Foundations of Computational Decision Analysis

Mats Danielson

Foundations of Computational Decision Analysis

Mats Danielson

Sine Metu

Published by Sine Metu Productions, Stockholm, Sweden, 2023, 2025
Website: www.sinemetu.se

ISBN 978-91-531-0456-8

First edition, September 2023
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. (1853/2009). The Positive Philosophy of Auguste Comte, Vol. II, (H. Martineau, Transl.). Cambridge University Press, Cambridge, MA.

Contents

Part I: Fundamentals of Computational Decision Analysis

Part II: Methods of Multi-Criteria Decision Analysis

Part III: Software for Computational Decision Analysis

Preface

Decision methodology, mainly in the form of decision theory and decision analysis, has been studied for quite some time. Not least a number of Nobel Prize winners in economics have worked in the field, including Simon (1978), Allais (1988), Kahneman (2002) and Hurwicz (2007). Most of them have contributed to *normative* theory, that is, the study of how we should rightfully choose. This is usually in an idealised and theorised form that can't be used in 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're deciding in an optimal way" but they say nothing about how to get there. It's as unhelpful as a theoretical description of how to ride a bicycle. You can't just read the description and then peddle off. Or reading a couple of books on how to swim. Thrown in at 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 fooled by all sorts of information and disinformation. This can be very amusing to read about and Kahneman's book, *Thinking, fast and slow*, is recommended for both entertainment as well as thought-provoking reading. But what we really need is perhaps not a catalogue of cognitive missteps, but rather a method that carries us in a reasonably safe way from decision problems to decisions. Descriptive research, therefore, is unfortunately not much help to us either when we're going to make real decisions of good quality. Continuing the cycling analogy: reading about bicycle accidents and how riders fell off their bikes, and how large their grazes were, will not help us much either. We will still not be able to peddle off after reading about them. The same holds for stories of life guards and swimming incidents.

Fortunately, there's a third research direction, prescriptive decision analysis, which focuses on procedures 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 the book do not rely on unrealistic assumptions about the decision-makers' abilities to present precise information and achieve flawless decisions. Rather, they aim at providing useful support in real-world decision situations of various kinds.

Thus, a focus on the foundations of prescriptive decision analysis means studying real-life applicability, which by necessity implies computer software. There is today a wide gap between decision theory and real-world needs, which unfortunately leads to decision analysis not being used and leveraged in society to the extent it could and ought to, given the many hard decisions that have to be made every day. Writing more large theoretical books will not help – there are already good ones available and most of the stuff in them is unfortunately not applicable to practical decision analysis.

Originally based on the author's Ph.D. thesis from 1997 (Part I), the text has been extended, rewritten, and pruned more than once but not published until now. It is hoped to serve as a basis for rejuvenated research interests in decision-analytic methods solidly based on sound and well-established theoretical research results, both from within decision theory and outside from adjacent theories such as mathematics, statistics and microeconomics. 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 also encompassing the larger process from data collection, over elicitation, analysis, presentation and possibly negotiations, all the way through to making a decision. Thus, Part II is concerned with the core of MCDM, namely MCDA.

Happy reading!

The author, Stockholm, September 2023

Preface to the Second Edition

Decision analysis is, as almost any management or professional support technology, dependent on computer power to be 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 procedures and ad hoc formalisms, 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 mission 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 the software, this book 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, quantitative, and effective approach to making decisions under uncertainty. 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.

The formal study of decision analysis dates back to the mid-20th century, with contributions from pioneers such as von Neumann and Morgenstern, who introduced utility theory, and Savage, who developed subjective expected utility. These foundational ideas have since evolved into various probabilistic and deterministic methods used in fields such as economics, operations research, artificial intelligence, and policy-making.

Bayesian inference was introduced into decision theory, initially through the work of Savage in the mid-20th century. 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 (a key 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 key 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.

Uncertainty is perhaps the most defining characteristic of classic decision analysis. In virtually every decision, there are elements that 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 is either known or 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.

The implications of uncertainty are vast. In high-stakes decisions, such as those faced by healthcare providers, financial institutions, or government policymakers, the ability to manage uncertainty can be the difference between success and failure. Probabilistic 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.

At the core of most probabilistic decision models is the concept of expected utility, a foundational idea in decision theory. 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 vital 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 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.

The power of probabilistic 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 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). Expected utility theory is the most widely used framework in probabilistic decision analysis. Expected utility theory 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. Conversely, a risk-seeking individual may prefer alternatives that offer the potential for higher rewards, even if those alternatives come with higher levels of uncertainty.

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 process is crucial for understanding the robustness of decisions, particularly when there is uncertainty in the underlying 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.

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 16th and 17th 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* (1202) 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) raises 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 they are interrupted and cannot continue when one player has won five rounds and his counterpart has won three (David, 1962, p. 37).

Later, Cardano (1501–1571), an Italian mathematician, physician, and gambler, tried to answer the question in *Liber de ludo aleae* (A book about gambling, 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.

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 17th and 18th 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 19th 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, which describes the probability of a given number of events occurring in a fixed interval of time or space. His work had wide-ranging applications in areas such as physics, finance, and insurance (Poisson, 1837). 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 a crucial 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-19th century, probability theory had developed into a mature mathematical discipline with growing applications in science, economics, and industry. 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 power 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 20th 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 Ω, 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 $\{A_1, \ldots, A_n\}$ and B with outcomes $\{B_1, \ldots, B_m\}$ 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 $P(A_i \mid B_j, C_k) = P(A_i \mid C_k)$.

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 Ω consist 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 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 fact that the distinguished Bernoulli family was in many ways connected to St. Petersburg. In this paradox, Nicolaus Bernoulli considered a fair coin, defined by the property that the probability of heads is ½. This coin is tossed until *head* appears. The gambler is rewarded with 2^n ducats if the first head appears on the n:th trial. The expected monetary value of this game is

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

It is a bit hard to believe that any gambler would be willing to pay an infinite amount of money to participate in such a game. Bernoulli concluded therefore that the expected monetary value is inappropriate as a decision rule. Bernoulli's solution to this paradox involved two ideas that have had a great impact on economic theory. Firstly, he stated that the utility of money cannot be linearly related to the amount of money; it rather increases at a decreasing rate.

Bernoulli identified the value of the consequences of a choice as being different from the objective economic outcome, commonly referred to as the idea of diminishing marginal utility. Bernoulli's second idea is that a person's valuation of a risky prospect is not the expected return of that prospect, but rather the prospect's 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 is the mathematically expected value, when subjective utility is taken into account. In the St. Petersburg Paradox, the value of the game is finite due to the principle of diminishing marginal utility. Originally Bernoulli employed a logarithmic utility function, $u(x) = \alpha \log x$, where the α is dependent on the gambler's wealth prior to the gamble itself, and x is the outcome.

Substituting this value for x in (EMV) yields a finite number. Consequently, people would only be willing to pay a finite amount of money to participate, even though the expected monetary value of the game is infinite.

The term utility can be regarded as a measure of some degree of satisfaction, and a utility function is a mapping from outcomes, i.e., losses or gains, to real numbers representing this degree of satisfaction. The logarithmic utility function defined by Bernoulli was in itself considered adequate 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 on the basis of 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 important 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 β 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 by means of positive affine transformations. Apart from ratio scales, interval scales do not have an absolute zero (e.g., zero length); 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 degrees has a meaning, while the gap between two ratios does not.

In general, people are willing to pay more money for what they consider to be more desirable. In this respect, a monetary scale can at least be expected to be an ordinal scale, i.e., a scale measuring preference ordering without the possibility of stating the strengths of desires. For a majority of business decisions, the use of monetary scales is considered a reasonable and acceptable measure of utility. However, it is not uncommon that monetary values are used to scale non-monetary outcomes, such as public health and environmental damage. In many cases, this problem is due to a lack of means and usable tools for representing and evaluating intangibles and vague valuations. This is particularly troublesome when aggregating ordinal information and can be severely misleading.

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) takes 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 var-*

iable 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 for a short while, 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-20th 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 decision 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., 1983).

Prescriptive decision theory is a more recent perspective, brought about by the inadequacies of the two earlier perspectives in real-life decision situations. It focuses on identifying the discrepancies between how decisions are made (descriptive) and how the normative theory claims they should be made. One purpose of this prescriptive theory is to bridge the apparent gap between older decision analysis and actual decision making.

This body of theory contains approaches that deal with mechanising (rather than automating) the structuring and analysis of decision situations. Presuming the decision-maker has a desire to be rational, the mechanical model can aid by devising suitable courses of action given supplied (elicited) information. A decision-analytic tool handles a finite number of alternative courses of action and supports the decision-maker in the evaluation of and selection among those alternatives. Stated differently, such a tool aids decision-makers in their search for a preference order of a set of alternatives and in their strive for rationality. The remainder of this book will have a prescriptive process perspective, aiming at a foundation for tools and procedures usable in real-life decision situations.

03. A Thorough Foundation 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 something is 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. From each category, some common models are presented.

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_i, C_{ik}, v_{ik}, and s_i are known with certainty, and there is a known deterministic relationship between the choice of an a_i and the corresponding C_{ik}, 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** (wicked)
  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 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 that a decision-maker can have preferences on.

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: $\max(\Sigma_{j\leq n} v_{ij})/n$, where $1 \leq k \leq n$, and where 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 $r_{ij} = \max\{v_{sj}, s = 1,...,m\} - v_{ij}$.
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.

	s_1	s_2	s_3	s_4	**Rule picks alternative**
a_1	2	2	0	1	Laplace
a_2	1	1	1	1	Wald
a_3	0	4	0	0	Hurwicz ($\alpha > ¼$)
a_4	1	3	0	0	Savage

Table 1. Milnor's 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 this book.

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, 1990):

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 $\Sigma p_i V_a(c_i)$ 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 axiomatisation are given in (Fishburn, 1981).

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 normative 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*.

$a \geq_p b$ 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 $u(a) \geq u(b)$ if and only if $a \geq_p b$, 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 $\langle p_1 \cdot v_1, \ldots, p_i \cdot v_i, \ldots, p_r \cdot v_r \rangle$, 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 $a \geq_p b$ or $b \geq_p a$, and if $a \geq_p b$ and $b \geq_p c$ then $a \geq_p c$.

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

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

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

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.

Shifting our attention to the system of Savage, he 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 $\varpi_\iota(c_i)$ 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 primitives: i) the binary preference relation $\geq_p$, ii) a set $S = \{s_1, s_2, \ldots\}$ of states, iii) a set $C = \{c_1, c_2, \ldots\}$ of consequences, and iv) a set $F = \{f: S \rightarrow C\}$ 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. $c_i \leq_p c_j$ if and only if $f \leq_p f'$ when $f(s) = c_i$ and $f'(s) = c_j$, for all $s \in S$.
3. B is null ($B = \varnothing$) if and only if $f \leq_p g$ given B, for all $f,g \in F$.
4. A is not more probable than B ($A \leq B$) if and only if $f_A \leq_p f_B$ or $c_i \leq_p c_j$, 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. $f \leq_p c_i$ given B ($c_i \leq_p f$ given B) if and only if $f \leq_p h$ given B ($h \leq_p f$ given B), when $h(s) = c_i$, for all $s \in S$.

To clarify some of the concepts: 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 *not* A, and f_B given *not* B is preferred to f_A given B, then if f_B is preferred to f_A this means that a decision-maker holds event B more probable than event A (and vice versa). Then Savage proposes the following seven axioms:

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

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

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

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

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

Axiom 6: *Continuity*: Suppose $h \leq_p g$, 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 $h \leq_p g'$ or $h' \leq_p g$.

Axiom 7: *Dominance*: If $f \leq_p g(s)$ given B ($g(s) \leq_p f$ given B) for every $s \in B$, then $f \leq_p g$ given B ($g \leq_p f$ 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 $L_c = \{L_1, L_2, \ldots\}$ be a set of lotteries on C (alternatives with uncertain outcomes in the consequence set C), then they showed that the utility function $u:L_c \rightarrow R$, has a representation $u(L_i) = \Sigma p_i(c_i) \cdot u(c_i)$ 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 vein in the approach of Savage, his theory is often referred to as subjective expected utility.

Descriptive decision theory

Human decision-makers tend to, under given circumstances, behave inconsistent 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 his conclusion is 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 it is impossible for any normative theory of decision making to embrace all inherent peculiarities in a free world of heterogeneous decision-making inhabitants.

However, this perspective has been heavily critiqued. A common counter-argument is that the axioms of utility theory are 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 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, 1988). 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 crucial 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 inconsistent 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., $u(x_{ce}) = \Sigma\, p_i u(x_i)$, where $u(x_i)$ is the utility of gaining the monetary value x_i. The risk premium, **p**, 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., $p = \Sigma\, p_i x_i - x_{ce}$. With respect to the risk premium **p**, a classification of decision-makers into three classes can be made: a decision-maker is risk-averse if **p** > 0; risk-prone if **p** < 0; and risk-neutral if **p** = 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 be in need of 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 her last funds that will give her 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 **p** is probably negative.

However, some argue that it will never be possible to formalise the decision process with all reasonable risk attitudes by a utility function and an associated risk premium. Many critics emphasise that a majority of the mathematical models of decision analysis are oversimplified. Consider, for example, the reasons for gambling. Most people would agree that there is a pleasure involved in the pure act of participating in a game with uncertain outcomes. If mathematical expectation were the only criterion for gambling, no games would ever be arranged by rational beings, since when the rules of the game would make it rational for the gambler to bet, then the arranger should be irrational to offer the bet. However, people do still arrange and participate in games, although either the gambler or the bookmaker will be on the irrational side. Furthermore, it has also been argued that humans tend to disregard very small probabilities, even in games with finite mathematical expectations (like nation-wide lotteries), and that, in the case of very high probabilities, a gambler is not willing to risk arbitrary amounts (Menger, 1934). Such arguments cross the border to descriptive decision theory, and while both important and interesting, they do not aid in the formulation of viable prescriptive decision-analytic theories, models or procedures.

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

04. Decision-Analytic Evaluations

As seen in the previous chapter, much thought has been given to which rules correctly rank alternatives under strict uncertainty. Among the first suggestions were maximin and minimax. Both are in some sense extreme strategies. A generalisation using the weighted average is the Hurwicz rule (Hurwicz, 1951). Another evaluation rule, the minimax-regret (Savage, 1951), is the same rule as minimax but applied to alternative losses instead.

As noted in that chapter, Milnor suggested some reasonable criteria by which evaluation rules could be judged and demonstrated that no rule could comply with all of the requirements (Milnor, 1954). Hence, none of the rules can be universally agreed on as being *the* rule. However, Milnor's requirement of invariance under the addition of a new state by column duplication is too strong. If the decision matrix changes, the entire decision problem has changed into a new one. Then it is not surprising that decision rules run into trouble. Excluding that requirement, Laplace's rule does indeed satisfy all the other requirements. The rule is based on the implicit assumption that uncertainty is the same as assigning equal probabilities to all states. This assumption constitutes a link between methods for making decisions under strict uncertainty and methods for making decisions under risk. Laplace's rule is similar to maximising the expected value and when all probabilities are assigned the same number, the expectation turns into a simple sum.

Consider instead a situation where in addition the decision-maker has some estimates of the probabilities of the states involved. 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 ($A_1,\ldots,A_m$), 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 group 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 $f(p_{i1},\ldots,p_{im_i},v_{i1},\ldots,v_{im_i})$, 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 in spite of 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 $s_{ik} = \sum_{j=1}^{k} p_{ij}$. 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_{ij} is as above.

These suggestions come from a normative standpoint. As noted in Chapter 3, Malmnäs (1996) shows for those above and for other proposals that their performances can at best be equal 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. 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 (1994) is to supplement the expected value with qualitative evaluations. An example is the qualitative sorting function, which has the basic form

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

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 $\langle\{\{C_{ik}\}_{m_i}\}_m,P,V\rangle$, δ_{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 $i,j,k \in N+$ as appropriate. The α and β 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 $a,b,d \in [0,1]$, $M_{ij}[a,b]$ is the set $\{(K_i,K_j)(d) \mid d \in [a,b]\}$.

$M_{ij}[a,b]$ 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 Δ-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 Δ*[I]-dominates* C_j **iff**

$\forall$ $(K_i,K_j)(d) \in M_{ij}[I]$ $\sum_{k \in K_i} f(p_{ik}, v_{ik}, \alpha) - \sum_{k \in K_j} f(p_{jk}, v_{jk}, \beta) \geq 0$ **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 Δ-dominance. The first subdivision of the Δ-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 Δ-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 Δ[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.

Note: C_i *1[I]-dominates* C_j **iff** $\forall$ $(K_i,K_j)(d) \in M_{ij}[I]$ $\sum_{k \in K_i} g(p_{ik}) \geq \sum_{k \in K_j} g(p_{jk})$.

The note points out the resemblance with some familiar dominance concepts. One further specialisation of the first order Δ-dominance is the first order stochastic dominance, a well-known concept. To reach there, the general first order Δ-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 Δ-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 *1SE-dominates* C_j **iff** $\forall (K_i,K_j)(d) \in M_{ij}[I]$ $\sum_{k \in K_i} p_{ik} \geq \sum_{k \in K_j} p_{jk}$.

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 Δ-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 $K_j = \{k \mid v_{jk} \geq t\}$

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

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 Δ-dominance.

The Δ-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 ⟨C,P,V⟩ and functions f and h, C_i *2[I]-dominates* C_j **iff** C_i Δ[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_j **iff** $\forall$ $(K_i,K_j)(d) \in M_{ij}[I]$

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

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 Δ-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 ⟨C,P,V⟩, C_i *2S-dominates* C_j **iff** C_i 2[0,1]-dominates C_j. C_i *2SE-dominates* C_j **iff** C_i 2S-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 Δ-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** $\forall$ $(K_i,K_j)(d) \in M_{ij}[0,1]$ $\sum_{k\in K_i} p_{ik} \cdot v_{ik} \geq \sum_{k\in K_j} p_{jk} \cdot v_{jk} \cdot$

Another second order Δ-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 2Δ-dominance.

Definition: Given a decision frame ⟨C,P,V⟩, 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 $(K_i,K_j)(0)$ $\delta_{ij} \geq 0$.

Also note that $\delta_{ij} \geq 0$ is not applicable 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 Δ-dominance, which acts as a unifying concept in comparing and discussing evaluation rules.

This chapter builds on (Danielson, 1997, Ch.5)

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 0.3 to 0.7 (Shapira 1995). 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 judgments, 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, for example a utility function. For instance, (Rowe, 1994) incorporates more qualitative sources of uncertainty, acknowledging that uncertainty due to communication difficulties and divergent values is an unavoidable aspect of policy decision making. 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. Finally, Wynne (1992) addresses uncertainty in the foundations of information and knowledge, as well as in processing information.

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, 1997). 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 judgments 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 $\Sigma\, p_{ij} = 1$ 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 $\langle\{\{c_{ij}\}_{j=1,\ldots,h_i}\}_{i=1,\ldots,m},P,V\rangle$, 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 $\langle\{\{c_{ij}\}_{j=1,\ldots,h_i}\}_{i=1,\ldots,m},P,V\rangle$, the *expected utility $E(A_i)$* of an action A_i is $E(A_i) = \Sigma_{k\leq h_i}\, p_{ik}\cdot u_{ik}$, 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 $\langle\{\{c_{ij}\}_{j=1,\ldots,h_i}\}_{i=1,\ldots,m},P,V\rangle$, let a and b be two vectors of real numbers $(a_{i1},\ldots,a_{ih_i})$ and $(b_{i1},\ldots,b_{ih_i})$ 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 $\{p_1,\ldots, p_k\}$. A list of numbers $[n_1,\ldots,n_k]$ is a *solution vector* to a base L if the substitution of n_i for p_i, for all $1 \leq i \leq k$, 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 $(a_{i1},\ldots,a_{ih_i})$ and $(b_{i1},\ldots,b_{ih_i})$, 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 “never eliminate or disqualify an action that might be the best one”. The only option then becomes “never 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 $\langle\{\{c_{ij}\}_{j=1,\ldots,h_i}\}_{i=1,\ldots,m},P,V\rangle$, the *difference in expected utility* δ_{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 $\langle\{\{c_{ij}\}_{j=1,\ldots,h_i}\}_{i=1,\ldots,m},P,V\rangle$, A_i is *at least as good* as A_j iff ${}^{abcd}\delta_{ij} \geq 0$, for all a, b, c and d, where 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, $\{u_{i1} = b_{i1}\}$ & ... & $\{u_{ih_i} = b_{ih_i}\}$ & $\{u_{j1} = d_{j1}\}$ & ... & $\{u_{jh_j} = d_{jh_j}\}$ 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 *π-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 *π-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 *π-contraction* of X.

The π-cut is a linear procedure, but non-linear procedures are plausible as well. In addition, the procedure can work from either side ((L)π-cut and (R)π-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 in 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, where only local, simpler intervals may be available.

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 evaluating imprecise decision situations, several approaches have been proposed, each with its own strengths and limitations. While interval-based models have been extended to represent imprecise probabilities and utilities, the challenge remains to effectively incorporate belief strength and handle qualitative aspects of decision making. Another suggestion that extends the classical analysis is the application of fuzzy set theory, which also relaxes the requirements for numerically precise data and purports to provide a fairly realistic model of the vagueness in subjective estimates. These approaches also allow the modelling of problems in vague linguistic terms, and membership functions can be defined in accordance with the statements involved. Fuzzy set theory is quite a widespread approach to relaxing the requirement of numerically precise data and providing a more realistic model of the vagueness in subjective estimates of probabilities and values, however with known problems in the evaluation of the set membership. One major disadvantage of such formalisms is the problem of communication between analysts and stakeholders. While sometimes possessing attractive mathematical properties, the basic concepts are most often not known to the decision-makers, creating a knowledge gap hard to bridge.

By contrast, the aforementioned 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 real-life decision problems within standard decision analysis. Moreover, the possibility to state, for example, that one consequence is inferior to another is useful, particularly when handling qualitative information. Therefore, in addition to allowing interval statements, some modern decision methods allow statements containing comparisons between probabilities or between values, or even between differences between them, a feature lacking in older approaches.

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-20th century with the development of utility theory and the axiomatic foundations of rational choice, most notably embodied in 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 normative ideal of rational behaviour under uncertainty. Probabilistic 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. Methods and tools built on this tradition have been useful whenever the problem specifications fit the framework.

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

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 evaluation 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 unescapable 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 PROMÉTHÉE, among others. Each class of method carries its own assumptions, strengths and weaknesses. Value-based methods often rely on compensatory aggregation rules and require strong 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.

As a consequence, most present-day developments in computational decision analysis occur within MCDA rather than single-criterion probabilistic 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 researches, 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, conflicting objectives. In these 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 (MOO) 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 Harold Kuhn and Albert 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 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 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 developed methods to analyse trade-offs between criteria in the 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.

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 transparent 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).

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, 1980). 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 the systematic 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 in applying it to real-world decision problems.

The traditions in multi-criteria decision analysis (MCDA) are quite different from those in "traditional" probabilistic decision analysis (PDA). While PDA traditionally has a more theoretical and axiomatic approach, focusing on well-foundedness, MCDA has been more concerned with processes, procedures and calculation schemes. There is nothing inherently wrong in any of the two sets of approaches, rather they 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 industrial engineering (TOPSIS) and civil engineering (VIKOR). While an engineering approach to a research problem is not per se better or worse than a mathematical/theoretical one, they yield vastly different outcomes. This book aims at unifying both views by first presenting the important theoretical results of PDA and then trying to map them onto MCDA while keeping the engineering perspective of such methods intact and adding computability as a third pillar of a modern, real-world view of decision analysis. See, e.g., (Greco et al., 2016) for an overview of the current state-of-the-art in the field.

In (Danielson, 1997), only PDA is treated in detail, due to the general research focus back then being on probabilistic models. Since then, and characterising the 21st century, multi-criteria decision problems have been much more in focus. Luckily, most results from PDA carry over to MCDA, albeit with some modifications and exceptions. 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. Keeney and Raiffa (1976/1993) present four adequate examples of decision situations where the decision-maker cannot hide from the fact that there are multiple objectives in conflict with each other. 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 $C_i = (c^1, \ldots, c^n)$. Some literature uses the terms *criteria* or *perspective* instead of attribute, however, from a theoretical point of view these terms may be used interchangeably.

A number of approaches to aggregate utility functions under a variety of attributes have been suggested, such as (Keeney and Raiffa, 1976/1993; Keeney, 1980; 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 in order 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 $\sum w_j = 1$, $w_j \in [0,1]$, 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 u_i:s map to zero for the worst possible state regarding the i:th 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

$$1 + KU(x_i) = \prod \left[K k_i u_i(x_i) + 1 \right]$$

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 u_i:s 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

$$1 + K = \prod \left(1 + K k_i \right)$$

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. 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 outcomes 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:

- Utility independence of attributes
- Monotonicity in attributes
- Decomposability (e.g., additive or multiplicative form)

When these are satisfied, an additive utility function of the form can represent preferences. Unlike vNM, KR allows the modelling of preferences without uncertainty, making it foundational for MCDA. 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 a utility function that is linear in the appropriate domain. 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 special case of a multi-attribute 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. The 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 independence 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 accommodates 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 are unfortunately over-engineered and must be either reformulated to be used as proper decision analysis frameworks or disregarded as theoretically motivated tools and methods. To properly discuss them, we introduce ten desiderata that are derived from vNM, KR, and the general theory of multi-attribute utility.

Desideratum 1 (Ordering): The preference relation is complete and transitive. For all A and B, either $A \succ B$, $B \succ A$, or $A \sim B$. If $A \succ B$ and $B \succ C$, then $B \succ C$. vNM assumes completeness and transitivity as axiomatic to ensure coherent preferences. KR carries these over to deterministic multi-attribute models.

Desideratum 2 (Dominance): If for all i, $s_i(A) \geq s_i(B)$ and for some i, $s_i(A) > s_i(B)$ then $A \succ B$. Strong dominance is compatible with both vNM and KR. It ensures that if one alternative is objectively better, it must be preferred.

Desideratum 3 (Monotonicity): If $A \succ B$, and A' is such that $s_i(A) \geq s_i(A')$ for all i (with strict inequality for at least one i), then $A' \succ B$. A standard assumption in both vNM and general MAUT. In KR, this corresponds to increasing value functions: improving an attribute must not worsen utility.

Desideratum 4 (Independence of Irrelevant Alternatives, IIA): If $A \succ B$ in set X, and $C \notin \{A, B\}$, then $A \succ B$ in $X \cup C$. 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 5 (Score Independence): The preference between A and B depends only on the vector of scores $\{s_1(A), \ldots, s_n(A)\}$ and $\{s_1(B), \ldots, s_n(B)\}$. It does not depend on the scores of other alternatives in $X \setminus \{A, B\}$. This follows directly from utility independence and separability in KR. It ensures that preferences are not context-sensitive to unrelated alternatives.

Desideratum 6 (Criteria Transparency): For any preference $A \succ B$, 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 7 (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 $s_i(A) \geq s_i(B)$ should not reduce A's preference over B. This follows from sensitivity analyses in MAUT and reflects the principle that weights encode preference intensities and must affect final utility accordingly.

Desideratum 8 (Criteria Independence): If criteria C_i and C_j produce identical scores for all alternatives, then swapping or merging them is permissible only if the weight structure is adjusted accordingly. Related to the independence of attributes in MAUT: duplication of identical attributes without adjusting weights violates utility independence and will overstate their importance.

Desideratum 9 (Scale Invariance): For any criterion C_i, if a monotonic transformation $f : \mathbb{R} \rightarrow \mathbb{R}$ is applied to all $s_i(\cdot)$, then preferences remain unchanged. In both vNM and MAUT, utility functions are ordinal up to a monotonic transformation. As long as the transformation preserves order, preferences must remain stable.

Desideratum 10 (Rank Preservation under Deletion): If $A \succ B$ in X, and C is a third alternative not affecting the scores of A or B, then removing C does not alter the ranking $A \succ B$. Follows up on Desideratum 4 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.

These ten desiderata form a requirements system called DAMS (Decision-Analytic Methodologic System) which guarantees well-behaving and well-functioning

MCDA methods if the ten are all adhered to. From these ten desiderata, some consequences follow:

Proposition 1 (Utility Representability): If DAMS Desiderata 1–7 are accepted, then there exists a utility function $U: X \rightarrow \mathbb{R}$, representable as a weighted additive model

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

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

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 4, 5, and 10 are accepted, then the decision method is immune to rank reversal caused by irrelevant alternatives.

Desideratum 4 ensures rankings are stable under expansion of the alternative set, Desideratum 5 ensures no dependence on unrelated scores, and Desideratum 10 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 6 and 7 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 desideratum as well. However, it is concerned with the understanding of the underlying process elements rather than the internal consistency of the calculation steps.

Desideratum 11 (Explanatory transparency): It must be possible for the users to make and maintain a requisite mental model of the analytic process as a whole, including but not limited to the calculation steps. In this book, however, the concern is with the computational steps of the process, which should be possible to understand on a conceptual level, yielding auditability and replicability.

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, Usually, the process is facilitated by a designer or an expert, which makes the users not question the traceability of the output from the input, often relying on the expertise of the facilitator. If MCDA methods are to become widespread, however, there is a need for transparency in unguided sessions to build trust in the output results. Desideratum 11 is different from 6 despite sharing the word 'transparent'.

While the desiderata are formulated to be conceptually independent, some exhibit logical or functional overlap under classic utility theory assumptions. The following discussion considers potential overlaps and candidates for a possible reduced minimalist core.

DAMS Desiderata 1 (Ordering), 2 (Dominance), 3 (Monotonicity), and 5 (Score Independence) form a foundational core. These suffice to guarantee transitive, rational preferences that respect utility dominance and maintain independence from unrelated alternatives.

Desideratum 10 (Rank Preservation under Deletion) can be viewed as a corollary of Desiderata 4 (IIA) and 5 (Score Independence). If preferences are independent of irrelevant alternatives and based solely on score vectors, the deletion of an irrelevant third option should not affect pairwise comparisons. Desideratum 8 (Criteria Independence) implicitly relies on Desideratum 6 (Criteria Transparency) and Desideratum 7 (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 6 (Transparency) and 7 (Weight Sensitivity) are not strictly necessary but conceptually desirable since they ensure interpretability. Desideratum 9 (Scale Invariance) stands largely independent from the others but supports robustness under unit changes. It is justified on theoretical rather than logical grounds and could be omitted if all inputs are pre-normalised or internally transformed. However, they aid in the understanding of what should be required from a modern MCDA method.

To sum up, a minimal core could consist of Desiderata 1–3 and 5. The remaining ones either reinforce practical robustness (4, 6, 7, 9) or follow logically under standard assumptions (8, 10). However, for pedagogical reasons as well as argument's sake, all ten are kept in the DAMS system as beacons in the ensuing discussions of MCDA methods.

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 offer practical tools but lack coherence. As a case in point, take rank reversal, 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 the desiderata.

In line with the desiderata, preferences are supposed to be constructed to be invariant under irrelevant changes. This is encoded in 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 Score Independence (Desideratum 5) and Rank Preservation (Desideratum 10), 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 utility-theoretic foundation.

This makes rank reversal more than a nuisance. It is a diagnostic signal of representational failure. It flags methods that do not preserve rationality across varying contexts and undermines the claim that such methods are decision-theoretically grounded. Accordingly, testing for rank reversal should be a litmus test in evaluating any MCDA method's compliance with sound decision-theoretic principles. Rank reversal is not a minor technical flaw – it is a symptom of a deeper incompatibility with core desiderata. As such, it remains an accessible and informative indicator of a method's performative adequacy since methods exhibiting such reversals fail to satisfy decomposability, independence, and context-invariance, making rank reversal not only observable but also diagnostic.

The DAMS framework with ten desiderata will be employed to evaluate six different well-known MCDA methods: SMART (representing the SAW class of methods evaluating the alternatives using a sum-of-weighted-values approach), VIKOR, TOPSIS, ÉLECTRE, PROMÉTHÉE and AHP. These methods were selected because of their spread and reach. The observed patterns of usage and citations suggest that method popularity often reflects branding success rather than demonstrable methodological superiority. The prominence of certain techniques appears to be driven less by empirical performance and more by factors such as catchy acronyms, compelling narratives, and academic network effects. Additionally, being early to the methods scene has played a significant role, allowing some methods to establish a dominant position before competing approaches emerged, further reinforced by snowballing citation effects. In classic marketing theory, users are locked in to a product or a service by branding and terminology, creating a mental barrier to switching contexts. The proliferation of MCDA methods resembles a form of implicit marketing, where name recognition and institutional affiliation influence uptake, often independently of rigorous comparative validations or theoretical coherence, circumstances one could wish academia were mainly devoid of.

07. SMART

The SMART family of methods originated in the context of decision making 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 of decision theory.

The SMART (Simple Multi-Attribute Rating Technique) family of methods constitutes a group of approaches developed within the field of multi-criteria decision analysis (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 proposed 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 with Swing weights), 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 provides more meaningful 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 criteria importance and employs statistical techniques, most notably rank-order centroid (ROC) weighting, to derive approximate cardinal weights from the rankings. This approach trades off some theoretical rigour for increased ease of use and has been studied as a practical compromise in settings 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. Hybridisations with other MCDA techniques, such as the Analytic Hierarchy Process (AHP), have also been developed, particularly in contexts where qualitative criteria are present or when pairwise comparison techniques are deemed preferable for elicitation.

Overall, the SMART family of methods reflects an ongoing trajectory within MCDA to balance normative soundness with practical applicability. Its various forms (SMART, SMARTS, SMARTER, and others) share a common structure rooted in additive value models but diverge in their assumptions, elicitation procedures, and treatment of uncertainty. The evolution of SMART methods has been closely tied to developments in decision analysis, cognitive psychology, and applied statistics, reflecting broader methodological trends in decision support system design.

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 for each criterion are reflected in the final decision, and the process is computationally efficient.

In response to the limitations of the original SMART method, the SMARTER (SMART Extended) method was developed by von Neumann et al. (1981) to handle more complex decision making scenarios. The primary aim of SMARTER was to provide a more flexible framework that could accommodate interactions between criteria and handle more sophisticated preferences.

The SMARTER method retains the basic structure of SMART while it allows for the inclusion of non-linear relationships between criteria and their respective weights. While SMART relies on a linear aggregation of scores, SMARTER includes the possibility of applying non-linear functions to the scores of alternatives, which allows the method to better reflect the diminishing returns that might be encountered in real-world decision making situations. This extension of the SMART framework makes it better suited to situations where the decision-maker's preferences are not perfectly linear and where the marginal utility of improving performance on a criterion may decrease as performance improves.

SMARTER also allows for greater flexibility in handling interdependencies between criteria. In complex decision problems, it is often the case that the performance of an alternative on one criterion can affect its performance on another. For example, in an environmental decision making context, improving the economic efficiency of a project may negatively impact its environmental sustainability. SMARTER provides a mechanism for modelling such trade-offs between criteria, often through the use of more sophisticated utility functions or through pairwise comparisons of criteria that capture their relative importance and interdependencies. The increased flexibility of SMARTER also introduces additional complexity. The method requires more detailed input from decision-makers, including the specification of non-linear functions and interdependencies between criteria. This can increase the time and effort required to implement SMARTER and may limit its accessibility to decision-makers without strong analytical backgrounds.

SMART in its basic form complies with the desiderata of DAMS, not least by its additive structure, monotonicity, and transparency. Scores and weights are transparent, independence and dominance are preserved, and rank stability is guaranteed. It is sometimes argued that it struggles with Desideratum 8 (Criteria Independence) in that duplicating a criterion inflates its influence. This is a misunderstanding, though, since additive methods must always adjust their weights when the criteria set changes. As for the rank reversal litmus test, methods like SMART never display such behaviour.

08. VIKOR

VIKOR is a method developed in the late 1970s by Duckstein and Opricovic (1980) and further disseminated in academic literature during the 1980s and 1990s. It was designed as a method for ranking and selecting alternatives in the presence of conflicting criteria, based on the concept of compromise programming and the use of an aggregating function that incorporates individual regret and group utility.

The development of VIKOR emerged from earlier work in multi-objective optimisation and compromise solutions. Its formulation is mathematically related to a metric used in compromise programming, where the distance of each alternative to an ideal solution is computed. However, unlike traditional compromise programming which uses a parametric family of distance functions, VIKOR applies a discrete compromise approach involving a specific form of aggregating function that balances the majority utility (group benefit) and the individual regret (worst performance). The basic VIKOR method uses two measures: the S measure (representing the aggregated group utility) and the R measure (representing the maximum individual regret). These are then combined into a single ranking index Q using an external parameter v.

The computational procedure of VIKOR involves the identification of the best and worst values for each criterion among all alternatives (known as the ideal and anti-ideal solutions), 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 certain acceptability conditions that involve both rank consistency and a threshold for closeness between top-ranked alternatives.

The first step in VIKOR is the identification of the decision criteria and alternatives. The decision-maker must define the relevant criteria for evaluation, which can be both quantitative and qualitative, and select the set of alternatives that will be assessed. The next step 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 actual MCDA process begins, and ideal and regret 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. These solutions serve as benchmarks for evaluating the alternatives.

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 $A_1, A_2, \dots, A_n$) and m criteria (denoted as $C_1, C_2, \dots, C_m$) 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_j.

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

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

while for non-beneficial criteria, it is

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

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:

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

where $y_{\max,j}$ are the maximum values for each normalised criterion j. For each alternative A_i, the total distance to the ideal solution is then calculated using the following reversed formula:

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

rather than a more traditional standard formula

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

where w_j represents the weight of the j-th criterion. Thus, the normalised scores are now reversed and reinterpreted as the multi-dimensional closeness to the worst outcome instead. Next, the regret is computed by the 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 value of any of the constituent criteria.

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

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

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 $v = 1$ and also 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, the standard accounts of the method do not seem to require $0 \leq v \leq 1$, which opens up even stranger possibilities. This calculation procedure yields three ranking orders of the alternatives based on their performances S_i, R_i and Q_i. A complex set of rules determine which of the rankings take precedence, with the Q_i-ranking being the primary to consider first.

To sum up, the calculations of the VIKOR method involves normalisation, determination of ideal and regret solutions, distance measures, and the final calculation of closeness coefficients to rank the alternatives. The beneficial and adverse criteria scales are just standard and reversed scales, respectively. 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 pre-normalised values y_{ij}, i.e. their ranges are [0, 1]. Then A^+ becomes the vector (1, 1, 1, …) and all S_j become $\Sigma_i[w_i \cdot (1 - y_{ij})]$, 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_j mostly resembles a standard MCDA measure. However, as pointed out, with a reversed scale where lower numbers are better, a measure of distance from the synthetic optimal alternative. Still, this is in line with DAMS since all operators are linear and thus there exists a 1–1 transformation. 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, although it can 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 Desideratum 11.

Further, VIKOR departs from DAMS Desiderata 2, 3, 4, and 5. It ranks alternatives based on regret measures in addition to strict dominance or monotonic utility. Its results depend on the full set of alternatives and the definition of ideal/worst points, violating independence. The method lacks decomposability and transparency (Desideratum 6), further undermining its normative robustness.

The use of compromise ranking regret measures similarly deviates from standard MAUT. Its aggregation formula includes a balance parameter, which lacks a clear normative grounding in utility theory. It fails decomposability and is sensitive to dataset composition, violating utility independence. While it introduces trade-off modelling in spirit, it lacks the mathematical axioms needed for rational trade-off representation. The use of compromise programming makes it structurally prone to rank reversal when the ‘best’ or ‘worst’ alternatives change upon set modification. It violates IIA, Score Independence, and Rank Preservation (Desideratum 10) by design.

09. TOPSIS

TOPSIS (Technique for Order Preference by Similarity to Ideal Solution) is a method developed by Hwang and Yoon (1981). It was introduced in the book *Multiple Attribute Decision Making: Methods and Applications* and was intended to 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 of TOPSIS 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) in a multi-dimensional attribute space.

The method operates, as is common, on a decision matrix consisting of a finite set of alternatives and a set of evaluation criteria, which are assumed to be of either benefit-type (where higher values are preferred) or cost-type (where lower values are preferred). 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 comprises the best values for each criterion (maximum for benefit-type, minimum for cost-type), and the NIS comprises the worst values. 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 alternatives are ranked based on these coefficients, with higher values indicating greater proximity to the ideal solution.

The evaluation principle of TOPSIS 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. In decision analysis, this type of approach is not uncommon in some MCDA methods, where the ideal solution represents the optimal choice across all criteria, and the 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 Euclidean distance is, however, controversial.

In more detail, the first step, after forming the traditional two-dimensional matric 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 / Σ 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).

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

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

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

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) distances to the vectors of the best and worst solutions calculated by

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

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 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 α 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 α + α = 2α units to become equal to A^{+}. But TOPSIS would consider the required improvement to be $\sqrt{2}$ which is clearly wrong. The distance in a weight space is measured by a city block (or Manhattan) measure, not a Euclidean 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 clearly not the case in TOPSIS which assigns the length $\sqrt{(w_1^2 + w_2^2)}$ to the traversal while DAMS, requiring a city block metric, will have 1 for every conceivable traversal.

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 ¼ of the cases, TOPSIS' results differ from the linearity-based standard DAMS model.

To continue, TOPSIS violates several of the desiderata in DAMS. It fails Desideratum 2 (Dominance) since alternatives closer to the ideal can be outranked. It violates Desiderata 4 and 10 due to reliance on dataset-dependent reference points, causing rank reversal. It also violates Desideratum 5 (Score Independence), as rankings depend on the presence of other alternatives. While transparent in its formulation, its foundation lacks utility representability.

TOPSIS ranks alternatives by their distance from an ideal and anti-ideal solution. While intuitive, this method violates utility decomposability since the reference points are dataset-dependent rather than criterion-level constructs. The presence of context-sensitive anchors introduces violations of utility independence. Moreover, score transformation is not monotonic in all cases, especially under min-max normalisation. These features depart from MAUT's normative framework.

Let alternatives A and B be evaluated based on two criteria. A is closer to the ideal solution than B. Introducing a third alternative C, with extreme values in one criterion, shifts the position of the ideal point. Now, A appears farther and B closer to the new reference, potentially reversing their ranks. This clearly breaks Desiderata 4 and 10 (Rank Preservation), showing that preferences depend on context rather than intrinsic performance.

To sum up, adding or removing alternatives shifts the reference points, violating IIA (Desideratum 4) and Score Independence (Desideratum 5). Rankings can change even if the added alternative is irrelevant, manifesting classic rank reversal.

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 Société d'Économie et de Mathématiques Appliquées (Benayoun et al., 1966; Benayoun and Sussmann, 1966). 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. An 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 for 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 into account criteria weights. 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 (Roy, 1991).

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, each differing in the treatment of assignment rules and model structure.

The ÉLECTRE family of methods follows a systematic 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 help to 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, ÉLECTRE also uses 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 with respect to 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:

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

where x_{ij} is the original utility of alternative A_i with respect to criterion C_j. As for TOPSIS, it is easy to see that all x_{ij}^2 / Σ 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. The TOPSIS method copied this RMS rescaling, which is a vector space metric rather than a DAMS-compliant metric, from ÉLECTRE without reflecting on the consequences of adopting it.

But 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 d_{ik} is calculated as

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

using the concordance set membership function

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

In a similar but not mirrored way, calculate the discordance index

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

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, yielding one final ranking. So while the ÈLECTRE 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

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

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

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

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

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

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 ÈLECTRE family contains many methods that differ in various respects. All those modifications take the ÈLECTRE set of methods even further from the DAMS model. The final ranking is 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 preferable.

ÉLECTRE violates Desiderata 3–5 and 10. 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 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, violating Desideratum 10 and 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.

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) founded 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 for group decision making, accommodating 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 ÈLECTRE 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 crucial 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

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

while for reversed scales, it is instead

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

where, as usual, 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 φ 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 π_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_j for alternatives A_i and A_k.

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

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

and

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

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 on the separate positive and negative flows. The alternative with the highest Φ_i is the most preferred, and the one with the lowest Φ_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 4 and 5, as the net flow scores depend on the entire set of alternatives, not just pairwise comparisons. It also violates Desideratum 2 (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.

PROMÉTHÉE 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 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 change. This violates Score Independence (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 undermines IIA and Score Independence and leads to frequent rank reversals when alternatives are added or dropped.

12. AHP

The Analytic Hierarchy Process (AHP) is a method developed by Saaty in the 1970s, with its theoretical foundations first formally presented in 1980. AHP was introduced to support complex 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 judgment 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 judgments 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 susceptibility 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):

- 1 means equal importance.
- 3 means one element is slightly more important.
- 5 means one element is significantly more important.
- 7 means one element is very strongly more important.
- 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 pairwise 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 (i.e., if A > B and B > C, then A > C should hold, also called 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:

$$CI = \frac{\lambda_{\max} - n}{n - 1}$$

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 CRCRCR is below a 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 is the most non-compliant of the methods discussed in this book. It fails or partially violates nearly every desideratum: Desideratum 1 (due to tolerated inconsistency), Desideratum 2 (dominance ignored), Desideratum 3 (monotonicity not guaranteed), and Desiderata 4–5 and 8–10 (due to scale sensitivity, context dependence, and rank reversal). The eigenvector approach further obscures criteria transparency (11) and utility interpretability.

Let alternatives A and B be evaluated in an AHP framework with pairwise comparisons indicating A $\succ$ B. Now introduce C, which is strictly worse than both A and B across all criteria, i.e. A $\succ$ C and B $\succ$ 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 $\succ$ A might occur. This violates Desideratum 4 (IIA), Desideratum 5 (Score Independence), and by extension, the separability required in utility theory, and 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 occurs when a new, even dominated, alternative is added. AHP violates IIA, Score Independence, and Criteria Independence due to its normalisation and weighting procedures being context-sensitive. It is perhaps somewhat of a stretch to call AHP flawed, as in (Abbas, 2018, Ch.3), but it is certainly the least compliant with DAMS.

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 normative ideals 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.

Desideratum → **Method ↓**	1. Ordering	2. Dominance	3. Monotonicity	4. IIA	5. Score Indep.	6. Transparency	7. Weight Sens.	8. Criteria Indep.	9. Scale Invar.	10. Rank Del.
SAW	OK	OK	OK	OK	OK	OK	OK	OK	OK	OK
TOPSIS	OK	NO	OK	NO	NO	OK	OK	NO	OK	NO
VIKOR	OK	NO	NO	NO	NO	NO	OK	NO	OK	NO
PROMÉTHÉE	OK	NO	OK	NO	NO	OK	OK	NO	OK	NO
ÉLECTRE	OK	OK	NO	NO	NO	NO	OK	NO	NO	NO
AHP	OK	NO	NO	NO	NO	NO	OK	NO	NO	NO

Table 2. Five MCDA methods compared using the DAMS desiderata

As demonstrated in this book through analysis, classification, and counterexamples, many popular MCDA methods fall short of satisfying these desiderata – especially where rank reversal is concerned. Such failure points reveal deeper inconsistencies with utility-analytic rationality. The literature is plagued with comparisons of MCDA methods where all the compared methods point in different directions. Rather than trying to appoint a "winning" method, which is most often the aim, they should be taken as signs of the overall health of the research field as a whole.

As an example, in (Opricovic and Tzeng, 2004) VIKOR and TOPSIS are compared. The comparison illustrates clearly the ad-hoc nature of both methods and also 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 point out different alternatives as the "best" one, but the reasons for or against either variant are hard to grasp for a reader. In fact, the 18 variants together succeed in ranking the three alternatives in all eight (!) possible permutations of the ranking order. Imagine how impossible it is for a layman user to understand the pros and cons of each method. None of the methods contain any means of sensitivity analyses, but rather present the results with three decimals. In (Opricovic and Tzeng, 2007), the four methods VIKOR, TOPSIS, ÉLECTRE and PROMÉTHÉE are compared. There are six sets of 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. Within the top set, the best alternative changes frequently or is undetermined.

Further, in (Zlaugotne et al., 2020), five methods are compared of which three are VIKOR, TOPSIS, and PROMÉTHÉE. For the four alternatives in the article, 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 four 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 "better" than any single method. The substantial efforts required notwithstanding, there is no theoretical proof that such averaging should lead to a better analysis. If that were the case, it would be possible to construct one optimal giant "the-more-the-merrier method" consisting of all the world's known MCDA methods (or perhaps all variants of all methods), weighted by some mysterious all-encompassing meta-weighing scheme.

What all the methods (except SMART) fail to do is to separate the calculus of 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, processes dealing with negotiation, bargaining, regret, etc. should be kept and handled in an orderly fashion. If that handling requires additional calculations, they can 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 why and subsequently suitability verified 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 obtain a ranking with a favourite alternative on top – is surely a contributing factor to the prevailing mistrust and underutilisation of MCDA in society.

This book provides both a diagnostic and prescriptive perspective: identifying logical weaknesses in existing methods, while also pointing at a route toward greater decision-theoretic coherence. This is not a plea for process conformity. The differences in philosophy and the different brandings of "schools" of thought are invigorating and should influence the elicitation processes, the presentation formats, the group decision mechanisms, and much more – as long as the methods stand on established scientific mathematical ground. Substituting a since-long well-established and sound axiomatic computational core for homemade calculi is not a way to gain trust, it is a way to opaqueness. Neither is mixing descriptive and psychological factors with an axiomatically grounded computational core. The need to stand out by branding and uniqueness should be satisfied in other ways and by other means, less detrimental to the MCDA field.

Despite the logical clarity and mathematical rigor of the unified utility framework grounded in von Neumann–Morgenstern and Keeney-Raiffa (vNM/KR) axioms, a wide range of popular MCDA methods persist that violate these principles. This raises the question: are there any compelling mathematical or logical reasons to prefer these methods? The answer, in short, appears to be a resounding no.

None of the well-established MCDA methods that violate vNM/KR axioms, such as VIKOR, TOPSIS, ÉLECTRE or PROMÉTHÉE, are grounded in a rigorous theoretical foundation. These methods often make heuristic or procedural sense but fail when held to standards of decomposability, independence, and consistency. Some of the problems are:

- No representation theorem supports the forms of aggregation used in these methods.

- Rank reversal and reference dependence violate basic tenets of rational choice.
- Non-decomposability in scoring means there is no underlying utility function being maximised.

Despite these shortcomings, the abovementioned methods proliferate and are widely used in practice. There are several reasons for that:

1. **Software Availability:** Many are 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 or KR and thus default to 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 toward 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 (e.g. completeness, independence, decomposability).
- It supports additive utility representations and generalises both von Neumann-Morgenstern's and Keeney-Raiffa's theories.
- It yields decisions that are transparent, defensible, and logically justified.

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

- It promotes heuristics and approximations over consistency.
- It focuses on cognitive ease, group dynamics, and stakeholder inclusion over correctness.
- It most often adopts methods that fail DAMS but are thought to be easier to explain.

Part II of this book demonstrates that methods like 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 normative coherence while retaining formats familiar to prescriptive users. This convergence is essential for elevating decision analysis from plausible heuristics to justifiable practice.

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

1. **Loss of Interpretability:** Methods that violate decomposability, transitivity, or independence lose any claim to be preference-preserving. Their rankings are artefacts of procedure, not reflections of rational preference.
2. **Rank Reversal and Context-Dependence:** As discussed above, violations of key axioms produce arbitrary 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 critical to draw a boundary between two distinct 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.
- Procedural methods, such as outranking, voting mechanisms, or pairwise flows, are implementation strategies. While they may offer heuristic appeal or operational ease, they are not substitutes for foundational coherence.

Confusing these two levels leads to mistaken beliefs, for instance, that a visually compelling ranking procedure is comparable to a DAMS-based decision analysis. 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 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 and modified in countless ways – at the end of the line, but only after the core results according to established theories have been calculated, and the post-processing can be shown still to comply with the axioms and desiderata of well-established scientific theoretical bodies instead of arbitrary made-up formulas – arbitrary seen from a decision-theoretic point of view.

To move the MCDA field forward in a scientific direction, and to unify rigor with usability, a set of guiding principles is necessary. Such 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 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 in rational defensibility or practical clarity. 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 (any many others with them) violate simplicity.

- Outranking methods that require multiple thresholds and preference functions across criteria.
- Multi-phase methods where the output of one arbitrary step becomes the input to another, often without clear justification.
- Methods that produce partial orderings through procedures that cannot be linked to any utility representation.

The desiderata proposed in DAMS are supposed to lead naturally to models that are both simple and normatively sound. Additive utility models, dominance-based comparisons, and weighted sums need not be 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.

14. Three Notes

A Note on 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 doesn't necessarily make sense.

In the additive model of MCDA, you sum up the weighted differences in performance across various criteria. In other words, you're aggregating the differences in scores or performance metrics, which is typically associated with the difference scale.

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

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're combining the differences between each alternative's performance across criteria, not their ratios.

AHP, on the other hand, explicitly requires that the pairwise comparison scale 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_1has 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 judgments 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 sum it up: AHP's ratio scale means that when you compare alternatives (or criteria) pairwise, you're 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 don't require the ratios between them to be meaningful, but rather just their relative differences (additive increments).

A Note on 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, synthetic criterion. 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 | 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 DA 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 $O(\mathrm{n})$ task and turning it into an $O(\mathrm{n}^2)$ 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 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 criteria (Howard, 2009, p.52). While that is an extreme case, 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.

A Note on 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 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 separately handle decision problems in which certain criteria represent thresholds or veto points that cannot be offset, regardless of performance elsewhere.

Outranking methods achieve this by embedding thresholds, calling them names 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 should 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 judgment *inside* the calculation core, outranking methods inadvertently produce black-box-like outcomes. 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 judgmental subtleties, in reality they offer greater normative 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 methodological tension: 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 the perceived 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 large 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, turning an outranking-based process into a less rational exercise. So the real-world process requirements turn the tables against opaque calculation methods.

15. Probabilistic MCDA

The structural similarities between the von Neumann–Morgenstern (vNM) and Keeney-Raiffa (KR) utility models suggest that they are not competing frameworks, but rather special cases of a more general probabilistic multi-criteria decision analysis (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 outcomes, where each attribute has multiple probabilistic outcomes. The general utility of an alternative is: This nested structure expresses vNM utility as the special case where there is only one attribute and only uncertainty exists, and KR (and MAUT) 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 multi-criteria models. The desiderata serve as a scaffold for such synthesis, and their expansion into this domain may mark the next frontier in normative decision theory.

In real-world decisions, alternatives often involve uncertainty in addition to multiple criteria. A natural extension of MCDA 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 classical 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 sort opens the door to richer, securely grounded models. It allows decision-makers to accommodate both subjective probabilities and value trade-offs. It also 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 subtle tensions 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 or define a stronger axiom that encompasses both.

Decomposability Tensions KR assumes additive or multiplicative 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 Adjustments To align vNM and KR under the expression: one must assume 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.

Implicit Normative Shifts 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 are formally equivalent in the role they play within the utility aggregation. This requires a commitment to 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 fragments of a more general system.

To reconcile and extend vNM and KR within a general probabilistic multi-criteria decision framework, we propose the following unifying desiderata. They are designed to support utility representations of the form

$$U(a) = \sum_i w_i \sum_j p_{ij}\, u\big(x_{ij}\big)$$

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

Desideratum 2 (Continuity): For any alternatives A, B, and C, with A $\succ$ B $\succ$ C, there exists a $\lambda \in (0, 1)$ such that B $\sim \lambda$A $+ (1–\lambda)$C This applies both to probabilistic mixtures (as in vNM) and to attribute trade-offs (as in KR).

Desideratum 3 (Probabilistic Independence): For all alternatives A, B, and C, if A $\succ$ B, then for any $\lambda \in (0, 1)$: $\lambda A + (1-\lambda)C \succ \lambda B + (1-\lambda)C$. This ensures linearity in probabilities.

Desideratum 4 (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 5 (Additive Decomposability): If utility independence holds for all attributes, then the overall utility function is additive across attributes and linear in probabilities: $U(a) = \sum_i w_i \sum_j p_{ij}\, u(x_{ij})$.

Desideratum 6 (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 7 (Weight Normalisation): For all weights and probabilities, $w_i \geq 0$, $\Sigma w_i = 1$ and $p_{ij} \geq 0$, $\Sigma p_{ij} = 1$ respectively for each *i*.

Desideratum 8 (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 successfully unify the vNM and KR theories into a single coherent foundation for probabilistic MCDA. They allow trade-offs across attributes and beliefs while preserving normative coherence and a clear interpretative structure. They form the basis for the UNEDA software platform which handles 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. The open-source software platform is described next in Part III (Chapter 16).

16. Computational Evaluation

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

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.

The 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 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, $x_{i1}, x_{i2},\ldots,x_{in}$ and the m children of the node x_{ij} are denoted $x_{ij1}, x_{ij2},\ldots,x_{ijm}$, 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 ${}^{X}\max(f(x)) =_{\text{def}} \sup(a \mid \{f(x) > a\} \cup X \text{ is consistent})$. In a similar way, 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*, ${}^{X}\text{argmax}(f(x))$ is a solution vector that is a solution to ${}^{X}\max(f(x))$, and ${}^{X}\text{argmin}(f(x))$ 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 $\left\{\langle {}^X\min(x_i), {}^X\max(x_i)\rangle\right\}_{i\in I}$ is the feasible box (orthogonal hull) of the set and is denoted $\langle {}^X\min(x_i), {}^X\max(x_i)\rangle_I$.

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 $\Pi:n_j\rightarrow[0,1]$ over C_i. Let p_j denote the function value $\Pi(n_j)$. Π 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 $\langle a_k,b_k\rangle$, $k\in C_i$, which represents the user and default range constraints $\forall k \in C_i : p_k \in [0,1]$.
- 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 = ⟨I,N,E,r⟩ 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 = ⟨I,N,E,r⟩, 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 ⟨T,P,V⟩ containing the following representation of the alternative:

- A rooted tree T = ⟨I,N,E,r⟩ 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 $\{p_i\}$, $i \in I \backslash \{r\}$, 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_i. 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.

$$EV(A_i) = \sum_{i_1=1}^{n_{i_0}} p_{ii_1} \sum_{i_2=1}^{n_{i_1}} p_{ii_1i_2} \cdots \sum_{i_{m-1}=1}^{n_{i_{m-2}}} p_{ii_1i_2\cdots i_{m-2}i_{m-1}} \sum_{i_m=1}^{n_{i_{m-1}}} p_{ii_1i_2\cdots i_{m-2}i_{m-1}i_m} v_{ii_1i_2\cdots i_{m-2}i_{m-1}i_m1},$$

where $p_{...i_j...}$, $j \in \{1,\ldots,m\}$ denote probabilities in P_i and $v_{...i_j...1}$ 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 $x_1,\ldots,x_n$, $\pi \in [0,1]$ is a real number, and $\{\pi_i \in [0,1] : i = 1,\ldots,n\}$ 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 $\bar{k} = (k_1,\ldots,k_n)$ is a consistent point in X. A π-contraction of X is to add the interval statements $\{x_i \in [a_i+\pi\cdot\pi_i\cdot(k_i-a_i), b_i-\pi\cdot\pi_i\cdot(b_i-k_i)] : i = 1,\ldots,n\}$ to the base X. $\bar{k}$ is called the contraction point (or focal point).

By varying π from 0 to 1, the intervals are decreased proportionally using the gain factors in the π_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 textbook on the subject, e.g. [BHM77, C83]. The LP problem is the following optimising problem:

$$\max f(\mathbf{x})$$
$$\text{when } \mathbf{Ax} \geq \mathbf{b}$$
$$\text{and } \mathbf{x} \geq \mathbf{0}$$

where $f(\mathbf{x})$ is a linear expression of the type $k_1x_1 + k_2x_2 + \ldots + k_nx_n$, $\mathbf{Ax} \geq \mathbf{b}$ 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 $y_1,\ldots,y_k$, with $k = 2 \cdot m$, to the consistency problem, it can be reformulated as

$$\min (y_1 + \ldots + y_k)$$
$$\text{when } \mathbf{Ax} \geq \mathbf{b}$$
$$\text{and } \mathbf{x} \geq \mathbf{0},\ \mathbf{y} \geq \mathbf{0}$$

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 $a_{i1}x_1 + a_{i2}x_2 + \ldots + a_{in}x_n + y_{2i-1} \geq a$ and $a_{i1}x_1 + a_{i2}x_2 + \ldots + a_{in}x_n - y_{2i} \leq b$. 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 $\langle {}^{X}\min(x_i), {}^{X}\max(x_i)\rangle_n$, 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 ${}^{PV}\max(\Delta_{ij})$ and ${}^{PV}\min(\Delta_{ij})$. The optimisation of general Δ_{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 ${}^{PV}\max(\Delta_{ij})$ turns into ${}^{PV}\max(\sum_k p_{ik} - \sum_k p_{jk})$ for first order Δ-dominance such as 1SE and security levels, and into ${}^{PV}\max(\sum_k p_{ik}\cdot v_{ik} - \sum_k p_{jk}\cdot v_{jk})$ for second order ones such as 2SE or NE.

First Order Dominance

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

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

for some index sets K_i or index set pairs $(K_i,K_j)(d)$ 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 $\mathbf{Ax} \geq \mathbf{b}$ 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 ±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

$$\text{(QPI)} \qquad \max\ (\mathbf{x}^T\mathbf{Q}\mathbf{x} + \mathbf{c}^T\mathbf{x}) \quad \text{when } \mathbf{A}\mathbf{x} \geq \mathbf{b}$$

where $\mathbf{A}$ is a $m \times n$ matrix with linearly independent rows, $\mathbf{Q}$ is a symmetric $n \times n$ matrix, and $\mathbf{c}$ is a vector in 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 Danzig 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 null-space methods, depending on matrix sparsity, stability requirements, and other criteria. The BLP1 problem maps well onto QPI since there is one second-order non-linear 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 $\mathbf{Q}$ 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 δ_{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 stack.

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 $\langle C,P,V\rangle$,

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

${}^{V}E_j^{min}$ is $\sum_{k=1}^{m_j} p_{jk} \cdot b_{jk}$, where $b_{jk} = {}^{V}\min(v_{jk})$.

${}^{V}\delta_{ij}$ is ${}^{V}E_i^{max} - {}^{V}E_j^{min}$.

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 $(b_{i1},\ldots,b_{im_i})$ and $(b_{j1},\ldots,b_{jm_j})$ be as in the definitions of ${}^{V}E_i^{max}$ and ${}^{V}E_j^{min}$ above. For all feasible vectors $(p_{i1},\ldots,p_{im_i})$, $(p_{j1},\ldots,p_{jm_j})$, $(v_{i1},\ldots,v_{im_i})$, and $(v_{j1},\ldots,v_{jm_j})$ ${}^{V}E_i^{max} \geq \sum_k p_{ik} \cdot v_{ik}$ 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 $1 \le k,l \le m_i$, 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) $v_{il} \ge a + v_{ik}$ and ii b) $v_{il} \le b + v_{ik}$. 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 $(b_{i1},\ldots,b_{im_i})$ is a feasible vector as is $(b_{j1},\ldots,b_{jm_j})$ 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 $(b_{i1},\ldots,b_{im_i})$ and each b_{jk} in $(b_{j1},\ldots,b_{jm_j})$ 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 ${}^P\max({}^V\delta_{ij})$ by means of an LP algorithm such as Simplex. The proposition then guarantees that ${}^{PV}\max(\delta_{ij})$ can be determined by calculating ${}^P\max({}^V\delta_{ij})$ instead provided the precondition is met. Similarly, ${}^{PV}\max(\sum_k p_{ik}\cdot v_{ik})$ 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:

$$\mathrm{EV}(\mathrm{A_i}) = \sum_{i_1=1}^{n_{i_0}} p_{ii_1} \sum_{i_2=1}^{n_{i_1}} p_{ii_1i_2} \cdots \sum_{i_{m-1}=1}^{n_{i_{m-2}}} p_{ii_1i_2\cdots i_{m-2}i_{m-1}} \sum_{i_m=1}^{n_{i_{m-1}}} p_{ii_1i_2\cdots i_{m-2}i_{m-1}i_m} v_{ii_1i_2\cdots i_{m-2}i_{m-1}i_m1},$$

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

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

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 book has been developed in parallel with the software platform UNEDA (Universal Engine for Decision Analysis) over a long period of time. A fair amount of material connected to UNEDA has never been published other than on the author's university webpages. Those documents describe aspects of prescriptive decision analysis that have been incorporated into the UNEDA computational engine library.

The UNEDA software platform is an open-source library, free to use for any purpose, academic and non-profit as well as (from June 6, 2025) commercial. The software library is found at the GitHub repository

github.com/uneda-cda/UNEDA

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

doi.org/10.5281/zenodo.15114623.

DOI 10.5281/zenodo.15114623

References

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

Ackoff, R. L. (1962). *Scientific method: Optimizing applied research decisions*. New York, NY: 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*. Dordrecht: D. Reidel Publishing Company.

Bayes, T. (1763). An essay toward 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.

Benayoun, R., Roy, B., & Sussmann, B. (1966). ELECTRE: 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). Paris: 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, Transl.). *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). *Liber de ludo aleae* [Book on games of chance]. (Reprinted 1953, Princeton University Press, Princeton, NJ).

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.

Dantzig, G. B. (1947). *Linear programming and extensions*. Princeton, NJ: Princeton University Press.

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

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*. Stanford, CA: AAAI Press. Republished as: Doyle, J., & Thomason, R. H. (1999). Background to qualitative decision theory. *AI Magazine, 20*(2), 55–68.

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

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

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

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

Fischhoff, B., Slovic, P., Lichtenstein, S., Read, S., & Combs, B. (1983). How safe is safe enough? A psychometric study of attitudes towards technological risks and benefits. *Policy Sciences, 14*(2).

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*. Chichester: Ellis Horwood.

Horwood, B. (1988). Planning theory: Consistency and convergence. In J. Friend & A. Hickling (Eds.), *Planning under pressure: The strategic choice approach* (pp. 243–260). Oxford: Pergamon Press.

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

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

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

Gauss, C. F. (1809). *Theoria motus corporum coelestium in sectionibus conicis solem ambientium*. Hamburg: 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, CA: Stanford University Press.

Greco, S., Ehrgott, M., & Figueira, J. R. (Eds.). (2016). *Multiple criteria decision analysis: State of the art surveys*. New York, NY: 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, UK: Cambridge University Press.

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

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). De ratiociniis in ludo aleae. In D. E. Smith (Ed. & Transl.), *A source book in mathematics* (pp. 75–83). New York, NY: Dover. (Original work published 1657)

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

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

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

Keynes, J. M. (1921). *A treatise on probability*. London, UK: Macmillan and Co.

Kolmogorov, A. N. (1933). *Grundbegriffe der Wahrscheinlichkeitsrechnung*. Berlin, Germany: 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*. Singapore: World Scientific Publishing.

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

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

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

Laplace, P.-S. (1816/1952). *Essai philosophique sur les probabilités* (3rd ed.). Paris: Courcier. (English translation published by Dover, 1952)

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*. New York, NY: John Wiley & Sons.

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

Malmnäs, P. E. (1994). 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*. New York, NY: Wiley.

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). New York, NY: Wiley.

von Mises, R. (1957). *Probability, statistics and truth* (2nd English ed., H. L. Cambel, Transl.). New York, NY: Macmillan. (Original work published 1928)

de Moivre, A. (1718). *The doctrine of chances: A method of calculating the probability of events in play*. London: W. Pearson.

de Moivre, A. (1730). *Miscellanea analytica*. London, UK: Tonson & Watts.

de Moivre, A. (1738). *The doctrine of chances* (2nd ed.). London, UK: Woodfall. (Reprinted 1967, Cass).

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

Morgan, M. G., & Henrion, M. (1990). *Uncertainty: A guide to dealing with uncertainty in quantitative risk and policy analysis*. Cambridge, UK: Cambridge University Press.

von Neumann, J., & Morgenstern, O. (1944). *Theory of games and economic behavior*. Princeton, NJ: 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.

Opricovic, 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.

Opricovic, 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.). Florence: Aboca Edizioni.

Poisson, S. D. (1837). *Recherches sur la probabilité des jugements en matière criminelle et en matière civile*. Paris: Bachelier.

Quetelet, A. (1846). *Sur l'homme et le développement de ses facultés, ou Essai de physique sociale*. Paris: Bachelier.

Raiffa, H. (1968). *Decision analysis: Introductory lectures on choices under uncertainty*. Reading, MA: Addison-Wesley.

Ramsey, F. P. (1931). Truth and probability. In R. B. Braithwaite (Ed.), *The foundations of mathematics and other logical essays* (pp. 156–198). London: Kegan Paul, Trench, Trubner & Co. (Original work written 1926)

Rowe, W. D. (1994). Understanding uncertainty. *Risk Analysis, 14*(5), 743–750.

Roy, B. (1991). The outranking approach and the foundations of ÉLECTRE methods. *Theory and Decision, 31*, 49–73.

Saaty, T. L. (1980). *The analytic hierarchy process: Planning, priority setting, resource allocation*. New York, NY: McGraw-Hill.

Savage, L. J. (1954/1972). *The foundations of statistics* (2nd ed. 1972). New York, NY: 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, NJ: Princeton University Press.

Shapira, Z. (1995). *Risk taking: A managerial perspective*. New York, NY: Russell Sage Foundation.

Simon, H. A. (1955). A behavioral model of rational choice. *The Quarterly Journal of Economics, 69*(1), 99–118.

Slovic, P., & Tversky, A. (1974). Who accepts Savage's axiom? *Behavioral Science, 19*(6), 368–373.

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*. Cambridge, MA: 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*. New York, NY: Wiley.

Walley, P. (1991). *Statistical reasoning with imprecise probabilities*. London, UK: Chapman and Hall.

Weichselberger, K., & Pöhlmann, S. (1990). *A methodology for uncertainty in knowledge-based systems*. Berlin: Springer-Verlag.

Wynne, B. (1992). Uncertainty and environmental learning: Reconceiving science and policy in the preventative paradigm. *Global Environmental Change, 2*, 111–127.

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.

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 in the software industry for 20 years.

Other books from Sine Metu in the same series

Transcending Business Intelligence, Third Edition, 2022
K. Borking, M. Danielson, G. Davies, L. Ekenberg, J. Idefeldt, A. Larsson
ISBN 978-91-978-4505-2

A Decision-Analytic Manifesto, Third Printing, 2024
L. Ekenberg, M. Danielson
ISBN 978-91-527-5306-4

The Rise and Fall of Philips Data Systems, Fifth Printing, 2024
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 toward a critical assessment of multi-criteria decision analysis (MCDA). This book is grounded in the conviction that decision-analytic methods must rest on solid scientific foundations, logical coherence, empirical testability, 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 reconstructs the theory with both philosophical care and real-world decision problem relevance.

The second part turns to MCDA, the rapidly expanding family of methods designed to guide the analysis of complex decisions with multiple objectives. Rather than treating these as a neutral 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 in the face of uncertainty.

The final part deals with computations and a software platform for enabling applications for 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 scientists, analysts, and anyone serious about the logic of choice.

UNEDA

Sine Metu