). PROMÉTHÉE I falls short because of two rankings as output which may differ, declaring the alternatives incomparable. ÉLECTRE is explicitly built to *not* require completeness or transitivity. It is more aligned with partial and ambiguous preferences. AHP is formally complete and transitive if judgement matrices are consistent, but that is a big if in practice since a consistency index of zero is notoriously hard to obtain.

**Desideratum 3** (Dominance): If for all i,  $s_i(A) \ge s_i(B)$  and for some j,  $s_j(A) > s_j(B)$  then A > B. This is often referred to as Pareto dominance or the strong dominance rule. This desideratum demands that if alternative A is at least as good as B in every criterion, and strictly better in at least one, then A must be strictly preferred over B. It is a straightforward principle in spirit, *better is better*, and important for rational consistency.

Method	Fulfils D3	Comment
VIKOR	Partly	The Q-index blends utility (S) and regret (R). Be-
		cause $Q_i$ is monotonically increasing in both $S_i$
		and $R_i$ , it follows that $Q_A > Q_B \Rightarrow B > A$ if it had
		not been for the C1/C2 threshold rules.
TOPSIS	OK	TOPSIS compares relative closeness to the ideal
		vs. anti-ideal. A dominant alternative will always
		have equal or better distance metrics, leading to a
		better closeness coefficient.
ÉLECTRE	NO	ÉLECTRE is built on outranking, not dominance.
		It allows for thresholds, veto effects, and incompa-
		rability. A dominated alternative (in the Pareto
		sense) might appear above a dominant one in the
		final ranking, due to indifference or thresholds.

PROM I	NO	PROMÉTHÉE I does not always respect dominance due to thresholds. A dominating alternative might get the same preference flows. Further, incomparability is still possible, which can obscure strict dominance.
PROM II	Partly	PROMÉTHÉE II does not have incomparability in its output but still suffers from thresholds. Depends on which stage you call "the relation". Once reduced to a single real score, dominance is respected (so the <i>final</i> ranking is consistent).
АНР	Only under perfect consistency	AHP is vulnerable to dominance violations due to inconsistencies in pairwise comparisons. If the decision-maker's judgements do not reflect dominance, the eigenvector weights may still assign a higher rank to a dominated alternative. Dominance is not structurally enforced.

Only TOPSIS robustly preserves the dominance principle. That is, if one alternative is strictly better in at least one criterion and no worse in the others, it will be preferred. Others may violate this property due to compromise calculations or judgement inconsistencies. ÉLECTRE, by design, allows outranking contradictions and veto thresholds to override dominance, reflecting its commitment to modelling partial and conflicted preferences rather than idealised rational consistency.

**Desideratum 4** (Monotonicity): If A > B, and A' is such that  $s_i(A) = s_i(A')$  for all  $i \neq j$  and  $s_i(A) = s_i(A') + \varepsilon$  for some small  $\varepsilon > 0$ , then  $A' \geq B$ . This says that if A is better than B, and an alternative A' is created that is worse than A in one criterion, then A' should not be worse than B (i.e., B should not leap ahead just because A' got slightly worse than A in a differential sense). This is a stability condition: weakening a better alternative should not reverse an established preference.

Method	Fulfils D4	Comment
VIKOR	OK	VIKOR preserves monotonicity: worsening an al-
		ternative increases both the utility loss (S) and po-

	ı	
		tentially the regret (R), which raises the compro-
		mise index (Q). Even if the worsening shifts the
		maximum-regret criterion from a low-weighted to
		a high-weighted one, the regret term cannot cause
		the ranking to flip. And C1/C2 can at worst make
		A and B equally preferred (incomplete ranking).
TOPSIS	OK	When an alternative is worsened, its distance to
		the ideal increases and its closeness coefficient de-
		creases. The method ensures that preference is
		preserved unless the worsening is large enough to
		fully reverse the closeness relation, making it con-
		sistent with monotonicity.
ÉLECTRE	NO	ÉLECTRE can violate this rule due to veto thresh-
	not neces-	olds and discordance. If A' is slightly worse than
	sarily	A in a vetoed criterion, it might lose the outrank-
		ing status, even if B is globally worse. Also, in-
		comparability may replace a previous strict prefer-
		ence.
PROM I	NO	PROMÉTHÉE I respects this property when A'
		and B are comparable because of monotonic pref-
		erence functions. If A' is strictly worse than A, it
		will score slightly lower but still above B if A did.
		However, it can collapse the order into incompara-
		bility due to how flows are separately compared.
PROM II	OK	PROMÉTHÉE II assigns net flow scores via pair-
		wise comparisons. A' being slightly worse than A
		reduces its score while still being comparable.
AHP	NO	AHP depends on subjective pairwise judgements,
	no guarantee	not direct performance scores. If $A > B$ because
	g	of a particular judgement and then A' worsens in
		objective terms, it does not automatically follow
		that $A' > B$ . The pairwise matrix might yield dif-
		ferent eigenvalue-based rankings, especially if in-
		consistencies are present.
		consistencies are present.

The monotonicity property that worsening a preferred alternative should not reverse its dominance over a clearly inferior one is fully respected by three Big Five methods. ÉLECTRE and PROMÉTHÉE may switch preference status based on thresholds or veto conditions. AHP, based on subjective judgements, has no structural mechanism to enforce this kind of ordinal stability.

**Desideratum 5** (Independence of Irrelevant Alternatives): If A > B in set X, and  $C \notin \{A, B\}$ , then A > B in  $X \cup C$ , provided that criteria weights are automatically adjusted to preserve the relative importances of one unit on each original scale if C caused any scale renormalisations.

Method	Fulfils D5	Comment
VIKOR	OK	VIKOR normalises criterion values based on the best and worst in the current set, which can alter the loss profiles and rankings when a new alternative is introduced. If the weights are automatically rescaled, the relative preference between two unchanged alternatives remains stable.
TOPSIS	OK	TOPSIS defines ideal and anti-ideal reference points based on the full set of alternatives, making its normalisation sensitive to the presence of new options. If weights are rescaled, the preference between any two unchanged alternatives remains stable.
ÉLECTRE	NO	ÉLECTRE is based on pairwise outranking with thresholds, and adding a new alternative C can change the concordance/discordance matrices, especially if C introduces new veto situations or changes credibility scores. It satisfies IIA at the outranking level but violates it at the final ranking level.
PROM I	NO	PROMÉTHÉE I uses pairwise preference flows, and adding C introduces new comparisons (A vs. C and B vs. C), which can affect overall flow values. Thus, the preference between A and B may change.
PROM II	NO	Same as PROMÉTHÉE I: net flow scores are recomputed based on all pairwise comparisons.

		Adding an irrelevant third option can redistribute outranking flows, which may alter $A > B$ .
АНР	NO	AHP uses pairwise comparison matrices, so adding C creates new comparisons (A vs. C, B vs. C, etc.). The derived priority vector can shift even if $A > B$ held before.

Two of the five benchmarked methods satisfy the condition of independence of irrelevant alternatives. Three of them define preferences in relation to the full set of alternatives, whether through outranking flows (PROMÉTHÉE), concordance matrices (ÉLECTRE), or pairwise judgements (AHP). Consequently, adding or removing an option not directly involved in a preference relation (e.g.  $C \notin \{A, B\}$ ) can still cause a reversal of A > B, making the methods context-dependent rather than strictly ordinal.

**Desideratum 6** (Rank Preservation): If A > B in X, and C is a third alternative not affecting the scores of A or B, then removing C from X does not alter the ranking A > B (allowing for automatic weight adjustment to preserve per-unit criterion meaning). This is an instance of the rank reversal property in its removal form. It is a consistency condition under contraction. If C is irrelevant to the comparison between A and B, then removing C should not disturb that comparison. Violation of this desideratum is a hallmark of context-dependent or relativistic MCDA methods where scores are based on entire sets of data.

Method	Fulfils D6	Comment
VIKOR	NO	Even if alternative C does not affect the scores of A or B, its removal changes the threshold DQ, which increases and might obliterate $A > B$ .
TOPSIS	NO	Its use of context-sensitive extreme artificial solutions (ideal and anti-ideal) means that removing an irrelevant alternative can shift the reference points. Even with automatic weight adjustment, the non-linear geometry shifts and closeness can change.

ÉLECTRE	Partly	Although its pairwise outranking relations $A \gtrsim B$ are computed independently, the final rankings depend on the entire set of alternatives. Removing a third alternative C can alter the structural dominance relations or the set of outranked alternatives. However, the many variants of the method have different sensitivity to rank reversal.
PROM I/II	NO	In PROMÉTHÉE, the net flow scores are aggregate constructs over the entire set. Removing an irrelevant alternative C can change the flow balance, and alter the ranking between A and B, despite their pairwise scores being unchanged.
АНР	NO	Its ranking outputs are based on eigenvectors or geometric means of a full matrix. Removing a third alternative changes the dimensional structure, and can shift relative priority values, leading to preference reversals between unchanged pairs.

The desideratum of invariance under irrelevant removal asks whether removing a third alternative leaves the ranking A > B intact. None of the five methods reviewed satisfy this condition. VIKOR trips on the finalising rules C1 and C2. TOPSIS depends on reference points derived from the entire alternative set. Removing C can shift the ideal or anti-ideal positions, thereby altering A's and B's closeness or compromise scores. ÉLECTRE shows some resilience, thanks to its pairwise structure. But also here, changes in the concordance or veto dynamics can distort the ranking. PROMÉTHÉE, relying on net preference flows, also fails: any contraction of the alternative set alters the flow landscape, making outcomes sensitive to seemingly irrelevant options. AHP, built on global pairwise matrices, responds to removal with complete recalibration, often changing weights and ranks. It has no internal model of what a score means and only sees the ratios as directly meaningful.

**Desideratum 7** (Criteria Transparency): For any preference A > B, there exists a representable and decomposable justification based on the contribution of each criterion to the total evaluation. This is a prescriptive rationality requirement. It implies three things: i) that the decision process should be transparent, i.e. not a black

box, ii) the contribution of each criterion to the ranking must be explicit and traceable, and iii) a decision-maker (or stakeholder) should be able to understand and explain why A > B, broken down by criteria.

Method	Fulfils D7	Comment
VIKOR	Partial	VIKOR produces scores S (utility) and R (regret), which are aggregated into the Q-index. While S is decomposable (weighted sum of distances to ideal per criterion), R is non-compensatory, taking the maximum deviation. This makes full decomposition asymmetric and harder to explain. Moreover, conditional decision rules further obscure interpretability.
TOPSIS	OK	Although the closeness coefficient is not additive, it is constructed entirely from per-criterion terms that are geometrically and algebraically interpretable. This allows preferences such as $A > B$ to be justified based on specific criteria that favour A in relation to the ideal solution.
ÉLECTRE	Partial	ÉLECTRE builds an outranking relation using concordance (supporting criteria) and discordance (opposing criteria). While it can be described why A outranks B, veto rules and thresholds make explanations non-additive and discontinuous, and thus hard to justify in scalar terms.
PROM I	Partial	PROMÉTHÉE uses criterion-wise preference functions and produces positive/negative flow contributions. One can break down $A > B$ by examining how much each criterion contributes to A's net flow. The final ranking is based on the net flow $\phi(A)$ which is traceable but blends information about A's performance vs all others.
PROM II	Partial	PROMÉTHÉE II extends this with a complete ranking. Since it is still based on criterion-wise prefer-

		ence functions and weighted net flows, one can generate a step-by-step breakdown of why $A > B$ . Same reservation as I.
АНР	Partial	AHP computes local priorities per criterion and aggregates them into global weights. In theory, the preference A > B can be decomposed. However, because it is based on subjective judgements, and due to eigenvalue-based weighting, explanations may lack clarity. Also, when inconsistencies exist in the matrix, the decomposability decreases.

The desideratum is fully satisfied by VIKOR and partly satisfied by all others. ÉLECTRE, relying on thresholded outranking logic, produces qualitative preference relations that resist scalar decomposition, making justifications difficult to articulate in terms of continuous contribution. AHP, while decomposable in principle, lacks interpretive clarity due to reliance on subjective pairwise judgements and potential inconsistency.

**Desideratum 8** (Weight Sensitivity): Let  $w_i \in [0, 1]$  be weights summing to 1. A change in  $w_i$  that increases the influence of criterion  $C_i$  in which  $s_i(A) \ge s_i(B)$  should not alter the preference A > B. This is a monotonicity property with respect to weights, an important consistency criterion in weight-sensitive MCDA methods. This expresses directional weight monotonicity. If A is already as good as or better than B under criterion  $C_i$  and the weight of  $C_i$  is increased, then A > B should be preserved. It assumes that methods are sensitive to weights in a directionally consistent way.

Method	Fulfils D8	Comment		
VIKOR	NO	VIKOR normalises each criterion into [0, 1], mak-		
		ing weights operate proportionally on commensurable scores. However, the regret term (R) uses a max function, which is sensitive to weight shifts if		
		the criterion is the worst for an alternative.		
TOPSIS	OK	TOPSIS uses vector normalisation, but this does		
		not guarantee that the transformed scores span the		

ÉLECTRE	OK	full [0, 1] interval, often falling within a narrower subrange. Thus, increasing w <sub>i</sub> may have a dampening or unpredictable effect. However, it preserves ranking in the ordinal sense as requested.  ÉLECTRE's concordance relations use weights, but because normalisation is not standardised to [0, 1], and thresholds/veto rules apply, weight multiplication is not cleanly interpretable. But at worst, the preference is not strengthened.
PROM I/II	OK	PROMÉTHÉE apply weightings to preference degrees derived from non-linear, threshold-based functions. These preference functions introduce regions of flat sensitivity (indifference thresholds), where increasing the weight on a favourable criterion has no impact, so at worst the preference is not strengthened.
АНР	NO	AHP works on ratio-scale judgements, not normalised performance data. The weights are derived, not applied, and the final ranking is influenced by the entire matrix, not marginal criterion values. Increasing $w_i$ directly (e.g. through matrix adjustments) does not reliably preserve $A > B$ , and this behaviour is not interpretable as scalar weight application at all.

When evaluating weight sensitivity, it is important to distinguish between the aggregation logic and the normalisation method used by each technique. PRO-MÉTHÉE maintains directional consistency under weight changes when preference functions are well-behaved and monotonic. TOPSIS, despite applying weights to vector-normalised scores, exhibits partial robustness in an ordinal sense. VIKOR, while based on [0,1] scaling, suffers from its max-regret term  $R_i$ , which is insensitive to most weight shifts unless the criterion is dominant, failing to meet the desideratum. ÉLECTRE further complicates weight interpretation through threshold and veto structures, making weight effects non-transparent and context-dependent. AHP, relying on eigenvector-derived weights from subjective matrices, provides no

direct or interpretable mechanism to adjust or monitor criterion impact, and fails the test of rational, directional weight response.

**Desideratum 9** (Criteria Independence): If criteria  $C_i$  and  $C_j$  produce identical scores for all alternatives, the results should be cardinally invariant under merging them into one criterion with a combined weight  $w_i + w_j$ . This is a structural rationality desideratum concerned with weight integrity and redundancy handling, testing method invariance under model equivalence. The desideratum assumes that a method i) should not allow redundant criteria to artificially inflate influence, ii) ensures that merging two identical criteria into one does not distort the result, provided weights are added, and iii) reflects the additivity or sensitivity to criterion structure, essential for model parsimony and usability.

Method	Fulfils D9	Comment
VIKOR	NO	VIKOR's S and R scores are based on weighted $L_1$ (sum) and $L_{\infty}$ (max) distances from the ideal. If two criteria $C_i$ and $C_j$ have identical scores across all alternatives, they contribute twice the same deviation. Regret takes only the largest weighted shortfall, having two identical columns means that the lesser of the two weights risk being ignored.
TOPSIS	NO	The method calculates weighted Euclidean distances to the ideal and anti-ideal points. If two criteria have identical scores across all alternatives, treating them separately introduces half their contributions twice. Treating them separately vs. merged produces different geometric (root-mean-square) L <sub>2</sub> distances.
ÉLECTRE	NO	ÉLECTRE's concordance and discordance matrices rely on weights across all criteria. Redundant criteria each independently contributes to concordance, but the RMS distances ruin the calculus the same way it does for TOPSIS.
PROM I/II	OK	PROMÉTHÉE's preference flows are additive and criterion-wise. If two criteria are identical, their individual preference functions will be identical as well.

		Merging them with adjusted weights produces identical net flows.
АНР	NO	The weights are not user-controlled, they are computed from a pairwise comparison matrix of criteria. Merging two criteria is not a matter of adjusting weights. It means removing a row and column from the criteria matrix and recomputing the entire weight vector. This may change all other weights due to renormalisation and eigenvector sensitivity.

The desideratum of redundancy sensitivity is important for avoiding hidden bias and ensuring model parsimony. Of the five methods reviewed, only PROMÉTHÉE reliably supports this principle: if two identical criteria are merged and their weights recombined, the net flows remain unchanged. VIKOR is also compliant under similar conditions, though its maximum-regret term (R) introduces asymmetry. In contrast, TOPSIS and ÉLECTRE rely on geometric RMS distances that do not reproduce under these conditions. AHP, structured around subjective pairwise comparisons, does not support criterion merging at all, and redundancy is treated as legitimate additional information, a problematic stance in analytical modelling.

**Desideratum 10** (Scale Invariance): For any criterion  $C_i$ , if a positive affine transformation  $f: \mathbb{R} \to \mathbb{R}$  is applied to all  $s_i(\cdot)$ , then the preference relation A > B should remain unchanged. This means that any such transformation that does not change the direction of scores should not alter the ranking of alternatives. Methods that rely on non-linear magnitudes or ratios may violate this desideratum.

Method	Fulfils D10	Comment			
VIKOR	OK	VIKOR calculates deviations from the ideal			
		point, using weighted $L_1$ (S) and $L_{\infty}$ (R)			
		measures. It applies min-max normalisation to			
		each criterion. This transformation is invari-			
		ant under strictly affine functions.			
TOPSIS	NO	Because it relies on vector normalisation and			
		Euclidean L <sub>2</sub> distances, it is sensitive to the			

		magnitude of scores. Even if a criterion's relative distances are preserved, a transformation can change the final ranking.
ÉLECTRE	NO	The method is purely ordinal and thus invariant under affine transformations. But if veto thresholds are used, they rely on numerical gaps which can be distorted under scale changes.
PROM I/II	NO	Relying on preference functions defined over numerical differences, affine transformations of a criterion's scale might trigger thresholds and non-linear preference functions, thereby altering the overall ranking.
АНР	NO	AHP uses subjective pairwise judgements, not score functions. An affine transformation of raw scores changes the subjective ratio, leading to different matrices and altered results.

The desideratum of affine transformation invariance tests whether methods rely solely on the difference of criterion values. ÉLECTRE does not satisfy this condition, even though its outranking logic depends only on ordinal comparisons, due to veto thresholds. PROMÉTHÉE could have respected this property, but only if its preference functions had been truly affine to begin with. In contrast, VIKOR has such a truly affine-respecting construction. TOPSIS is sensitive to value transformations and violate the desideratum. AHP, grounded in subjective comparisons and eigenvector derivation, is scale-dependent and non-invariant by construction.

This concludes the discussion of the Big Five methods. Most of them display some strengths and some weaknesses, while AHP does not do well regarding any desideratum. This is due to it being the farthest away from the well-established and validated axiom systems of von Neumann-Morgenstern and Keeney-Raiffa that the DAMS desiderata are built on and reflect. There is no other coherent set of axiom systems that could optionally be adhered to, thereby invalidating the "smorgasbord approach" sometimes advocated for as a replacement for rigorous foundations.

All desiderata are summarised in Table A1, which is structurally identical to Table 2 but pivoted and more nuanced since desiderata can be partially fulfilled. The judgement in the new table is not always black and white. SAW methods such as SMART are not included in Table A1 since they are not discussed in the appendix due to their conformance with the DAMS desiderata set.

		70	CTRE	II/I	
$\mathbf{Methods} \rightarrow$	VIKOR	FOPSIS	ECI	PROM I/II	<b>a</b>
<b>Desiderata</b> ↓	<b>N</b>	TO	ÉLEC	PR	AHP
D1. Ordering	OK	OK	NO	PART	OK
D2. Transitivity	OK	OK	NO	PART	COND
D3. Dominance	PART	OK	NO	PART	COND
D4. Monotonicity	OK	OK	NO	PART	NO
D5. Indep. Irrelevant Alt.	OK	OK	NO	NO	NO
D6. Rank Preservation	NO	NO	PART	NO	NO
D7. Transparency	PART	OK	PART	PART	PART
D8. Weight Sensitivity	NO	OK	OK	OK	NO
D9. Criteria Independence	NO	NO	NO	OK	NO
D10. Scale Invariance	OK	NO	NO	NO	NO

Table A1. The Big Five methods reassessed using the DAMS desiderata

The appendix ends with a summary in Table A2 of some major pros and cons of each of the five methods discussed. For more detailed strengths and weaknesses, refer to the discussions on each desideratum in this appendix and also the methods' respective chapters in Part II. The last column indicates whether the methods can use the UNEDA open-source platform with some modifications.

Method	Strengths	Weaknesses	UNEDA
VIKOR	Scale-robust, value- sensitive, compro- mise-aware, optimisa- tion-style	Violates context and transformation invariance; R component obscures monotonicity	YES
TOPSIS	Transparent, stable under scaling, geometrically explainable	Fails on IIA, context dependence, rank re- versal	YES
ÉLECTRE	Ordinal reasoning, soft appearance, par- tial comparability	Lacks transparency, threshold-sensitive, context-dependent, partial ranking	YES
PROM I/II	Decomposable, monotonic, stable under value transformation	Fails rank reversal, IIA, version I has only partial ranking	YES
АНР	Conceptual simplicity, consistent when judgements are	Sensitive to inconsistency, scale, and context; non-additive, lacks interpretability	NO

**Table A2**. Some major pros and cons of the Big Five methods

To conclude, there are essentially three major lines of development within the MCDA method spectrum. The first is the classical tradition, grounded in established theoretical frameworks. The second is the ÉLECTRE lineage, which four of the five methods in the appendix belong to, and which accepts additive utility but discards much of the rest. The third is AHP, the least compliant of the three lineages and consequently labelled as fundamentally flawed (Abbas, 2018, Ch.3). Unfortunately, the latter two categories have attracted the most attention in the last decades, diverting focus and resources away from real progress in the field. Brand recognition, arguably one of the most important success factors, is not addressed in this appendix. A 2023 ranking of brand name visibility among the Big Five methods lists: (1) AHP, (2) TOPSIS, (3) PROMÉTHÉE, (4) ÉLECTRE and (5) VIKOR. No other methods came close to their levels of recognition. With AHP on top, the ranking closely resembles an inverse of their DAMS compliance, highlighting both the importance and effectiveness of branding efforts.

### References

Abbas, A. E. (2018). *Foundations of multiattribute utility*. Cambridge University Press.

Ackoff, R. L. (1962). *Scientific method: Optimizing applied research decisions*. John Wiley & Sons.

Allais, M. (1953). Fondements d'une théorie positive des choix comportant un risque et critique des postulats et axiomes de l'école américaine. D. Reidel Publishing Company.

Bayes, T. (1763/1963). An essay towards solving a problem in the doctrine of chances. *Philosophical Transactions of the Royal Society of London*, 53, 370–418.

Belton, V., & Gear, T. (1983). On a short-coming of Saaty's method of analytic hierarchies. *Omega*, 11(3), 227–230.

Belton, V., & Stewart, T. J. (2002). *Multiple Criteria Decision Analysis: An Integrated Approach*. Kluwer Academic Publishers.

Benayoun, R., Roy, B., & Sussmann, B. (1966). ÉLECTRE: Une méthode pour guider le choix en présence de points de vue multiples. *Note de Travail No. 49*, Société d'Économie et de Mathématiques Appliquées, Direction Scientifique.

Benayoun, R., & Sussmann, B. (1966). Manuel de référence du programme ÉLECTRE. *Note de Synthèse et Formation No. 25*. Société d'Économie et de Mathématiques Appliquées, Direction Scientifique.

Bernoulli, D. (1738/1954). Specimen theoriae novae de mensura sortis [Theory on the measurement of risk] (L. Sommer, Trans.). *Econometrica*, 22, 22–36. (Original work published 1738)

Brans, J.-P. (1982). L'élaboration d'instruments d'aide à la décision. In R. Nadeau & M. Landry (Eds.), *L'aide à la décision: Nature, instruments et perspectives d'avenir* (pp. 183–214). Les Presses de l'Université Laval.

Brans, J.-P., & Vincke, P. (1985). A preference ranking organisation method. *Management Science*, *31*(6), 647–656.

Cardano, G. (1663/1953). *Liber de ludo aleae* [Book on games of chance]. (Reprinted 1953, Princeton University Press).

Choquet, G. (1953). Theory of capacities. Annales de l'Institut Fourier, 5, 131–295.

Danielson, M. (1997). Computational decision analysis (Doctoral dissertation).

Department of Computer and Systems Sciences, Royal Institute of Technology.

Danielson, M., & Ekenberg, L. (2016). The CAR method for preference strength in multi-criteria decision making. *Group Decision and Negotiation*, 25(4), 775–797.

Dantzig, G. B. (1947). Linear programming and extensions. Princeton Univ. Press.

David, F. N. (1962). *Games, gods and gambling: The origins and history of probability and statistical ideas*. Charles Griffin & Company.

David, L., & Duckstein, L. (1976). Multicriterion ranking of alternative long-range water resource systems. *Water Resources Bulletin*, 12(4), 731–754.

Debreu, G. (1952). Definite and proper preferences. Econometrica, 20(2), 295-309.

Dempster, A. P. (1967). Upper and lower probabilities induced by a multivalued mapping. *Annals of Mathematical Statistics*, *38*, 325–339.

Doyle, J., & Thomason, R. H. (1997). Background to qualitative decision theory.

In J. Doyle & R. H. Thomason (Eds.), *Qualitative preferences in deliberation and practical reasoning*. AAAI Press. Republished as: Doyle, J., & Thomason, R. H. (1999). Background to qualitative decision theory. *AI Magazine*, 20(2), 55–68.

Duckstein, L., & Opricović, S. (1980). Multiobjective optimization in river basin development. *Water Resources Research*, *16*(1), 14–20.

Dyer, J. S. (1990). Remarks on the Analytic Hierarchy Process. *Management Science*, *36*(3), 249–258.

Edwards, W. (1977). *How to make good decisions in the face of uncertainty*. RAND Corporation.

Edwards, W., & Barron, F. H. (1994). SMARTS and SMARTER: Improved simple methods for multiattribute utility measurement. *Organizational Behavior and Human Decision Processes*, 60(3), 306–325.

Ellsberg, D. (1961). Risk, ambiguity, and the Savage axioms. *Quarterly Journal of Economics*, 75, 643–669.

Fibonacci, L. (1202/1857/2002). *Liber abaci* (B. Boncompagni, Ed. 1857). Rome: Tipografia delle Scienze Matematiche e Fisiche. Republished as: *Liber abaci* (L. E. Sigler, Trans. 2002). Springer-Verlag. (Original work published 1202)

Fischhoff, B., Goitein, B., & Shapira, Z. (1981). Subjective expected utility:

A model of decision-making. *Journal of the American Society for Information Science*, 32(5), 391–399.

Fishburn, P. C. (1981). Subjective expected utility: A review of normative theories. *Theory and Decision*, 13, 139–199.

Fishburn, P. C. (1983). Transitive measurable utility. *Journal of Economic Theory*, *31*, 293–317.

French, S. (1986). *Decision theory: An introduction to the mathematics of rationality*. Ellis Horwood.

Friedman, M. (1953). The methodology of positive economics. In M. Friedman (Ed.), *Essays in positive economics* (pp. 3–43). University of Chicago Press.

Funtowicz, S. O., & Ravetz, J. R. (1990). *Uncertainty and quality in science for public policy*. Kluwer Academic Publishers.

Føllesdal, D. (1984). Values, context and justification. In W. Krawietz, H. Schelsky, & D. H. Scheffler (Eds.), *Wertproblem und Wissenschaft* (pp. 151–163). Duncker & Humblot.

Gauss, C. F. (1809/1871/2011). Theoria motus corporum coelestium in sectionibus conicis solem ambientium. Perthes et Besser.

Good, I. J. (1962). Subjective probability as the measure of a non-measurable set. In P. Suppes, P.T. Nagel, & A. Tarski (Eds.), *Logic, methodology, and the philosophy of science* (pp. 319–329). Stanford University Press.

Greco, S., Ehrgott, M., & Figueira, J. R. (Eds.). (2016). *Multiple criteria decision analysis: State of the art surveys*. Springer.

Gärdenfors, P., & Sahlin, N.-E. (1982). Unreliable probabilities, risk taking, and decision making. *Synthese*, *53*, 361–386.

Gärdenfors, P., & Sahlin, N.-E. (1983). Decision making with unreliable probabilities. *British Journal of Mathematical and Statistical Psychology*, *36*(2), 240–251.

Hacking, I. (1975). The emergence of probability. Cambridge University Press.

Herstein, I. N., & Milnor, J. (1953). An axiomatic approach to measurable utility. *Econometrica*, 21, 291–297.

Howard, R. A. (1966). Decision analysis: Applied decision theory. *Proceedings of the Fourth International Conference on Operational Research*, 55–71.

Howard, R. A. (2009). Foundations of decision analysis revisited. *Decision Analysis*, 6(3), 216–246.

Huber, P. J. (1973). The case of Choquet capacities in statistics. *Bulletin of the International Statistical Institute*, 45, 181–188.

Hurwicz, L. (1951). Optimality criteria for decision making under ignorance. *Cowles Commission Discussion Paper No. 370*.

Huygens, C. (1657/1929/1959). De ratiociniis in ludo aleae. In D. E. Smith (Ed. & Trans. 1929), *A source book in mathematics* (pp. 75–83). Dover Publications.

Hwang, C. L., & Yoon, K. (1981). *Multiple attribute decision making: Methods and applications*. Springer-Verlag.

Jaynes, E. T. (2003). *Probability theory: The logic of science* (G. L. Bretthorst, Ed.). Cambridge University Press.

Keeney, R. L. (1992). *Value-focused thinking: A path to creative decisionmaking*. Harvard University Press.

Keeney, R. L., & Raiffa, H. (1976/1993). *Decisions with multiple objectives: Preferences and value trade-offs* (Rev. ed.). Cambridge University Press. (Original work published 1976)

Keynes, J. M. (1921). A treatise on probability. Macmillan and Co.

Kolmogorov, A. N. (1933). *Grundbegriffe der Wahrscheinlichkeitsrechnung*. Springer-Verlag.

Kolmogorov, A. N. (1938). On the analytic methods of probability theory. *Uspekhi Matematicheskikh Nauk*, *5*, 5–41.

Kotz, S., & van Dorp, J. R. (2004). Beyond beta: Other continuous families of distributions with bounded support and applications. World Scientific Publ.

Krantz, D. H., Luce, R. D., Suppes, P., & Tversky, A. (1971). *Foundations of measurement: Additive and polynomial representations*. Academic Press.

Kuhn, H. W., & Tucker, A. W. (1951). Nonlinear programming. In J. Neyman (Ed.), *Proceedings of the Second Berkeley Symposium on Mathematical Statistics and Probability* (pp. 481–492). University of California Press.

Laplace, P.-S. (1812/1820/1825/2012). *Théorie analytique des probabilités* (3e éd., R. J. Pulskamp, Trans. 2012). Veuve Courcier. (1st ed. 1812)

Laplace, P.-S. (1816/1952). Essai philosophique sur les probabilités (3rd ed., F.

W. Truscott & F. L. Emory, Trans.). Dover. (Original work published 1816)

Levi, I. (1974). On indeterminate probabilities. *The Journal of Philosophy*, 71, 391–418.

Loomes, G., & Sugden, R. (1982). Regret theory: An alternative theory of rational choice under uncertainty. *The Economic Journal*, 92, 805–824.

Luce, R. D., & Raiffa, H. (1957). *Games and decisions: Introduction and critical survey*. John Wiley & Sons.

Malmnäs, P. E. (1994a). Towards a mechanization of real-life decisions. In D. Prawitz & D. Westerståhl (Eds.), *Logic and philosophy of science in Uppsala* (pp. 231–243). Kluwer Academic Publishers.

Malmnäs, P. E. (1994b). Axiomatic justification of the utility principle. *Synthese*, 99, 233–249.

Malmnäs, P. E. (1996). *Evaluations, preferences and choice rules* (Research Report). Department of Philosophy, Stockholm University.

March, J. G., & Simon, H. A. (1958). Organizations. John Wiley & Sons.

Markowitz, H. (1952). Portfolio selection. The Journal of Finance, 7(1), 77–91.

Menger, K. (1934). Das Unsicherheitsmoment in der Wertlehre. *Zeitschrift für Nationalökonomie*, *5*, 459–485.

Milnor, J. (1954). Games against nature. In R. M. Thrall, C. H. Coombs, & R. L. Davis (Eds.), *Decision processes* (pp. 49–60). John Wiley & Sons.

von Mises, R. (1928/1957). Probability, statistics and truth (2nd English ed., H.

L. Cambel, Trans.). The Macmillan Company. (Original work published 1928)

de Moivre, A. (1730). Miscellanea analytica. Tonson & Watts.

de Moivre, A. (1718/1738/1967). *The doctrine of chances: A method of calculating the probability of events in play* (2nd ed. 1738). Chelsea Publishing Company.

de Montmort, P.-R. (1708). Essay d'analyse sur les jeux de hazard. Le Conte.

Morgan, M. G., & Henrion, M. (1990). *Uncertainty: A guide to dealing with uncertainty in quantitative risk and policy analysis*. Cambridge University Press.

von Neumann, J., & Morgenstern, O. (1944). *Theory of games and economic behavior*. Princeton University Press.

von Neumann, J. (1928). Zur Theorie der Gesellschaftsspiele. *Mathematische Annalen*, 100(1), 295–320.

Oddie, G., & Milne, P. (1990). Act and value. *Theoria*, 57, 42–76.

Opricović, S. (1998). *Višekriterijumska optimizacija sistema u građevinarstvu*. Građevinski fakultet Univerziteta u Beogradu. [In Serbian, Cyrillic script].

Opricović, S., & Tzeng, G.-H. (2004). Compromise solution by MCDM methods: A comparative analysis of VIKOR and TOPSIS. *European Journal of Operational Research*, 156(2), 445–455.

Opricović, S., & Tzeng, G.-H. (2007). Extended VIKOR method in comparison with outranking methods. *European Journal of Operational Research*, 178(2), 514–529.

Pacioli, L. (1494/2008). Summa de arithmetica, geometria, proportioni et proportionalità (Facsimile ed.). Aboca Edizioni.

Poisson, S. D. (1837). Recherches sur la probabilité des jugements en matière criminelle et en matière civile. Bachelier.

Quiggin, J. (1982). A theory of anticipated utility. *Journal of Economic Behavior and Organization*, *3*(4), 323–343.

Quetelet, A. (1846). Sur l'homme et le développement de ses facultés, ou Essai de physique sociale. Bachelier.

Raiffa, H. (1968). *Decision analysis: Introductory lectures on choices under uncertainty*. Addison-Wesley.

Ramsey, F. P. (1926/1931). Truth and probability. In R. B. Braithwaite (Ed.), *The foundations of mathematics and other logical essays* (pp. 156–198). Kegan Paul, Trench, Trubner & Co. (Original work written 1926)

Roy, B. (1968). Classement et choix en présence de points de vue multiples (la méthode ÉLECTRE). Revue française d'automatique, d'informatique et de recherche opérationnelle. Série verte, 2(V1), 57–75.

Saaty, T. L. (1977). A scaling method for priorities in hierarchical structures. *Journal of Mathematical Psychology*, 15(3), 234–281.

Saaty, T. L. (1980). The analytic hierarchy process: Planning, priority setting, resource allocation. McGraw-Hill.

Saaty, T. L. (2001). Ratio scales are critical for modeling neural synthesis in the brain. In R. Trappl (Ed.), *Artificial neural nets and genetic algorithms*. Springer.

Savage, L. J. (1954/1972). *The foundations of statistics* (2nd ed. 1972). Dover Publications.

Schoemaker, P. J. H. (1982). The expected utility model: Its variants, purposes, evidence and limitations. *Journal of Economic Literature*, 20, 529–563.

Shafer, G. (1976). A mathematical theory of evidence. Princeton University Press.

Simon, H. A. (1955). A behavioral model of rational choice. *The Quarterly Journal of Economics*, 69(1), 99–118.

Slovic, P. (1974). Assessment of risk taking behavior. *Psychological Bulletin*, 81(4), 330–333.

Smith, C. A. B. (1961). Consistency in statistical inference and decision. *Journal of the Royal Statistical Society: Series B*, 23, 1–25.

Stigler, S. M. (1986). *The history of statistics: The measurement of uncertainty before 1900*. Harvard University Press.

Tversky, A. (1981). The framing of decisions and the psychology of choice. *Science*, 211(4481), 453–458.

Tversky, A., & Kahneman, D. (1986). Rational choice and the framing of decisions. *Journal of Business*, 59(4), Part 2: S251–S278.

Wald, A. (1950). Statistical decision functions. John Wiley & Sons.

Walley, P. (1991). *Statistical reasoning with imprecise probabilities*. Chapman & Hall.

Weichselberger, K., & Pöhlmann, S. (1990). A methodology for uncertainty in knowledge-based systems. Springer-Verlag.

von Winterfeldt, D., & Edwards, W. (1986). *Decision analysis and behavioral research*. Cambridge University Press.

Yaari, M. E. (1987). The dual theory of choice under risk. *Econometrica*, 55(1), 95–115.

Zlaugotne, B., Žihare, L., Balode, L., Kalnbalkite, A., Khabdullin, A., & Blumberga, D. (2020). Multi-criteria decision analysis methods comparison. *Environmental and Climate Technologies*, 24(1), 454–471.

### **About the Author**

**Mats Danielson** is a Full Professor of Computer and Systems Sciences at Stockholm University. He also holds a UNESCO Chair Professorship in Higher Education within Computer Science and is a Senior Research Scholar at IIASA, the International Institute for Applied Systems Analysis in Vienna, the birthplace of formalised MCDA. He has previously served as Vice President of External Relations, Innovation and Information Technology and as Dean of the Faculty of Social Sciences at Stockholm University.

He holds a PhD in Computer and Systems Sciences from KTH Royal Institute of Technology as well as dual degrees in Computer Science and Engineering from KTH and in Economics and Business Administration from the Stockholm University School of Business. Before entering academia, he worked for more than 20 years as an IT and management consultant.

#### Other books from Sine Metu in the same series

Transcending Business Intelligence

K. Borking, M. Danielson, G. Davies, L. Ekenberg, J. Idefeldt, A. Larsson ISBN 978-91-978-4505-2

A Little Book on Error-Free Software

M. Danielson

ISBN 978-91-531-0456-8

A Decision-Analytic Manifesto

L. Ekenberg, M. Danielson

ISBN 978-91-527-5306-4

Computational Intelligence

M. Danielson

ISBN 978-91-527-6234-9

The Rise and Fall of Philips Data Systems

M. Danielson, A. Läppinen

ISBN 978-91-527-6233-2

Foundations of Computational Decision Analysis begins with the established basics of classic probabilistic decision theory and builds towards a critical assessment of multi-criteria decision analysis (MCDA) methods. The book is grounded in the conviction that decision-analytic methods must rest on solid scientific foundations, logical coherence and conceptual clarity to enable transparency.

The first part revisits the roots of decision theory, examining subjective probabilities, utility, and the fundamental role of value in rational choice. It disentangles common confusions, highlights core assumptions, and presents the theory with both philosophical care and real-world decision problem relevance.

The second part turns to MCDA, the expanding family of methods designed to guide the analysis of complex decisions with multiple objectives. Rather than treating these as a set of tools, the book examines them as scientific constructs and potential guides, asking not just how they work, but also why, when and whether they should be trusted as support tools.

The final part deals with computations and an open-source software platform for enabling applications of effective modern decision analysis, with special attention to real-world imprecision and the need for systematic sensitivity analyses.

This is a book for readers who want more than procedural knowledge. It is for those who seek conceptual depth, methodological clarity, and logical reasoning in the design and evaluation of decision-analytic methods. A must-read for decision analysts, students, and anyone serious about the logic of choice.



Sine Metu