

TRANSCENDING
BUSINESS
INTELLIGENCE

Transcending Business Intelligence

Third Edition

Kjell Borking Mats Danielson Guy Davies Love Ekenberg Jim Idefeldt Aron Larsson

Sine Metu

**Sine Metu Productions
Skönviksvägen 256
SE-122 42 Enskede
SWEDEN**

E-mail: info@sinemetu.se

Website: www.sinemetu.se

**Editors: Mats Danielson and Love Ekenberg
Cover and graphic design: Love Ekenberg
Typesetting: Jim Idefeldt
Third revision editor: Mats Danielson**

We support open access as far as possible. This work may therefore be copied and distributed anywhere in the world in any form, freely and without economic gain, for non-commercial use only, as long as the book is unaltered and correctly reproduced in its entirety, including the copyright text below being preserved (CC-BY-NC-ND 4.0).



**© 2011, 2019, 2022 Kjell Borking, Mats Danielson, Guy Davies, Love Ekenberg,
Jim Idefeldt and Aron Larsson**

Third revised edition

ISBN 978-91-978450-5-2

Printed by Elanders Sverige AB, Stockholm, Sweden

By thinking about the causes of, and the consequences of our attitudes and actions, we can affect and change our decisions. We try, for example, to notice which facts provoke us, to cast them aside. We name what we desire and what we fear, thereby gradually becoming better acquainted with our feelings. In this way, our views and values about ourselves and whatever the issue is change. We expand our freedom. If, on the other hand, we never look at ourselves as objects, never ask ourselves if we really desire the desires we have and never try to process our inner selves, our freedom shrinks and we slowly become turned to stone.

Harald Ofstad: Vi kan ändra världen. Hur bör vi ställa frågorna?

[We can change the world. How should we pose the questions?]

Prisma, Stockholm 1990. (Translated from the Swedish edition.)

This book is dedicated to our dear friend, esteemed colleague and co-author Professor Love Ekenberg, who passed away in September 2022 during the preparation of this edition. We mourn his untimely death, and that he is not here with us to see the end result.

Foreword

Sometimes, it would be rather nice if we could take for gospel all that talk about us being the pinnacle of creation with free will at our command, that we are the ultimate manifestation of rationality. But we can't. 'Cos we ain't. In most regards, we are rather limited. Yes, we can invent one or the other nifty widget but in total, our prospects are meagre.

At the same time, the flow of information, dependency on the world around us, the global economy and the continually increasing pace of change mean that old business models or service structures are no longer effective, or even relevant, and therefore need continual reworking. This does place heavy demands on organisations and individuals to work with change. Decisions need to be made continually – and on the whole, things just turn out as they do. Sometimes well, sometimes not so well. But there are methods and tools available to handle and counter a passenger-seat worldview when it comes to support for decision making. Let us show you.

The main problem with human decision making is that **intuition** is seldom sufficient. In many cases, making rational and well-considered decisions unaided is simply too difficult. An even greater problem is that most people don't realise this. People believe themselves to be rational but really and truly their capacity for rationality is rather limited, a situation that's made worse by the delusion that, without any kind of help, we can make wise decisions.

You might think that organisations and people should just get on with things as well as they can and leave it at that, but it's often the case that decisions affect other people, sometimes quite a lot, sometimes crucially, and in some cases they affect many other people or organisms, born and unborn. This applies, for example, to public decisions on policies, not least during the Covid-19 pandemic, and on large-scale infrastructure projects. Not uncommonly, we get the impression of an unacceptably large element of chance, even in very costly or controversial matters. This is quite unnecessary.

We who've written this book find it both striking and worrying that large investments are effected without using more qualified methods of decision making in which priorities, values and weights are openly expressed and taken into account. Feelings and political components are often all mixed up with facts into an impenetrable quagmire. Feelings are certainly important components, not least in the making of decisions, but they shouldn't be confused with facts. That's when things go awry and nobody understands why.

In response to such situations, Business Intelligence (BI) systems have emerged (formerly decision support systems). Sometimes they're hailed as a saviour, sometimes they're beheld like something from outer space and sometimes coveted only in order to keep up with the Joneses. But a large number of all BI projects don't run as expected. Their objectives are seldom met, not because they're technically difficult when conducted more or less in the right way, but rather because... well, read on and you'll see for yourself.

All these misgivings and deficiencies are what we've tried to do something about, using plenty of knowledge and ideas about how to handle decisions and risk. Ideas that we have used in a great number of both large and small projects, both in the private and public sectors. These projects have varied in size between 6 and 1,800 consequences and with budgets from € 10,000 to € 800,000,000 and have included everything from large acquisition bids covering site planning before extending infrastructure to choice of national insurance systems for economic and social risks associated with river flooding. Often, many stakeholders are involved. This means that large amounts of opinion need siphoning off. If you know what's what, then you can distil from an analysis what should be done.

This tends to go fairly well and now we believe that we know *what* should be done in most cases. A nice side effect of our methods is that we can now also see *why*. Sweet!

So now we're going to let *you* know why too.

Some of us have been working with decision analysis for a couple of decades, as university professors and consultants. One of us at least as long with data warehousing and other areas of decision support systems as well as data-

bases, long before the concept of BI had arisen. All of us have considerable experience as consultants in one form or another and together we possess the sharpness of vision of a hawk soaring over the field.

And with our hawk's perception, we have seen that BI has wandered off course. Like many others, we felt, and are now convinced, that supporting tools for decision making is the way to go, though perhaps not for all BI projects. Read on and you will understand why.

We would like your BI projects to transcend the organisational, architectural and implementational difficulties they may have encountered (or may be headed for) and be launched as fully functional systems for the benefit of your business.

For this reason, we've taken wing in writing this text in order to offer you our hawk's view over the field with the addition of the odd bit of advice that we hope might help you hunt for a satisfying decision support solution.

The authors, Winter 2011 (first edition)

Even though the terminology has shifted around a bit, the confusion around and the need for a better understanding of decision support (systems as well as analytic methods) not only remain but are steadily increasing. This third edition has been carefully updated while aiming at keeping the book's style and messages intact.

The authors, Winter 2022

Contents

<i>Foreword</i>	<i>I</i>
-----------------	----------

<i>Introduction</i>	<i>IX</i>
---------------------	-----------

1 <i>Business Intelligence – a Panacea?</i>	1
What Exactly is BI?	2
What BI Ought not to Be	3
What BI Ought to Be	5
Monitoring and Follow-up.....	7
Management by Objectives.....	9
Analysis	11
Decisions	13
Systematically Deploying BI	14
Implementing BI Systems.....	19
Architecture for an Enduring Data Warehouse.....	21
Why do Things So Often Go Wrong?	26
Future Indicators	29
Finally	33

2 <i>Clear Decision Processes – Let’s Sing It</i>	35
The Basics of Decision Processes	36

The Anatomy of Decision Processes.....	37
The Decision Process in a Nutshell	41
Building an Organisation's Decision Competence	42
Screening	43
The Decision Management Process	48
Progress Factors	55
Decision Quality	55
Decision Context	57
A Sound Decision Basis	57
Communicating Decision Problems and their Bases.....	59
To Round Off	66

3 <i>What Underlies an Actual Decision?</i>	67
The Importance of Making Decisions	68
Decision Rules	69
Criteria for Decision Rules.....	73
The Search for Good Decision Rule Criteria	75
The Great Dilemma	79
We Need Sound Methods	80
What Does a Sound Method Look Like?	81
Various Decision Models	85
Determinism	85
Strict Uncertainty	86
Uncertainty and Risk	89
Tree Models	91
Problems with the Expected Value	94
Expected Utility	96

Precise and Imprecise Data	97
Intervals and Comparisons	100
Comparing Alternatives	101
The Gangrene Example.....	104
More on Uncertainty	106
More Advanced Models	110
Several Goals and Stakeholders	110
Multi-Criteria Analysis	112
Sensitivity Analyses	120
Risk Management.....	123
A Final Reflection	125
<hr/>	
4 <i>How it Can be Done & How it Should be Done</i>	127
Bypass Stockholm.....	128
The River Svartå, Bathable	136
Final Words	142
<hr/>	

Introduction

How on Earth did this happen?

The question is not infrequent, but more to the point, it's unnecessary. It arises because in many cases, people haven't thought things through beforehand. Yes, sometimes people can be unlucky, but far from every time that things go awry. Bad luck is never that consistent – by definition.

Thinking things through properly beforehand might seem a fairly obvious requirement for making good decisions, but startlingly often, even critical decisions are made without any in-depth analysis. Well, there may have been some background data lying around somewhere, but not much is done with it. Though it's easy to blame authoritarian leadership or the like, poor decision making is far more usually due to not really knowing what to do with the available information, whether it's sufficient and what else needs to be known.

The situation today is that most organisations don't have any structured procedure at all for handling decisions. Lacking are both methods and the knowledge by which to identify and analyse even fairly simple problems. Universities have been offering training in these methods¹ since the mid-90s and a number of companies in the US have begun to introduce a more structured approach to decision making but in many other countries, this doesn't seem to have gained much ground yet.

But you don't need to do a whole course to become a better decision maker. All you need is to know a few basic principles. And those are what we want

¹ For instance, Stanford University developed a training programme called the "Stanford Strategic Decision and Risk Management Certificate Program" aimed at decision makers in business.

to convey in order to help you understand. Our ambition is that this book will result in better decisions not only by your organisation, be it small or large, private or public but, in the long-term, also by you as an individual.

We aren't going to devote much space to how people actually make decisions or why, although this is interesting, and rather than dwelling on poor practices and irrationality, our focus is on what people ought to do. The key point is that by using well-defined processes, it's possible to systematically condense bodies of information into an adequate basis for a decision. And happily, there are some splendid tools for doing this.

Making wise decisions *is* difficult but can be facilitated greatly by a methodical approach. In spite of this, many decision makers react negatively to the idea that decisions should or even can be made with the help of precise methods. How can any method possibly subsume experience, intuition and common sense? they argue. How can any impersonal tool possibly surpass human judgement? The answer is that human judgement, in many cases, is quite simply insufficient but the human should always stay in the loop.

We're not proposing to replace or eliminate the human decision maker. On the contrary, our aim is to strengthen their position by improving their abilities to make well-founded decisions. We'll be pointing out the common pitfalls, as well as showing how a precisely defined working process brings out the basic components needed for well-made decisions. When decision makers gain a good overview of the data that describes the decision to be made, the final choice can be made with a better outcome and with improved confidence. With a better decision basis, it also becomes clear in which areas more information would provide for an even more well-founded decision. Furthermore, a methodology helps to document the basis for each decision and provides for ensuing assessments and amendments.

There's absolutely nothing weird or idiotic about today's decision makers, but there often is about the tools and methods available to them. It would be more accurate to say that many decision makers have been led astray regarding their remit and capacity, but fortunately, with structured decision processes they can be guided back on track far more easily than had they been as dim-witted as some of their decisions would indicate.

As our starting point, we will be using what is known as decision support systems or Business Intelligence. We'll describe what this is, its strengths and weaknesses. It's through Business Intelligence that many organisations first come into contact with structured decision processes, or at least what is touted as such.

Chapter 2 runs through a concrete example that integrates the collecting of information with actual decision making. There, we tackle how to develop good decision structures and how proper processes can enter into and salvage an organisation. Chapter 3 continues by describing various decision rules and pitfalls when making decisions, as well as what we and others have found out about what to do in such cases. Finally, Chapter 4 provides a couple of real-life examples to illuminate it all and hopefully finally convince you.

This book is about strategic business intelligence, not its rather boring particular instantiations regarding this or that software package or implementation routines. Our quest is much more general – what is business intelligence, decision support systems or whatever we choose to call it really about? How can this technology make us wiser? How can we make wiser business decisions? Do we need support for that? Welcome to our business decision trip. Buckle up and ride along!

*Deciding can leave you perplexed
'cos complexity gets people vexed,
but don't fall for feigning
best choice; we're explaining
decision tools here in this text.*

1 Business Intelligence – a Panacea?

It's a sad fact that up to 50% of a dreadful number of projects for building BI systems are dreadful, that is, they end up in limited use or even total project breakdown¹, even though the ideas behind these projects are splendid, well usually. So perhaps it would be useful to take a closer look at this thing called BI. If nothing else, out of either curiosity or despair.

So, let's begin our decision trip with the concept of **Business Intelligence**, BI. This is an awfully trendy concept, one that spices up business banter in all manner of contexts. What people mean by BI is often broadly unparticular, which well matches most people's equally vague perception of decision support, hence the utility of the BI expression and its stardom as a buzzword.

Whatever the term may denote, Business Intelligence is still, at the time of writing the third edition in 2022, a scalding hot area even if it more recently has been rivalled by concepts such as data mining (which is a much narrower concept). All organisations still need to make informed decisions based on vast amounts of data residing in their information systems. Maybe the area will change names in the coming years but the needs will not go away. On the contrary, they will continue to increase. Enormous disparate databases reside behind the walls of any reasonably sized organisation. So allow us to offer a somewhat less indefinite concept by saying that the core of BI includes most of that which supports the controlling and appraisal of an enterprise – and all the jollies that it involves. Still too vague? Well, read on.

¹ Some say up to 70%. We don't go that far. Whatever the figure, though, it's a lot.

What Exactly is BI?

The most common meaning of ‘intelligence’, and the first entry in most dictionaries, refers to mental prowess; a human quality that’s very useful. This meaning is more recent than the older sense of the word, still prevalent in police and military contexts, concerning the gathering and interpretation of information. This sense lies closer to its etymology² in the Latin *intelligere* to understand, from *inter*, between and *legere*, to read or gather.

Self-appointed business gurus with freshly-coined ownership of recent buzzwords benefit greatly from potential customers’ misinterpreting the ‘I’ in BI for its modern meaning because this makes businesses feel compelled to embrace BI since none would wish to be seen as unintelligent. And so they set up whatever is touted as BI as part of their operation, thus securing their profile as an intelligent business. But being intelligent and practicing intelligence in the second sense of the word are quite independent.

We believe that all organisations have sufficient intelligence for their tasks, in the first sense of the word, just insufficient methods to accomplish the second sense of the word, which is what *we* mean when we talk about BI – the gathering of information, reading between the lines in order to understand it and to make the best use of it through well-deliberated decisions.

BI systems are cool,³ not least thanks to the first sense of the word ‘intelligence’. Organisation after organisation has pumped huge amounts of money and work-hours into BI projects. What is it they’re all after that they so often fail to attain? The promise is to introduce into an organisation something that will improve its decisions and strategic advantage over other organisations.

2 The science of the historical origin, development and relationships between words. A rather interesting science, by the way. Did you know, for instance, that Eskimos have only two words for snow – fewer than for patronising Caucasians? This is contrary to urban myths, which are usually...myths.

3 ... in 2011 at least, as we write this (first edition). In 2022 (third edition), it’s all about the latest buzzword, AI (by which is meant machine learning). However, the need for decision support in general, and wiser decision making in particular, certainly isn’t going to decrease. And for sure, AI cannot solve that. The core methods of BI and decision support will simply in time due be given new names and paraded along the sales catwalk in the latest haute couture.

That's a really strong sales argument indeed. Everyone wants to make better decisions and outdo the competition. Right?

None of this would be a problem if BI really wasn't needed. If businesses could trundle along as they did in the good ol' days passing from father to son with little change. But in today's mean and lean streamlined existence, those old ways will lead a business to the sideline, to watch in consternation as the competition gallops ahead, perhaps offering to buy up and revamp the occasional bystander.

Those who've succeeded at introducing good IT support for gathering **business information** are already halfway there. They have established a good platform on which they can later build proper **decision support** processes, something we'll be looking at from Chapter 2 and onwards.

Returning to the question of what BI is, and the fact that there are a frightful number of differing opinions, perhaps it's easier to start by describing what BI isn't – or at least what it *ought* not to be. Some disturbing quotations below testify to its abuse. We were truly shaken to our bones hearing them.

What BI Ought not to Be

Nobody doubts the fact that **operational systems** constitute a central part of the everyday workings of a business, such as ordering, stock keeping, billing and other such scintillating essentials.⁴ These simply have to work, otherwise business would grind to a halt. At the same time, the *processes* in the business have to support the needs of the operational system. If the two are well-matched then the business will probably run fairly smoothly.

A commonly misconceived attitude towards BI systems is that they're for solving all the issues that the operational systems cannot. BI – the jack of all trades! Perhaps this is at least partly due to suppliers of ERP systems tending to want to solve the report part of their system deliveries by referring to BI instead.

⁴ In many companies, these functionalities and sub-systems are combined into large monoliths called Enterprise Resource Planning (ERP) systems.

“The few reports we can’t produce with our system we solve by adding on a little BI module and that solves all our requirements.”

Such are the self-comfortings of the uninformed...

Equally comical as such underrating of BI is the over-confidence in what a BI system can achieve.

A sales manager at a bank, who also looks after payment card systems, launched a sales campaign promoting credit cards by offering new customers a € 10 starting bonus if they signed up for a new credit account. However, he failed to check whether the operational system would actually be able to cope with the increased clientele. When it transpired that it couldn’t, he approached a BI systems rep with:

“This must surely be a simple matter of fixing the data repository. That’ll solve everything.”

Hmmm, well actually...

Another rather odd outlook is that BI should make as much data as possible available to as many people as possible.

“We collect all data from our operational systems and make it available to our entire staff, they’ll surely get something out of it.”

Whoops! Maybe not the most structured approach.

Others seem to be quite convinced that they’re BI users if they can generate reports that cross departmental boundaries.

“We can now get reports with data both from the financial system and from the sales support system, so we’ve introduced BI.”

Might seem that way to them. You see what we’re up against?

So why do all these misconceptions persist? Woolly **requirements**, that’s why – not to mention the unfortunate tendency to confuse the responsibility for functionality (**processes**) with putting it into practice (**implementation**).

Some ill-fated BI projects begin by choosing the software. There are even organisations that openly proclaim their BI **strategy** to consist of their use of a certain software supplier – hardly a conceptual leap forward, more like starting to build a house by choosing who will supply the workmen’s tools before even considering an architect or drawings.

Although **BI tools** in themselves don’t contain a drawing, an architect *is* needed nonetheless. Some projects choose the most expensive software tools and construct a grand palace of a solution, largely without any input from an architect, with the consequence that eye-popping sums of money (as much as € 5–10 million) are wasted. The projects fail, only to have to start again from scratch. Out of compassion, we shan’t shame the downtrodden; suffice to say that these figures are real.

In many cases, the frustration of the users was so great that the organisations swapped out the grand tool for another one, even though it wasn’t the tool that was at fault. You can guess how effective that extra expense was. In kindness, we venture to comment that this might possibly have been just an ever so slightly sub-optimal way in which to optimise resources. With the right **architecture**, the original tool would have worked.

Investments such as these are driven by highly skilled software salespeople who with fabulous presentations and slick demonstrations show how easily 100% control and regulation can be gained over every part of the business. Sounds appealing, doesn’t it? Their sales pitches cleverly concentrate on the most superficial part of a BI project but the part that is most easily understood, namely the user interface – that’s right, just as you would expect from a salesperson, the focus is on cuddling the client.

That is not where the focus should lie in a BI project. And the tool is one of the last things that should be chosen.

What BI Ought to Be

OK, enough fretting. But where to start?

Well, the first thing is to look at the functionality of the desired BI system, not some software package.

A suitable point of departure is a clear division between what should be included in the operational systems and what should be in the BI system. In order to decide where a function should be housed, there are some simple questions you can ask:

Does the function support the business processes and make them more efficient? If so, then it should belong to an operational system.

Does the function solve tasks related to control and **follow-up**? If so, then it should belong to the BI system.

Operational information systems (IS) directly support **operational processes** and BI systems support follow-up and decisions. A BI solution is complementary to an IS and it should be a support for managing the business at several different levels.

In short, a BI system should constitute an integrated part of an organisation's control mechanism, its steering wheel. As such, it's very much the responsibility of senior management.⁵

Like most software projects, BI projects also contain components of specification and implementation. Ensuring that the specification is correct, that is, that it'll provide proper support to **decision processes**, is the responsibility of, not to mention the interests of, senior management. This is something that they must be prepared to give plenty of time and thought to. The implementation, on the other hand, is the IT department's responsibility. They must ensure that the BI system correctly fulfils senior management's specifications whilst providing them with efficient interfaces for their needs.

Few people are skilled in both of these areas of responsibility and this should be considered when manning a project.

⁵ In consensus-ridden cultures like Sweden, there is a tendency to abdicate from this responsibility which can leave co-workers uncertain as to what is actually expected of them.

As well as describing more or less all support for control and follow-up, with widely different demands on functionality, the BI concept sometimes includes a good deal of other stuff that lies outside the operational systems. In order to gain some control over this and gain a conceptual structure, we separate **BI functionality** into the following sub-categories, each of which we'll discuss in turn.

- ♦ Monitoring and follow-up
- ♦ Management by objectives
- ♦ Analysis
- ♦ Decisions

Monitoring and Follow-up

Monitoring is senior management's opportunity to oversee the function of the whole organisation. It is the collective picture of how the business is working and how this differs from the way in which it *should* be working. This large and important task needs dissecting so that it can be dealt with in several steps.

A good place to start is to introduce monitoring of the business' **key performance indicators** and thereafter use **management by objectives** as a method for managing change. Here, for example, **balanced scorecards**⁶ is a good method, but there are others too. Just select one, but whatever you do, make sure not to control the organisation using financial information alone.

This provides only one, and perhaps not even the most interesting, of all the various perspectives on a business. And as we will see in the coming chapters,

⁶ Wikipedia states: "The balanced scorecard is a strategic performance management tool – a semi-standard structured report supported by proven design methods and automation tools that can be used by managers to keep track of the execution of activities by the staff within their control and monitor the consequences arising from these actions. It is perhaps the best known of several such frameworks and was widely adopted in English-speaking Western countries and Scandinavia in the early 1990s."

most interesting decision problems are **multi-criteria** problems. This means that the problem must be scrutinised from more than one perspective. Since BI systems are decision support systems, it is fitting that they be designed so that information from several perspectives is available.

A small example: a consultancy company in the computer industry (or any industry for that matter) that only manages according to their financial results will single-mindedly prioritise a high level of paid visits to customers.⁷ This makes it all too easy to forget important activities like ongoing training of staff and enabling new employees to gain experience and develop their roles. The usual course of events in such cases is that after a number of years, the company will stagnate, losing energy and competence, while almost all development will be stifled; dissatisfaction will fester and misery reign. In this light, it seems like a better idea to ensure that monitoring and follow-up deal with *multiple* perspectives.

Try to resist including every desired functionality in the very first version of the BI system. Sometimes expectations are so high that it's difficult for even a fairly successful project to live up to them. It is, after all, rather complicated to collect all the important indicators for an organisation and present them simply in an overview.

A desire often expressed by management is a simple and clear control panel where they can see precisely and clearly what's going on in the business. For this to work, fully connected and functioning sub-structures are needed that can deliver the information.⁸ Most tools intended to effectuate such control panels are worryingly simplistic and are hard-pressed to satisfy the need for consolidated and concise information. Simple arithmetical operations are insufficient for reducing an entire business down to a few indicators.⁹

7 ... or as they say in the trade: a high consultant utilisation rate.

8 It's not much use having a dashboard in a car unless the thing is properly connected to the rest of the vehicle. Clip a few wires and see what happens. Errr, on second thoughts...

9 In contrast, the theory for decision support covered in Chapters 2 and 3 is more sophisticated while still being relatively simple, but above all useful. Hang on for some revelations.

Management by Objectives

Management by objectives can most simply be described as a tool for creating the desired behaviours in a business. This can be achieved by defining precise goal-measuring tools that provide individuals in the business with the means and incentive to continually follow up the degree to which goals are being achieved. Traditionally, management by objectives has been achieved by producing a large number of reports on paper at regular intervals. These were then distributed to evaluators with the intention that they should not only read them but also act upon their content. Often these reports would be tailored by the evaluator to target a specific problem.

When evaluators leave their positions or if the problem no longer exists, reports often continue to be written and distributed just the same. That's right! Repetition begets routine and routines beget self-perpetuating traditions, to be respected and upheld; ours not to reason why. So much for effective use of resources – and we're not primarily thinking of the paper.

A strong indication that there's potential for improvement in a business' goal management is a rich flora of reports and forms, many spontaneously coming straight from the operational systems.

Each co-worker should quite simply be able to see what's expected of that part of the business for which she's responsible and can influence. This has to be done in a clear and comprehensible way. Usually by the system presenting what's expected, as well as current and earlier results, in a way that's crystal clear for the co-workers. However, the number of indicators must be greatly reduced. Preferably no more than 3 or 4, otherwise they lose clarity. At a pinch that could be extended up to 7 but anything above that becomes vertiginous.

Seven is the magic number at which the human cognitive capacity takes a nosedive when trying to hold chunks of information in mind. Without explicit support, people simply cannot cope with more and still understand their meaning, especially if they're interdependent and thus correlated.

One grievous situation we've witnessed was a number of middle managers in an insurance company who were expected to assess almost 50 indicators. A non-starter, to be sure.

The part of the BI system that supports management by objectives is the part that has the greatest number of users. Unfortunately, this is the part that receives the least attention during acquisition and deployment. When the various parts of the system are worked over, the most attention is nearly always afforded to the analytical tools, in the narrowest sense too, despite the fact that difficulties in analysis nearly always originate in the bad use and quality of the underlying data.

Tuning the engine is not going to improve performance if the air filter is clogged or you're using the wrong kind of fuel. Key perspectives first.

A well-deployed BI system should support co-workers, enabling them to follow up on how the business and they themselves are performing in relation to set objectives. The interface should be simply presented like the dashboard of a car. There you find continuous indicators of primary information such as speed, remaining fuel, etc. and some parameters are monitored with alarms such as unfastened seatbelts, low oil pressure, etc. It's most unusual in modern cars to distract the driver with oil pressure so long as it stays within normal bounds. Likewise, if any operational part of your business loses pressure, an alarm must immediately signal this to whoever is responsible.

Management by objectives can take on many forms. Whichever they may be, the most important effect is that they create the desired behaviours right the way through the entire organisation and that individuals are provided with the tools they need to see precisely and clearly what's expected of them, as well as how well they're fulfilling set goals.

In principle, it's all quite simple, just problematic to implement efficiently.

Analysis

The simplest form of analysis is the principle that an indicator shows when something is off course and that the relevant area of activity is broken down in order to find the part or parts of the system that are causing the deviation. The obvious thing to do then is to drill down into the underlying structure in order to find the core of the deviation. If an indicator consists of a simple and easily understood structure this is usually doable, especially if it also coincides with how the business is organised. Otherwise not. The users must know and have confidence that the information is always there to be inspected and that it reliably provides the same answer for the same set of parameters. Trust in output data is always key.

If such monitoring systems reveal that the business isn't achieving set goals that are more complex, then more complicated analyses need to be available. Most BI systems are sold on the basis of this very possibility, to carry out advanced analyses of data. The techniques for dealing with this include the use of **star structures**¹⁰ in relational databases or **cubes**¹¹ where data has also been aggregated.

The main problem, however, is that most people just aren't analytically minded. That's just the way it is, however we might feel about it, and those that do have a smidgen more such talents may still not understand the data model in any meaningful way. Thus, one of the underlying assumptions for BI analysis begins to crumble. And continues to do so.

Actually, not even all business analysts really understand data models. BI systems must therefore be able to present the **data space** in a way that's

10 A star structure is a table with the facts to be stored. Around these facts, there are tables of dimensions that represent aspects according to which it's desirable to extract or collate further facts. Thus a star structure makes it simpler for a user to understand and refer to data.

11 A cube is a further development of a star structure that's often generated from a relational one. So why call it a cube? One way of looking at a star structure is to think of facts in three dimensions. For example, sales measured in Euros above have the dimensions product, time and geography. Facts with these three dimensions then form a data space in the shape of a cube. The sales of a particular product can easily be chosen for a given time and place. In order to see total sales on a given day, all products and all locations are simply aggregated. In order to speed this process up, facts are accumulated across dimensions so that answers can be supplied quickly.

understandable to them too, or else run the risk of a wrong analysis based on some incorrectly delimited set of data. This places rather large demands on a BI system. It must also be apparent how to navigate through the system and choose data suited to each kind of analysis; otherwise, you might as well scrap the whole thing. In some cases, that'd be the way to proceed.

Even from the perspective of analysis, there's an important difference between an operational IS and a BI system. The operational system is held together by the processes in the business that it supports. These range from creating bills to buying in. For a BI system, there are seldom such processes and it's often difficult to predict just what needs analysing. In the worst case, there's only data and perhaps some data about data, so-called **metadata**.¹²

So, analyses demand of a BI system that data is:

- ♦ always available
- ♦ sufficient
- ♦ correct
- ♦ persistent over time

In addition, analyses demand a number of things from the user too:

- ♦ an analytical faculty
- ♦ comprehension that matches the level of analysis
- ♦ an understanding of how the data space represents the business information

There are rather well-developed tools that build on various mathematical models. However, it's no less dangerous to rely on the results of these

12 'Meta', by the way, is a prefix that actually means 'between', 'after' or 'over'. Compare here in particular Aristotle's *Metaphysics* which followed his *Physics*. This has been mystified a number of times here and there.

analyses if the user doesn't understand the mathematical background. So it's important to ensure this.

For example, we encountered a salesman who, in desperation over a sales slide that looked like a hail storm, ran a regression analysis on the data, chose a curve that pointed up and included it in his PowerPoint presentation. The fact that this curve had not the slightest connection with the meaning of the data didn't faze him one jot. The system had calculated it, so...

It's also possible to use *data mining*. This is a method, or rather a class of methods, for finding connections and patterns that our brains just aren't capable of picking out from large data sets. However, these tools place even greater demands on the ability of users in order to deliver meaningful and correct results. And ten years on from the first edition of this book, not much has happened, at least not to the abilities of the users.

Don't despair, though. In the next chapter, we'll see how to deal with these issues in a systematic way in connection with an actual decision process.

Decisions

Uncertainty and risk – these concepts are ever-present when we make decisions in a changing world. This often gives rise to difficult questions. How can we best deal with uncertainty? How can we use **probabilities** without statistical data? Can probabilities from data analyses be mixed with other types of probabilities? Which uncertainties are critical in a given situation?

Today, an appalling number of enterprises have no computer support for decision making, unless you count spreadsheets.¹³ And if that weren't enough, they don't even have well-thought-through processes for dealing with decisions. A good manager should be able to make good decisions – all alone. That is, after all, what being a manager is really about. Well, maybe not entirely. Many boards of directors and managerial bodies are required as part of their duties to make decisions that are *also* rational and objective. Whoops!

¹³ Almost every type of application in computer history has been described as able to support decisions, even spreadsheets. In a way it's true. Decisions *can* be made based on data from a spreadsheet, but this requires a meretricious definition. One might just as well call pen and paper decision support.

Unfortunately, we humans are not especially good at making objective and **rational decisions**. We would like to think we are, but our decision abilities are merely an expression of our genome, which evolved to survive in quite savage and primitive conditions. You might, of course, object that such an allusion, replete with wild animals and jungles, is a rather refined metaphor for today's **business environment**, but today the fact remains that in order to make rational and objective decisions we need support. On that point, there is no debate.

A first step is to introduce requirements on processes for the decision work. This is simply about preparing for the situation in which the decision needs to be made and the demands that will be made on the different types of decisions. There is, after all, a rather large difference between decisions about how much paper should be bought in for the printers and the decision about whether to invest in a new production plant. Got that?

The structure in the business' various decision situations should be mapped in order to ensure that it is available together with correct base data for making effective decisions at the strategic, tactical and operational levels. Do not forget either to include different enterprise and business risks. By making proper demands on the decision process and its basis, the system will contain the possibility of an augmented design that will support more or less all types of future decisions. Don't let anyone tell you otherwise. Just don't.

Chapter 2 deals with this very kind of process in businesses. In this chapter, we'll stick to system support for these processes.

Systematically Deploying BI

Deploying BI in a systematic way means having control over the project work from the earliest stages. We assume here that your organisation has the capacity and the experience to follow through a traditional IS project. We'll just top up with some particulars for BI (or decision support) projects.

To begin with, let senior management think the project over and not just give the green light to a low-level readymade proposal. What kind of support do

they actually want to provide to the organisation? Analyse this within the context of the four main user areas that we discussed earlier:

- ♦ Monitoring and follow-up
- ♦ Management by objectives
- ♦ Analysis
- ♦ Decisions

Then when the requirement specification is put together, it usually yields a project that's far too extensive, one that would take an unreasonable amount of time to introduce in its entirety. It's wisest then to introduce one stage at a time, rather than to pressure the project leaders to outdo themselves and deliver everything at once. The latter course of action will only cement deficiencies in one sub-system in all the dependent ones. As the first stage, start preferentially with monitoring and follow-up.

The design and introduction of a BI system are greatly facilitated if the business has an explicitly stated strategy¹⁴ as well as a tactic for how it should be developed. It goes without saying that there are no ready-made BI tools that can solve that. If this foundational work isn't done, the BI project will just make the situation worse. Be assured!

Once the strategy and the tactics have been set by senior management and communicated to the project group, it must be decided how to ensure that the project is advancing in the right direction. This work must not be limited by the information presently available, but rather it should be a *wish list* for how the overarching goals are intended to be secured. It should be obvious that the presently available information should not be allowed to limit the future strategy. Don't confuse the goals of the business and those of the deployment strategy. Start deployment using existing data as far as possible, data that should initially be quality assured. Continue then by supplementing with the rest of the data desired.

¹⁴ ... and a vision and a mission and all that stuff that management consultants promote.

In other words, complete a small success story first. It's an old guile, but it works. Always "fail fast and easy, redo it, then do it right". Always.

Management consultants can be heard dyseloquating this as 'starting with the golden mean', which in this context they've chosen to mean that the goal is to educe, from the available data, everything needed from the business' perspective, see Figure 1.¹⁵ Deployment starts with all the data that is easily accessible and of sufficient quality.

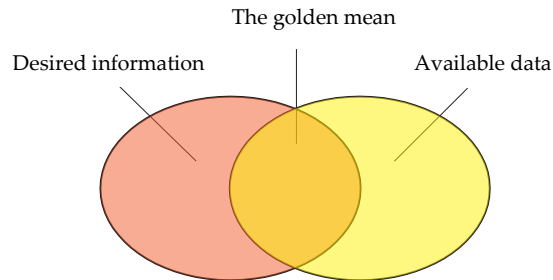


Figure 1. The 'golden mean' of management consultants

It's wisest to start on a small scale, particularly if the business hasn't used BI systems before. The world is both large and complicated and the requirements are guaranteed to change when users realise what BI is and what it can do for them. At that point, to start with a gigantic project is seldom wise, nor does it save time. It's just costly.

Make sure senior management is continually involved. If the BI venture turns into a pure IT project then it's almost guaranteed to fail. The project group should always include a strong devotee or passionate spokesman who's deeply seated in the business and who knows its needs.

The choice of products for deploying BI is, on the other hand, not a direct strategic decision. It might feel like an assertive action but it only puts the

¹⁵ The golden mean is actually an irrational number that can be expressed as an infinite fraction of only 1's. This converges more slowly than all other irrational numbers expressed as infinite fractions and in this sense, it's the most irrational number there is. There now. What management consultants miss out on...

cart before the horse. Such a back-to-front plan is often due to pressure from an impatient management team – *Let's get things moving!* – whose impetuosity is gleefully driven on by slick salesmen on provision-based salaries from suppliers of BI software. Do not allow yourself to become ensnared.

Unfortunately, it's for the organisation as for the individual – that you usually can't buy yourself out of problems, at least not structural ones. Buying the most expensive and sophisticated BI software on the market won't solve a thing unless senior managers know what they want out of BI – and this they *must* know – though they may need help expressing it in precise terms.

We feel it's worth pointing out that there's sometimes a tendency amongst IT people to underestimate the abilities of the management. Managers are considered to have less ability to deal with modern tools and everything must therefore be *so simplified that even the boss can cope with it*. This may have been the case for previous generations of managers, but it's hardly so nowadays. Even for senior management, using IT in support of daily tasks goes without saying.¹⁶ Stay firmly on guard to ensure that senior management maintains responsibility for the general BI functionality and that the IT department sticks to the implementation. It's akin to dividing up household chores.

If the information presented by a BI system hasn't been well-thought-through, but rather become based on the capabilities and structure of a predetermined tool (and possibly also on dubious data), then it's all too easy to become steered in the wrong direction.¹⁷ All too many BI projects where the goal was to introduce a certain tool have fallen apart. The initiative for such introductions is worryingly taken by IT departments in order to reduce the number of reports, especially spontaneous ones, and admirable though a reduction in the flora of reports may be, this is not a primary criterion for the

16 This has not always been the case. Abstaining from the technicalities of IT support used to carry an air of loftiness. Trivialities like dealing with e-mail were handed down to a secretary who printed out messages and laid them neatly on the manager's desk or faxed them to wherever he or she was expected to be found next. But today, even directors have to be able to handle modern tools in order to manage their organisations.

17 This is like relying on an unswung compass while coasting the eastern seaboard of the Atlantic. Experience has it that America rather than Africa then becomes the next port of call.

Footnote to footnote: A swung compass is one that has been corrected for magnetic interference from the ship itself.

choice of tools and certainly not a point of departure for the whole project. The work must start with the business and how senior management wants to develop the business. This is the major difference between BI systems and operational IS. The development of BI demands a completely different kind of involvement from the management during the development phase.

By the way, it's rather common early in a BI project to want to involve people who have good insight into which data is available in the various operational systems. This is admirable. But to allow this to steer the way the system is built up isn't admirable, for it locks the business into the present, rather than embracing the opportunity to transcend the current state – something BI can do, and will do when developed in the proper way. Note: the proper way.

Often the scope of the project is straight-jacketed too, quite unnecessarily, by looking only to the business' own information sources. Broaden your field of view. For effective control and follow-up, data is needed from sources outside the organisation, from the so-called environment.

For instance, there's enormous potential in also using and analysing text-based information and linking this with purely numerical information. When things happen in the environment, they're often first visible in textual information and it's often not until the opportunity for an effective reaction has been lost that the event is reflected in numerical data.

A small example: Today, there's a large flow of information in the environments of businesses, which isn't only directly readable but also provided with XML tags, making it simple to process text for analysis. The days when companies subscribed to press cuttings and announcements excised from newspapers, pasted up on board are long since gone – aren't they?¹⁸

As we've already stated, operational systems should be aligned with the processes of the business. If the systems support the business, it will function better. Unfortunately, there are often no such systems with which to align a BI system, unless, of course, decision processes are already well-documented processes in the business, something that all too few businesses have at

18 A couple of emails a day from the communication and information department doesn't constitute a structured monitoring of the environment, either.

present. Without such documentation, the only alignment that can be carried out is to ensure that the structure in the **data warehouse** correctly reflects the business in general terms.

Without documented decision processes, the data space must be created so that it's comprehensible as is. Having no supporting processes making sense of the data space is analogous to looking for galaxy clusters with a magnifying glass. Such an approach also makes designing all the more difficult.

In many implementations of BI systems, the work that consumes the most time and money involves creating continual flows of high-quality and validated data. The design and architecture of the data warehouse are therefore considerably more important than which specific tool is used. A well-designed data warehouse can survive several generations of operational IS and can be adapted gradually to the changes in the business. Therefore it's worth spending the extra time and resources on the design of the data warehouse.

Finally, we would like to alert you again to the fact that at present, 50% of all BI projects fail because the resulting systems don't deliver what's expected of them. This failure rate is probably tolerated only because the nature of the businesses is such that they can continue to survive without a BI system. Even if that is the case, apart from wasting resources and losing the promised improvements, the whole of the enterprise is subjected to considerable **operational risk** with no return whatsoever.

However, in those cases where a well-functioning BI system has been deployed, our experience shows that it becomes a critical asset that the business would neither wish nor dare to operate without.

Implementing BI Systems

This is not a course textbook on implementation so we shan't be immersing ourselves in technical details. However, since many BI projects have severe problems with implementation a few words about its difficulties are in order, not about implementation so much as about data architecture. This is definitely something senior management should know about.

The fiddliest thing with the implementation of a BI project is often the underlying parts – collecting, quality-assuring and cleaning up the information about the business so that it really can be used in a meaningful way for control and follow-up. If this is successful then it's relatively easy to find tools that can present the information. However, the collection of data represents a formidable difficulty and this is the crux of the matter.

Unfortunately, in many organisations' systems, the quality of data is as solid as a barrel of woodworm. We've encountered sales support systems where 10% of reported sales were for product numbers that weren't even part of the product range. Yes, this really is a true case – another that we'll spare naming and shaming. Aren't sales statistics important?

If 10% of the transactions in a financial system referred to non-existent account numbers, would this be ignored? Hardly! Relying on false data like this yields a distorted picture of reality, to say the least.

In a successful BI project, up to 70–80% of the implementation costs arise from the capture, cleaning up and quality assessment of data. These must be stored in a structure that in a sensible way reflects the business activities of the organisation – a structure that must also be *designed* for the purpose of analysis and follow-up. It's not entirely surprising that this requires starting with the business and its needs. This must be done without exaggerated consideration of the operational data available as the design is begun.^{18B}

Also, developing *only* the operational systems is unlikely to yield a better basis for overview and control. Operational information often lacks a temporal dimension and the underlying data structures are adapted to dealing with transactions, not with follow-up and analysis.

The greatest part of the work with implementation, however, is to provide these functions with correct and quality-assured data; to build a data ware-

18B The concept of ERP (Enterprise Resource Planning, from footnote 4) is sometimes interpreted as meaning that a single system should solve all of the IT needs in a business. Strangely, this understanding prevails in quite a number of international organisations that buy and sell subsidiary parts of their enterprise ten times faster than the time taken to introduce an ERP system. This is another misconception that has spread, not least in the Swedish public sector, which has had what might politely be called "some rather interesting consequences".

house that will constitute the organisation's amassed memory¹⁹ over time. A data warehouse like this also constitutes a stable set of data for analyses covering a given period and one that can be updated at frequent and regular intervals.

Operational systems in a company come and go. The same is true of staff and the formal structure of the organisation. But a well-developed data warehouse endures – at least its architecture does – and this can provide a reasonably correct basis for monitoring and follow-up, as well as management by objectives and not least analysis and decisions.

Herein lies part of the justification for the high-flown discipline of BI – its structural endurance over time. The world at large changes, often faster than we expect. If opportunities are not to be lost, and threats are not to turn injurious, it is essential for a business to move quickly, preferably to be proactive rather than reacting after the event. A well-designed BI system will provide for this ability, increasing not only the organisation's resilience to change but much more importantly, its ability to take advantage of change.

Having said all this, it shouldn't come as a complete surprise that building up a well-functioning data warehouse is a bit complicated, sometimes bordering on difficult. So how to go about it?

Architecture for an Enduring Data Warehouse

A well-functioning data warehouse must be built in several layers, where each layer contributes to the process of transforming data into information, rather like the production stages of a factory, where each stage contributes to transforming raw materials into an end product.

Because of its sheer quantity, dealing efficiently with the information in a BI system and describing its data space to users, requires a high-level yet meaningful description. Enter KPI.

¹⁹ This 'memory' also needs a well-developed security function since it contains information that can severely damage the business if it were to fall into the wrong hands. Security issues are dealt with in many books. Read at least one to educate yourself on what is a truly crucial matter.

Whoa! ... what does this mean?

KPI is that well-established TLA²⁰ for Key Performance Indicator. A KPI is always a quantified unit, such as ‘sales in Euros’. This unit can be further broken down by time period, by product, by region and so on. Each ‘by’ defines a **dimension** of the indicator. By setting up a list that consists of columns with all the desired KPIs, and a column with all the desired dimensions applicable to that indicator, you can quickly gain an overview of what the data warehouse should contain.

In addition, you’ll need to look at the various layers, or depth, of the data. The time dimension might find expression as years, months or days. When building the data warehouse, it’s important to see which KPIs share dimensions and at what depth. All KPIs with common dimensions should be viewed in the same context.

In conjunction with this work, it’s often appropriate to sort out the problem of definitions. This usually requires more work than you’d imagine – whatever you imagined! All the people involved are usually convinced that their definition is the one and only, and that everybody else should see it their way.²¹

Quality assuring the underlying data is essential for assuring the quality of indicator values. It is only in the financial system that data is relatively well-defined, owing to the accounting structure. In all the other systems, both the quality of the data and the definitions are far worse. For example, in sales systems, there may be no uniform definition of what a customer is. A product number that doesn’t exist may happily be sold with not a twinge or whinge of discord from the sales system. What are those salespeople doing when 10% of sales are for hot cakes straight out of a vacuum? It is rather less likely that the financial system has such glaring anomalies since entries are subject to balancing, checking and the scrutiny of auditors.

20 TLA is a well-established three-letter acronym for... well, you guess!

21 This attitude has given rise to many grievances. Allow us to cite Bertrand Russell: “The whole problem with the world is that fools and fanatics are always so certain of themselves, and wiser people so full of doubts.” (Nobel Prize laureate in 1950, but not for this quote.)

Suppose that we want to measure the number of purchases. These are the kinds of questions that emerge: What is a purchase? When is it completed – when the order specification has been filled out or at some other point? Where can this information be gleaned?

By working methodically, indicator by indicator and dimension by dimension, common definitions for the whole business are created.

And now to the physical data warehouse.

Not so surprisingly, the steps for transforming data into a data warehouse are numerous, so here we merely offer a short overview of how they can be phased.

- ♦ Extracting from the source system
- ♦ Archiving batches of extracted data
- ♦ Staging data transfer with quality feedback to the source
- ♦ Standardising formats and normalising data
- ♦ Building dimensions and performance measures
- ♦ Cubes and views
- ♦ Final user tools and reports

Extracting from the Source System

Once KPIs have been defined, the work of finding the data in **source systems** can begin. Here it's important that the people responsible for each source system are also responsible for the **extraction** process. There are two reasons for this. First, considerable knowledge about the source system is normally needed in order to find the desired data. Secondly, it's important to consider the flow of data into the data warehouse when making any subsequent changes to the source system. This step must also be included in tests of any new versions of the system, something that's all too easily overlooked.

Archiving Batches of Extracted Data

When data is collated from several different sources, discrepancies may be seen that weren't previously apparent. In order to rectify these, the cause of each discrepancy must be found. Therefore, when data is extracted from its original source in the operational system and transferred to the data warehouse, it must be **tracked** in order to render it traceable. This tracking history is often called the data's **lineage** and it has great importance because data without lineage undermines users' confidence in the system. Users need to dispel any doubt about the validity of information and its origin before they can base important decisions on it. A simple way to track data is simply to timestamp it. If data is loaded from a file, then the file name and row number in the file should be saved together with each record. This is because source systems are often purged and changed. By archiving extracted data in its initial format, it can be reloaded if need be, even if the original source has gone up in smoke.

Staging Data Transfer with Quality Feedback to Source

From the source system, the data is read. While the transfer is carried out, checks on underlying properties are carried out, such as correct data typing and column counts. If errors are detected here, it's best to leave the offending records where they are so that they can be dealt with later. Often it's enough just to note the error in the data warehouse's log file, and to send out an error message – by e-mail or text message for example – to the system manager, but if errors affect important content, the data warehouse should not be “opened” until this has been corrected. During loading the number of records processed should also be logged as well as the time spent loading. Everything in the name of traceability.

Standardising Formats and Normalising Data

The term ‘normalise’ is used in database contexts to mean that data is stored only in one place.²² This isn't what it means here. Here it means that data is adjusted to a standardised model that's used in the data warehouse. For instance, different source systems might use different codes in order to

22 Well, sort of. Greatly over-simplified.

denote the same customer. Product catalogues might come from different systems with different 'layouts'. The most important part of normalisation, however, is to ensure that the relationship between data is consistent. So, no sales of non-existent products, no empty order lines and no orders placed by phantoms, ghouls or deities. All such details should be adjusted at this level in order to render the right structure for the right business model and the right business rules. It's really quite straightforward, but surprisingly all manner of cheating and diddling goes on within this step.

There are a lot of traps to fall into and mistakes to make during this step. At the same time, they are rather easy to avoid if only some BI commandments are adhered to. Let's make a short enumeration: i) Avoid over-normalising the data. This will lead to overly complex and fragmented BI data structures. Let the complexity reside in the source databases instead. ii) Take extra care in identifying non-obvious relationships. This can really only be done by communicating with the owners of the datasets involved. No guessing from specifications or documentation. Walking the walk is required. iii) Consider scalability from the outset. Your BI application will either be a hit (hopefully) and then attract a lot of new users with new requirements, or it will be a dud that (hopefully) will be closed down at the earliest possible point in time.

Building Dimensions and Performance Measures

From standardised data, we can now begin to build star structures which are the foundation of all data warehouses. Since star structures consist of tables of dimensions extending facts, and since KPIs are stored across those dimensions, it's important to determine and set the system's dimensions first. After that, values for updated KPIs can be stored. Here too, all changes need to be logged, not only for maintaining traceability but also for the purpose of understanding how data storage is adapted to the structure of the data warehouse, as that structure develops. Describing the structural development of the data warehouse also permits analysis that can help anticipate possible problems with capacity due to growth – remember data just keeps on rolling on in. The description also permits analyses that can help anticipate possible problems with capacity.

Cubes and Views

Often, the user is not permitted to work directly with the relational database. Instead, a so-called cube²³ is created. The advantage of cubes is that they contain metadata (data about data) which means that the tool chosen by the user can work directly with the data space and present **aspects** of it to the user. The cube also stores aggregates of the data, which often provides for really swift responses to users' queries. But great care is needed here because there's no general standard for metadata for defining cubes. A well-functioning security system must also be in place since the data warehouse contains information that should not fall into the wrong hands. Sometimes a few keystrokes suffice to steal the whole of a business' history. Dolesome!

Final User Tools and Reports

Too often, formal procurement and evaluation are only applied at the level of user tools. There's no question that the user interface is important but far more important is the underlying structure. No tool can compensate for shoddy data warehouse architecture, nor for missing **data quality**. Remember, too, that different users require different tools of varying functionality. No one needs everything. And tools with greater functionality are often more difficult to learn. It's a little like giving a smartphone to an octogenarian.

Why do Things So Often Go Wrong?

The most common mistake is the lack of overarching architecture. This isn't so strange if there's no obvious structure to stick to. Clear decision processes are often absent and then the only thing that really helps is experience and knowledge from having previously implemented a BI system and having learnt from one's mistakes. Therefore, it can be wise to bring in outside help the first few times but without relinquishing the initiative or responsibility.²⁴

23 ... which is a storage form prepared for presentation and analysis, see footnote 11.

24 Don't let the huskies direct the sled – an old Inuit proverb we just made up.

The problem of low data quality in the underlying systems is nearly always underrated. Correcting this problem is nearly always a long process that often involves communication right down to the grassroots of the organisation where the information is generated. Various (un)intelligent objections to **quality assurance** arise in this context. Does all this really need doing? Surely a few wrong references to a product are of little consequence? As long as the price is right, what difference does it make?²⁵ Such avoidance is rife.

Data quality is an issue that concerns the whole business and is attained in an iterative, often time-consuming process. Those who contribute data must also have access to it in the refined form so that they understand the importance of good quality. If they don't understand what the information is to be used for, and nobody demands correctness, then carelessness leading to errors is a pretty predictable outcome.

It's also often the case that definitions in the business are unclear, contradictory or simply wrong. This occurs both on a large and small scale. It can concern fundamental concepts or trivialities. We've come up against production and sales departments that have had completely different definitions of a day. Clearly, production days must start at 07:00 hrs with the first shift, whereas for the sales teams it was just as obvious that the day should start at 00:00 hrs – not so strange then that they could never agree on what was coming out of production during a day. A BI project must ensure that a common and non-relative set of concepts is formed over time as the BI system is extended to include large parts of the enterprise.

There are often problems collating the operational system with the data warehouse. If the BI system doesn't contain dependable data that its users know they can trust, then it won't succeed. Users' mistrust will be greatest in the beginning, so gaining their faith in the data is particularly important when the system is least mature. That's why it's so very important from the outset to ensure efficient traceability so that the source of discrepancies is quickly

25 You have no doubt been shopping on occasions when several different but similar items had the same price. Maybe three pizzas with different toppings of the same size from the same manufacturer. The cashier scans the first and then presses $\times 3$ on the till since "*they do have the same price after all*". Well, yes, the sum total on the receipt will be the same. But what about the quality of the stock level data for those three products? The devil always swims around in a pool of details...

and easily found, allowing it to be rectified at the earliest possible stage, preferably *before* being discovered by users.

BI system data should consist of a never-ending sequence of snapshots of the business's life cycle that can be played back – one per day, usually. These should be consistent with all the other systems.

It is also necessary to decide how deep down in the structure data should be stored and what degree of granularity it should have. Once data has been aggregated, it's no longer possible to see what has happened at the levels below the aggregation. So the system should be constructed in such a way that any discrepancies that arise can easily be traced and completely explained. This will uphold the confidence of users.

It's easy to create too many and too complicated goals, indicators and reports. So, be sensible, especially during the first version of the system, and focus on clarity and structure. The analysis part cannot be solved solely with the help of indicators. Analysis requires other tools. The senior management must not just envisage a system that visualises information and possibilities, they must decide what's important and ensure that the business runs along those lines.

Another common problem is to fetch information directly from many different source systems and sometimes even from different levels in the data warehouse. This creates an unpalatable spaghetti structure²⁶ which after a while can neither be dissected nor maintained. We've seen some very large projects that have had to be abandoned halfway, in order to start afresh because nobody was able to see through the tangle. It might seem simple to quickly correct an error from a source system by correcting the report where the error can be seen but then, when the source system is corrected, the report will be wrong. Such corrective iterations are tiresome and benefit no one. Instead, already from the outset, there must be a clear structure with precise accountability for each step in the refinement of data.

Inappropriate forms of presentation and **visualisation** of data are also worryingly common. It's tempting to fall for fabulous forms of presentation.

26 ... all knotted and completely indigestible.

However, these often only mess things up. Beware! Never forget that it's the information that's important. Simple forms are preferable; be suspicious otherwise.

Oh yes, and before choosing the form of the presentation, it might be in order to conduct a usability test of a mock-up of the tool. Does it support the users in efficiently achieving their goals or is it just pretty? Three-dimensional effects and cinematic fly-arounds almost never enhance understanding of the data. It's frighteningly seldom that any kind of usability study is conducted before a tool is purchased. However well-designed the underlying system is, if users find the interface to obstruct their workflow, they'll be reticent to use it.

Future Indicators

Benchmarking is a relatively simple way of comparing how well a BI effort has succeeded. With only a few simple types of measurement, a comparison can be made with those who have done similar things. But adequate measurements can be difficult to define. The possibilities include:

- ♦ Development costs and initial licensing costs
- ♦ Total annual running costs of the system
- ♦ Total annual costs of the system per user

These offer a first indication of which to take a position, both internally and in comparison with other businesses. A little benchmarking never hurts.

We've seen such simple and concrete types of measurements as the annual renewal cost of a licence per user vary from almost zero to € 10,000 – per head!

There is no assured way to determine unequivocally whether a BI investment has been successful or not. It is usually a matter of degree.²⁷ However, there are a number of symptoms that indicate possible problems. The absence of

²⁷ In some cases, though, the failure is so obvious that the project is abandoned.

these indicate success just like the absence of symptoms of ill health in you, dear reader, indicates a degree of good health. So what are these loathsome symptoms?

- ♦ The cost is well over twice the planned
- ♦ The BI system isn't yet launched
- ♦ There are only a few users
- ♦ Dependency on consultants and individuals is high
- ♦ Problems with collation and low credibility
- ♦ Manual efforts are required to adjust data loading
- ♦ Continued demand for reports from operational IS
- ♦ Attempts to correct problems by changing hardware or software
- ♦ A tendency to throw good money after bad

Let's look at each of these symptoms in turn.

The cost is well over twice the planned. The cost of a BI project is often difficult to judge for external consultants too, because so much of the work involves quality-assuring data. Though often the problem is more one of over-confidence that the tools will perform miracles. We've seen estimates from large software companies where the introduction costs are set at 10% of the software licence. This may cover the installation costs, but hardly much more than that. You'd be wiser estimating the introduction costs at *ten times* the licensing costs. Akin to the good ol' belt-and-braces principle.

The BI system isn't yet launched. In many cases the BI system becomes relegated to a lonely and forlorn existence outside the comforting embrace of the usual maintenance team, often never becoming accepted by them. This can be due to ongoing changes but also a lack of clear separation between development- test- and maintenance environments. Sometimes there's no defined procedure either, for how to progress from development to testing and then

on to production. If the BI system is not properly launched, then that's not so surprising and clearly indicates that all's not well. Often some kind of external intervention is required in order to get things going.

There are only a few users. Monitor this well and closely. Read the logs and follow up on how many users of the system there really are compared with the number initially planned. Obviously, unless the support system delivers what people need, they'll stop using it. **Degree of use** is therefore a good measure of success.

Dependency on consultants and individuals is high. Many BI introductions are conducted by consultancy companies. There is nothing wrong with that – on the contrary, you may well need help from somebody with experience. But it does become a problem when the consultants remain because the business has become dependent on their continuing support. Control and responsibility must *always* be transferred to the system owner, which in turn means that the internal IT department must accept the system. We stated earlier that a BI project should be initiated and controlled by senior management and not by the IT department, but that is not to say that they shouldn't be involved from the outset. They should. Otherwise, you risk the IT department viewing the project as an outside project that they are unwilling to admit into their maintenance regime. That would be the end of it.

Problems with collation and low credibility. If the credibility of the BI system is questioned and the users don't trust the data from the operational systems, you have a **big** problem. In this case, the whole extract, transform and load (ETL)²⁸ process for **eliciting** and quality-assuring data, including ensuring traceability, must be overhauled. If the users don't have confidence in the BI system, it is utterly useless.

Manual efforts are required to adjust data loading. If problems in loading the data warehouse often arise, then the loading process is probably incorrectly designed. Perhaps this is because of carelessness in control and corrections. Unfortunately, this isn't unusual. A common mistake relates particularly to the E in the ETL process. The **extract function** is used to educe data from the source systems. However, this seldom works well after the passing of time.

28 The sequence of work phases when building data warehouses.

The responsibility for extracting data from the source systems must be shouldered by the respective owners of the source systems and be regulated by a Service Level Agreement (SLA).²⁹ These system owners, therefore, need to know what data they must provide the data warehouse with, and when any local alterations to their local systems are made, they must ensure that they continue to provide data to the data warehouse with the agreed format and frequency.

Continued demand for reports from operational IS. If users continue to solicit new or updated reports from the operational systems, then the BI system has failed to meet these users' requirements. The reports reflect an information or analytical need that should be covered by the BI system. Users should be able to find what they need without special services directly from the IT department. A high degree, or at least the potential of a high degree, of self-service is one of the hallmarks of a well-functioning BI system.

Attempts to correct problems by changing hardware or software. Many businesses try to correct problems they may experience by acquiring new and more powerful hardware or 'better' software products, often just more aggressively marketed – and more expensive. In many cases businesses are strung along by the unscrupulous yet all too common software 'development' strategy whereby poor design and deficient testing of prematurely released products is covered up with the old adage: "If you upgrade to our latest version everything will work beautifully". Yeah, sure. Unfortunately, these products are often innately flawed. But punters will still pay to have their cages gilded.

A tendency to throw good money after bad. Sometimes the unsettling phenomenon occurs of not terminating poor investments even when they don't work. "But we've spent soo much time and money on this, so we must continue with it". That would signify a good time to bring in an independent consultant to evaluate the whole venture.

29 This is often used to describe exactly which level of service should apply. Often, this stipulates the conditions of delivery and also the quality of the deliverables.

Finally

So the whole undertaking is really not so terribly difficult, even though the pitfalls are many and at times deep, particularly for the data warehouse and the specification of the BI system, and not least in relation to the decision processes that it should support. In conclusion, one should, as usual, think a little before acting. Thinking rather a lot would of course be even better. Acquire the services of somebody whose hand you can hold if that feels safer, but retain the initiative and, above all, the command.

Be sure to distinguish between responsibility for the specification and for the implementation. Both are important but their natures are quite different. The ultimate responsibility for ensuring that the specification effectively supports the decision processes rests with senior management. That doesn't mean that they should sit daily and edit the latest version of the BI specification. Responsibility for daily work should, of course, be delegated – and with clear directives. Since there won't be much routine in **structuring** decision processes, clear managerial guidelines are needed in order to carry this out efficiently.

The world isn't evil, even though gremlins sometimes seem to thwart human endeavours³⁰, but it is complex. The introduction of a BI system may be spared the antics of those gremlins by being done in stages. We suggest beginning with the business functions for monitoring and follow-up. When these work well, the BI system can become a platform upon which to build, as well as a medium through which to aid management by objectives. If the underlying data warehouse has been correctly and wisely designed, it will constitute an excellent basis for various analyses.

But what is it we really want – ultimately?

Assuming we're sound of mind, what we want is to be able to make good decisions at every turn. All this gathering of information, storing and analysis is really about providing a basis for decisions, whether they be small daily

30 ... not least when you are stranded with a BI system that isn't deliverable by the deadline and the budget has already exceeded the famous π factor.

ones or ones of greater strategic nature. Here lies another hurdle. It's more difficult to get a BI system to work well in an organisation that doesn't have a structured way of working with decisions. For this reason, next we'll be looking at something that's at least as difficult as building a BI system – namely establishing a sound *method* to support decision processes.

A *complete* decision system consists of a supporting BI system *and* decision makers. These two components must function together as a synergistic whole, like two symbiotic beings, where the whole system is greater than the sum of its parts. It would be most convenient if this could be achieved only by carrying out the technical construction and loading up quality-assured data so that all interesting views of all the information were easily accessed from anywhere. If that alone produced synergy it would be wonderful – but it doesn't.

This chapter was never about teaching you the details about implementing a BI system. It is rather about the strategic decisions that have to be made and how to go about implementing them. There is always someone who can to the nitty-gritty details, but why? What's the purpose? Does it serve the goals of the organisation or is it just for vanity, flair, perceived needs? Some soul-searching at the organisational level will do the bottom line very good here.

The methods and knowledge in this chapter, important though they are, are *not* enough to be able to introduce efficient and modern decision support because decision *processes*, at least well-formulated ones, don't just self-assemble or spontaneously result from clicking through umpteen different data views. For decision systems to work properly, the decision makers must be able to handle and take advantage of the information that the supporting BI system delivers, and to do so in a structured and integrated way – which brings us to the next chapter. Hang on!

2 Clear Decision Processes – Let's Sing It

Chapter 1 established that a BI system should support a business' decision processes in a similar way to how general information systems (IS) should support a business' operational processes. This is one of the simpler ways of understanding the role of BI systems in an organisation. So let's take a look at how to orchestrate decision processes.

Decision making is the major key around which the whole enterprise is orchestrated. The orchestration works together with some kind of methodical accompaniment or performance routine, even if a methodical accompaniment is not always precisely scored or even consciously formulated. For each playing unit, be it the whole orchestra, a section or a solo player, each must make decisions of the most varied kind in order to express and fulfil the goals of the work to be performed. The interpretation of how to adapt performance to a particular context depends on the interpretation of the alternatives and how they find expression. The simpler the choice, the simpler the manner of resolution and the more routine and practice can be relied upon. Even if the preferred style of performance is improvisation, common structural patterns and methods are still essential for players to perform well in a coordinated way, with different aspects and periods of the performance delegated. Sharing responsibility for harmonised activities ensures that the flow and dynamics of the work stay on track. And on track they need to be.

Decisions don't just spontaneously arise, they come from somewhere. Being aware of the origins of a decision problem is the first step towards attaining quality in decision processes, which brings us to the process itself.

Early decision situations are often classified as decision *problems*. This perspective views the situation as resulting from external conditions beyond our direct control that force us to act and choose a specific path or strategy. Another way to view it is through its potential. This means that we see all the possibilities the situation has to offer of replacing the current well-trodden path with one or more of a host of newly identified ones. However, it is not entirely obvious that it is necessary to change the path. Also, surprisingly often the perspective from which we view a decision situation steers the form that the problem assumes. Therefore it's preferable that problem formulations should be built into a systematic process – a decision process.

That's why this chapter dips into a discussion of the underlying theory of decision analysis and its related processes. The chapter also shows what a well-structured and efficient process for handling decisions can look like, as exemplified by the Preference Decision Management Process – PDMP. It deals really well with complicated issues in decision situations, both organisational as well as analytical. It's as good a managerial tool as you could hope to find. But first, let's tune in to some basics.

The Basics of Decision Processes

So how should decisions actually be made? And what are the processes and methods to do this? The cross-cutting issue for both of these questions is *why* we need to make a given decision at all, and if we do, what it is that makes that decision complex?

The risk of making the wrong decision using **intuition** rises with the complexity of the decision problem and is usually far greater than over-confident managers realise. Often, supporting analyses are carried out *after*

a decision has been made. This is a screaming example of what's known as **conformational bias**¹ and it's a scary realisation.

The Anatomy of Decision Processes

It's a consolation that there's a variety of models that can be of help in making decisions, well, at least for delineating the conditions that provide the basis for making decisions. There are also several generations of **decision models**. There are the slightly older, traditional ones, and the more modern ones which are better suited for use in processes.

A characteristic of most of the traditional decision models is that they require relatively well-structured problems, but don't let that scare you off from using models for analysing **ill-structured**² strategic problems. This can be done too, as long as analyses are included in the process.

By using a clearly defined working process, the basic elements of the decision situation become more evident and the holistic perspective on the problem improves. A structured working process provides decision makers with a good overview of the material on which the decision is based, what we call the **decision basis**. When the underlying material is clearer, it'll be easier to see in which areas any further information is needed as well as the nature of that information. This is important for resource allocation. During the process, everything can and should also be documented, clarified and evaluated.

Seems a bit woolly? Don't worry. It'll soon become clearer. We hope.

There are several schools of thought on how decisions should be approached. Let's start with the intuitive approach, that feels a bit comforting.

1 This is the term used by cognitive scientists to describe our human bias for seeking only to confirm our beliefs and not to look for, or to ignore, indications that we might be mistaken.

2 An ill-structured problem is characterised by one or more of the following: unstated constraints, insufficient information, conflicting information, vague goals, more information alters the problem, there's no solving formula or algorithm, multiple solutions exist, solutions are not comparable so none is optimal, multiple criteria for evaluating solutions mean choices need justifying. In contrast, well-structured problems have enough information about the problem and solutions to derive the latter from the former using systematic procedures that find the best solutions.

Intuitive Decisions

The word comes from the Latin *intueri* which means to observe, look at or regard. Intuition is the brain's amazing ability to form an immediate understanding or judgement without consciously perceiving all the necessary facts. Intuition offers a crucially important survival advantage when wandering through jungles or savannas. Imagine it – the trees are fruiting, the birds are singing, the monkeys calling, and suddenly a tiger, crocodile, snake or some other less-than-friendly creature appears. This is perhaps not the best occasion to deliberate over decision alternatives, their consequences and respective probabilities. Those individuals of our developing species who showed any such propensities were greedily gobbled up long before they identified the best course of action. The ones who acted swiftly and decisively without thinking were the ones who survived, the descendants of whom, many thousand generations later, came to populate today's businesses and management teams.

That's the underlying problem with decisions. That is why we need help – lots of it – all of us – to make rational decisions.

It's not only the hastiness of intuitive decisions that haunt management but even more detrimentally, the false confidence that managers often have in their decisions. Were it only the hastiness, it would be easy to avoid mistakes with a little reflection but overconfidence dispenses with reflection and urges us to act now. Most managers think they're good at using a later more rational refinement to review an initially intuitive judgement. This is an ignorant and worrying belief because it is *completely wrong*.

Even knowing about this weakness doesn't help much without a methodology to bypass the bias toward one's own intuition. Worse still is that this false belief is a particularly unwelcome fact for a person in authority to be reminded of, just when he's filled with the satisfaction of having made a 'wise' decision. We authors suffer from this too. We're pre-programmed in just the same way. None of us like to admit our weaknesses, nor to backtrack, but we have to accept our innate imperfections and light a candle for those whose cogitation genes became tigers' tacos or crocodiles' crunchies.

If we're to say anything positive about intuition, then one advantage is that we can be sensitive to factors that perhaps can't be dealt with through rational analysis or that we find difficult to evaluate algorithmically. The question is which information intuition really can capture, over and above that which we can supply consciously, and to what extent we have reason to put faith in that. Perhaps it's closer to hand to believe that decision makers resort to intuition due to a lack of background knowledge, time or competence – or some combination of these.

Intuition is characterised by an ability to capture a situation in its entirety. It builds on two fundamental factors: information and experience. The never-ending stream of information and percepts are all processed to some extent, with no attempt to limit either, both of which greatly exceed what we can deal with consciously. But in order for intuition to work reasonably well, a decision maker needs to be extraordinarily well-informed and extraordinarily experienced with similar situations. Here, self-assessment is yet another opportunity for misjudgement. As is being prone to grandiosity.

When we find ourselves in a decision situation, we perhaps try to grasp the situation as well as we can, at least to gain an overview. We try to use earlier experiences and knowledge in order to reach a good decision, or at least a not-too-abysmal one. If we've never faced similar situations the result is often consternation or decision paralysis unless we let intuition take over – our untrustworthy prehistoric autopilot. **Strategic decisions** often suffer this fate because they're often one-time decisions of which almost nobody has experience.³

The Need for Rational Decisions

If the informational basis for the decision is also ill-structured, then the decision is even more difficult for a decision maker to grasp. We'll say it again. We *need* decision procedures. We *need* methods. We *need* models. All else is an unwise risk, a foolhardy gamble, and all the more so the more complex the problem.

3 It has been shown that many decision makers when faced with complex problems and new situations still place greater faith in their own intuition than in rational analyses.

If you don't believe us, then believe the research results that emanate from work on decision processes done over many, many years in many, many countries by many, many researchers. Huge numbers of decisions have been observed, of varying quality, from the dubious to the catastrophic. The majority of these decisions were bad because valuable information was ignored, dispelled or misunderstood. The very least one can conclude from this work is that the methods for collecting and organising information for making decisions can always and should always be improved and that a more systematic approach results in better decisions.

But we need to adopt and *follow* a good process in order to achieve this. Some adopted methods aren't used the way they should be: there isn't the time; there isn't the manpower; there's too much confusion; or people simply forgot. It just didn't happen. Denial mode – a far too common illness.

Even though decisions constitute an important part of our lives, very few of us have any training to speak of in making them. Many decision makers consider themselves to be sufficiently good at making decisions, that they don't need any training or methods to help them. Imagine how well a professional football player would fair who considered he didn't need a coach or training programme. Protected by loyalties, the dependent reputations of colleagues and stakeholders, outside observation and often a lack of accountability, decision makers can consistently make poor decisions at low personal cost. If the internal affairs of organisations were as transparent and accountable as those of footballers on the field, the incentive to train methodically in order to improve the quality of their performance would undoubtedly improve decision makers' track records. Apart from the importance of an incentive, it's important to note that help in improving performance is usually needed from an outside source. An individual gains from advice from other individuals and an organisation from other organisations.

Those who deny the need for external help, have no choice but to learn from experience. Experience is an expensive and inefficient teacher, and unreliable too since decision situations are so incredibly varied. Often there are few similarities between decisions in different situations. But with methodology, it *is* possible to become good at making decisions because the similarities lie

not so much in *what* is to be decided as in *how* the decision process is constructed.

For the individual, this entails learning about a process that leads to the best decision with the least possible cost in time, money and effort. For the organisation, it's about actively supporting this process.

A decision method is used in order to structure the underlying material and to deal with uncertainty. It's a way to make the decision process more effective. The process is not a formula that automatically leads to the correct solution. It works as a tool to systematise and process the underlying material, and the methodology works as a tool selector and sharpener.

Some managers consider it unnecessary to work according to a formalised decision process. Perhaps they feel that it's embarrassing not to act assertively and immediately. For people with such profiles, action tends to yield all the more satisfaction when framed in a risk scenario of pressing importance with almost no time for deliberation.⁴

But this approach is topsy-turvy. Experience from many professional areas shows that the less time there is to make a decision, the greater the need to work towards a decision in a structured way and *not* to fall for the temptation to act as the heroic Cro-Magnon. Any time gained from an intuitive decision is often lost in the time (not to mention the money) spent trying to compensate or correct for poor choices based on incorrect evaluation norms, omitted parameters, unnoticed alternatives, you name it...

The Decision Process in a Nutshell

First, a model is built. Then it's analysed. Refinement and reanalysis of the model may then iterate.

⁴ This behaviour is a throwback to jungle survival. Overt displays likely helped females select fitter males, which might explain some of the antics of modern males in soccer shorts as well as business suits. Experiments in various settings show that in general, men's risk aversion falls notably in the presence of attractive women, even when this reduces their overall success. No doubt most of us have witnessed bravado of this kind.

To build the model, first identify the decision process. Try to comprehend the goals, and take into account various perspectives on the decision, which we will call the aspects of the problem. Then identify the alternative actions that address the problem. Then identify calculations of the consequences for each alternative action and estimate their probabilities and utilities. Then begin the evaluation of all these parameters.

When an alternative has been identified that appears better than others, a **sensitivity analysis** can be done. This tests whether the best alternative indicated is still better when uncertain factors are varied. If there is considerable uncertainty about which alternative is best then more information is needed. With that in hand, return to and modify the model. When you're fairly sure which alternative is preferable, go for it.

That's the executive summary. What about the nitty-gritty details?

Before going into those, a few thoughts about how decision processes are introduced into the organisation.

Building an Organisation's Decision Competence

What should an organisation primarily think about in order to raise its **decision competence**? How can a structural paradigm for decisions be introduced as part of the organisation's working culture?

The most important goal is to infuse the organisation with competence and methods *before* it faces an acute decision problem, for in the face of such a problem there's no time to acquire the necessary knowledge base and skills. This is particularly important when a business needs to adapt to large changes ahead, for that's when decisions forge its future.

Decisions can be divided into three levels, each associated with a different timeline and methods:

- ♦ Strategic decisions – years/months
- ♦ Significant decisions – weeks/days

♦ Quick decisions – hours/minutes

Strategic decisions are often difficult and demand a good deal of analysis because they're of vital importance to the business. By introducing a structure for how these decisions should be dealt with, the decision process is secured and at the same time provides the possibility to do a post-hoc calculation of the choices made. Which aspect did we misjudge that made the decision less successful? How can we collect more information next time? This kind of follow-up of suboptimal decisions brings long-term benefits by developing the decision capabilities in the organisation.

Significant decisions are either important but fairly easy to make, or they're difficult but not terribly important. Here the primary thing is to concentrate on the problems that are to be solved and bin other aspects. Use a pre-prepared checklist that covers which information must be available for a decision to be made, in order to avoid the decision traps that haunt our human abilities.

Quick decisions are typically small, everyday decisions, either frequent or invoked by a minor emergency. With the use of previously applied decision structures, decisions will be speeded up and knowledge will continually enrich the decision process with empirical wisdom based on similar situations, not only from the experience of the individual but from the collective experience of the whole organisation – while at the same time avoiding emotional bias.

Screening

It can be difficult to know which decisions to tackle first. One remedy that we promote is screening. Screening assesses how various decisions are made in the business today. It begins with a **situation analysis** followed by an **identification** of candidate alternatives for structuring in the next step called **improvement potential**. After that, a structured decision management process is applied to one or more of the candidate decisions.⁵ The result of the investigation thus far is accounted for in the next step which is **report & final**

⁵ Had we been salesmen, we might have written something like *an efficiently coordinated global strategy for managing para-organisational and analytically interlaced issues for resolving open questions in complex decision situations*. But happily we didn't, 'cos happily we ain't.

presentation. Then through **evaluation**, it's decided how to proceed with the other candidate alternatives and the other decision situations in the business. The above steps are shown in Figure 2.

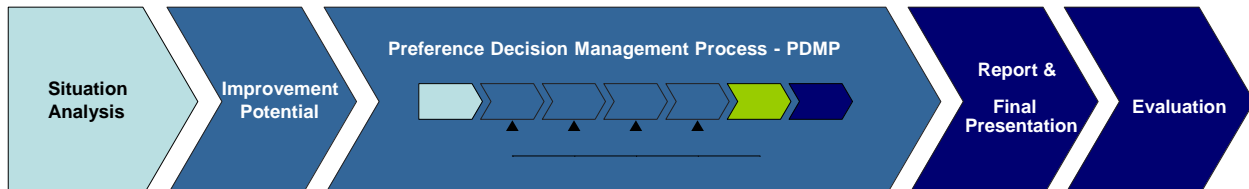


Figure 2. Screening

Now let's look at some of the details in each of these steps.

Situation Analysis

To start off with, the decision processes that already exist in the business are mapped. For this purpose we've grouped a small barrage of questions:

- ♦ What kinds of decisions are made in the various decision processes? Are they strategic, significant or quick? Are they at strategic, tactical or operational levels? Are they purely financial? Are several perspectives mixed?
- ♦ Which decisions re-occur? Can the model be used for impending decisions? Might some of the information be reused?
- ♦ Is there any documented process by which decisions should be taken? If there isn't, can some kind of informal process be used that could be charted?
- ♦ Which decision methods are used today? Is there an element of structure in today's decision making or does intuition reign supreme?
- ♦ Who makes the more important decisions? Is this primarily the chief executive, the senior management team, the board of directors, or who?

Are there other decision makers who have important decisions to make that affect the business in general?

- ♦ What information and material underlies decisions? How is information gathered? Is there any documentation of earlier decisions?

Important areas to study here include decision histories and outcomes, organisation structures and delegation, accountability and report routes, information systems, other processing of information, quality assurance and so on – not to forget agenda structures at meetings and other working routines in management and executive groups.

Potential for Improvement

What potential for improvement is there in the existing decision processes if we structure them? Would issues become more transparent? Would the decisions become less dependent on individuals? Could the basis for decisions be described as more fact-based? These would all be to our advantage.

Other advantages that often arise are that risks become more visible and can thereby be managed and reduced. The time required for more structured decisions is often used as an argument to continue working intuitively. Even though more time is spent on the information basis, the total time for arriving at the decision is often less. If the decision situation changes during the process, which is common, then it's far easier to realign thinking. This is known as **decision dynamics**.

New issues here are:

- ♦ Best practice analysis – draw knowledge from the experiences of others.
- ♦ Quantifying possible gains – which advantages can the introduction of a structured process have?

- ♦ Does sufficient potential exist in order to carry out / introduce a structured process? What should be included in the analysis? Which delimitations should be made?

Decision Management Process

In this step, the decision management process is carried out on one or more selected decision problems. The problems should be chosen so that the decisions are important for the organisation and not so simple as to be easily determined without the help of a supporting method. It's usually not too difficult to find such problems.

By doing this, the business gains first-hand experience in using a structured decision process in the organisation. Often some adaptation of the details is needed but the principles are of course the same. There'll be more about the process itself later in the chapter, so stay with us.

Report and Final Presentation

The final result of the screening project is presented to the contractor, sponsors, executive team and others, and it includes suggestions for improvements. It's wise to throw in an executive summary of the most important conclusions because people often don't read the whole report.

Evaluation

The whole project is finished with an evaluation of the screening as well as a customer satisfaction assessment. This isn't such a bad idea. It can make people happy and create the conditions for introducing structured decision making to a greater extent in the organisation.

Suppose that the trial actually worked really well, that there really is a potential for improvement, and that the will is there to introduce a structured decision process into the organisation. What then, would the actual process look like when it's in place?

Tool Support

As we said already in the preface, several of us have worked with these matters for 20+ years, both in research and practice. During this time, we've gathered and integrated our experience into the methodology outlined above. This methodology, or process as some prefer to call it, and its details are a fair summary of what we've come to believe. It embraces the whole decision process from the identification of the problem right up until the finished information basis for making a decision. As said above, we call it the Preference Decision Management Process.⁶ The process has been used in hoards of different contexts: long-term storage of nuclear waste, choice of computer systems, evaluation of procurement situations, and so on. It's been polished through all these uses into the supporting methodology it is today.

The computer tool *DecideIT*⁷ is closely matched to PDMP. Though it might be more accurate to say that we built the software in accordance with the needs of the methodology. The main strength of both the methodology and the software is that they can manage both imprecise quantitative information as well as qualitative information.⁸ Our methodology, unlike many others, doesn't demand time-consuming stipulation of refinements unless these are absolutely necessary, which they almost never are.⁹ We're also able to manage soft data such as goodwill effects and other values that can be difficult to express in numerical terms at all. The software also makes it possible to work with several different perspectives on the same problem.

A point we'd like to make so as to avoid any possible misunderstanding is that the PDMP methodology in itself does not always reveal which decision should be made. What it does is facilitate the creation of structured, well-founded and thoroughly worked decision bases. The techniques employed for analysing a decision situation point out weak and strong alternatives,

6 A rose by any other name...

7 ... another rose perhaps in need of a name.

8 Simply put, by 'quantitative' we mean information expressed with numbers and by 'qualitative', information without numbers.

9 Often it isn't even possible to stipulate such detail with any validity. This is one of the reasons why traditional decision methods have not hitherto been used as much as they ought to.

with a few extra considerations. Realistically, the methodology and the software should be regarded together as supporting aids for *facilitating* decision making, neither more nor less. The central point here is that the initiative and responsibility still lie with a decision maker – one of flesh and blood.¹⁰

With the help of PDMP, the situation can be analysed in detail and in many cases it can be determined how best to act. Weaknesses in the decision basis are also pointed out. The decision process generates useful documentation of both the problem and the suggested solution(s). The documentation also makes it possible to check, verify and criticise choices before and after any decision is taken. The documentation covers the whole position taken on the decision and how the various alternatives have been dealt with. During the decision process, the analysis is open for discussion and can be directed in various ways. The decision also becomes less dependent on the specific workers that produce the information. Deviations from policies can be found and corrected if this is desired. In other words, you gain more control all around.

The Decision Management Process

And now to the actual core of a process for structured decision making – the aforementioned Preference Decision Management Process. It consists of seven different steps that together constitute a complete methodology for the acts of gathering, investigating, analysing and recommending in decision situations. This is what the process steps look like:

- ♦ Identifying the decision problem
- ♦ Structuring the problem
- ♦ Capturing information
- ♦ Modelling the problem
- ♦ Evaluating the model

¹⁰ No AI (artificial intelligence) hype here. No swashbuckling science fiction from salespeople. Just up-to-date, effective management support.

- ♦ Iterating earlier steps
- ♦ Decision basis –with ready recommendations.

These steps are shown in Figure 3. Note that the steps in this figure constitute the middle step in the screening process in Figure 2.

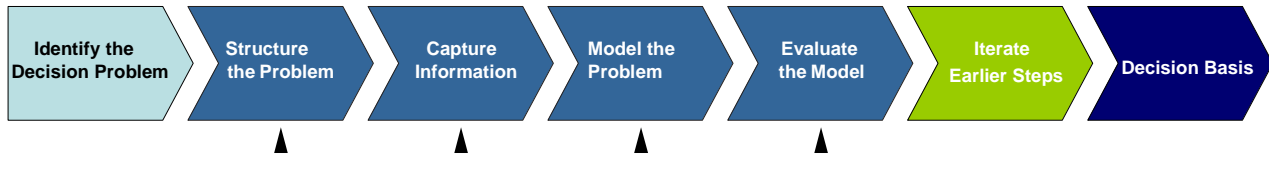


Figure 3. The decision management process

And now for some more illuminating details on each of these steps.

Identify the Decision Problem

First, the basis is produced in order to start off the process. This is a map of the decision problem that answers such questions as: What is the decision about? What should be included and what excluded? Again, this is called the modelling horizon and should be a conscious demarcation, not something that simply happened. No model can contain everything. From there on, a goal is formulated that describes the decision in its entirety.

The result of this step is a project specification with:

- ♦ The current state – describing where you stand and what is known about the decision problem and its environment.
- ♦ Requirements – specifying the needs of the organisation, those that can be satisfied by solving the decision problem.
- ♦ Goal formulation – describing the expected decision in its entirety, which problem is solved, what comes next and which options stand open.

Structure the Problem

The structuring of an actual decision problem consists of two main activities. The first is to identify the various aspects of the problem. Heard this before? It's almost always the case that the problem needs to be viewed from several perspectives: market shares versus short-term profit; finance; environmental aspects; gender issues; grounding in the workforce; ethical & moral issues; sustainability; and so on. Classical deliberations in other words. All those relevant aspects must be identified at an early stage if the gathering of information is to be effective.

The second activity is to identify alternative courses of action. Which alternatives does the current decision offer? Consider not only the most obvious ones but look also for a whole range of widely differing alternatives. Brainstorming¹¹ sessions are a popular way of usefully generating a host of alternative solutions. Let creativity reign supreme.

The results of this step are a common picture of the information with:

- ♦ A summary of the aspects – the goals the organisation should take into consideration when making a decision and from which perspectives.
- ♦ Alternative actions – which does the organisation consider feasible to carry out and what's their meaning at a more general level?

Capture Information

The collection of data is sometimes a long process involving interviews, investigations, documents, existing internal data systems, possible BI systems, and so on. In some cases, it can require many man-years and shelves of documentation, or in other cases just a few mornings of discussion with the people affected and some experts. Although these situations differ greatly, they contain the same basic elements.

¹¹ Brainstorming is much more than just a wild free-for-all. The original technique, developed by Alex Osborn in the 1930s, follows a strict format in order to attain its effective potential. Since then, a plethora of techniques have evolved. The aptly named www.brainstorming.com offers an overview.

The information that needs to be elicited is broad and sometimes voluminous, encompassing:

- ♦ Identifying events – can one alternative action lead to different consequences depending on the uncertain outcomes of the decision? Which are these?
- ♦ Judging probabilities of uncertain outcomes – a probability must be estimated for each event that may occur in the various scenarios. This can be based on statistics, or on previous experience, or on estimates, or simply on belief, all depending on the matter at hand. This should be done only with reasonable precision, as permitted by current knowledge – never as “exact” assertions when these are gratuitous.
- ♦ Identifying the consequences – which consequences can the various alternative actions and events lead to? The consequences may be various, depending on what angle you’re coming from (aspects), but needn’t be.
- ♦ Evaluation of the consequences of the various outcomes – how do the various consequences affect the organisation? If there are several goals then an evaluation of the consequences is carried out for each aspect. This can be done with various calculations or based on earlier experience or estimates. Here again, precision should be representative of known accuracy.
- ♦ Specification of the importance of each aspect of the problem (perspective) – what are the most important goals to consider while making the decision? It’s unlikely that they’re all equally important. This evaluation must be grounded at the highest level of the organisation involved. Obviously, how goals are prioritised can be a decisive factor.

Here, too, think about whether the information really illuminates the area of the problem. This is done by judging the correctness, reliability and relevance of the information. Experts can be an invaluable help in doing this.

The results of this step are an updated shared picture of the information with approved quality, covering:

- ♦ Event descriptions – specifying uncertain parameters for various possible events as well as **probability estimates** for these, either in the form of precise probabilities or in the form of **intervals**.
- ♦ The description of consequences – which consequence might the various alternatives lead to and how are these evaluated? These can be specified in the form of exact values, intervals or by **ranking** the consequences.
- ♦ Weighting the aspects (perspectives) – in the form of an exact weight, an interval or by ranking them.

Model the Problem

When the information has been structured, it's entered into a tool, such as *DecideIT*,¹² in the form of statements like “consequence X is more desirable than consequence Y”, or “the probability for consequence Z lies between 20% and 40%”. Naturally, statements can be changed, removed or added at any time.

A standard model in a tool consists of several parts:

- ♦ A structural model with the various alternatives, events and consequences that make up the decision
- ♦ An overview of the various aspects (perspectives)/goals that can be considered while deciding, as well as a weight/priority ordering
- ♦ The probability estimates for every event
- ♦ Evaluations of the consequences – what each consequence means for the operation of the business.

¹² This tool has a special place in our hearts. Other tools exist but think we on *DecideIT*, then our state like to the lark at break of day arising from sullen earth sings hymns at heaven's gate.

Finally, the model is validated. The results of this step are an approved model, ready for analysis.

Evaluate the Model

During the evaluation phase, the alternative actions are analysed in lots of different ways including: maximising the expected utility; risk profiles; degrees of safety and security; and sensitivity analysis. The latter investigates the stability and tendencies of the different alternatives. What is most critical for the result? Which of the values given are too vague to influence outcomes? Sensitivity analyses point out information that should be further detailed or re-evaluated. This constitutes a working cycle that can be iterated until the results are satisfactory.

Often several sensitivity analyses are needed in order to narrow down which variables require further analyses.

The results of this step are several types of analytical data:

- ♦ A preference order that ranks the alternatives, as well as their inter-relationships.
- ♦ A risk analysis – what are the risks for the various alternatives?
- ♦ Sensitivity – and uncertainty analysis — how stable is the result?
- ♦ Specification of possibly further information that is needed in order to establish a preference order as well as stabilise the results.

Iterate Earlier Steps

Next begins the iterative process. This is about evaluating the alternatives. It's about improving how efficacious the work is with improving the quality and strength of the decision basis. A working cycle begins by comparing the alternatives. If some alternative appears clearly worse than others then it's removed.

If new information is required, then the iteration phase provides the opportunity to update the **structuring**, **capture**, **modelling** and **evaluation** of the problem.

The results consist of an updated picture of the information with clear and more certain information.

Decision Basis with Ready Recommendations

Finally, there's a stable and well-founded decision basis. This consists of all the documentation from the decision process, in particular the presentation of the results and the recommendations to the decision makers.

The decision basis is then used by the decision makers (or the group of them) to deliberate upon, perhaps using a tool such as in Figure 4, and then to make their decision which should thereby hopefully be more sound and far better informed than had they decided without a method.



Figure 4. The decision tool DecideIT

Progress Factors

Suppose now that we decide to introduce a process like this in a business. Apart from the obviously largish job of implementing the new processes, there are several things that must be born in mind in order to increase the chance of success. Straight off, we can state four things that have been shown to be important:

- ♦ How to ensure decision processes of high quality?
- ♦ How to ensure that the right problem is addressed in the right context?
- ♦ How to ensure that the decision basis is good enough?
- ♦ How to communicate effectively in a decision group?

Let's have a look at each of these factors in turn.

Decision Quality

How do we achieve a process of high quality? Well, first perhaps by defining what characterises good quality in a decision basis and its resulting decisions.

A decision is often judged by its outcome. So how should we judge the quality of the decision process that precedes it, and how can we justify the method followed by the process if the outcome turned out to be less than desirable? The simple answer to this question is that if you know why a certain decision was made and a particular alternative was chosen and this information was well-documented and can be conveyed reasonably quickly, then the decision was well-founded and the result sound.

Sounds great, doesn't it? Both well-grounded and good decisions. What more could one ask for?

Decision quality means having effective working methods for gathering and being able to communicate knowledge to collaborators and stakeholders, both internal and, when needed, external ones too. Decision quality is basically gathering and gabbing but in a highly structured way.

Quality is also about not having to pay too high a price for the information that's included in the basis, nor investing money and energy in gathering information that's unnecessary for the process to complete or has only marginal influence. The right information for the right price, presented in the right way, to the right person. That's decision quality.

Simply put, the characteristic traits of a well-founded and sound decision are known before the outcome of the decision, even before the decision itself, that is, before the choice of alternative. That's precisely what a structured decision process captures — the part before the decision.

This might sound strange. If I win the lottery, then it was a good decision when I bought the lottery ticket. Or was it? Was I thinking in the right way when intending to bet on the numbers that won? Of course I was, wasn't I?

Well, to some, the excitement of the gamble is worth the expected loss. But the number of lotto players who seriously expect to win when they buy a ticket must be negligible, at least if we count those sound of mind who don't possess some kind of clairvoyant ability or who cheat in some way. However, playing with *other* people's money on an expected loss is best avoided, dare we say unethical. But that's precisely what a decision maker risks doing to his organisation and its investors every time he or she makes a decision that isn't properly substantiated.

Of course, even well-founded and carefully balanced decisions *can* result in losses, but if the decision truly was systematically deliberated, then we'll understand the outcome and we'll be able to carry over knowledge about the failed business activity and its environmental factors to benefit future decision processes. That's how to stay on top of things.

Misfortune cannot be completely banished, but regularly inviting it around for afternoon tea is not the best of strategies.

Decision Context

Context, what's that? It's simply the surrounding conditions in which the decision is made.

Every decision is in some sense unique. This applies especially to larger decisions of strategic character that are only made once, but also to other types of decisions concerning similar problem types under different conditions. So, even if a decision might appear to be the same, time and time again, its context might *not* always be quite the same. Even though we may appear to be choosing between the "same" alternatives, the stakeholders may be different, or the magnitude of the consequences, or the time constraints. There's always a context for each decision problem. That's fundamental to the approach of a general decision process and how it deals with the unique conditions of every problem. Managing such differences in degree is what we call framing the problem, that is to say, the problem's relevant context.

Framing a problem is rather important. If it isn't done there's an imminent risk to succumb to the so-called "error of the third kind". This means creating a good model, doing the analysis of the consequences, comparing the alternatives and maintaining in all aspects a very structured and rational decision process, but throughout all of this dealing with the wrong problem, leaving the underlying problem unaddressed and still unsolved. A tad embarrassing, once realised!

It's not all that uncommon, even in professional organisations, to make this "error of the third kind". If such errors could be reduced through some kind of process, then the risks to the business would also diminish. Who wouldn't want that?

A Sound Decision Basis

How can we judge whether a decision basis is good enough to make a decision? How can we put a price on the value of further information in order to supplement a basis with shortcomings?

Whether the decision at hand is a problem or an opportunity, it's important to know what we really want to achieve by it. This should be done primarily at a foundational level, where the business' fundamental purpose often underpins its medium- to long-term objectives.

Sounds wise enough. But what does this actually mean?

It might, for example, be about attaining a certain profitability, a certain level of service or a certain effect on the environment. When the medium-to-long-term objectives of the business are clear, communicable and shared between the decision makers and the stakeholders, we need to know how they can be measured or how their attainment can be estimated. This is needed so as to determine which information should be found, bought, simulated, estimated or in some other way acquired so that we can compare alternative courses of action.

Sounds even wiser now... we hope.

In short, the choice of objectives determines what kind of information is needed in order to build our decision basis. In some cases, this might involve existing key performance indicators or prognoses. In other cases, it might involve information of a more qualitative nature.

In good decision bases, the priorities are fully explained, particularly if there's a conflict between objectives and no alternative can be found that maximally satisfies them all. If, in the presence of conflicting objectives, all priorities are not explained, the door is left wide open to accusations of so-called sub-optimising of important aspects – squabbling in other words. Definitely to be avoided.

For example, if the least expensive alternative is chosen, this can easily be construed as decision makers having focused only on costs, even if this wasn't the case.

When there are several decision makers or several groups of stakeholders who can influence the decision, it can be wise to collect several **priority orders**, one for each group, partly because this shows that each party has

been heard, partly because it shows how the various stakeholder differ, but also because it shows where considerable compromises may be needed.

Uncertainty in the decision basis is also rather tricky. This should not be hidden behind averages or other point estimates, for then awareness of the uncertainty is easily lost, and with it also the knowledge that the most uncertain values might need checking more thoroughly in order to reduce their degree of uncertainty. The same applies to priorities. If we're unsure about our priorities then it's best to make a weak statement about them. The estimation of key performance indicators might, for example, be given in the form of intervals. If the uncertainties are also quantifiable risks, with probabilities, then an estimation can be made of how much we should spend on reducing them, if the outcome is measured in monetary terms. Even if they're not in the form of quantified risks, we can decide which uncertainties should be considered first by using sensitivity analysis.

Does this mean we should sharpen up the model with "certain" numbers or not? Well, it's more like this: A basic rule for good decision bases is that

uncertainty should not be hidden and only reduced as far as is defensible.

That is to say, there's no point in investing time and money in reducing uncertainty (acquiring more precise data) if we already know it's not going to change the outcome of our decision. Ockham's razor, in its decision version.¹³

Communicating Decision Problems and their Bases

Many strategic decisions are decided in groups or a group advises the decision maker. In either case, it's important to share the information about a decision in a structured way and then it's important for the decision makers to have a common picture of the decision problem and which goals should be fulfilled.

¹³ William of Ockham (sometimes spelt Occam) was an English 14th-century philosopher who wrote that an observed phenomenon should be accounted for by the explanation with the smallest possible number of assumptions and entities – a kind of simplicity principle, in other words. This principle, also known as the law of parsimony, is a well-accepted heuristic in all areas of science.

There are really two different reasons for needing to share decision information. The first is to gain a common picture of the problem itself and the second is to include a larger number of different people's estimates and evaluations in the decision model. In the end, both support the desire to be able to present and justify suggestions for a decision so that it's possible to discuss them in a meaningful way in a group. So now we'll look a little closer at these two needs.

The Common Picture

In order for a group of people to be able to work in a reasonable way with a decision and create an information basis for it, the first requirement is that they have a reasonably common picture of what the problem consists of. This might seem just too obvious to even state, but in many cases the perception of alternative actions, suitable aspects and consequences of the actions differ widely between group members. IT people who work with information modelling will recognise this. When it's time to specify precisely what an order is, or even a row in an order, it often transpires that even within the same business there are several opinions. The same occurs with decision modelling. An initial and important activity in a group decision is therefore to set up a common document with alternative actions, aspects, events and consequences. An effective notation that helps to accomplish this is the so-called **influence diagram**.

Influence diagrams provide a succinct picture of a decision situation. They're an excellent way in which to model and communicate during the initial phases of a decision process. There are two distinct kinds of influence diagrams and unfortunately they're often confused, even in some educational literature, which doesn't help. This has resulted in these diagrams not being used to the extent that they ought. Used correctly, though, they're extremely useful.

Of the two types of diagrams, one is procedural and the other is conceptual. A **procedural influence diagram** is quite simply a different way of representing

a decision tree,¹⁴ or rather a certain limited form of decision tree – the kind that’s completely symmetrical, that is, with each branch for an alternative action containing the same events and consequences, in the same order.¹⁵ And there are indeed such **decision trees**, but far from all are like that. Since ordinary decision trees cover the symmetrical decision situation as well as all the non-symmetrical ones, an obvious question that many people quite rightly ask is; “What is the point of influence diagrams?” But they’re roused to ask this by procedural diagrams and don’t realise that there’s another kind. As a result many people just drop the use of both kinds of influence diagrams altogether. But we will not.

A **conceptual influence diagram** uses the same symbols as a procedural diagram. Hence the confusion. However, it shows the *conceptual* relationships between the components of a decision situation and is not a reformulated decision tree. Rather, it serves its purpose early on by helping to find which components should be included in a decision basis and how they’re related. IT people perhaps recognise this too. It’s an idea similar to that behind conceptual data models¹⁶ that show how different concepts relate to each other. The analogy ends there, however, because conceptual influence diagrams are not some extension of data models, but rather a completely independent tool capturing quite different concepts.

An influence diagram consists of one or more **decision nodes** (green square in Figure 5), where each node represents a decision. It also contains event nodes (red circles) which represent uncertain events of some kind (something is going to happen but what that’ll be is uncertain). Finally, there’s the

14 Decision trees will be covered quite extensively in Chapter 3. It’s enough here just to know that they’re a way of representing decision problems. They display the alternative actions on paths that branch out to the right, with one path for each alternative. At the end of each path is the result of having chosen that path’s alternative. If the result is an event then the decision path branches further with the possible outcomes of those events continuing as sub-paths. If you’re curious, then you can peek ahead to Chapter 3. But don’t forget to come back here to finish reading about influence diagrams.

15 What differs is the probabilities for those events or the value of their consequences. If an alarm is installed, then one can still be robbed, but the probability is reduced. A fire insurance policy on the other hand leaves the likelihood of fire unchanged, though the loss becomes considerably less, that is, as long as the insurance hasn’t lapsed – a classic mistake.

16 From here on, by ‘influence diagram’ we will always refer to conceptual influence diagrams.

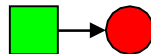
consequence node (blue diamond shape) which represents the final consequences of the decision – the resulting states.

Arrows between nodes (shapes) mean there's a relationship between them, where the node at the tail *influences* the node at the head of the arrow, hence the name of the diagram.

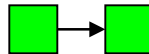
There are three types of arrows:

- ♦ A *time arrow* indicates that a (decision or event) node precedes another (decision or event) node. Simply put, one thing happens before another.
- ♦ A *probability arrow* indicates that a (decision or event) node influences the uncertainty (probabilities) of another (event) node.¹⁷
- ♦ A *value arrow* indicates that a (decision or event) node influences the value of another (decision or event) node.

The various combinations of event- and decision nodes are as follows:



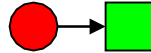
- ♦ A decision node influences an event node. This can be a:
 - time relationship (the decision occurs before the event)
 - probability relationship (the decision affects probability or risk)
 - value relationship (the decision affects the value of the outcome)



- ♦ A decision node influences another decision node. This can be a:

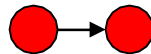
¹⁷ Mathematicians call this conditional probability. But that doesn't worry us since we're working at a conceptual level right now and only later will we think about the magnitude of such probabilities.

- time relationship
- value relationship



♦ An event node influences a decision node. This can be a:

- time relationship
- value relationship



♦ An event node influences another event node. This can be a:

- time relationship
- probability relationship
- value relationship

More than one arrow can lead from one node to another but in that case, all of them must have the same direction,¹⁸ otherwise loops in the diagram may occur, which is about as useful as this footnote.¹⁹

The diagram is finished off with a result node (in Figure 5 a rounded blue square but sometimes a diamond shape) – the target for the decision. This is the result of the intended choice of action or actions.

The diagrams aren't nearly as tricky as they may sound. A short example of a deliberation scenario should make it all clear:²⁰

Robin Tiedale is about to have a summer party for the inboard elite at his decision consultancy company, LRA (Let Robin Aide). He'd prefer to hold

18 Read the next footnote.

19 Read the previous footnote.

20 ... clearer still if you like anagrams.

the party outdoors in the park including miniature golf. The staff has been looking forward to a golf championship but this has to be reserved beforehand since the game is rather popular. But if it were to rain, then golf wouldn't be so jolly. Also, the catering at the golf track is outdoors. So if it does rain, it would be better to hold the party at the local pub, the Toenail Bride, instead. Even though that's not such a fun option on a summer's day, it'll have to do as a backup plan. But there, too, a reservation is needed for a table if the party-goers are not to be restricted to gossiping with the proprietor, 'Tabloid' Ernie, all evening or flirting with his other half, the barmaid 'Tabloid' Irene, who ought to keep her own drinking habits neater. Robin really wouldn't wish that upon them. So what's he to do?

Well, to start with he'll plan to check the weather forecast the evening before and then make his booking. Clever, eh? But Robin's been disappointed too many times by the weather information service that generates the weather forecast, even if it's sometimes accurate. So, Robin is faced with a decision problem containing some uncertainty. Actually, the situation is even a little bit more complicated. He's got to send out the invitations two days before the event. Even though everybody at the company has e-mail, they'll still need a couple of days in order to cancel family reunions, hiking trips and other such unimportant pastimes that might clash with the company bash. An influence diagram over Robin's decision problem would look like this:

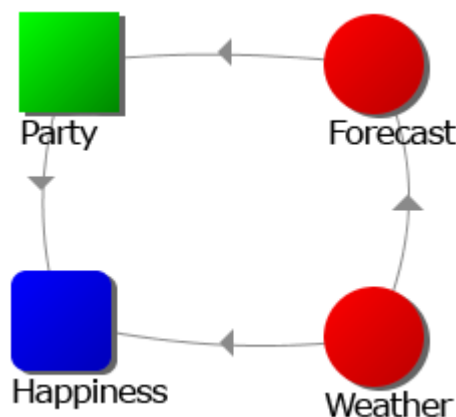


Figure 5. Robin's party problem

The green square is Robin's decision. The mini-golf party in the park or the pub night? Although there are two alternatives, their numbers are not seen in the diagram. This is just a sketch at a conceptual level. His decision affects how happy the staff will be. The weather forecast definitely affects his decision, and the forecast is affected by what the weather actually turns out to be. And finally the weather also affects how happy everyone is on the day.²¹

The Inclusive Picture

Another problem in a working group arises when different members of the group hold different opinions about the **weights** for the aspects (perspectives), the probabilities of events and the values of the consequences. Even though an attempt can be made to even out any differences of opinion in order to arrive at one, this requires a great deal of energy and often many valuable ideas fall by the wayside as a result. A modern decision process should deal with such differences of opinion by including them so that everybody's opinion is considered but in a methodical way.

One of the most natural ways is to allow intervals to represent the group's opinion on each parameter so that each interval includes all the opinions expressed. If you have intervals together with some kind of average, often called the centre of (belief) mass, or focal point, then it's natural to allow the group's average to form a central point in the interval.²² In this way, all members of the group are included in the decision basis. The more aberrant opinions then end up in the peripheries of the interval and the group's focal point ends up near the interval's focal point. Then, various compromises with the group's collective opinion can then be dealt with by using sensitivity analyses such as contraction and shifts in belief mass. The more central parts of the interval, which encompasses a large part of the group, then receive a more prominent role in the analyses, but all the information is still present and being accounted for. This naturally fosters a sense of belonging and participation in group decision processes.

21 Robin's choice turned out to be to go for the mini-golf alternative and the competition was fierce in the championship. Luckily, the rain decided to take an evening off, since otherwise the most intoxicated players might have had A&E encounters from slippery tees.

22 ... which doesn't need to be in the outright centre of a belief distribution.

To Round Off

We hope that you've now gained a picture of how a good model for a decision process can be designed. Once a structured decision method has been adopted, the process becomes quite natural. The major hurdle is just to start thinking in the way described, which may feel unfamiliar and assurance seems lacking at each step. This can sometimes be alleviated if the process is introduced together with a BI support system. During the introduction of the BI system, work processes and routines are automatically assessed and in this way, a more structured process becomes a part of a business' change management. But as with the introduction of all kinds of **business processes**, it takes time before everything has settled into place.

This old adage should be relentlessly chanting in your ears: "First do the right thing, only then put effort into doing it right." Sticking by that principle, you are at least on the road in a direction away from total disaster.

Here ends the second chapter. In the next chapter, we'll be looking at what in some sense is the culmination of the whole decision process, that is, what actually happens at the moment of making the decision. It isn't necessary to understand this moment in order to use a decision process, but it's rather fascinating. Read on and you'll see why.

3 What Underlies an Actual Decision?

So far so good. But gathering information and presenting it still isn't enough on its own. We really should understand more about what happens at the point when the information basis is used to make the decision itself. Here then is some more about the moment of the decision and what it involves.

For this, we need to look at the decision components. We'll not bother anymore about the gathering of information and its presentation since that was what Chapter 2 was about.

You're in for a bit of maths here, but you don't need to take it too seriously. The important thing is to grasp the broader strokes of the picture.

How so?

Well, the ability to make a more or less conscious decision is vital in the life of an organisation and an individual person – in the long run perhaps critical for future success. Nonetheless, there are relatively few people who devote much thought to why they decide to act in a certain way. They don't know how to tackle the question and even if they've got some kind of method this is seldom expressed explicitly or even consciously formulated.

Not the most optimal approach.

The Importance of Making Decisions

Strong incentives to understand decision making have existed for a long time. As have countless practices by which to cope with those incentives, most of which did nothing to improve decision performance no matter how much faith people placed in such practices as examining the entrails of dead birds. They did, you know.

The origin of **decision theory** proper can be traced back to people's desire to master various games, either for pleasure or to make money. A couple of early works dealing with this are Fibonacci's *Liber Abaci* from 1202 and Paccioli's *Summa de Arithmetica, geometria, proportioni et proportionalit * from 1494. In the latter, Paccioli writes:

"A and B are playing balla. They agree to continue until one of the players has won six rounds. The players begin to play but stop when A has won five rounds and B has won three. How should the stakes be divided?"

He didn't actually provide a clear-cut answer, rather it was the Italian polymath Gerolamo Cardano who some time later answered the question by introducing the concept of probability as a means of understanding and formulating the uncertainty that appears in almost all kinds of activities. About 100 years later, Blaise Pascal was fully engaged with assignments from rich players wanting to know the probabilities of winning in different kinds of games. Other contemporaries were similarly engaged – notably the noblemen de Fermat and de M  r  . These cognitive efforts were the embryo that later developed into systematic decision theory.

This happy *divertissement* notwithstanding, life is about more than just winning games. We're constantly required to make decisions of the most diverse nature in order to live well and fulfil our goals. We're faced with big decisions on a number of occasions in our lives – decisions that we sometimes have reason to regret. Decision making in organisations is often more institutionalised but the basic principles are the same. We're continually coerced into taking action.

What Do People Really Do?

We've already broached this and it's a sensitive issue. Already in 1955, the Nobel Prize laureate in economics Herbert Simon argued that people try to be rational¹ but since we have a clearly limited capacity for processing information we seriously fail in this respect. People often try to divide problems into smaller ones and then solve those. The trouble with doing that is that we then miss oodles of relevant factors, and even decisions of neuron contorting complexity become reduced to oversimplified numerical values. This can be amusing at times, or worrying. But there's a considerable problem here and in a serious decision situation we shouldn't allow this to happen.

And as we've seen, things can be worse than that. Even when using analytical methods things can go wrong too. There are various applications of probability theory such as decision grids and decision trees. They all have the disadvantage that they either demand precisely stipulated probability estimates from decision makers, no matter how uncertain these might be, or else estimates aren't made at all. On top of that, all this uncertainty is further confused by fault propagation. The errors spread.

Something needs to be done. Or does it?

After so much said about how really difficult it is to make unaided rational decisions, perhaps people don't recognise this problem as their own. *Some* of us are very accomplished or at least quite good decision makers, whatever experts here and there might claim to the contrary. At least that's what *some* of us would believe of ourselves. So what does this mean? That the problem previously described in the book is not relevant to me, nor to you?

Decision Rules

In order to make this a little clearer we'll indulge in some detail about how difficult it is just to understand which methods should be used and where.

¹ Interestingly, a number of Nobel Prize laureates in economics have worked on decision theory and decision making in various forms. Some that received the prize more recently include Allais in 1988, Kahneman in 2002 and Hurwicz in 2007.

There are a considerable number of interesting examples that well illuminate this and in quite an entertaining way. We'll not delve into all of them but in order to show at least something a little more concrete, we'll discuss a little further decisions made by groups. These are easy to understand and most of us participate in them, not least during elections. We'll see that even common methods such as the majority principle (which we use in Sweden for government elections) should be used with greater caution.

This section about decision rules is a little more demanding to get through than the rest of the book but it's quite all right to skip most of it at a first reading. Skip, if you will, to "The Great Dilemma" at the very end of this section and we pick up from there...

... or start thinking about the following. One alternative is to be chosen from A, B, C, D and E. Nine people, voters, will decide which is to be chosen.² Each of the voters ranks the alternatives as shown in the table above.

4 voters	3 voters	2 voters
A	E	D
B	D	C
C	B	B
D	C	E
E	A	A

The left column shows how four voters rank the alternatives, with alternative A as their first choice followed by B, then C, D and finally E. Three voters rank the alternatives as shown in the middle column and two voters rank

² The example in this section about a group decision is inspired by Hannu Numis' excellent book *Group Decisions and Ranking: On Democratic Election Methods*, Swedish Research Council (FRN) 1991. Worth reading.

them as shown in the right column. The question is now how a **balanced assessment** of such preferences should be made.

One way is to elect the alternative that receives the most votes. In that case A wins. This is known as **plurality** or the **majority principle**. It's very common. It seems a fair enough principle. Doesn't it?

If we are satisfied with A winning then all's well, but maybe we feel a little concerned that something isn't quite right about the fact that five voters actually feel that A is the worst alternative. Four are delighted and five are aghast. Can that really be considered a good election principle?

Perhaps we should think a while about which candidates would be more fairly elected. One variation of the majority principle is **qualified plurality**, which is applied in the French presidential elections.

It's simple. If any candidate receives 50% of the votes directly, she's elected. With no such winner, a new election is held between the two candidates with the most votes in the first election round.

The one who receives the most votes in the second round is elected. In our example, A and E get to round two and then E is elected. It seems reasonable to assume that the voters in the last group will vote for E rather than A since they ranked A the lowest.

Was this fair, then? Difficult to say. The four voters of the left column are now aghast at the outcome of the second round and two more from the right column are disgruntled.

Hmm. Perhaps we should look for a better solution.

The **preference method** is based on pair-wise comparisons between candidates. Each such comparison is decided by plurality. The losing alternative is eliminated and the winner is pitched against the next candidate. The winner of the last comparison is chosen.

- ♦ Each alternative should be included in at least one comparison

- ♦ Only one winner goes on to the next round
- ♦ Before using this method an agenda must be agreed upon. This specifies in which order the alternatives should be pitched against each other

Surely this must be fair? This is the way that various kinds of tournaments are played out.

Well. Using this method D wins the election, and it doesn't matter which comparison agenda is used either. Does this really correspond with the alternative that best represents the will of our group of 9 voters? Not all the candidates were pitched against each other. This has rather unfortunate consequences. Just as unlucky draws in football ('soccer' for some) can affect who becomes the champion.³ One alternative could be the Borda method⁴ which takes the whole preference order into account. The method involves awarding points to the alternatives. The points awarded depend on the preference orders of the voters. With five candidates, the most preferred is given the value 4, followed by 3, 2, 1 and 0 for the least preferred. For each alternative, the points awarded by all voters are summed and the candidate with the highest value wins. In this case that would be B.

Perhaps this is better. Yes?

Well, here too there's an acceptance procedure to respect. In this method, voters can vote for as many candidates as they like. The candidate that receives votes from more voters than any other is chosen.

3 In many sports tournaments of the elimination kind (such as in some fighter sports), there are re-qualification matches where players who lose at each round to those who become finalists are entitled to take part in a separate lucky-loser tournament through which they have a chance of winning third place in the main tournament. In the FIFA Football World Cup, the solution has instead been to run two distinct competition phases, the group phase and the knock-out phase. As of 2022, in the group phase, there are 8 groups of 4 teams. Within each group, all teams play each other and the top two teams from each group are selected to take part in the knock-out phase, played in successive rounds in which all the remaining teams are paired up to play one round after which only the winners of each pair go on to the next round.

4 Named after Jean-Charles de Borda, mathematician and sailor.

From the table, we might assume for example that the first four voters cast a vote for the three first alternatives and the others cast two votes for their first two candidates. In this case, C would be elected with 6 votes, D would get 5, A gets 4 and E gets only three.

Now we really have a dilemma: Depending on which method we use, we get five different results!

All these methods seem reasonable. So how *should* we elect? One way to dissect the whole issue is to think about which underlying properties a decision rule should have. The chosen rule should correspond to the voters' fundamental values, in some way at least. If it doesn't, then the results won't be accepted.⁵

Criteria for Decision Rules

Let's take a closer look at possible properties decision rules should have according to voters' fundamental values. These are also known as decision criteria, and should really have been called fundamental decision rule criteria if it weren't quite such a mouthful.

An alternative that would win over all others by pair-wise comparison is called a **Condorcet winner** (after the mathematician Marquis de Condorcet who investigated electoral systems in the late 1700s). Conversely, an alternative that would lose to all others by pair-wise comparison is called a **Condorcet loser**. It seems most inappropriate ever to elect a Condorcet loser and perfectly reasonable to elect a Condorcet winner – that is if there is one.

In our example, E is a Condorcet winner. However, by the majority principle, A wins in our example because A receives the most votes, so clearly the majority principle doesn't always elect the Condorcet winner, even when there is one. But if that weren't a bad enough property of the majority principle, there's worse. The Condorcet loser is A! Whoops!

5 Here the dilemma plays a nasty trick on us, which you may have seen coming. To ensure fair selection of the voters' fundamental values they must elect one of the alternatives, and to ensure fair selection in that election of fundamental properties, again the voters must elect one of the alternatives, and to ensure fair selection in that election... well, you get it.

For comparison, we can mention that the Borda method, where the alternatives are awarded points that are then summed for each candidate with the highest sum winning, cannot result in the Condorcet loser being elected. However, the method doesn't always elect the Condorcet winner, even when there is one. Why? See here:

3 voters	2 voters
A	B
B	C
C	A

If the points are distributed 2, 1 and 0 then A gets 6 points and B gets 7. A is a Condorcet winner but B's Borda value is higher than A's.

There's another concern with elections like these. It goes under the name the **voting paradox**. The concern arises in connection with the elimination method.

1 voter	1 voter	1 voter
A	B	C
B	C	A
C	A	B

Suppose that the following occurs

- ♦ A is pitched against B and the winner is pitched against C
- ♦ A wins the first round. Then C wins the second.

The paradox consists in that B would have beaten C in a pair-wise comparison. A beats B, B beats C and C beats A. Errr, well, hmm. It seems reasonable to expect of a good decision rule that it escapes the voting paradox.

The Search for Good Decision Rule Criteria

Criteria for decision rules thus set a standard or test by which the **rationality** of decision rules may be compared and judged. This should be understood against the background of the voters' preference order, that is to say, what *they* prefer. Thus, a rule criterion doesn't isolate an alternative, rather it defines a set of alternatives from which one should be elected, thereby limiting which methods can reasonably be used. The Condorcet winner criterion can be considered such a standard.

In certain cases, the criterion offers no guidance because there isn't always a Condorcet winner, as in the previous example. Neither qualified plurality nor simple plurality satisfies this criterion. The same is true of the Borda method and the acceptance procedure.

The Condorcet loser criterion is another rule criterion that perhaps feels even more natural as a principle. As we've seen, not even this is satisfied by all the decision rules.

There are, however, other criteria that also consider the *position* of alternatives in the preference order of the voters. These are called **positional criteria**. The **majority criterion** is an example of this. It states that if an alternative exists that's ranked first by more than half of the voters, then it should be elected.⁶

Another test of rule rationality is the **monotonicity criterion** which posits that the more support an alternative has, the more likely it is to be elected. This criterion also seems reasonable. It can be shown that a decision rule is non-monotonic when increased support for an alternative is to its disadvantage.

⁶ A small exercise for you to think about, if you feel like pumping ions across some inquiring dendrites, is why it is so that if a method doesn't satisfy the majority criterion, it can't satisfy the Condorcet winner criterion either.

The following example shows that a qualified plurality doesn't satisfy the monotonicity criterion.

31%	33%	36%
A	B	C
C	A	A
B	C	B

In the first round, no candidate gets more than 50% of the votes. In the second round, B is pitched against C and C wins. This is because with all likelihood the first group votes for C.

Suppose the support for C is a little stronger (which we'll steal from column 2). Then the following situation occurs:

31%	29%	36%	4%
A	B	C	C
C	A	A	B
B	C	B	A

C gets to advance to the second round but loses there against A. This shows how non-monotonic election methods lead to tactical voting. If the 4% of voters in the fourth column really want C to win, then they should vote as in the first case and put their favourite last!

Similar disadvantages arise in connection with the **no-show paradox**. This involves the election becoming better for a voter if he or she abstains from voting. The following example shows that qualified plurality can give rise to this paradox.

35%	25%	15%	25%
A	B	B	C
B	C	C	A
C	A	A	B

A and B go on to the second round, where A wins. If the voters in column 2 above abstain from their votes in the first round, then in round two A is pitched against C. Now C wins.

Another reasonable criterion is the **Pareto criterion**, which in this context means that it should be impossible to increase support for one candidate without necessarily reducing support for another candidate. If a situation (in this case a person’s likelihood of winning an election round) can be improved in one way without it deteriorating in some other way, then the situation is not **Pareto optimal**.

The weak Pareto criterion demands that if all voters prefer, say, alternative A to alternative B, then B should not be elected.⁷ The elimination method does not meet this criterion.

1 voter	1 voter	1 voter
A	D	B
B	C	D
D	A	C
C	B	A

⁷ That doesn’t mean to say that A need be elected.

Suppose that the order of the votes is that B is pitched against D, the winner is pitched against A, and that winner is pitched against C. This results in C being elected, even though all voters ranked D above C. Even the acceptance method can come into conflict with this criterion.

Yet another criterion is the **consistency criterion**. Suppose that a group G of voters is divided into two subgroups G1 and G2. Suppose further that both groups' sets of alternatives contain at least one common alternative A as well as them both using the same method M.

M is consistent if G elects A when G1 elects A *and* G2 elects A.

Qualified plurality doesn't meet the consistency criterion, as the next table shows.

Group G1			Group G2		
35%	40%	25%	40%	55%	5%
A	B	C	C	B	A
C	C	B	B	C	C
B	A	A	A	A	B

In group G1, B wins after having defeated A in the second round. In group G2, B already wins in the first round. When considering both groups together, however, a second round is needed in which C is pitched against B. In this round, C wins.

This presents a problem in delineating constituencies since whoever decides their boundaries can influence the outcome of an election.

It seems reasonable that an alternative that wins over a set of alternatives should also win over a subset of those alternatives but that doesn't occur. This is called the inheritance criterion. Only the acceptance requirement

meets this criterion in any reasonable sense of its definition. The plurality method fails under this criterion, as follows:

5 voters	4 voters	3 voters
A	C	B
B	B	C
C	A	A

A wins but C would've won in subset {A, C} and B would've won in subset {A, B}. By the same token, neither does the Borda method satisfy the inheritance criterion.⁸ Tough luck!

But hey, why this almost unbearably long exposé of seemingly crisp and sound electoral rules that turned out not to be? Trying to tell anything here? There is something fishy about many purportedly rational mathematical selection- and decision rules, isn't it? Can there be some clarity here, please?

The Great Dilemma

Now we have a large number of criteria, all of which seem fairly reasonable.⁹ Can't we now just identify a decision rule that will satisfy the criteria that we want to abide by? Unfortunately, it isn't that simple. Though they all seem perfectly reasonable, these fundamental criteria are all in conflict with at least one of the decision rules, which also all seem perfectly reasonable. Kenneth Arrow¹⁰ showed already in 1951 that in general, it is impossible to find an election method that meets several very reasonable criteria.

This doesn't mean that we should abandon all decision analysis in favour of throwing dice, it just means that it's extremely important to understand what

8 To see this, use suitable groups with the preference orders [A, B, C], [B, C, A] and [C, A, B].

9 G'day all you who have skipped to here from the beginning of this chapter. Read the skipped section next time around. A short summary of the skipped part is that we *tried* to find some good fundamental criteria for choosing decision rules so that the rules don't yield unreasonable results. We tried.

10 ... also a Nobel Prize laureate in economics, 1972.

you're doing. That's the main point here – that intuition cannot cope, not even with fairly simple decisions, because *none* of the intuitively reasonable criteria yield reliably rational results. Many decision makers don't even understand the rule criteria involved here. Even what appears simple and obvious can't be dealt with in a simple way if you scratch the surface a little.

So now what?

There are solutions to all of this. Our point is merely that nothing about decisions is as simple as it seems. It really isn't.

Did you have any idea things were this dire?

That's why we need, really, really need, well-thought-through methods – methods that can encapsulate the difficulties and lead to reliable results without having to ruminate at every decision, on the detailed reasoning surrounding the pitfalls accounted for above. Most people are happy to be able to avoid having to slog through such reasoning. Admit it was a teensy bit tiring – or did you simply skip forward?

We Need Sound Methods

One important thing to note is that in the reasoning above, we didn't even deal with uncertainty in the background data. By uncertainty, we mean that it is not 100% certain that a particular consequence occurs, but rather that it may occur with a somewhat lower level of probability. If we mix in such aspects, then everything becomes even trickier, and this is nearly always the case in reality. Since things are difficult enough without involving probabilities, this further strengthens the argument that methods really, really, really are needed.

Ok, ok. So now what do we do?

Since we have to make decisions and we assume that it's desirable that most decisions, at least those made in organisations, should be well-founded and rational, but neither individuals nor businesses use sensible methods to do this and don't even...

Stooooooooop, we've got that point – for the fifty-eleventh time. Let's have a solution instead.

Alright, then we have to find a medium to help deal with decisions and to understand criteria for decision rules...

That's not a surprise and it's not a solution? Come on, stop beating about the bush and tell us how.

Many decision makers doubtlessly react negatively to the thought of being replaced by an automatic method and a computer program...

Said that already too. Yaaaaaawn!

In spite of all these pitfalls, most decision makers still think they make good decisions. This doesn't correspond to reality, as we have seen, and we should learn to make decisions in a structured way. As with many things, it is possible to learn this. However, it seems a trifle foolhardy to expect to become a better decision maker through a do-it-yourself approach, when it's taken hundreds of years of the best human minds to identify and prove the existence of the many pitfalls.

So what are these principles for decisions?

What Does a Sound Method Look Like?

To answer that question, we need to reflect on some overarching issues. Which the right decision is, that is actually intertwined with the fundamental values of the decision maker.

What, in truth, are the foundations of our position relative to the values of different actions? Why is abortion right or wrong? Why is euthanasia right or wrong? Can it ever be defensible to kill a person? What constitutes an inviolated life? Is it right to revolt against authority? How much can we allow ourselves to profit at the expense of others? Should we tolerate the desperation of the needy? How much may altruism cost us? Is an individual human life ever worth less than the survival of a non-human species?

Uncomfortable but revealing questions.

But taking such questions seriously and thinking for a while leads straight to the overarching theories for moral decision making. A preliminary and coarse characterisation can be made through theological and deontological theories.

Theological theories consider that only the *consequences* are of importance to the decision. Nothing else. The only thing that counts is what an action results in. The likes of origin, motives, circumstances, state of mind or intentions have absolutely no significance at all.

Deontological theories, on the other hand, balance various factors in a decision situation. The morality of the Old Testament is to a large degree rule-based. Its point of departure is for the most part the Ten Commandments, *regardless of the consequences*. Even such varied modern philosophers as Kant, Prichard and Sartre can be said to have built on deontological theories to a large extent, and thus take into consideration a great deal more than just the consequences of actions.

Here we must stop and think about what's important. We've taken the position of propounding a theological view through and through. The consequences of actions are all that matters. The point of departure is, in short, that if someone hit you on the head, it hurts just the same even if the blow was brought about with the best of intentions. In organisational decision making, we should ask ourselves if it really is possible at an organisational level to uphold any view other than the theological one.

If we assume that it's only the consequences of actions that we should focus on, then the question immediately arises as to how we should judge these. How should they be valued? We do so many things for so many reasons, seldom for one reason alone. So the problem here seems to be understanding what we really want to achieve.

Phew, heavy stuff.

Of course, lots of savvy people have thought about this too and a great deal has been said about it. Some consider that pleasure is the only thing that counts in our lives. This is usually labelled **hedonism**. Others consider this to be vulgar and prefer to weigh in other issues like knowledge, meaning and other more subtle feelings which can be loosely labelled ‘happiness’ or more precisely **eudaemonism**.

There is a plethora of different schools of thought both within hedonism and eudaemonism and we’ll not go into these. However, it might be interesting to reflect on what your ultimate values actually are: pleasure, happiness, safety, money, love, power, pure fulfilment of your preference orders or other idealised values?

These questions are deeply existential and have probably haunted human-kind throughout our history. At least many thinkers throughout the centuries have had a great deal to say about such things. A few of the prominent figures who have expounded on these questions are Epicurus, Hobbes, Nietzsche, Bentham, Mill, Moore and Harsanyi.

Oo aahr, that be them there phiphopholers, that be.

Supposing you do have an understanding of which values are important, who then should benefit from them? Should you carry out actions that benefit only yourself (**egoism**) and your family (**nepotism**), or your organisation (**corporate enslavement**) – or should you aim to ensure that they’re spread out for the benefit of many (**utilitarianism**)? It’s therefore also important to think about distribution issues. Most people think far too little about this.¹¹ Naturally, a lot’s been said about this too. Millions and millions of people have been tortured and murdered, and states have been both created and destroyed in various “interesting” social experiments, but still we haven’t found the values leading to an even remotely perfect model for societies.

In the light of history, therefore, we urge gentle contemplation when selecting values and beseech as much caution and prudence as possible, so as avoid affronting the sensibilities of others with too much genocide, torture, rape and ethnic cleansing. There are no objective truths. This we guarantee. So

11 Whoever dies with the most toys wins. And all that jazz. Maximalism rather than minimalism.

when you've finished thinking, whatever your mission turns out to be, keep it reined.

And again, in organisations, we assume that most employees usually work toward the visions and goals that the owners and senior management have set. If this is not the case then management definitely has a problem, though one completely different in nature from those under scrutiny in this book.

Methodological Issues

So we're back to methods again. The first thing to get comfortable with is thinking about the values of our organisation. Not only that but all the awkward issues that arise when creating a decision basis. Some of them are ungraciously unpropitious.

- ♦ Which are the relevant alternatives? How can we be sure that we've included all the important alternatives?
- ♦ Which are the consequences and which of them are relevant?
- ♦ Which opinions or people are important to take into account?
- ♦ How important are the various people and their opinions?
- ♦ Which perspectives are important to consider?
- ♦ What values are important to consider in the current situation?
- ♦ How can we estimate these values?
- ♦ How can we correctly estimate the probabilities?
- ♦ Do we need more information in order to reach a decision?
- ♦ Which decision rules should we use in which situation?

Enough, enough. Phew. One thing at a time, please.

Various Decision Models

Decision models exist to fulfil many different purposes. It might be enlightening to take a closer look at some of the different models and how they can serve us. One of the ways in which they differ is the assumptions about how they can predict the future.

- ♦ **Determinism.** In a deterministic world, there is no doubt about future events.
- ♦ **Strict uncertainty.** In a strictly uncertain world, on the other hand, there are many possible scenarios but their relative probabilities are unknown or simply not considered.
- ♦ **Uncertainty and risk.** In a risky world, both the various outcomes and their probabilities are considered.

Let's have a look at each of these kinds of models in turn.

Determinism

Many decision rules are based on there not being any uncertainty in the decision situation. In basic finance, Gordon's model for stock valuation is a good, though perhaps less than inspiring, example.

The model assumes that the value of a stock can be wholly determined as a function of the expected future dividends. There's no doubt about the future of the organisation. It's expected to live until the end of the universe when time ceases to exist. Furthermore, the politics of business and commerce is expected to be the same each year for the duration of the universe. Similarly, a risk-free interest rate (and thereby the required rate of return) is assumed constant.¹² These conditions are rather unreasonable and Gordon's formula seems useless. However, it does have its uses as long as the conditions are

¹² If you care about the details, then under these conditions the fundamental value of the current share (P_0) is calculated as $P_0 = D_0 \cdot (1+g) / (r-g)$ if $r > g$, where D_0 = dividends year 0, g = growth rate in the dividends and r = required rate of return.

kept in mind. It's no coincidence that the formula has the formulation it does. In many modelling situations, especially before computers became generally available, closed analytical expressions were sought. That was the only kind of expression that could be solved numerically in a short space of time. As a basis for general decision making, however, they offer little flexibility. We can abandon this kind of model without much grief.

Strict Uncertainty

In decision theoretical contexts, strict uncertainty is present when the decision maker has no idea about the probabilities of the various consequences. Decisions made under strict uncertainty are difficult. The uncertainty can be partial or total. Sometimes it's possible to attribute probabilities to some of the outcomes but not all, and sometimes it's impossible to attribute probabilities to any at all. That leaves the poor decision maker with the only option of considering the values of the consequences.

Here we can introduce a way of modelling decisions. Consider a decision situation where our potential actions are the two alternatives A_1 and A_2 . The future world can exist in various states, one kind of which is a national obsession in the UK – the weather. If we wonder whether to take along an umbrella on a walk, then there are two possible future states of interest:

- ♦ It will rain
- ♦ It won't rain

We assume here that all else is equal. For instance, when making such daily decisions, we don't consider the possibility of a nuclear attack. The effects of common or garden rain in that case would pale in significance and render the umbrella problem irrelevant.¹³

Now imagine two states T_1 and T_2 . These are common to both alternatives in a model like this, which yields different consequences depending on what

13 This is called the modelling horizon. But if you hear of a nuke-resistant brolly, let us know.

we do, and depending on which state the world ends up in. If we choose the alternative not to take along an umbrella and if it begins to rain, then we'll get wet. If on the other hand, we choose to take along an umbrella and it doesn't rain, the exertion of carrying such a heavy item¹⁴ to no avail might cause us to become over-heated. In this model, we'll not consider the knock-on consequences risking to produce the following bad effects: enervation of the fleshy parts, impotence of the nerves, torpor of the understanding, haemorrhages, deliquia and, along with these, death.^{15,16}

To each of the four consequences, a value is assigned that that outcome represents (v_{ij}). As we've seen, the value can be defined in various ways though here we'll use some kind of numerical scale. If we put all this together, we end up with the small grid shown in Figure 6.

	State T_1	State T_2	
Alternative A_1	v_{11}	v_{12}	$\text{MAX}(v_{11}, v_{12})$
Alternative A_2	v_{21}	v_{22}	$\text{MAX}(v_{21}, v_{22})$

Figure 6. Decision grid for problems like taking a walk with or without a gamp

Swoosh. The hawk made a swoop. Here it is at last – a model. But it's still without a method and before we reach there, we need finally to decide upon a few decision rules.

A few Decision Rules

Over the course of many years, much thought has gone into which rules correctly distinguish alternatives. Several more or less nifty rules have been suggested. We'll briefly describe a few candidates. But these are too simplistic so we shan't dwell on them for long.

¹⁴ ... especially one that's nuclear resistant. We've been on the lookout for those for ages.

¹⁵ You'll be relieved to know that it's perfectly possible to model such chains of events, in case we got you worried. But that's beyond the scope of this book.

¹⁶ In case your worry instead was the progression, find some solace in that Hippocrates' prognoses were published without peer review.

- ♦ **Maximin** is a *defensive* rule that chooses those alternatives whose worst consequence is the least bad. In the decision grid, the row chosen is simply the one with the highest least value. In this way, the decision maker can ensure that the outcome receives the best possible guaranteed outcome. For decisions where organisations are sensitive to backlashes, this can be the preferred approach.
- ♦ **Maximax** has precisely the opposite goal and chooses the alternative where the best consequence is the very best. The rule maximises the best outcome with an *offensive* approach. For decisions where organisations are in great need of an advantage or some form of progress, this can be a possible strategy.

Both maximin and maximax are in some sense extreme decision rules. An attempt to choose a golden medium between the two is the Hurwicz rule, which weighs the least advantageous (x) and the most advantageous (y) outcome of each alternative, yielding a resulting value $r \times v_x + (1-r) \times v_y$. Here v_x is the value of outcome x , and r is a weighting factor between 0 and 1 usually called a **risk index**.

The smaller r is, the more hope we place in the future, since then the positive factors are weighted more heavily than the negative ones.

- ♦ **Minimax** compares the maximum outcome values of all the alternatives and chooses the alternative with the lowest maximal value.
- ♦ **Minimax-regret** is a variant on the minimax rule which evaluates the various losses (f_{ij}) incurred by each outcome j for a given alternative i . Each combination of alternative and outcome has a loss value. The loss value f_{ij} for an alternative i and outcome j is the difference between the value of the best outcome for alternative i and of outcome j . So the smallest value that f_{ij} can assume is 0 when outcome j is the best outcome for alternative i ; and the largest value that f_{ij} can assume is when outcome j is the worst outcome for alternative i . The loss values can be entered into the decision grid as shown in Figure 7. Minimax-regret compares the maximum

loss of all the alternatives and chooses the alternative with the lowest maximum. This minimises exposure to uncertainty.

	State T_1	State T_2	
Alternative A_1	f_{11}	f_{12}	$\text{MAX}(f_{11}, f_{12})$
Alternative A_2	f_{21}	f_{22}	$\text{MAX}(f_{21}, f_{22})$

Figure 7. Decision grid with alternative forfeits under uncertainty

Don't fret if this isn't entirely clear. These are rubbishy rules anyway. We've only included them to put our later offerings into perspective.¹⁷ It isn't credible in a real-life decision situation that you really haven't got any information or opinion about the probabilities that would justify using these rules. Therefore, it seems a reasonable extension to the grid above to attempt to incorporate more of the information that's available in a decision process.

Uncertainty and Risk

Said is done. A decision under risk occurs when a decision maker has estimates or opinions about the probabilities of the various consequences. This is the case for example in Black Jack, Russian roulette or dice games, but also in decisions about marketing and investments. This might seem strange when it comes to investments, but even though no exact probability can be assigned to the success of investments like the acquisition of a company, it's often advantageous to treat the investment as though it were a decision under uncertainty.

Even these things can be modelled using a grid. The grid looks much the same as before but for the addition of a new row. This new row contains a probability for each outcome.

Again, we're back to a decision situation with two alternatives and two outcomes. Each of the four consequences is assigned a value. Each of the two outcomes receives a probability p_j and you might ask where these come from.

¹⁷ ... and in an attempt to be pedagogical by building up your knowledge step by step.

A probability for an outcome can be based on a series of experiments by *observing* a large number of times what happens after a certain alternative takes effect, and noting how many times the various consequences occur. Based on the proportion of occurrences, each outcome is assigned a **frequency probability**. Observation of this kind may seldom be a realistic prospect, but that's how frequency probabilities are set.

Another way to go is to let the decision maker *estimate* the probability. Such estimation is usually used when for some reason or other it's expensive, difficult or outright impossible to carry out experiments in order to discover probabilities based on frequency. This is generally all too often the case.

Either way, the resulting grid looks like that shown in Figure 8.

	State T ₁	State T ₂	
	Probability p ₁	Probability p ₂	
Alternative A ₁	v ₁₁	v ₁₂	$E(A_1) = p_1 \times v_{11} + p_2 \times v_{12}$
Alternative A ₂	v ₂₁	v ₂₂	$E(A_2) = p_1 \times v_{21} + p_2 \times v_{22}$

Figure 8. Decision grid for decisions made under risk

The rule for deciding the best alternative is now not quite as extreme as before. The information about the probabilities is weighted in an expression that's usually called the expected value, which for alternative A₁ is calculated as

$$E(A_1) = p_1 \times v_{11} + p_2 \times v_{12}$$

and similarly for the other alternatives. The alternative with the highest expected value is that which can be expected to be the best (or at least most advantageous for the decision maker if she's been wise). The expected value rule has the great advantage of being unambiguous and simple to calculate as long as the probabilities and values are given as real numbers.

The theory for maximising the expected utility has dominated decision analysis for decisions under risk. It's been widely applied, even as a descriptive model in several fields, including economic behaviour, and become generally accepted as a normative model for rational choice.¹⁸

At least in those cases where the possible outcomes can be assigned a monetary value, it seems natural that the total value should be a balanced assessment of these. Apart from that it also seems natural to give different weights to the outcomes so that more probable outcomes gain more influence over the final value. So the approach isn't completely up the spout.

Tree Models

The classical way to represent a decision problem is to use the grid above. But in many situations representing the problem with a tree feels more natural, is more flexible, and offers a better overview. Decision grids are a static representation. They only take one decision level into account, where there is only one choice to be made between different alternatives and once the decision is taken the problem is solved. However, in reality, decision problems often have more levels and are more dynamic. One decision problem may lead directly into another and perhaps another, and so on. Choices may need to be made in a particular order or the decision maker may be able to influence the probabilities of events occurring. For these situations, a tree representation is definitely better than a grid. Actually, it's better in most cases.

Decision trees are very useful for complex decisions where the tree structure provides the decision maker with a graphical representation of the core of the problem, and shows all the inter-relationships between the possible choices and uncertain factors. In its graphical form, it describes the architecture of the problem methodically and fairly objectively.

So how do trees relate to the grids we looked at just now?

¹⁸ 'Descriptive' means trying to describe what decision makers *actually do*, rather than what they *ought to do*.

Well, the tree contains the same information as a grid, at least when the tree only contains one decision with one level of events. But the more levels there are to the decision, the more advantageous a tree becomes as a supporting medium because the overview of the problem isn't lost. At the same time, it's important to try to limit the number of decision nodes so that the overview of the tree isn't lost either. If many nodes are needed then a summarising tree can be created for the overview accompanied by several expanded sub-trees which contain the details.

A tree consists of three types of nodes: In **decision nodes** the decision maker chooses one from several alternatives. Then there are **event nodes** where the outcome is outside the immediate control of the decision maker. Nature quite simply tosses a coin and chooses one of its available courses of action.¹⁹ Finally there are **outcome nodes** that show the outcomes of the various paths through the tree. Figure 9 offers an example of a decision tree.

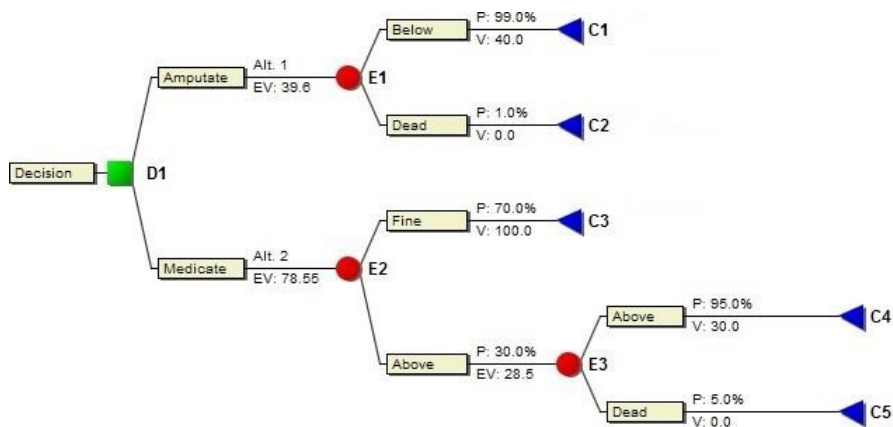


Figure 9. Decision tree for decisions under uncertainty

A decision tree consists of a **root** (green square in the diagram), which represents a decision.²⁰ It also contains a number of **event nodes** (red circles) that represent various kinds of uncertainty. Finally, there are **outcome nodes**, also known as **consequence nodes** (blue triangles) which represent final

¹⁹ Pure animism!

²⁰ We don't want to complicate things too much, but a tree can perfectly well contain more than one decision node.

outcomes. Each path through the decision tree from the main decision node to a final outcome represents an alternative course of action that may be taken in the decision situation modelled by the tree.

In trees like this, measures of uncertainty, such as **probability distributions**, are assigned to the event nodes. This is interpreted quite simply as if one alternative has been chosen then each possible outcome has an uncertainty. These are often expressed with a number that denotes the respective probabilities for the various events. An event can result in a final outcome or in a subsequent event. Based on this structure a decision rule can be applied, such as maximising the expected value or minimax-regret.

In Figure 10 there are two alternatives. Each of these can lead to various events and outcomes. Outcomes don't necessarily occur but may do so with a known probability as indicated in the tree. The values in the blue outcome nodes denote the monetary values if the outcome occurs. Here's a small example of this in the form of a decision tree.



Figure 10. Example of an investment decision tree

Here's the decision problem:

- ♦ to invest, or
- ♦ not to invest...

... that is the question.

An investment brings with it a 40% risk that the outcome will be a loss of € 40. On the other hand, the probability is 60% of gaining € 30. If the investment is not made, there is still a risk of losing € 50 but the probability of this is only 20%. The remaining 80% pertains to the possibility of gaining € 10.

Here we can make a decision by calculating the **expected monetary value** (EMV). This is done by multiplying the probabilities for each outcome with the monetary value it yields.

In this way, alternative A yields $0.4 \times (-40) + 0.6 \times 30 = € 2$

0.4 is the probability of losing € 40. The remaining 60% is the probability of gaining € 30.

Alternative B yields $0.2 \times (-50) + 0.8 \times 10 = -€ 2$

The probability of losing € 50 is 20%. The remaining 80% is the probability of gaining € 10.

Here we see that the EMV for investing is a gain of € 2 and for not investing is a loss of € 2. The alternatives are chosen by the highest value which is to invest, because it's € 4 better. Nerve-racking decisions, these.²¹

Problems with the Expected Value

But what value does a given decision have? In those cases where the possible outcomes can have a monetary value, it seems natural that the total value should be a balanced assessment of these. It also seems natural to provide different weights to the outcomes so that more credible outcomes affect the values more. Sounds quite astute!

A rule that's often used, just as in the example above, is always to try to maximise the EMV. If a decision maker is **risk neutral**, then it seems like a reasonable rule. A risk-neutral decision maker would consider that twice the monetary gain is always twice as desirable. However, this rule has been

21 For even more adrenaline, add nine zeros (referring to digits here, not decision makers).

questioned and relentlessly debated amongst philosophers, mathematicians and economists. It was one of the last century's most skilled mathematicians, John von Neumann,²² who together with Oscar Morgenstern²³ laid the foundations for **utility theory**. Utility theory is covered in more or less yawn-invoking, eyelid-encumbering freestanding works, including big fat books on modern finance theory which are more or less guaranteed to bring terminal boredom eons before bounty. Suffice to summarise in a short example.

Suppose that old Professor Frege has been given a choice between receiving € 50,000 or playing a one-off game in which the probability of winning € 100,000 is 50%. The probability of winning € 0 is thus also 50%. Now, according to the expected value rule, he should be indifferent about choosing one or the other alternative. Since the expected value for both is € 50,000. But of course, this isn't so. Suppose further that our dear Professor is suffering from a life-threatening condition and has an immediate need of € 40,000 for an operation that has to be performed within a week in order to save his life, and furthermore that this money isn't available from any other source. In that case, € 50,000 in hand would be worth much more than a 50% chance of winning € 100,000 but equally possibly of ending up with zero. However, if the operation were to cost € 80,000 then the € 50,000 alternative isn't going to save him. But the game alternative offers a 50% chance of winning € 100,000 and surviving. Thus the expected monetary value of an alternative doesn't necessarily correspond to the real value for all people or organisations in all situations.

The conclusion is that in many situations, the EMV is not a sufficiently good principle and should be replaced by a utility-based rule. In such a rule, each outcome is assigned a number on a scale, for example a decimal number between 0 and 1, where 0 is the worst imaginable and 1 is the best imaginable outcome. These numbers correspond to the perceived utility that a person or organisation gains from the outcome. If the monetary values are judged to correspond completely in proportion to their utility, then utility boils down

22 Who was such a rare beast as a party-going genius with a strong affinity for fast cars, women and booze, all without moderation.

23 Not known for his frivolity, though.

to the same thing as a monetary scale. But here again, this is an idealised situation. Something we'll come back to in a minute.

Expected Utility

Obviously, in the methods presented above the values were given. In reality, as you might suspect after a little more thought, things aren't quite that simple.

The methods discussed above can easily be improved, as the last example hinted at, by formalising the concept of utility. The **expected utility** (EU) is calculated in a similar way to the expected monetary value (EMV). The calculations are still easy to carry out and provide well-defined decision rules in every given situation. There are, however, a number of objections even to utility theory:

- ♦ Utility functions are uniquely individual in the sense that all individuals have their own unique utility function based on their preferences. For this reason, there's no simple way to compare different individual utility functions. It's difficult to determine whether an individual is happier than another when they enjoy the same benefit. From this, one can conclude that the utility function, for groups like organisations, is somewhat tricky to use.
- ♦ Imagine if a decision maker doesn't know any exact number for the utility of a certain outcome. In that case, it would be desirable to be able to express statements such as "the utility of outcome X lies between 0.3 and 0.5" or "the value of outcome Y lies between € 250,000 and € 400,000". None of the above methods accommodate this. They all require exact values, regardless of how uncertain decision makers actually might be. They therefore fail to provide the possibility of expressing degrees of uncertainty in any way. A clear and worrying deficiency.
- ♦ It also transpires in many practical situations that decision makers don't want to express their utilities or monetary value in terms

of numbers. Often they wish to relate the outcomes to each other in statements like “the value of outcome X is greater than the value of outcome Y” or “outcome Z is much less desirable than outcome T”. Even this is impossible using the traditional models described so far. Unfortunately, this sometimes leads to decision methods not being used at all. Sad but true.

- ♦ Many decisions are of a once-only nature, or they’re so important that an outcome that is too undesirable cannot be tolerated, even if it has a low probability of occurring. In this case, it seems reasonable to require that if the probability of a very bad²⁴ outcome is too high²⁵ then the alternative should be discarded, even if it shows a good expected value. Such safety reasoning can be interesting for insurance companies, for example. They seldom want to enter into an insurance contract where the possibility of profit may be good, but where there’s a risk that isn’t negligible of a loss so great that the whole company could be ruined.²⁶ Safety levels aren’t found in traditional decision models either and therefore aren’t regularly used where they should be.

These objections notwithstanding, in practice the expected value is often a good approximation²⁷ to the expected utility – which is not perfect anyway. Thus, we will use the terms interchangeably moving forward.

Precise and Imprecise Data

It would be good if everything were now more or less done and complete. We have one representation and we have a rule that works fairly well, except for a few minor points. Unfortunately, things still aren’t so simple, for there are new troubles brewing, an inkling of which the points above imply.

24 ... that is: under a certain safety level of “too dire an outcome”.

25 ... that is: above a certain safety level of “too much risk”.

26 Some companies meet with ruin in any case, but usually for other reasons – mismanagement, as (mis)management consultants would say.

27 Especially when augmented with the abovementioned safety (often called security) levels.

If you think a little about how things usually are, then you'll find that you almost never have access to such precise information and that any such requirement makes unreasonable demands on decision makers, with the exception of a few rare situations. How can you, dear reader, or anybody else for that matter, indicate in a meaningful way the exact probability of a future event? In real life, people usually can't even distinguish between probabilities of 30% and 70%. So the hope of finding meaningful support for exact values is a hobbling duck. And yet many methods work in this very manner.

In order to really bring this home, allow us to point out that in real decision situations we seldom have access to exact information such as "outcome X will occur with a 17% probability". It's seldom that we have access to frequency-based probabilities at all, rather we rely on subjectively estimated probabilities. It should be obvious that estimates don't have sufficient precision to be indicated by a single fixed number. It's not surprising then that methods based on fixed values are difficult to use in a meaningful way. And yet they persist in spite of the uncertainty of the estimates.

Cor! Uncertain uncertainty. That's really a twister.

Amongst the greatest problems with traditional methods is the difficulty of dealing with intervals and comparisons between probabilities. But intervals aren't only normal in real-life decision problems, they're the rule rather than the exception, so it should go without saying that they should be an integral part of any model worth its mettle.

The same applies, as we saw just now, to utilities. It's often extremely difficult to quantify these, even if, contrary to what you might expect, we've managed to find good measures for them. Uncertainty in the form of **confidence intervals** and other kinds of uncertainties always turns up.

What can we do about this?

Quite a lot, actually. We need a way to be able to indicate probabilities in intervals, where each interval's width is determined by how certain the decision maker is. The statements "the probability of outcome X lies between

20% and 50%” and “the probability of outcome Y lies between 30% and 40%” are more meaningful, since they can better reflect the decision maker’s beliefs. Even statements of the kind “the probability of outcome X is low” and “the probability of outcome Y is relatively high” are both reasonable and natural. Don’t you think?

In a similar way, we can broaden our methods to become something much more useful by indicating value statements in a better way. Try for example to estimate the value of a medium-sized genocide, a slightly catastrophic natural disaster, a triflingly dangerous nuclear meltdown or something closer to your business. All estimates of these things in terms of precise numbers fall foul of reality, to say the least.

So it’s indeed practical to be able to say things like “the value of outcome X is between € 500,000 and 700,000” or “the value of outcome Y is relatively high and definitely higher than the value of outcome Z”.

So in between those methods without probability estimates and those methods with fixed estimates, there’s masses of room (and need for) models that permit greater expressive power in the form of intervals and comparative orders as illustrated in Figure 11. Let’s call these more modern methods **interval methods**, in contrast to the traditional ones.

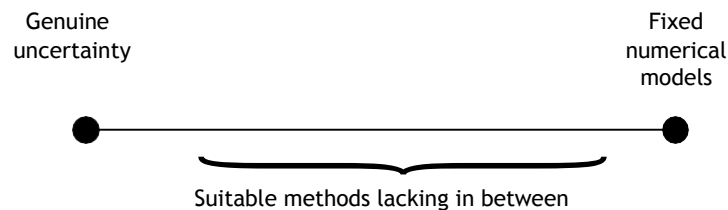


Figure 11. Traditional methods represent extreme points on a method scale

We’re about to go into some of the details about interval methods, but if we start to fray your brain, just skim through it all. You don’t need to appreciate the details but might like to see some of them at least. Fortunately, there are tools that can take care of them all and spare people recondite calculations.

Intervals and Comparisons

A different formalism in which to represent the problem is required and for this reason, we'll look a little closer at how alternatives can be represented and calculated.

Natural statements in this context seem to be the following:

- ♦ Event X is probable
- ♦ Event X is possible
- ♦ Event X is improbable
- ♦ The probability of event X is equal to p
- ♦ The probability of event X is greater than p
- ♦ The probability of event X lies between p and q
- ♦ Event X is just as probable as event Y
- ♦ Event X is more probable than event Y
- ♦ Event X is much more probable than event Y

Together, all such statements about the likelihood of events form what we call a **probability base**. This isn't a mathematically oriented book so we ask you to trust us when we say that all these statements can easily be represented with the help of simple expressions in logic and maths.²⁸

For the value (utility) of the current outcomes, there are the corresponding statements:

- ♦ Event X is desirable

28 A tad of maths for those whom it tickles. A probability base is said to be **satisfiable** if there's at least some numerical value that can be assigned to every variable (e.g. $p(X) = 39\%$, $p(Y) = 22\%$, etc.) so that all equations are satisfied at the same time. The idea behind satisfying the probability base is that without it, there's no way to calculate the value of the alternative since there are no probability assignments that can possibly satisfy the requirements.

- ♦ Event X is better than the worst outcome
- ♦ Event X is undesirable
- ♦ The value of event X is equal to v
- ♦ The value of event X is greater than v
- ♦ The value of event X lies between v and w
- ♦ Event X is just as desirable as event Y
- ♦ Event X is more desirable than event Y
- ♦ Event X is much more desirable than event Y

Together the statements about the desirability of outcomes form what we call a **utility base**. As for the probability statements, all these statements about the utility value of events are also easily represented by formal expressions, and the best news about this is that a tool can perform the **transformation** automatically.

Comparing Alternatives

Just like with traditional methods, the value of an alternative can now be given by the expected utility or the expected monetary value. Since both probabilities and utilities are now variables, calculations of the expected value give rise to a somewhat more complicated expression, so-called **multi-linear expressions**. A simple example of one of these is the ordinary expected utility

$$EU(A) = p_1 \times u_1 + \dots + p_n \times u_n$$

In the expression, each p_i is a variable that expresses probabilities. In the same way, u_i denotes the utilities of the various outcomes.

Expressions like these are considerably trickier to solve by hand and previously it was difficult to calculate them in a reasonable time even using

computers.²⁹ But today there are clever algorithms and programs that calculate comparisons between multi-linear expected values in a reasonable time.

Great. How then are alternatives compared?

Now that methods exist for calculating expected values for probability- and value bases using interval statements or comparative statements or both, new questions arise. How are the alternatives compared by using all this new stuff? First let's look at what are called **dominance concepts**, which establish when one alternative is better than another.

- ♦ Alternative A_1 is **at least as good** as A_2 in the probability- and utility bases if $EU(A_1)$ is greater or equal to $EU(A_2)$ for all permitted values of the probabilities and utilities.
- ♦ Alternative A_1 **dominates** A_2 in the probability- and utility bases if it is at least as good as A_2 and $EU(A_1)$ is greater than $EU(A_2)$ for some permitted values of the probabilities and utilities.
- ♦ Alternative A_1 is **permitted** if it isn't dominated by any other available alternative, that is, if no alternative exists that's better.³⁰

If there's only one permitted alternative then obviously it's also the best alternative and should be chosen. Usually, however, several alternatives are permitted. This is because if there's an obviously good or bad alternative, this will usually already have been dealt with manually before the decision process is set into full-scale motion, and therefore it wouldn't give rise to any decision problem. That there are several permitted alternatives means that, for different permitted numerical values of the variables (parameters) in the bases, different alternatives are the best.

29 Also necessary is the determination of what's called the centre of mass and with that, the complexity of the calculation rises exponentially. In really large problems, it would previously (before computers became widespread) have been necessary to wait several million years for a result, which would give a new meaning to the term bureaucracy.

30 This is explained less formally later on in the section about multi-criteria. Read there if this doesn't make sense yet.

Given the above, it seems natural to look at the parts of the bases in which each alternative is best and the size of those parts. In other words: how many values can be assigned to variables (different numerical values) so that an alternative remains at least as good as the other alternatives? This leads us to the concepts of **contraction** and **belief mass**.

Contraction and Belief Mass

Somewhat simplified, a contraction of the comparison between A_1 and A_2 indicates in how much of the information space (the bases) acceptable values of the variables lead to A_1 being a better alternative than A_2 . A contraction of A_2 compared with A_1 is determined in the same way. Now it might be the case that the first contraction becomes 40% and the second 80%. The rational conclusion would then be that A_2 is a better choice because in some sense it's more "robust". In a much larger proportion of all the possible cases, A_2 is then the better choice. In this way, there's a rational argument for choosing A_2 . If there are more than two alternatives, then many pair-wise comparisons arise between all the alternatives. These are then easily summarised in a single number for each alternative.

Another way to compare A_1 and A_2 is by studying how much the decision maker believes in the difference δ_{12} , where

$$\delta_{12} = EU(A_1) - EU(A_2).$$

It is relatively easy for a decision maker, using software, to express this belief for each piece of information in a decision situation. A total belief in each alternative and in differences δ_{ij} between them can then be calculated in order to study how the belief mass behaves under various sensitivity analyses and by doing so, a good sense can be gained of how much the alternatives differ.

This sounds a little complicated but it's all part and parcel of a computer tool. After a short while, people who work with such tools usually don't have any difficulty, neither entering statements into the system nor comparing alternatives. It's probably time for a small example, one that clearly makes the point that even small problems are already surprisingly tricky to solve without good methods.

The Gangrene Example

Petunia Croft has gangrene in her left leg and must quickly choose between amputation under the left knee (call this alternative A_1) and treatment only with medication (alternative A_2). She consults her doctor Rev Royce who offers information about the risks involved in the two methods of treatment. Dr. Royce in his turn reflects on the information and statistics from the National Health Service. Together they arrive at the following. The risk of death from amputation below the knee is less than 1%. When medication alone is used the risk is 20–30% that the infection will still spread and require amputation above the left knee. The risk of death from that more extensive operation is as high as 5–10%.

Which alternative should Ms. Croft choose? Think about this straightforward problem before reading on.

As you can see, even though the problem is structurally very simple, it's clearly quite difficult to give a reasonable recommendation justified by intuition.

The problems are:

- ♦ We don't have exact probabilities, since the statistical material is uncertain (because of the intervals).
- ♦ There's no reasonable scale of utility values. Obviously, we could try to set precise values, but if we take the problem seriously this can't be done anywhere near precisely enough.

But if we skip these difficulties, surprisingly enough we can advance a fair way with the problem just by ranking the outcomes.

- ♦ Full recovery is better than amputation below the knee.
- ♦ Amputation below the knee is better than amputation above.
- ♦ Amputation above the knee is better than dying.

Few would dispute these preferences. But are they enough to make a sound decision?

Figure 12 below (similar to Figure 9 but containing intervals) shows two alternatives and their outcomes, exactly as Ms. Croft and Dr. Royce perceived them. The outcomes (C1–C5) don't necessarily occur but rather occur with those probabilities indicated above the connecting edges in the tree. The numbers on the outcome nodes indicate their values if the events leading into them occur.

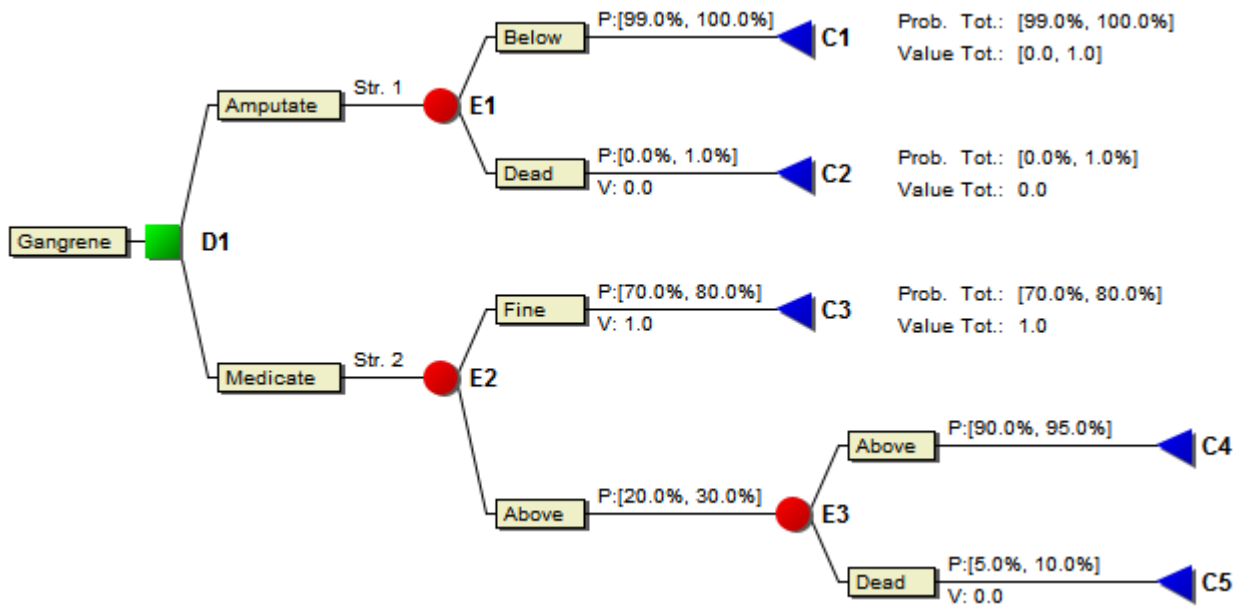


Figure 12. Decision tree for the gangrene case

Having modelled the problem, it's time for some calculations. Again the calculations are quite complicated to do by hand so we'll use *DecideIT* to solve the problem. The curves in Figure 13 below show the strength of the alternatives for various degrees of contraction. Without going into the details, we can say that the cone represents the difference in belief in expected value. 96% of the belief mass lies above the x-axis, favouring A_2 = medication. The rest,

4%, favours A_1 = amputation. Thus, alternative A_2 is considerably better, and the result is stable in relation to various possible changes to the input values. Note that this result was obtained by only specifying that complete health is the best outcome and death the worst, but without specifying the values of outcomes in between other than ranking them.

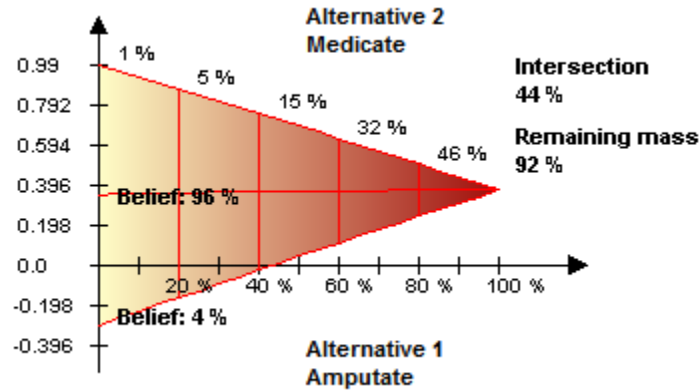


Figure 13. Recommendation graph for gangrene treatment

The patient can therefore begin the medicinal treatment with the knowledge that she made the wisest decision given the information available.

Using decision methodology in this way, problems like this, even considerably more difficult ones, are relatively simple to deal with. Not only that but the analysis can be further refined with little additional effort.

More on Uncertainty

In a complex world, nothing important is simple and decision making with the help of intervals is no exception. The advantage of methods that use intervals is that they avoid problems with exact probability distributions.

On the other hand, it sometimes becomes difficult to find a reasonable decision rule that recommends one out of a set of alternatives, while still fully reflecting the intentions of the decision maker. Since the probabilities and values are represented by intervals and comparisons, the expected values for the alternatives will also be intervals. The problem is now that they can overlap. This means that the best possible value for one alternative can be better than the worst possible value for another. A situation like this is illustrated in the **cardinal comparison** in Figure 14 below. The intervals of the expected values for the two left alternatives (1 and 2) are strongly overlapping, while both are clearly better than alternative 3.

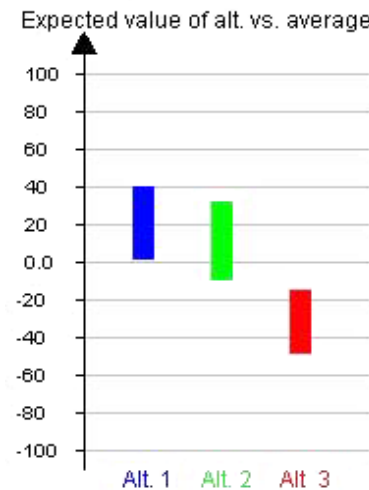


Figure 14. Cardinal comparison

Alternative 3 is so bad that it would probably have been discovered just by looking at the decision problem directly. But how should the others be tackled? Sometimes it's impossible to sufficiently distinguish between stronger alternatives, and pure interval analysis doesn't always provide enough advice in situations like that shown in the illustration above.

In this example, you might feel that this really isn't that important since the alternatives are actually rather similar. The difficulty in distinguishing them perhaps just reflects reality.

Wrong! It reflects the decision structure. The more **structural information** the problem contains, the narrower the EMV interval becomes which contains the "real" possible results. In this way, the structure provides information that can be meaningfully quantified. This is a useful property.

The idea builds on consideration of the decision maker's opinion in different parts of the interval, expressed or tacit. One result of this is that the alternatives' expected values can overlap but even so, it's still possible to determine which alternative is best. The non-overlapping parts can be seen as the core of the decision maker's evaluation of the decision.

This becomes somewhat cryptic, but the principles are actually quite simple.

It's like this. Suppose you estimate the probability that it'll rain tomorrow to be about 60%, but being an intelligent person, you don't suppose for a moment that the probability is exactly 60% so you want to indicate some form of uncertainty or confidence about this estimate. Let's say that after supplementing this, you draw the conclusion that the probability is not less than 50% and no higher than 80%. Furthermore, you believe a little more strongly in the lower probabilities than the higher ones. Since you originally assumed that the probability is about 60%, you probably tend to believe more in numbers around 60% than around 80%.

This opinion can be quantified and represented as in Figure 15. The higher the functional value (y-axis) the greater the belief in the respective probabilities. Let's call the diagram a **belief distribution** (or more formally a distribution of belief mass). Such distributions can have any form depending on the decision maker's various beliefs. The important thing is that they can be used to sharpen the analysis.

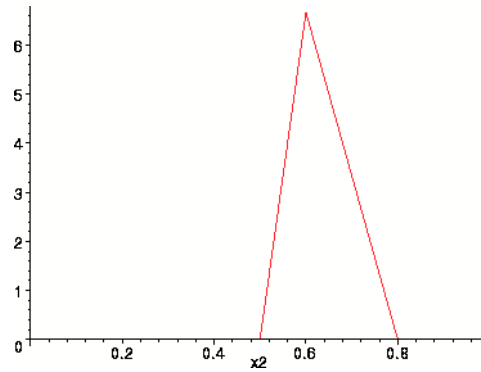


Figure 15. A triangle distribution

When a decision tree is evaluated, the final result, which is of course the expected value of each alternative, also gets a distribution of belief mass like this. The interesting and important observation now is that this actually leads to the distribution becoming more concentrated around a smaller partial interval, which strangely enough is fairly independent of the original distribution for the probabilities and values. In Figure 16, this phenomenon can be seen as the overlapping intervals of the two expected values. The two alternatives contain statements that make them impossible to distinguish by using interval analysis without either belief mass analysis or contraction.

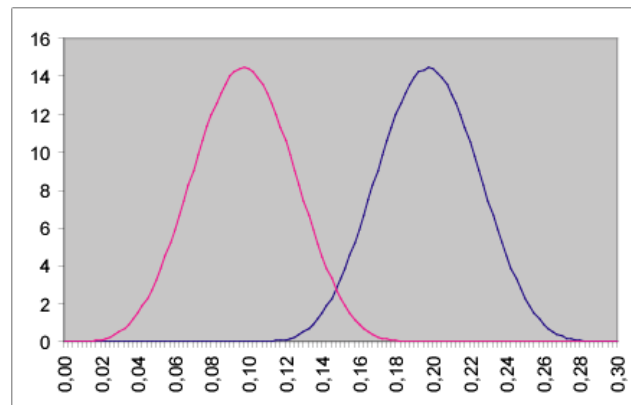


Figure 16. Overlapping belief mass distributions

Let's say that the lilac (left) distribution lies over the interval for the expected value of alternative A_1 . The blue (right) lies over alternative A_2 in a similar way. The interval for A_1 is 0.0–0.2 and the interval for A_2 is 0.1–0.3. Clearly, these intervals are overlapping and interval analysis alone can't distinguish between them. But now we also have the lilac and blue belief distributions. As you can see, the largest area under the curve for A_1 is between, say, 0.06 and 0.14. And the largest area under the curve of A_2 lies approximately between 0.16 and 0.24. So the aggregated belief in the alternatives differs considerably, indicating that A_2 should definitively be selected. This is the same phenomenon that underlies the result in Figure 13.

More Advanced Models

We've now gone through a little bit about representation and methods for decision problems. We've also braved a small in-depth jaunt (with those of you who bore with us) into more developed methods for calculating problems that naturally arise. It's important to remember that there are tools (that is, computer programs) that take care of all the troublesome calculations. But there are further aspects to this and some of them, in particular, deserve mentioning. We're thinking of multi-criteria models (several goals), sensitivity analysis and risk analysis. And here they come.

Several Goals and Stakeholders

There's often a prodigious number of opinions, frequently they clash, and frequently this causes trouble: we want to be free and safe; we want to maximise economic strength while minimising the effects on the climate; we want to minimise the risks and still attain the greatest possible profit; we want to maximise efficiency while maximising work quality; we want to juggle while riding a unicycle on a tight rope underwater.

This can be a tad toilsome as if we needed to point that out. On top of this, all parties expect to be satisfied while no important stakeholder should be completely dissatisfied, for that would cause toil indeed. Happiness is short-lived, and in our material world tends to require frequent increments in gain,

preferably greater gain than the people occupying neighbouring houses, driving lanes or career trajectories. Advertisers and salespeople use this human weakness to full effect. But because there aren't enough desirable niches and resources to satisfy everybody's or every company's desire for gain and a secure place in the world, conflicts result. These are often characterised by different people wanting different things but mostly having to settle for less of what they want and more of other things while, particularly in decision situations, it's about individual people or companies wanting to attain several goals and having to settle for less of them all because several may be incommensurable.

Boohoo, so what to do?

The problem is finding the optimum, but the optimum depends on what your perspective is. A decision can, and in most cases should be seen from several perspectives and the value of an alternative is then dependent on the perspective from which it's viewed.

Imagine a chemical producer that is about to invest in a new purification plant for sewerage. The decision to invest can be studied from several perspectives. The financial perspective is central for assessing both the cost of the investment and the possible reduced efficiency of production. Environmental concerns and risks to wildlife are both important for their public image as well as being an obligation to assess. Also, mistakes bring fines, litigation for damages, badwill, disrepute and dismal consequences. Staff welfare is another aspect, as are public health, noise levels and the distribution of tasks. The alternative that offers the best economic prognosis frequently leads to problems in natural and working environments.

How is it possible to take all these factors into consideration? Not least is the problem of how to communicate all of this to groups of stakeholders, decision makers, and, and, and. There are so many goals, so many criteria, so many perspectives.³¹

Tricky, so what do we do?

31 You can regard these concepts as a whole. The goals are seen from the perspective that generates aspects (which some schools of thought call evaluation criteria).

One way would be to evaluate the alternatives only from the most important perspective.³² A complication that's not all that unusual is the tendency then to ignore the information provided by the other perspectives. This doesn't seem terribly clever, not least because even if the information is available, there's a risk that it'll quickly be forgotten. That could be most disruptive. Believe us.³³

Imagine an alternative that with only a hair's width surpasses another from an economical perspective but which is clearly worse from an environmental perspective. You understand. A number of disquieting questions are likely to pop up further along the line, and the decision makers risk being portrayed as morality monsters and environmental marauders, which can be injurious as well as costly.

Another approach would be to separate the decision analyses by perspectives and later in some way try to tinker the results back together. But how can a balanced assessment of the whole be carried out? And what's the actual working process like?

Goal conflicts also force us to make various kinds of trade-offs between our stated goals. How is this done?

The first thing that's often missed, of the many things that are often missed, is to actually understand what our goals really are. Identifying these is a central issue, as is having a method that supports a priority order. These are the very essences of our next pet technique. Here it comes...

Multi-Criteria Analysis

Supposing that we *are* aware of our goals, then a common method to deal with them is to set up conditions, i.e. levels for **thresholds**. Alternatives that satisfy all these can then be considered good enough, and at least don't need

32 ... if such a perspective can be found.

33 The conditions surrounding the "most important" perspective can also change – as things often have a tendency to do.

to be discarded right away.³⁴ Creativity works better when it's focussed on satisfying goals rather than other matters. It works worst when firefighting.

Following this, the “dominated” alternatives are rejected. Those poor ostracised alternatives are those that *for all goals are at most as good as some other alternative, and worse for at least one of the goals*. On that token, they're worse than the ones allowed to remain. At least one other alternative therefore remains that's *at least as good at all goals and better than one or more* of each of the rejected alternatives. The alternatives that remain after this procedure are called, with poetic concinnity, **non-dominated alternatives**. And they, or some subset of them, can look forward to a more detailed analysis.

Just to make this really clear, if only one alternative remains and it's non-dominated, then the matter is resolved. The trouble is that this is seldom the case. How often do you find an example of a car that is best on all scores from price to performance, from comfort to luggage space, from fuel consumption to attractiveness? Even if a large number of cars fall by the wayside because of the pure boundary conditions, a number probably remain that do not dominate any other car – that is if you aren't a complete car fanatic.

Making this point again, often one alternative action is best for short-term gain, while another can be best for long-term sustainability. A third might be best from the point of view of the environment, and a fourth is best for the organisation's technical competence, and so on, and so on. In order to progress with the analysis we need some way by which to prioritise. There needs to be some way by which to assign aspects³⁵ different weights, afford them different sway in the discussion. If you think that everything is equally important then you should consider whether you really need all those aspects. In classical multi-criteria analysis, this means that our goals, or aspects, are given numerical weights that should reflect how important the aspects are.³⁶ The greater the weight, the greater the (assessed/perceived) importance.

34 If all alternatives are eliminated in the first step then the level of ambition needs to be lowered, or new alternatives found.

35 ... in contrast to utility values for consequences.

36 If you ever want to impress, the jargon is “relative significance of criteria”.

If you succeed with this then you have a model with aspects and weights. And now there's a plethora of methods you can use. Some build on classical decision- and utility theory, others build on decision makers providing preferences that for each aspect indicate when an alternative is better than another. In all these methods, some form or other of weighting is used in order to fairly assess the aspect.

Ok, so which methods are good at this?

A fairly simple and well-established method from 1976 is that invented by the American pioneers in decision analysis, Ralph Keeney and Howard Raiffa. The aspects and values are weighed together using something that resembles expected values,³⁷ like this: for one alternative with three aspects, you'd have the following expression of the alternative's total weighted value

$$V(A_1) = w_1 \times v_{11} + w_2 \times v_{12} + w_3 \times v_{13}.$$

where w_1 is the weight assigned to aspect 1 and v_{11} is the utility that alternative A_1 has when we look at aspect 1 of the problem. The expression yields weighted values for each alternative. This is really the same principle as for expected values, which we became acquainted with earlier.

The weights usually have the same requirement that they should be larger than or equal to 0 and the sum of all weights should equal 1. That isn't an absolute requirement, but it's really very practical when working them out. Now have a look at Table 1. The cheapest alternative is simply to dump the waste somewhere. That's irreverent of the environment, not to mention detrimental, and the competence of the staff probably won't develop much. Not the first time something similar has happened, though. Don't get us started on the multi-criteria problems and skewed utility functions in much of our modern society. The environmental issues we are dealing with today are a mere whisper of what's to come. But let us not wander astray from the example we are discussing, however doleful we might find it.

37 If you'd like the label given to this method, it's additive multi-attribute utility, a term that could conceivably provide a resonant ring suited to some occasions. Choose your occasions carefully, though.

Alternative	Cost	Environmental impact	Competence
Dump waste in the sea	€ 1 million	–10	10
Take care of waste in accordance with disposal directives	€ 2.5 million	–1	35
Dig and bury it	€ 1.2 million	–10	5

Table 1. What to do with environmentally harmful waste

Table 1 describes costs, environmental impact and competence development in terms of utility. The degree of this utility is represented here as a number. This might feel a bit odd at first, but this is what the method builds on.

The values here are analogous to how utility values were described earlier. So, a common method to do this is first, for each aspect, to let the worst utility be 0 and the best 1. Then set out each of the alternatives in-between, somewhere along the scale between 0 and 1, proportionately where this seems justified.

A cost of € 1 million yields a utility of 1 (lower costs are better) and a cost of € 2.5 million yields a utility of 0. The other values are worked out in the same way, so that they all lie on the same scale.

In the example, a cost of € 1.2 million is assigned the utility

$$(-12 - (-25)) / (-10 - (-25)) = 13/15 \approx 0.8667.$$

Immediately apparent is that the sea-dumping alternative dominates the dig-and-bury alternative, so digging is buried, R.I.P. The next step is to work out the weighted values for the remaining alternatives – in this case only two.

But here, there's more trouble – as usual. This time excessive precision. That should be avoided in decision modelling. All the problems with utility theory that we went harping on about earlier remain. The precision causes problems in quantifying and leaves no room for uncertainty or vagueness which means it becomes difficult to model well. The modelling space has been destroyed

through our desire to give the problem a simple mathematical dress-up. That is a common mistake, and a quite unnecessary one at that. Instead, we should stipulate the utilities as intervals or rankings, making use of the information available, but no more and no less. No artificial flavouring here, please.

What we mean is quite simply that instead of embodying the utility of a cost as a precise number like 0.8667, from the point of view of environmental impact we should be able to specify that the alternative that follows disposal directives is better than both the others. What a surprise! And the cost of € 1.2 million is much more desirable than a cost of € 2.5 million. Another surprise! Then we can apply the methods that we talked about earlier. The same story applies to the weights. If this seems convoluted, don't despair. It really isn't as fiddly as it might appear at first. We'll come back to this in Chapter 4 when we show a couple of real-life examples.

Eliciting Weights

It isn't easy to get people to say what they really think. This can be because they don't know themselves or for some reason they don't want to reveal what they think. Surprisingly often, though, people don't actually have an opinion or clear thoughts about something they're asked. This has implications for trying to elicit weights and utilities from those decision makers. The ability to avoid excessively precise numbers has proved very helpful in overcoming such problems with eliciting people's views.

Sometimes, it's common to compare aspects against what's called reference aspects (also **reference perspectives** or **reference criteria**). Cost is not an uncommon choice for this purpose – the aim is then to try to find a trade-off point that the decision maker seems willing to accept. A cost reduction of € 10,000 in some context is perhaps, round about, equivalent in value to a couple of units lower environmental impact. This is rather similar to a cost-benefit analysis but requires everything to be measured in well-defined scales and units, which isn't exactly simple as we have already established. Furthermore, the trade-offs aren't linear, a fact that rather messes things up.³⁸

38 Being non-linear means that there can be a big difference in usefulness between earnings rising from € 0 to € 100,000 compared with earnings rising from € 1,000,000,000 to € 1,000,100,000.

Weights can be gleaned from decision makers by having them talk about their preferences in relation to fictional decision alternatives. Somewhat faster methods can be based on interactively exploring weights using various graphical tools through which priorities can be found. Simply searching for a ranking is another method where not too much effort needs to be put into finding weights, except when these are actually needed. The needs usually become evident later in the process while conducting specific sensitivity analyses. Personally, we like the latter method because it's more time- and cost-efficient. Funnily enough, quite imprecise weights often suffice for differentiating decision alternatives if they're expressed in terms of intervals like "the weight lies between this and that value" or as ranked orders like "this is more important than that". Later on, if the weights need to be made more precise, it's more artful to delay³⁹ the effort of doing so until such precision is needed, and not earlier.

The method that we ourselves have developed offers, in addition to the above, lots of help with weighing together values. As before, the weights don't need to be stated in precise numbers. Interval statements work just fine. In those cases where people don't even want to work with intervals, aspects can be ranked from the most to the least important. Thereafter a similar calculation apparatus is used as before, yielding an analysis of the decision situation with all values weighed together.

Figure 17 shows how relationships between aspects can be expressed. The aspects' relative significances are displayed in falling order from the most important (in this case bathability) to the least important (business impact).⁴⁰ Larger distances indicate, perspicaciously enough, larger differences in weights (and thus importance) between two neighbouring aspects.

Expressed criteria importances can subsequently be changed by moving the respective arrows representing the criteria up and down the slider scale. The results from the changes are immediately displayed in several output windows if a software program such as *DecideIT* is being used.

³⁹ Welcome news for procrastinators.

⁴⁰ This was a county council rather than a company.

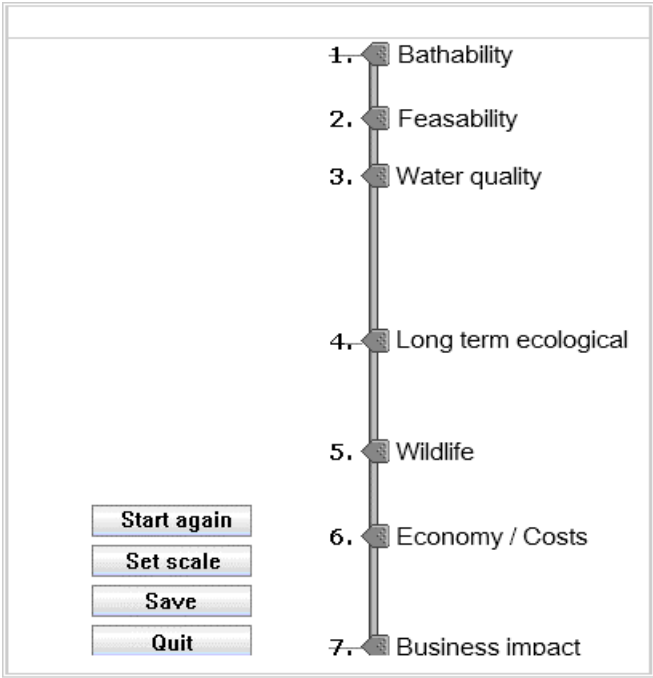


Figure 17. Aspect weighting on a slider⁴¹

There are various supporting methods that facilitate the discussion, **weight elicitation** and visualisation of priority orders. Such methods are suitable to include as part of a business decision process.

In *DecideIT*, there are two distinct ways of expressing rankings for decision situations where the decision maker does not want to express criteria importance (or probabilities or values) as numbers or intervals. One is a total ranking as in Figure 17, the other one is to specify relations between specific criteria or consequences, such as in the gangrene example in Figure 12, and leave other ranking relations unspecified as desired, possibly due to lack of information at the present stage of the analysis.

41 The elicited weights in the figure are taken from a decision case about the measures to be effectuated around a polluted wetland. The decision makers consisted of a group of politicians. Look out for this case in Chapter 4.

Tree Models for Multi-criteria Analysis

Having dealt with decision trees with probabilities and outcome nodes earlier, we'll now look into tree models for multi-criterion analysis. These are usually called aspect- or **criteria hierarchies**.⁴² These are based on the fact that there can be smaller facets or sub-aspects for some or even all of the aspects. This model is intended to help judge how well alternatives meet the various goals. Sub-aspects are often more concrete and in certain cases easier to estimate or measure.

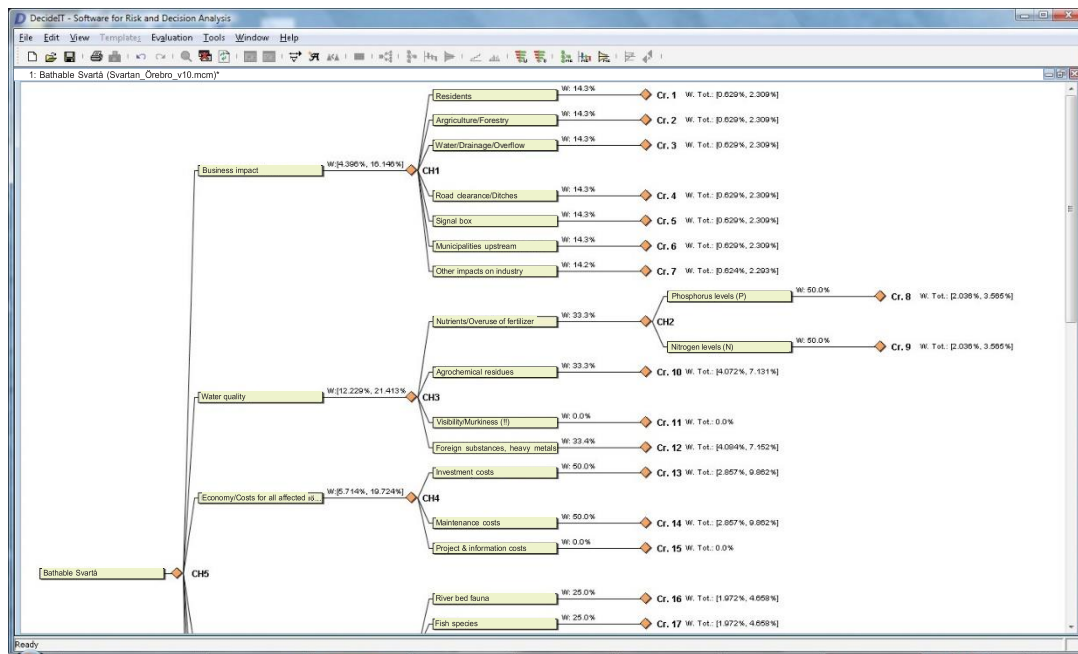


Figure 18. Criteria tree from the Svartå case in Chapter 4

In Figure 18, the more abstract main criterion *water quality* has been dissected into **sub-criteria**, of which two are very concrete: *phosphorus-* and *nitrogen* levels. In this way, the evaluation can be made easier.

⁴² In principle, these are structures similar to the decision trees already covered, except that the nodes here are aspects rather than probabilistic events and outcomes.

The tree model also makes it easier to explain the decision problem to others. At least it seems so whenever we've used it. The tree illustrates what's important for the decision while at the same time it's dynamic, helping stakeholders to work together on extending and modifying the important structural components that the tree consists of.

In a good decision tool, all the components that we've discussed in this chapter are present: aspects, sub-aspects, aspect weights, alternatives and utilities. In this way, we get a foundation for the decision method, just as with an ordinary decision tree, and the principles for evaluating it are completely analogous. Very practical. Very convenient. Very rational. And not particularly difficult, once you're used to them, which doesn't take long.

In really advanced tools these kinds of models can be combined.⁴³ Only when your tool can apply both decision trees and aspect trees to one and the same decision problem can you claim to have a tool with **complete coverage** of the decision process. That includes really complicated decisions, whether the problem has one or more aspects, whether it contains uncertain events or not, and whether the problem's parameters are well-defined or not. In all cases, a tool with complete coverage devours the whole problem, transforms it into a decision basis and delivers a suggested solution, or else shows that the information is insufficient and indicates which information needs sharpening.

Sensitivity Analyses

What happens if one or several of our judgements are not quite right, that is, not completely wrong, but perhaps a little off the mark? And what happens if the information basis changes? Maybe some outcome becomes more probable, less relevant or more expensive. These are all perfectly natural questions to have in mind. A good decision method leaves room for this and a good decision basis contains such information. Performing a sensitivity analysis is in other words an important component in making a proper decision analysis.

43 Need we add that *DecideIT* is a really advanced tool?

Part of the decision process is to continually question the structural elements it contains. Is this the best and most relevant alternative for solving the problem? Has some aspect been overlooked? Should any of the information be viewed in a different way? Are any essential outcomes missing?

If there are different opinions within the decision group, more alternative analyses can be carried out and pitched against each other in the form of workshops or the like. But the incomparably commonest sensitivity analysis is to vary the existent parameters, that is, the probabilities and outcomes. This kind of sensitivity analysis can be used to examine how resilient the choice of an alternative is to changes in the numbers that are used in the analysis. Usually, we try to solve for so-called **critical values**. These are the boundaries within which the probabilities and outcomes need to be confined in order for the decision not to be changed, given that the other values remain unchanged.

For example, if the stock keeping rate has been calculated at 35% then there may still be a degree of uncertainty about that figure. There could be a reason to find out how much an increase or decrease affects the final result. In order to do this, one or more variables are permitted to wander between the lowest and the highest credible values, while the other variables retain their original values.

This sounds quite easy.

Indeed it *sounds* so. Let's start with the apparently simplest model by fixing the values of the probabilities and utilities. First, we look at the values of the final outcomes. The value of each outcome should be varied up and down in a systematic way – one at a time, or several at a time. If the information has been structured in such a way that one outcome is worth more than another, then these must co-vary, otherwise you'd be looking at combinations of values that aren't permitted. There'll be rather a lot of combinations to test. Maybe not too overwhelming with a computer to do the work.

The next thing would be to look into the probabilities of the events. The probabilities of each outcome should also be varied up and down in a systematic way. One at a time. Let's start. Suppose that we have five possible outcomes for one alternative. The probability for these all together is 100%.

Call these P_1 to P_5 . Let's start with P_1 which we had assumed to be 25%. Now we increase it to 25.5%. But since the sum of the probabilities should be 100% we have to diminish the sum of the other P values. How? By reducing P_2 ? P_3 ? Both P_2 and P_4 ? We must do all of these and much more besides. Then we increase P_1 to 26%, and then to 26.5%, and then... Thereafter, we reduce P_1 to 24.5% and work our way down. For each step, first P_2 should compensate, then P_3 , then... ad nauseam.

Obviously, there'll be an effect on the alternatives in various cases, so how do we summarise all this? It's a lot of data and we've only looked at P_1 for one alternative so far! What if we'd had ten? Or a hundred?!? You see the problem. A huge number of calculations is needed to put this analysis together, an unmanageable number. This is what's usually called the combinatorial explosion, and all this results from a deceptively simple representation of probabilities, utilities and weights – *in the form of fixed numbers*.

However, if the probabilities, utilities and weights were represented as intervals and comparative rankings from the outset, then all you have to do is, within their ranges, to let them assume all the possible values that don't invert the comparative rankings that might exist between any of the outcomes. You've seen this in the context of contractions earlier. The sensitivity analysis is then working with the model proper, not some artificial add-on evaluation using hypothetical values that have no grounding in the stakeholders' views. In contrast, first building a model with fixed values and then afterward varying values around those points may find sensitivities that are completely irrelevant to the true values of the stakeholders. However, with intervals, such bogus results don't arise because the space for analysis is already in the model, and furthermore, owing to the process by which the intervals and rankings were elicited, that space has been fully validated.

One way to show how an alternative's expected value varies when its parameters vary is a **tornado diagram**. This shows how variation in things like probabilities in a decision problem affect the final result – the expected value. The parameters in the diagram are sorted so that the probabilities with the greatest effect are highest up followed in decreasing order of influence.

The breadth of each parameter indicates the degree to which it affects the expected value. The appearance of the diagram bears some resemblance to a tornado, see Figure 19, hence the name. But don't worry about these twisters devastating your business; on the contrary, the extra analysis they provide only increases your safety in the face of stormy decisions.

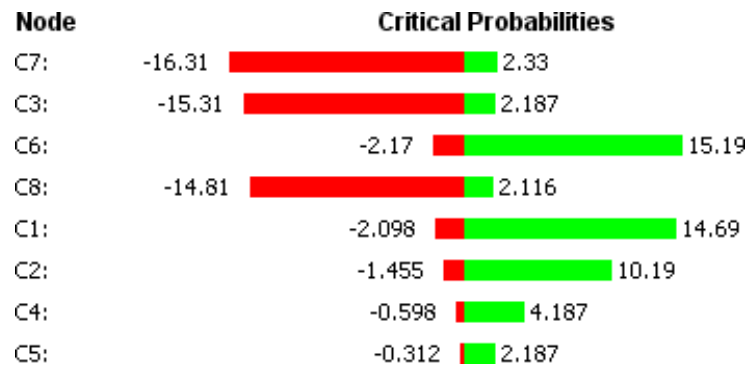


Figure 19. A tornado diagram

Risk Management

Risk management is an important part of many people's professional- and perhaps even personal lives and a plethora of methods exist for judging, communicating and quantifying risks. Unfortunately, many of these, even though they're diligently used, are inept.

You would think that disasters and the thought of them happening would be disagreeable enough to spur people to avert their risk, but surprisingly often, people responsible for avoidance are shocked when anything of magnitude really happens. Also, but perhaps less surprisingly, these people try to avoid accountability by making a huge effort to explain what they did or didn't do, often spending far more energy dodging blame than would have been necessary to avoid or mitigate the serious consequences of the event.

Notwithstanding what all these people say and do, the way that people handle risks is often a big problem and can result in huge losses for society

and businesses, not to mention death tolls – as we could witness during the Covid-19 pandemic. Risks must be taken seriously, and woolly analyses and intuitive guesses are extremely inappropriate in such contexts. Low probabilities of devastating events are far too serious to ignore, and although dismal outcomes are infrequent, it's often the very fact that they are so infrequent that results in them being ignored or put off, with catastrophic repercussions as a result, simply because they weren't managed in a systematic way.

Various risk management methods have arisen as a result. This is a lucrative business area and as usual, the quality of the work as well as the competence of risk managers varies greatly.

Nonetheless, successful risk management really can reduce the probability of undesired events and minimise their negative effects. Qualified risk management is needed, regardless of whether the risks pertain to extensive disasters or relatively small matters like project management, daily processes or the introduction of new standards.

Risk analysis can be said to be the most formal part of risk management and can be applied to decision making in all areas and at all levels. A risk analysis is intended to be a systematic identification of risk and risk-taking. It usually incorporates judgements about risks and vulnerability where all threats are analysed, at least to some extent. Such analyses are also used to control current security and protection, as well as to evaluate the effects of changing these.

The potential advantages of risk analysis are significant, for example adequate cost coverage, increased productivity, better focus on security and increased awareness in general. One principal function is also to create effective procedures for dealing with events when they happen. The lack of such procedures is blatant in all too many cases.

There are also different kinds of risk analytical methods. These range from simple checklists, which merely require a few hours of effort, to analyses that can take many months. By including analytical methods in the decision process, substantial problems can be solved. A very practical effect is that risk analyses can simply be included in those decision trees that are going to

be used anyway. Some may call this overkill, but one shouldn't scoff at such synergies.

A Final Reflection

There seem to be a great many decision methods and models, some of which appear to have more agreeable properties than others. Some can even model *real* decision situations. Woweee!

This calls for double celebrations. First, we have finally reached the end of Chapter 3. The roadmap to better decisions has been laid out before us. While there were some slightly difficult passages, the reward for reaching the end of this leg is that there will be no more mathematical formulas. Yeah, that's the second thingie. No more math.

And by sheer luck, we haven't met a single transcendental number despite the title of the book. A transcendental number (e.g. π) is non-algebraic, i.e. it is not the root of any non-zero polynomial of finite degree. Those numbers never made it here, except in a footnote discussing the π factor, but that's more folklore than math.

Phew! Time to kick back and relax...

Just a mo'. If all this is so great and not so terribly difficult, couldn't we please see a few examples? A couple of real-life, real-world, real-everything case studies?

Certainly. That is precisely what Chapter 4 is for.

Finally a small disclosure (or shall we call it an incentive?). For those that hang in there and clear the final leg, Chapter 4, thus making it to the end, there will be a little bonus waiting that can bring some enjoyment (dare we say excitement?). Well, let's get on with it.

4 How it Can be Done & How it Should be Done

Now for some real stuff. So far we've gone through this and that on various aspects of decision making with loads of talking and little walking, but finally in this chapter we're going for a stroll. For these perambulations, we've used our own PDMP methodology to impart order and structure to the case studies.

The first is from the Swedish Road Administration's¹ project proposal **Bypass Stockholm – Förbifart Stockholm**, which Stockholm's County Administrative Board² has promoted for some time. It's a ginormous project. The proposal is quite controversial for various reasons, but there's been little publicity concerning the decision methods employed. So we grew curious and made some inquiries.

The second study concerns water quality in the **Svartå**, a river of great beauty that runs through the Swedish town of Örebro.³ The quality of the river's water is so low that all the municipal bathing spots that dotted its banks have had to be closed. A bad state of affairs for bathers, not to mention wildlife. Measures needed to be taken to reopen the river to bathers. We know which

¹ At the time of the study, it was called *Vägverket* but has since coalesced with other bodies into *Trafikverket*, the Swedish Transport Administration.

² ... in Swedish *Länsstyrelsen*, we'll translate it to County Board.

³ One of Sweden's larger cities.

measures and others concur, now at least, following our use of PDMP.

Bypass Stockholm

As we have seen, for the most part perplexingly often, an unacceptably large element of randomness enters into decision making. This is frighteningly true even in highly costly and controversial decisions made by public bodies. The much-publicised controversy surrounding the project *Bypass Stockholm* is an illuminating example in this vein. The proposal has been controversial for various reasons, but we're not concerned here with which facts have been in focus and how relevant they are, even though this is also interesting. Here we're focussed on the decision methodology since it was deficient, to say the least.

The background is that car traffic in the Stockholm area was expected to rise dramatically by around 40% during the upcoming fifteen years. Thoroughfares were already chockablock, so there'd undoubtedly be problems ahead. The Swedish Road Administration conducted a study of how best to link up the north and south parts of the Stockholm region, which resulted in the publication of a report.⁴ Quite early on, we became curious as to the decision methodology used in this controversial project. In our eyes, it appeared already at that time that decisions were being made on dubious grounds.

In the Road Administration's report, two alternative constructions were described, *Bypass Stockholm* and *Diagonal Ulvsunda*, as well as an alternative called the *Combination Alternative* which focussed mostly on developing public transport and toll systems.⁵ These three alternatives were compared with each other as well as with a reference alternative, the *Null Alternative*. Depending on the alternative chosen, the cost of extending the road network

4 A 56-page report "Nord-sydliga förbindelser i Stockholmsområdet" that can be ordered from Trafikverket in Borlänge, Sweden [in Swedish].

5 We're actually talking about the revised report. The original one suggested the alternatives Bypass Stockholm, Ålstensleden and Diagonal Ulvsunda. Ålstensleden was rightfully removed as an alternative following large public protests led by the local NÅL initiative – No to ÅlstensLeden.

according to 2003 price levels was estimated at € 1.7–3.2 billion – rather a lot of bees and honey. At least we think so.

In August 2007, Stockholm’s County Administrative Board adopted the alternative *Bypass Stockholm* which is intended to become a link across Saltsjö-Mälarsnittet and stretch from the motorways E4/E20 at Kungens Kurva in the south, via Sättra, Kungshatt, Lovön, Vinsta, Lunda and Hjulsta, as far as the E4 at Häggvik in the north. The total length of the bypass road section will be about 21 kilometres (13 miles).

Our professional expertise lies in decision analyses and since rather a lot of money was involved, we thought it would be interesting to try to understand why this particular alternative was supposed to be the best – something we couldn’t comprehend. Fortunately, the decision basis consisted of about 80 different **valuations** of the various aspects of the alternatives under 19 different aspects. What the County Board claimed to have evaluated was how well the alternatives satisfied the aspects, including accessibility, environment, regional infrastructure, traffic safety and economic growth. Sounds fair enough.

Naturally, the decision was complicated and the intention was that the decision basis would provide a picture of the most essential factors. The evaluation was reported on two quite different scales, *Yes* and *No* as well as a scale from “– –” to “+ +” depending on how well the alternatives satisfied the goals. Furthermore, many evaluations yielded intermediate positions such as “+/-” or “0/+”. The best alternative for each aspect was marked by shading it. For example, for the aspect *Health*, the *Combination Alternative* was highlighted with shading even though all three alternatives were given a ‘+’. This was supposed to be interpreted as this alternative being somewhat better than the other alternatives in relation to this aspect. In relation to the aspect *Energy and Materials Accounting*, the *Combination Alternative* was given a “+/-” and *Diagonal Ulvsunda* a “–”. This was intended to be interpreted as the *Combination Alternative* being better than *Diagonal Ulvsunda* under this aspect. Figure 20, which is taken directly from the Road Administration’s own material, shows how things turned out.

How do the various alternatives satisfy the goals?				
Goal fulfilment	Null Alternative	Bypass Stockholm	Diagonal Ulvsunda	Combination Alternative
Accessibility				
Availability		Yes	Yes	No
Saltsjö - Mälar section		Yes	Yes	No
Central parts		Yes	Yes	Yes
Create bypass		Yes	Yes	No
Access roads		Yes	Yes	Yes
Regional structure				
Common market		Yes	Yes	No
Multi-centredness		Yes	Yes	No
Economic growth				
Basis for development		Yes	Yes	No
Traffic safety				
Fewer fatalities & seriously injured	0	+	+	++
Equality				
Improved for women	0	-	0	+
Environment				
Health	0	+	+	+
Safety	0	+	+	0 / +
Wildlife, heritage, recreation	0	--	0	0
Climate targets	0	-	-	+
Custody of land and water	0	+ / -	+	+
Energy & materials accounting	0	-	-	+ / -
Economy				
Investment	0	--	--	--
Maintenance	0	-	-	--
Social economics	0	0 / +	0 / +	0 / -

Goal fulfilment (best alternatives shaded)

Figure 20. Goal fulfilment for the Bypass project [Väggerket]

Already in this summary, certain difficulties became apparent to us since it was fairly unclear what was meant by these scales. It wasn't stated which

properties the scales had, something that's extraordinarily important for knowing which method to use for a balanced assessment of the values. Nor was there any account of, let alone reference to, any weighting having been done at all, so they probably hadn't even realised that particular problem.

Another, but worse worry was that the various aspects weren't assigned any proper weights. This resulted in not having the foggiest how important each of them was. There wasn't even an indication of what might be considered a reasonable priority order. Oddly enough, these issues were completely bypassed (no pun intended) even though this was absolutely pivotal for the decision since none of the alternatives were best for the majority of aspects of the problem. All this resulted in a decision basis that was even more difficult to evaluate. The actual evaluation wasn't done in any more qualified manner either and because of this, the complete lack of aspect weights wasn't revealed. The resulting information basis was the only decision support available to the politicians.

Could it have been the case that this decision process was *so* complex that reasonable methods for dealing with it really didn't exist? It would certainly excuse the Road Administration if that had been the case. It would also support the odd classification and structure that they used since in some sense it did at least function as a kind of support for the decision.

But no, this was not the case! All the components of the decision methods that we've described were readily available at that time too. The Road Administration stood within easy reach of decision methodologies fully capable of dealing with exactly these kinds of decision situations. They could easily have opted for a state-of-the-art methodology that could have dealt with the situation's lack of precise data; with opinion and subjective impulses exerting an essential influence; with opinions varying considerably over time and from person to person, or between organisations. However, from the material presented, we can only conclude that they didn't make use of any of the available apparatus.

Before the decision was made, we amused ourselves by looking into how the problem *should* have been analysed in a more reasonable fashion. In order to really highlight this, we allowed first-year students taking a basic course in

decision analysis at Mid-Sweden University⁶ to carry out analyses with the information that at that point was available from the Road Administration's own website. What was most surprising was that the students had no difficulty at all in producing very informative analyses that indicated which of the alternatives should have been promoted.

The students could easily establish that the alternative *Diagonal Ulvsunda* appeared the best in general, at least under a number of reasonable weightings, and we confirmed the students' correct application of the methodology. Furthermore, this alternative was at least equivalent to *Bypass Stockholm* even when the aspect 'accessibility' was given the highest priority. The environmental goals for *Diagonal Ulvsunda* were relatively well fulfilled, even though they weren't as good as for the *Combination Alternative*. But even so, the other weightings showed that *Diagonal Ulvsunda* was the best alternative when considered under the whole set of aspects. In comparison with *Bypass Stockholm*, it transpired that since the roads were to be lead across the island Lovön, the aspect 'environmental impact' made *Diagonal Ulvsunda* the better alternative.

Pivotal for the result was how the various aspects were weighted, something that didn't appear to have been done at all by the Road Administration, at least there's no indication of this from the publically available material. However, judging from the unexpected outcome, strong preferences seem to have been present. It's noteworthy that one of Sweden's all-time largest road projects could be decided without accounting for such preferences with any pivotal weightings to support the decision.

Having said that, we would like to add that we're not taking a stand on which decision should have been made. We're really not terribly bothered about that, but rather about the tragedy of the situation. We find it remarkable and worrying that such a huge investment is carried out without employing modern expertise and decision apparatus, where weightings, ranking orders and values are openly expressed and taken into consideration during the evaluation phase.

6 ... previously a university college, now proudly waving a 'university' title to all and sundry.

The whole business is rather sad, but the truly worrying thing is that this isn't an isolated case. The Road Administration isn't unique in this respect and this case doesn't constitute a single occurrence of a foreboding example. Everything would be much less disquieting if that were so.

We all expect, at least in Sweden, that public bodies be able to deal with costly and complex decisions with greater competence and transparency and to use effective, if not efficient methods, especially when these are freely and easily available. It's most unfortunate if the decision mechanism is not understandable by the public, who in the end foot the bill. This is even worse when demonstrably incorrect decisions are made.

The Decision Itself

Now to the details. The decision basis for this case needs a description of the alternatives and aspects. That, the Road Administration did succeed in compiling. We also have a (rather weird) scale for the alternatives. Not so good, but ok – we'll use it. What's missing, however, is the weighting of the aspects and this is where things really go awry.

Our intention was to collect everything relevant to the decision and enter it into a suitable tool. Unsurprisingly, we used *DecideIT*. We entered the aspects and represented the alternative scale in a suitable way so that we gave an interpretation of all these “++”, “+/-”... symbols in terms of intervals. Just how we translated these turned out not to be all that important. We set the intervals to be a subdivision of values between 0 and 1. “++”, for example, could then stand for all values between 0.8 and 1. An illustration of the aspect tree is shown in Figure 21.

Once this structure had been created and the available data entered into the model, we were able to carry out various kinds of evaluations and analyses. Figure 22 shows how the various alternatives were compared.

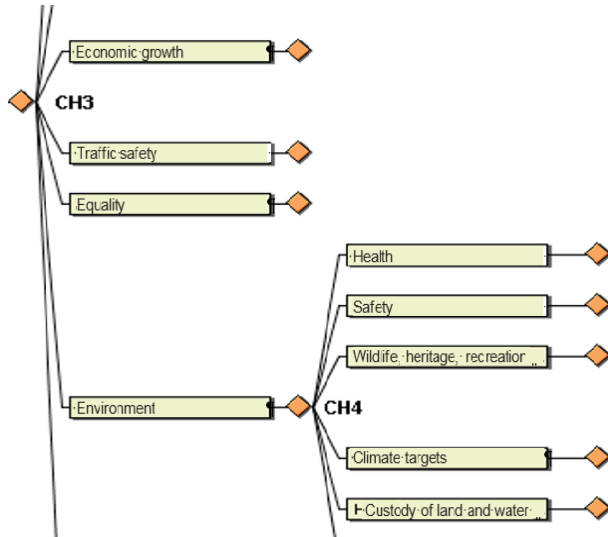
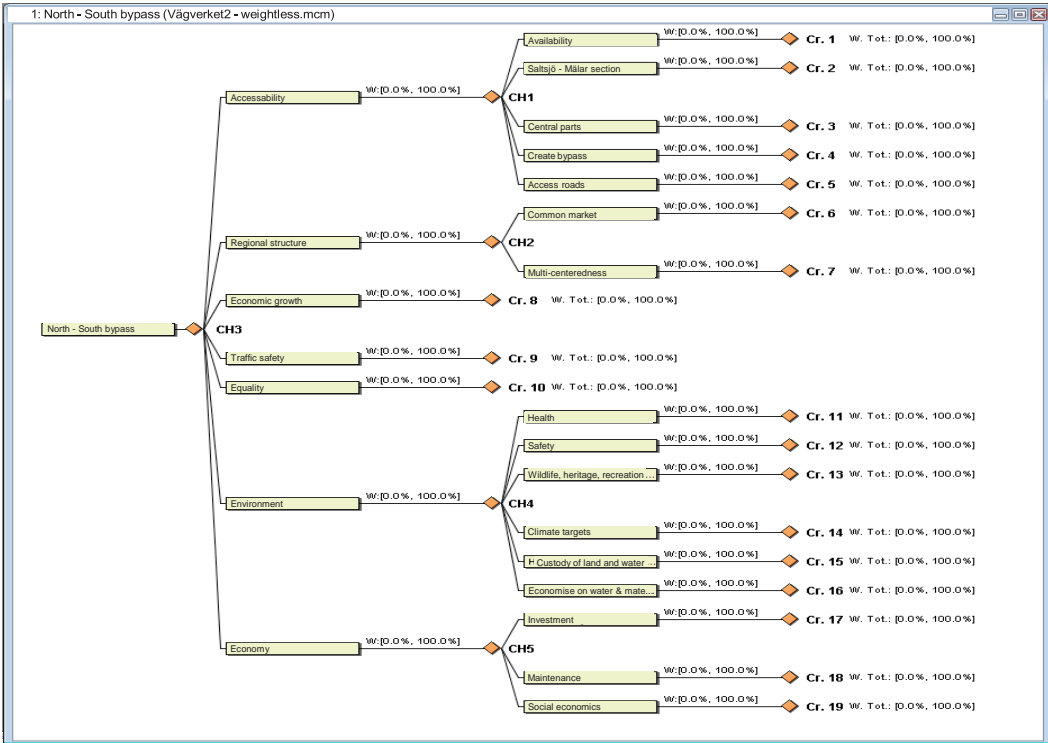


Figure 21. Complete aspect tree for the Bypass project plus a partial close-up

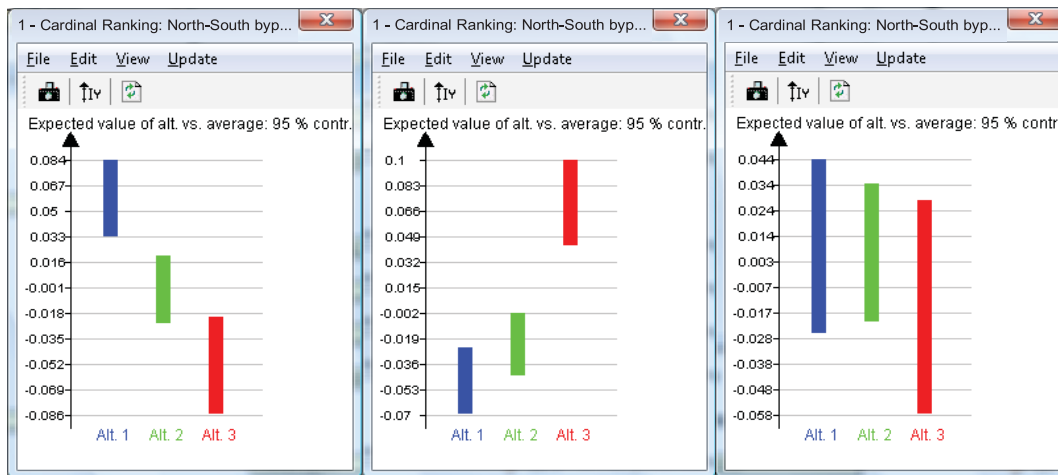


Figure 22. Results of comparisons between the alternatives for the Bypass project

The results of the comparisons are shown first without weightings (left), then where ‘accessibility’ is considered to be the most important aspect (middle) and last where ‘environmental impact’ and ‘traffic safety’ are considered the most important (right). Alt. 1 is *Bypass Stockholm*, Alt. 2 is *Diagonal Ulvsunda* and Alt. 3 is the *Combination Alternative*.

In these three graphs, we’ve used slightly different weightings for some of the central aspects in order to show how very different results are obtained depending on the weightings. This means that no well-founded decision can possibly be taken seriously without considering the weightings. Either these weightings are explicit as in Figure 22 or else they’re done unconsciously and indirectly. But when the decision is actually made, we have, whether we like it or not, weighted the various aspects against each other. From the Road Administration’s material, we weren’t able to discover any weightings, explicit or otherwise. A truly lamentable and disheartening observation.

A few comments about Figure 22. The evaluations shown are based on a method we call cardinal comparisons.⁷ This can be considered a generally sound evaluation method and we earlier explained it. For those of you

⁷ Oh, the joys of jargon! At least it’s an alliteration.

who've forgotten or skipped that, it basically involves working out the best and worst possible expected utility for each alternative. Here we can see that there's a rather large amount of overlap between the bars and that the decision basis needs to be further stipulated in order to really glean anything generally meaningful. So the uncertainty around this decision is still quite large and the poor data quality was of little help to the politicians. The Bypass Stockholm project commenced officially in 2016 and in 2020, four years into the project, it was delayed by *five* years (sic!) due to severely bad quality assurance. It is now scheduled to open in 2030 at the earliest and at a cost exceeding € 4 billion.

The River Svartå, Bathable

For quite some time, Örebro had a problem with water quality in their river Svartå, which flows right through the central part of the city. The quality of the water was so low that the public bathing areas located along its banks had been closed. Measures were required before they could re-open, ones that would secure the future supply of drinking water to the municipality and ones that would improve water quality in general. The problem encompassed neighbouring municipalities which would be affected in different ways depending on what Örebro municipality decided to do. As usual, there were a large number of stakeholders. Farmers and several other groups were those affected most negatively by the environmental measures, as is often the case.

The decision about priorities and what was to be done belonged to the municipal politicians. A long-term solution was desired and inescapable. The majority and the opposition were to be heard. Naturally, they had differing and not unexpectedly, mostly conflicting perspectives. 'Twas tricky.

We were brought in to advise on the decision process and decision support and naturally, we used our Preference Decision Management Process. Consequently, everything leading up to the decision was carried out as it should.

The opinions and ideologies of the various stakeholders and politicians had to be strictly separated. Civil servants and various environmental consultancy companies provided facts and expertise. The decision problem presented several aspects for evaluation. The problem also encompassed several decision

makers and groups of stakeholders. The statements made by civil servants and expert consultants about the outcomes were unavoidably uncertain, as were priorities. Quite a normal scenario, in other words.

It was quite simply impossible to ascertain beforehand which one was the best alternative, both for the experts with good knowledge about the circumstances and for the civil servants, indeed for anybody. Neither was it easy for the decision makers to express precise weights and trade-offs between the environmental and other goals. In addition, there were a large number of alternative solutions that had been suggested by three independent consultancy firms. It all needed structuring and grouping in order for the decision makers to be able to discuss the situation within a reasonable time frame.

In short, as is nearly always the case, the decision was difficult, it was messy and pressed for time by the desire for a rapid solution. How would Svartå water achieve its bathable status in accordance with EU norms?

The project was conducted in two stages. First, we needed to agree on a set of aspects for the politicians to take a position on. At this stage, we identified the decision problem and which goals were important to achieve.

We organised workshops. Municipal politicians and civil servants deliberated on the matter and seven aspects for evaluation survived.

- ♦ Economic costs
- ♦ Water quality
- ♦ Bathability
- ♦ Business impact
- ♦ Feasibility
- ♦ Wildlife
- ♦ Long-term ecological soundness of the constructed system

Thereafter, sub-aspects or facets were elicited, which helped the decision makers clarify what they considered important to this decision. Clearly, this wasn't just about citizens being able to bathe again. The issue was deeper than that.⁸ It soon became apparent what the issues were, which clearly helped the ongoing discussions with the civil servants.

Following that, the problem was structured and modelled.

The model was formulated by the politicians and civil servants together. Based on this model the politicians were able to rank the various aspects in order of priority, and the civil servants and consultants were able to judge the effectiveness and the impact of the various suggested measures. The elicitation was done with the help of sub-aspects for each main aspect. It turned out really well.

Weights were formulated by each decision maker in two steps. In order to support this well, we used the graphical tools of *DecideIT*, which is an integrated part of PDMP. Candidate sub-aspects were generated by the civil servants, whose suggestions were then subject to revision by all interested parties. The sub-aspects were more concrete and technical, which made passing judgement on them fairly easy.

Figure 23 on the next page shows the problem structure with several aspect levels. No, we don't expect you to grasp the details of the figure, it's mostly for showcasing the inherent complexity of the decision situation. There were six main criteria, each having 3-6 sub-criteria. Under all of these sub-criteria, each of the seven alternatives had to be assigned a value, either as a fixed number, an interval or a ranking position. This was done by civil servants from Örebro municipality. The division of work was such that the civil servants assisted by environmental experts were in charge of the values while the decision makers (the political level) were in charge of the importance weights and thus the prioritisation.

8 ... and we don't mean the depth of the river Svartå.

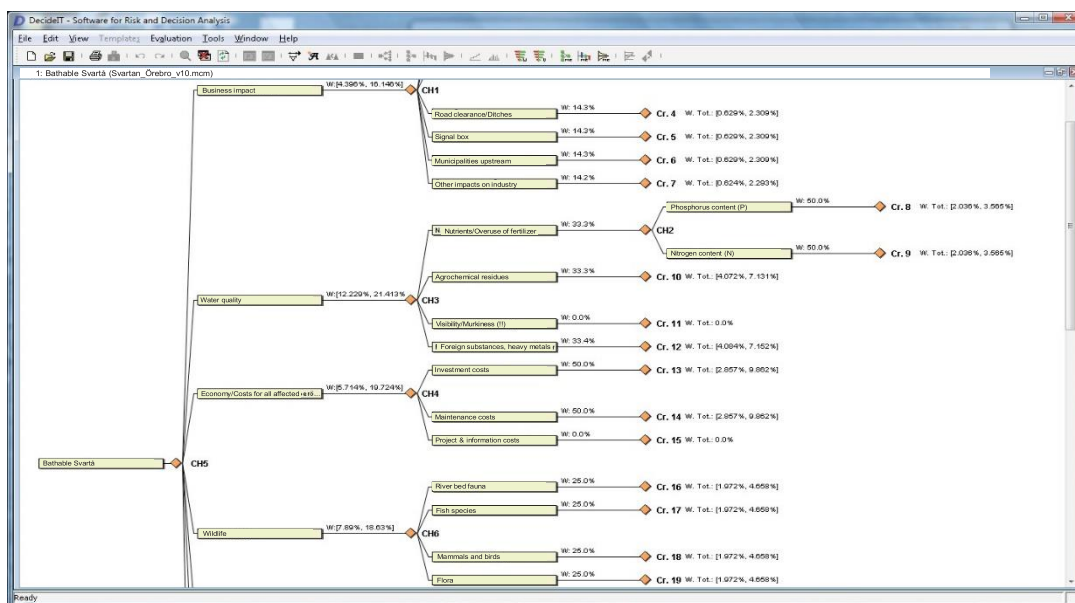


Figure 23. Aspect tree for the Svartå project

The next step was to formulate the alternatives. During the process of information capture, the experts among the civil servants judged how well various alternatives satisfied the goals in relation to the sub-aspects. Political opinion and ideals were separated from fact, and although everything was included in the model, it was clear what was what. A number of earlier proposals for solutions were eliminated when it became clear that they didn't contribute sufficiently to water quality and that their retention would have diminished the clarity of the decision process. Finally, seven *action packages* remained.⁹

1. Assess and rectify problems with private sewer outlets
2. Assess and rectify problems with municipal sewer outlets
3. Compost manure (biogas plant)

⁹ It's only a coincidence that there were both seven packages and seven aspects, holy though this number may appear.

4. Rectify city surface drainage
5. Establish wetlands
6. Rectify animal husbandry
7. Establish agricultural fields

These alternatives, numbered as in the figures, were those that were formally evaluated. Figures 24 and 25 show the results.

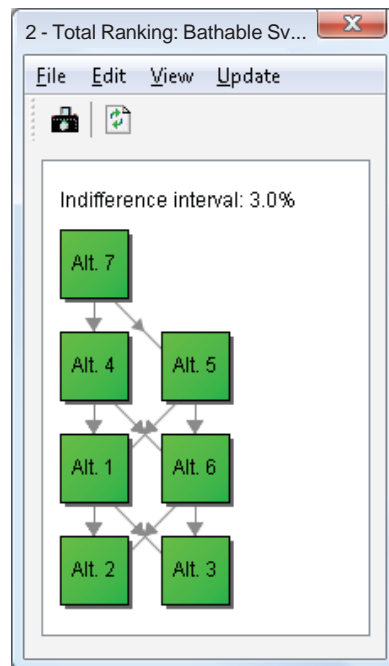


Figure 24. The Svartå alternatives are ranked. Alt. 7 is the best, followed by Alt. 4 and Alt. 5 which can be considered more or less equally good



Figure 25. Svartå comparisons in greater detail. Alt. 7 (agricultural fields) was quite simply the one that best corresponded to the preferences, weightings and expert opinions

How robust was this evaluation given changes to the decision basis? We're talking about sensitivity analyses. Alt. 7 (Establish agricultural fields), Alt. 5 (Establish wetlands) and Alt. 4 (Rectify surface drainage) were completely stable as the three highest-ranked alternatives. It would have required quite large changes in the judgements of the alternatives' goal fulfilment in order to change the suggested ranking, and such large changes didn't arise.

Bob's your uncle! All decided, justified and documented. Neat!

Well, at least if the recommendations were followed. Which they weren't. In 2014, environment inspectors still found *E. coli* bacteria and intestinal enterococci in the river from around 200 buildings that were not meeting the requirements and were given two years to rectify. In 2019, the river was still not bathable and the central Hästhagen bathing area had to be kept closed in order not to put citizens' health at risk, but in 2020 it could finally open.

The *DecideIT* tool can provide you with a wide range of sensitivity analyses. Those analyses will illuminate many aspects and angles of the basis material that underlies a decision made (or to be made). If you are interested in exploring those and getting a feel for a powerful real-world tool, have a look at the page after the index for the aforementioned bonus.

Final Words

There! That really wasn't so perplexing, was it? We've surveyed most of the territory, although we've done so rather quickly and generally, for this was only ever intended to offer our hawk's eye view. We would just like you to understand how reason and objectivity can win out when suitable techniques and methods are embraced. You don't need to shy away, ignore or wash over uncertainty and complexity. All you need to do is observe the terrain, follow its inherent structure and build a shared picture, before finally and accurately homing in on your target and swooping down for the kill. After that, sinking your talons into the choicest alternative is a fair reward, and above all, a true reward for all that flitting around in preparation. The result is visible and understandable. It really is possible to become better at understanding the background and the context and to make better decisions. You just need methods. And methods exist – and tools. Let the guessing games be over and done with. The hawk has landed.

*Deciding in dire situations
by instinct, without hesitations,
and shunning all tools
is a choice made by fools
in denial of the brain's limitations.*

Index

A

architecture 5, 19, 21, 26, 91
 aspect 26, 35, 42, 43, 50, 51, 52, 57, 58, 60, 65, 80,
 110–120, 127, 129, 131–139

B

balanced assessment 71, 91, 94, 112
 balanced scorecards 7
 bathable Svartå 119, 136–141
 belief distribution 108, 109, 110
 benchmarking 29
 BI functionality 1, 7, 17
 BI tools 5, 15
 business environment 14
 business information 3, 12
 business processes 66
 Bypass Stockholm 128–136

C

capture 20, 39, 50, 54, 56, 139
 capturing information 48
 cardinal comparison 107, 136
 communication 27
 complete coverage 120
 compromise 59, 65
 confidence interval 98
 consequence node 61, 92
 consultant dependency 30, 31
 contraction 103, 105, 109, 122
 corporate enslavement 83

critical values 121
 cubes 11, 23, 26

D

data quality 26, 27
 data space 11, 12, 19, 21, 26
 data warehouse 19, 20, 22–28, 31–33
DecideIT 47, 52, 54, 105, 133, 138
 decision basis 37, 48, 49, 53–59, 61,
 65, 84, 120, 129, 131, 133, 141
 decision competence 42
 decision context 57
 decision grids 69, 91
 decision models 37, 85, 97
 decision nodes 61, 62, 92
 decision processes 6, 18, 19, 26, 33–37,
 40, 42, 44, 45, 55, 56
 decision quality 55, 56
 decision rules 69, 70, 73, 75, 79, 81, 84–87, 96
 decision support 3, 8, 34, 131, 137
 decision theory 68
 decision trees 61, 69, 91, 120, 124
 degree of use 31
 deontological theory 82
 determinism 85
 dimension 20, 22, 23, 25
 dominance concept 102

E

egoism 83
 eliciting 31, 116
 eudaemonism 83
 evaluation 54
 expected monetary value (EMV) 94–96
 expected utility (EU) 54, 91–96, 101, 136
 extract, transform, load (ETL) 31
 extraction 23

F

follow-up 6, 7, 8, 15, 18, 20, 21, 33, 43
frequency probability 90

G

goal fulfilment 130, 141
golden mean 16

H

hedonism 83

I

identification 43, 47, 124
implementation 4, 6, 17, 19, 20, 33
improvement potential 43
influence diagram 60–64
interval method 99
intervals 9, 21, 52, 59, 65, 98–102,
104–110, 116, 117, 122, 133
intuition 36, 38, 39, 44, 79, 104
intuitive decisions 38

K

Key Performance Indicator (KPI) 7, 21–25,
58, 59

M

management by objectives 7, 9, 10, 15, 21, 33
maximax 88
maximin 88
metadata 12
minimax 88
minimax-regret 88, 93
modelling 48, 54, 60, 86, 115
monitoring 7, 8, 11, 15, 21, 33
multi-criteria 8, 110, 112, 113
multi-linear expressions 101

N

nepotism 83

O

operational processes 6, 35
operational risk 19
operational systems 3–9, 18, 20, 21, 31, 32
outcome nodes 92, 93, 105

P

perspectives 7, 8, 42, 44, 47, 50, 52, 65, 84, 111,
112, 116, 137
Preference Decision Management Process
(PDMP) 36, 47–54, 127, 128, 137, 138
priority orders 58, 118
probability 13, 51, 52, 62, 63, 68, 69, 80, 89, 90,
93–102, 106, 108, 121, 124
probability distributions 93, 106
probability estimates 52, 69, 99
probability intervals 52, 98, 104, 107, 122
problems collating 27
project management 124

Q

quality assurance 27, 45
quick decisions 43

R

ranking 52, 104, 116, 117, 122, 132, 141
rational decisions 14, 38, 39, 69
rationality 75
reference criteria 116
reference perspectives 116
requirements 4, 14, 16, 32, 49
risk 12–14, 19, 31, 36, 39, 41, 45, 53, 56, 57, 59,
62, 85, 88–91, 94–97, 104, 110–112, 123
risk analysis 53, 110, 124

risk management 123, 124
risk neutral 94, 95

S

screening 43, 44, 46
sensitivity analysis 42, 53, 59, 110, 120, 121, 122
significant decisions 42, 43
situation analysis 43, 44
source systems 23, 24, 28, 31, 32
spontaneous reports 9, 17
staging 23, 24
strategic decisions 39, 42, 43, 59
strategy 5, 15, 32, 36, 56, 88
strength 47, 53, 103, 105, 110
strict uncertainty 85, 86
structural information 108
structuring 33, 43, 48, 50, 54, 137
sub-criteria 119

T

theological theory 82
thresholds 112
tornado diagrams 122, 123
transformation 101

U

uncertainty 13, 41, 42, 53, 59–69, 80,
85–98, 106, 108, 121, 136, 142
uncertainty and risk 13, 85, 89
utilitarianism 83
utility intervals 122
utility theory 95, 96, 114, 115

V

valuations 85, 129
value intervals 52, 107
views 23, 26, 34, 36, 116, 122
visualisation 28, 118

W

weight elicitation 118
weight interval 122
weight intervals 52
weights 65, 91, 94, 113–119, 122, 131, 137, 138

DecideIT Software Licence

A licence for the decision-supporting and decision-analytical software *DecideIT*, seen in Figure 4, is included with this book and the software can be downloaded while stock lasts. The program is a user-friendly decision support tool developed for MS Windows by Preference AB and has been used in Chapters 2–4 of this book. It can handle various aspects of decisions with multiple criteria as well as event trees (probabilities). As you have seen in the book, there is no need to enter precise information in order to receive adequate decision support. Instead, rather vague information can still suffice to find out which decision alternative is the preferred one according to the available data. *DecideIT* has several remarkable properties and features which we now leave to you to explore on your own. There is a small guide on the webpage where you download the software and there is a large built-in help function.

How to install *DecideIT* on a Microsoft Windows PC (Windows 7 and higher):

1. Download the program from <https://www.preference.nu/digitrans>.
2. Follow the installation instructions on the screen.
3. Once installed, click on the *DecideIT* icon on your desktop or program menu to start the program for the first time.
4. The program will ask for your licence key.

<i>DecideIT</i> licence key 971F-B82-06B-1B8-FCF

This book comes with a one-year single-user licence for *DecideIT*. The licence is valid for one year from its entry into the program. The first time you start the program, you are prompted for a licence key. You should enter the key above into the relevant subfields. Once the key is entered, the expiry date is determined by the program and you are good to go for one year.

Happy exploring!

About the Authors

Kjell Borking has a long experience of both the Swedish and international IT industry. Kjell's competence areas are business development, management, project management, databases and decision support. He has been active in the IT industry since 1967 and has started several successful IT companies. In 1995, Kjell started a BI and data warehousing software company, and in 2002 he founded a consulting company that combined management consulting with computer experts in the field of business intelligence.

Mats Danielson is a Full Professor in Computer and Systems Sciences and former Vice President and former Dean of the Social Science Faculty at Stockholm University. He is also a UNESCO Chair Professor. He has a PhD in Computer and Systems Sciences from KTH Royal Institute of Technology as well as degrees in Business Administration and Computer Engineering. He has been working with decision and risk analysis academically and professionally for more than 20 years.

Guy Davies has a PhD from KTH Royal Institute of Technology in Computer and Systems Sciences focussed on specification and modelling. His preferred areas of consultancy include pedagogical and usability issues.

Love Ekenberg was a Full Professor in Computer and Systems Sciences at Stockholm University, a Full Professor of Information Systems at KTH Swedish Royal Institute of Technology, as well as a Guest Professor in Computer Science at Mid Sweden University. He had a PhD in Computer and Systems Sciences, as well as a PhD in Mathematics. He had been working with risk and decision analysis, i.e. development of products and methodologies within these areas, for more than 25 years when he passed away.

Jim Idefeldt is a specialist in methods and processes for risk and decision analysis. He has been responsible for the implementation of decision tools and techniques, and has been working as a consultant in the areas of BI and risk and decision analysis. He also has experience working as a Business Intelligence Manager. He has a PhD in Computer and Systems Sciences with emphasis on risk and decision analysis and has also an MBA in Industrial Management and Engineering.

Aron Larsson is a Full Professor in Computer and Systems Sciences at Mid Sweden University. He has extensive experience in facilitating structured decision processes including modelling, analyses and evaluations of decision problems with conflicting objectives, risks and several stakeholders in both private and public sectors. He is also working with development of decision software. He has a PhD in Computer and Systems Sciences and an MBA, and is also a researcher at Stockholm University.

How on Earth did this happen?

The question is not infrequent, but more to the point, it's unnecessary. It arises because in many cases people haven't thought things through beforehand. Yes, sometimes people can be unlucky, but far from every time that things go awry; bad luck is never so consistent – by definition.

Thinking things through properly beforehand might seem a fairly obvious requirement for making a good decision, but startlingly often, even critical decisions are made without any in depth analysis. Well, there may have been some background data lying around somewhere, but not much is done with it. Though it's easy to blame authoritarian leadership or the like, poor decision making is far more usually due to not really knowing what to do with the available information, whether it's sufficient and what else needs to be known

There's nothing weird or idiotic about today's decision makers, but there often is about the tools and methods available to them. It would be more accurate to say that many decision makers have been lead astray regarding their remit and capacity, but fortunately with structured decision processes they can be guided back on track far more easily than had they been as dim-witted as some of their decisions would indicate.

Businesses introduce various BI solutions both timely and untimely, but what most of them don't have is any kind of methodology by which to deal with decisions. With no processes, knowledge, techniques or tools, they soon come to grief. Happily this is fairly easy to remedy.

Decision processes and decision methods can be greatly improved. In this book we simply explain how to go about it.