1

Introduction

Welcome to Smart Decisions where you will learn about many ways to model and solve complex decision problems using Multi-Criteria Decision Analysis (MCDA). Each chapter will take you through a simple and common decision problem that most people face in their lifetimes, showing you how to solve it with a structured decision-making methodology in both Microsoft Excel and R. Microsoft Excel is the primary tool used in business and education for making spreadsheets and interactive formulaic templates while R is a popular tool used by specialists in analytics, data science and statistics. If you are not familiar with either Excel or R, don't worry! Chapter 2 is dedicated to telling you all about these software packages and how to use them. Trust me you will be an Excel and R whizz-kid in no time. That being said, this is not the primary aim of this book. The core of this book is to teach you about all of the wonderful and powerful decision-making methodologies out there which will help you justify and make better decisions at work and at home.

Let's start by discussing a hypothetical decision problem. Imagine you are a student who has just graduated and after years of partying and hard work you now have to decide what to do with your life. You are considering three distinct options; (1) find a job, (2) do further studies, or (3) start your own business. Your parents tell you to find a job related to your degree but your university professor suggests that you should do further studies before starting a professional career. To make matters more difficult, your friend, a recently established entrepreneur, has a great business idea and suggests you become a partner in their business. How do you decide what to do?

There are lots of interesting behavioural aspects to consider here, such as do you unconsciously favour the first option you discussed with your friend (this is referred to as anchoring) or do you prefer to stay in education as you haven't worked or started a business before (referred to as familiarity bias)? Although it is very important to acknowledge and understand the behavioural aspects of human decision-making, this will be rarely discussed in this book as the focus here is on the process of decision-making and not on the psychology of

Smart Decisions: A Structured Approach to Decision Analysis Using MCDA, First Edition.
Edited by Richard Edgar Hodgett, Sajid Siraj, and Ellen Louise Hogg.

decision-making. That being said, it is important to understand both areas of decision-making and therefore if you want to read more about behavioural decision-making we recommend reading *Thinking Fast and Slow* by Daniel Kahneman, *Psychology of Judgment and Decision Making* by Scott Plous and *Preference, Belief, and Similarity* by Amos Tversky.

Most people who study decision-making now agree that there are two systems of decision-making, system 1 which is based on intuition or gut, and system 2 which is controlled, conscious and requires considerable effort (and time). We will be focusing on the methods and techniques that can be used for important and complex decisions that fit within system 2, or as we like to call it methods for structured decision-making.

Going back to the example we discussed about the recent graduate choosing what to do with their life. This is a particular problem where alternative options are already known to the decision-maker and the decision problem is to evaluate each of these options in order to select one of these alternatives based on several different criteria. This is referred to as an evaluation problem where the ultimate goal is to obtain a ranking or rating for each alternative. Other examples of such problems would be a regulatory authority wishing to publish the ranking of all universities in a country or a food agency seeking to rate all food shops from 1 (inadequate) to 5 (outstanding).

Another category of decision problems would be those where we search for a feasible solution which is not explicitly known to us. This category of problems can be labelled as design problems, for example, product pricing or time-tabling problems. A product can be priced with any value on a continuum but the decision-maker is interested in finding the most feasible price on this continuum. Similarly, one may find it difficult to manually schedule a timetable with no conflicts whatsoever, and therefore seek to find a feasible solution. In these two examples, the former one involves selection of a single value while the latter can be visualised as the selection of multiple values. The problems involving multiple values (sometimes called multivariate) can also be referred to as allocation problems. Table 1.1 summarises these categories of decision-making problems along with an example for each possible type of problems.

Table 1.1 Examples of evaluation and design decision problems.

	Evaluation	**Design**
Selection	Hiring a candidate	Product pricing
Ranking	Ranking universities	
Rating	Rating hotels	
Allocation	Allocate money and possessions in a will	Time-tabling classrooms

Another important categorisation is to consider the number of decision-makers. In the case of two or more decision-makers, conflicts may occur, usually due to different preferences, interests, or each person's level of knowledge. By contrast, those decisions involving a single decision-maker may not face these kinds of conflicts, although conflicts among different objectives/criteria may still occur, for example, searching to buy the cheapest house but wanting to live in a nice area.

1.1 How to Structure Your Decisions?

Whilst there can be many ways to structure a decision problem, Hammond J.S., Keeney R.L., and Raiffa H. (2002) proposed a very nice and simple framework known as PrOACT, which is an acronym for <u>Pr</u>oblem, <u>O</u>bjectives, <u>A</u>lternatives, <u>C</u>onsequences, and <u>T</u>rade-offs.

Problem: The first element in this structure is the "Problem" itself. A good decision maker always questions the definition of the problem before identifying the ways to solve it. It is useless to find or suggest a wonderful solution to the wrong problem; therefore, it is important to identify the real problem before any further action.

Objective: After clearly identifying the real underlying problem, it is important to set objectives that can make us evaluate all the possible solutions. Some people may refer to these objectives as criteria. Linguistically, these two terms have slight difference, for example, a criterion can be "price", but the objective can be "finding cheapest price" or, if you are a rich person who wants to show off, the objective can be "finding the most expensive". In this example, both objectives are about the same price criterion, but one is trying to minimise while the other is trying to maximise. In most real-life problems, there are multiple criteria/objectives, and often these objectives are in conflict with each other. For example, one may dream of getting the cheapest item but with the highest quality.

Alternative: Consider a situation where there is only one possible course of action or choice. Will it be a decision problem? Not really, because the decision problem becomes a problem only when you have multiple alternative options, and you have to pick only one of these choices. Therefore, it is important to realise that identifying multiple alternative options is a very important step. One may miss a great (or best) solution without realising that it exists.

Consequences: After generating a good set of alternative options, it is important to evaluate the consequences of each of these alternatives. Usually this is done by creating a table where alternatives are placed in rows, objectives are places in columns, and then the consequences are penned down in each cell of this table. When describing the consequences in these cells, it is important to make sure they are as accurate as possible. Bad assessments will generate bad results, as they say "garbage in, garbage out".

Trade-offs: As we mentioned earlier, multiple objectives are often in conflict with each other, and therefore, you might end up with a table of consequences where there is no single clear winner. In this case, a trade-off analysis will be required where we need to sacrifice (or trade) something from one objective if we want to gain at least a little on another one.

In ProACT, there are three more elements that decision makers should consider in their decision-making exercise. These are Uncertainties, Risks, and Linked Decisions; but we will not discuss them here in detail. If you are interested in reading more about ProACT, we recommend reading the original book from Hammond J.S., Keeney R.L., and Raiffa H. (2002).

1.2 Different Stages in Decision-making

It is fair to say that every decision process starts with a need or desire to decide. The decision-maker realises this need when multiple alternatives appear viable to them whether known or unknown to them at the time of realisation. This realisation leads them to the two often-intertwined phases of structuring and exploration. For example, you have just moved to a new house and you realise you really need an air-conditioner. While exploring the various air-conditioning products on a popular shopping website, you make yourself aware of the criteria that must be taken into consideration when purchasing an air-conditioner. This is mostly true when the decision-maker has not taken a similar decision before. However, if the decision-maker is aware of taking similar decisions, it is likely they will have a decision structure in mind without needing to explore many options. For example, when buying lunch, you may consider the price, taste, travel time to get a particular food and the healthiness of the food. As you consider this decision very often in daily life you will already know many options available to you and the criteria that are important to you. This example is illustrated in Figure 1.1.

The two intertwined phases of exploration and structuring should give you enough information to construct a table of information, commonly referred to as a decision table where usually alternatives or options are listed as rows and the criteria as columns. A typical structure of a decision table is shown in Table 1.2. Some people prefer to format the decision table the other way around with criteria as rows and options as columns but in this book, we will stick to the format in Table 1.2.

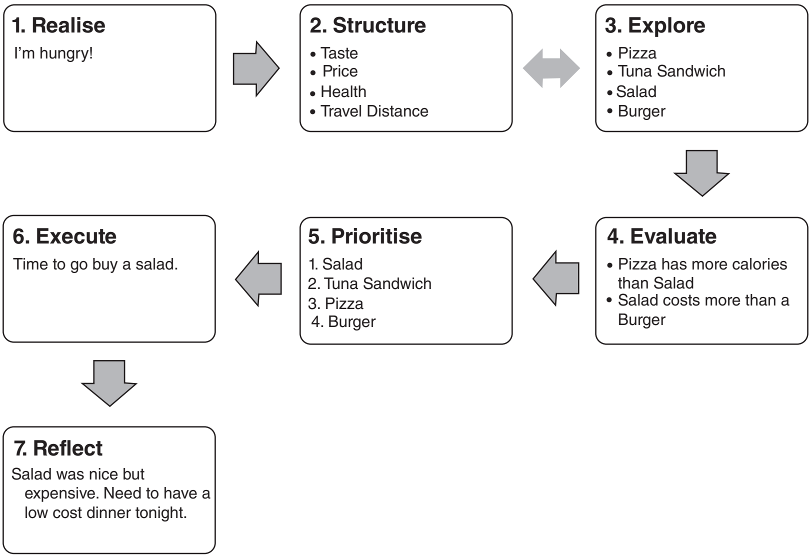

Figure 1.1 Steps in structured decision-making.

After the exploration and structuring phases, the next step in the decision-making process is the evaluation or assessment phase. Each cell in the decision table needs to be filled with information-bearing values. This information greatly varies from situation to situation. For example, in one situation it may be better to use a set of verbal values such as good, medium and bad for assessing each alternative with respect to each criterion. In a different situation, it may be better to look for factual values like battery life measured in years, engine capacity in horse power, product price or brand name. We included brand name here because factual values are not always quantitative and yet they bear information which may influence our decisions. Such non-numeric values will nevertheless need to be transformed into a quantitative value representing a

Table 1.2 The typical structure of a decision table.

	Criterion 1	**Criterion 2**	...	**Criterion** *m*
Alternative 1				
Alternative 2				
...				
Alternative *n*				

score for each alternative with respect to each criterion. For example, see Table 1.3 where the decision-maker has collected information about price, top speed, and brand for the three different cars they are interested in. The information about brand is purely categorical but he/she has assessed each brand to finally assign a percentage score representing their likeness for each brand.

Table 1.3 Example of preference elicitation.

	Price	Speed	Brand			Price	Speed	Brand
A	**£10k**	MEDIUM	Chipsy	→	A	**£10k**	70%	40%
B	£30k	**GOOD**	**Topsy**		B	£30k	**90%**	**80%**
C	£20k	BAD	Midsy		C	£20k	60%	60%

This process of converting raw information into quantitative scores is sometimes referred to as preference elicitation. Many decision analysts promote the use of direct ratings where the decision-maker is asked to assign a score between 0 and 100 to each alternative with respect to each criterion. We will discuss this in more detail in Chapter 3 with some practical examples. On the other hand, there is another school of thought which promotes the use of pairwise comparisons where a decision-maker has to make a number of one-to-one comparisons such as: do they prefer car A or car B in terms of Price and by how much? Some people say this way of decision-making is better as psychological theories suggest that the human mind can only focus on one comparison at a time. We will discuss this type of method in detail in Chapter 4 with some practical examples. Regardless of the method used for preference elicitation, the next obvious task is how to combine this information on different criteria in order to make (or aid) the final decision.

There are numerous methods which combine individual preference scores into an overall combined preference score (or a final ratings/ranking). One of the most straightforward and widely used approaches is the weighted sum method that is discussed in Chapter 4. As the name mentions, a weight is required for each criterion based on its importance, which is usually obtained by either the direct-rating or the pairwise comparison approach. Once the weights are obtained, the individual preference scores are aggregated using these weights.

Another class of aggregation methods are called ideal-point based methods. These methods create an imaginary ideal solution and in some cases an imaginary worst possible solution. The methods then attempt to quantify the distance of each alternative from the ideal or worst possible solutions. We will discuss this class of methods in Chapter 5.

All of the aggregation methods in Chapters 3–5 may need the individual preference scores in the decision table scaled appropriately before use. For example, you can't aggregate (or compare) preference scores ranging from 1 to 10 with preference scores ranging from 10 k to 100 k. For example, the price of cars in Table 1.3 may need to be converted into a 0-to-100 scale to overcome the problem of varied scales. This is achieved through a process called normalisation which is explained in Section 3.1 accompanied by a list of common techniques used for normalisation in Table 3.3.

Chapters 3–5 discuss the aggregation methods where overall scores are obtained by combining the individual preference scores. However, we often confront situations where someone vetoes against some of the alternatives. Another possibility is that some people may find minor differences to be insignificant. In such situations, it is difficult to use aggregation approaches. Another group of evaluation methods are outranking methods which, unlike the aggregation-based approaches, can still be used in these complicated situations with insignificance and vetoing. We will discuss outranking methods in Chapters 6 and 7.

In Chapter 8 we move on from evaluation methods to discuss methods for handling design problems where options are either infinite or unknown. These problems are mathematically modelled with an objective function (or multiple objective functions) along with a possible list of constraints that limit the possible range of solutions. This is discussed briefly in Section 2.2.9 then in more detail in Chapter 8. The goal of optimisation is to find a solution that maximises (or minimises) the value of an objective function without violating the constraints. There are situations in practice however where no solution can be found without violating the constraints. In such situations, the decision-maker may want to identify solutions that minimise these violations. In this respect, we will introduce Goal Programming in Chapter 8 where the constraints can be relaxed in order to find a feasible solution. The optimisation algorithm used in Chapter 8 is designed to work with problems that are represented with linear relationships. This method however will not work well with non-linear objective functions. To handle such non-linear problems, we will introduce the idea of evolutionary optimisation in Chapter 9 where we will show how Genetic Algorithms can be used to solve the travelling salesman problem. This is a practical problem of identifying a route between a number of different destinations that minimises distance or travelling time.

Throughout these chapters, we will introduce the problems with a decision table or objective function containing crisp quantitative values or equations – crisp meaning a precise number, not the snack food. However, we are often confronted with uncertainty in our data, be it the decision-makers assessments or the information we have captured from databases or other sources. In order to assess the robustness of solutions obtained from these methods, it is important to perform a sensitivity analysis which attempts to answer “what if?” types of questions. For example, what if the price of a particular car increases from £20 k to £22 k; will it still be the best alternative? Sensitivity analysis is a very important task to carry out post modelling a decision problem as it captures uncertainty in the obtained information.

This book will conclude with discussing this important topic of sensitivity analysis and will introduce a new technique called Simulated Uncertainty Range Evaluation (SURE) which has been practically applied in real-world applications.

References

“*Thinking Fast and Slow*” by Daniel Kahneman. Published 25 October 2011, publisher: Farrar, Straus and Giroux, New York, United States.

“*Psychology of Judgment and Decision Making*” by Scott Plous. Published 16 February 1993, publisher: McGraw-Hill Education, New York, United States.

“*Preference, Belief, and Similarity*” by Amos Tversky. Publication date: 16 December 2003, publisher: MIT Press, Massachusetts.

“Smart Choices: A Practical Guide to Making Better Decisions” by Howard Raiffa, John S. Hammond, and Ralph L. Keeney. Published 21 July 2015, publisher: Harvard Business Review Press, Brighton, Massachusetts.