1

Introduction

1.1 Decisions, decisions, decisions!

I really do hate making decisions. Life is so full of them. This evening I was fed up, and I told Thomasina how I felt. 'There are always choices to be made,' I complained. 'Whether to eat the meat that has been put out or to try the new crunchy things they've bought. Whether to go out mouse or shrew hunting. Whether to sharpen my claws on the oak tree or sycamore tree. And so on. Decisions. Decisions. Decisions.'

'You need to rest and relax,' said Thomasina.

'How?' I demanded.

'Go and lie down somewhere,' said Thomasina.

'Where?' I asked.

'Underneath the willow tree,' she replied. 'Or on the window seat. On the sofa. Or underneath the garden bench.'

Dear old Thomasina. She means well. (Vernon Coleman)

We make decisions all the time: whether to take the lift or the stairs; whether to buy a new car; whether to contradict our boss's latest edict; and so on. Some decisions have so little impact on our lives that we take them without much, if any, thought. Others have much greater potential impacts, and we reflect and deliberate upon the alternatives before choosing one. Some decisions are personal, some professional. How do we make decisions? How should we make them? Are we naturally good decision makers (DMs)? Can we learn techniques to improve our decision making? Can we develop computer programmes – *decision support systems* (DSSs) – that embody such techniques? These questions are essentially the ones that we address in the following chapters. We focus on the more significant of our decisions: whether to have a tea or a coffee can safely be left to whim.

No two situations that call for a decision are ever identical. They differ because a decision changes both the world and the DM in some small way, and neither can ever go back to the previous status quo. There are many other ways in which decisions differ, however: the context of the problem, the abilities, skills and dispositions of the people involved and the social context in which they find themselves (see figure 1.1).

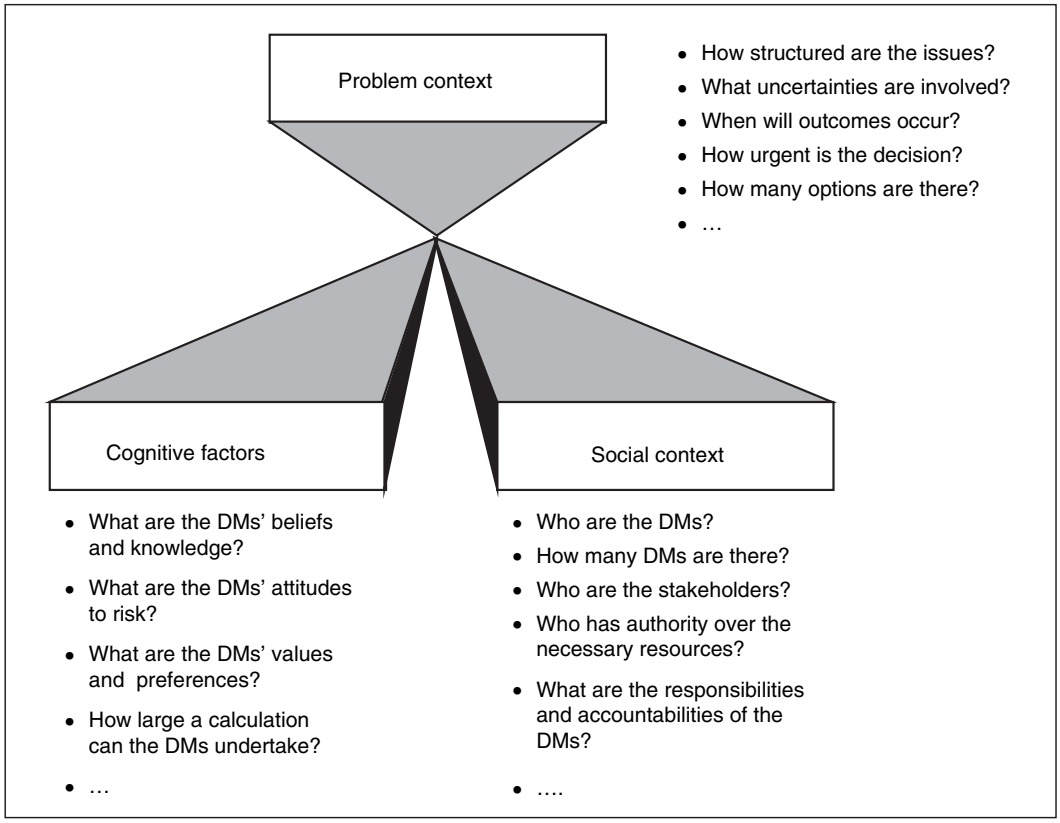

Figure 1.1 Factors that affect decision making
Source: Payne *et al.* (1993).

Our purpose in this chapter is to raise issues, introduce general terminologies and indicate topics that we cover in later chapters. We begin by considering a broad categorisation that will give shape to much of our discussion.

1.2 The strategy pyramid

You've got to be very careful if you don't know where you're going, because you might not get there. (Yogi Berra)

Perhaps the most commonly discussed distinction between decisions is that between strategic, tactical and operational decisions – the so-called *strategy pyramid* (see figure 1.2). *Strategic* decisions set the goals for an organisation or an individual. Mintzberg (1992) suggests that a strategy provides five P's: a *plan* for future action; a *ploy* to achieve some end: a *pattern* of behaviour; a

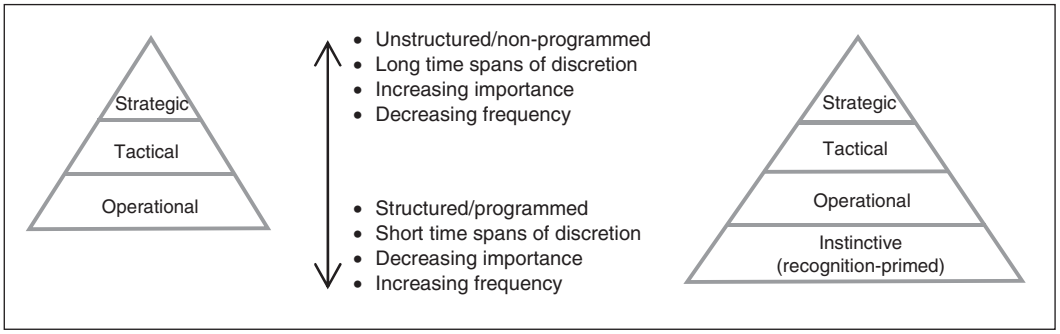

Figure 1.2 The strategy pyramid

position defined by goals and values; and a *perspective* on how to view the world. A strategy sets the direction and a broad framework in which more detailed decisions can be taken. *Tactical* and *operational* decisions fill in those details. Thus, a retail company might make a strategic decision to expand into a new region. It would then need to decide tactically in which towns and shopping malls it should establish itself and in what order it should open these. Operational decisions would develop and run the necessary supply chains, financial systems, staffing, etc. Similarly, a personal strategic decision might concern a career direction and be followed by operational and tactical decisions on where and for which company to work, how hard to strive for promotion, etc.

Simon (1960) notes that strategic decisions tend to be associated with *unstructured* or *non-programmed* problems. Seldom do DMs, such as a board of directors, come to a strategic issue with a straight choice between, say, various acquisitions. Rather, they first become aware that the company may need to grow. Through discussion, they formulate their objectives and the possible strategies they might follow. Only then do they have a strategic decision to make. In contrast, operational decisions are usually much more structured – for example, should an inventory level be increased to support a production plan or in what order should the production of various items be scheduled? Another concept, which correlates well with the unstructured/structured dimension, is that of the *time span of discretion* (Jacques, 1989). Roughly speaking, this relates to the length of time before the consequences of a decision have their full impact. The longer the time span of discretion the more unstructured and strategic the decision is likely to be.

The original 'three-level' strategy pyramid on the left of figure 1.2 misses an important type of decision. In many cases, DMs seem to match the current circumstances to something similar that has happened in the past and do roughly what they did then – or perhaps what they thought after

the event they should have done. In such *recognition-primed* decision making (Klein, 1993) there is little or no comparison of options, just an instinctive choice of action. Therefore, we extend the strategy pyramid to include a fourth level. The term 'programmed' fits well with the idea of instinctive decision making based upon recognising that the current situation is familiar and that the action proven to be successful in the past is appropriate. Situations are rarely identical, however, so DMs often simulate how the usual action will play out in the new situation and what small modifications are necessary to increase its effectiveness. This form of decision making is common among experts who regularly make very similar kinds of decisions, such as surgeons deciding on how to suture a wound, bank managers deciding whether to extend a loan or fire chiefs deciding how to tackle a fire in a building.

Within the discipline of artificial intelligence (AI) much effort has been expended on developing knowledge-based decision support systems (KB-DSSs), which seek to 'automate' decision making. These tools operate at the lower levels of the strategy pyramid precisely because they need training – i.e. they need to be provided either with a set of rules that tells them how to recognise and react to different types of situations or they need data on how experienced DMs reacted in the past. One of AI's research objectives is to develop KB-DSSs that need less training and operate at the highest levels of the strategy pyramid. For the present, however, machines able to think strategically and creatively in unstructured, novel situations belong to the realm of science fiction; we discuss this topic further in chapter 5.

We note that the importance of a decision increases the further up the pyramid we go – i.e. the potential consequences of strategic decisions are much more significant than those of instinctive ones. Conversely, the frequency with which a decision – or, rather, type of decision – is faced increases towards the base: operational and instinctive decisions are much more common than strategic ones.

Jacques (1989) argues that the tasks and decision making undertaken by staff at different levels within an organisation may be characterised by the longest time span of discretion required by their roles. Jacques' theory is a mixture of the descriptive and normative – i.e. it includes observations of how organisations *are* structured and reflections on how they *should* be. In many empirical studies he has shown that the concept of the time span of discretion provides a useful explanatory tool. He goes further, however, and argues persuasively that organisations are best able to achieve their objectives when members of the organisation work at levels with time

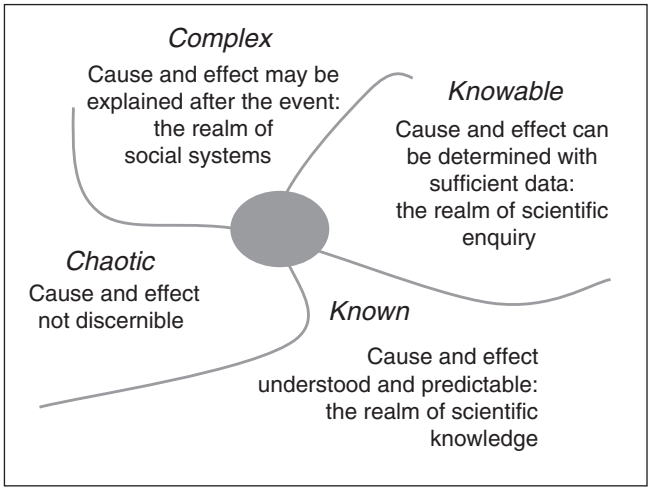

Figure 1.3 The cynefin model
Source: Snowden (2002).

spans of discretion within the limits of their ability to envisage the future. He terms such organisations *requisite*.

In his empirical studies, Jacques distinguishes four domains of activity:

- the *corporate strategic* domain, which sets the guiding values and vision and develops strategy to take the organisation towards these;
- the *general* domain, which develops an implementation plan for strategy;
- the *operational* domain, which organises the detailed delivery of the strategy; and
- the *hands-on work* domain, which delivers the work.

Note how these domains map onto the four levels (strategic, tactical, operational and instinctive) of the extended strategy pyramid (figure 1.2).

In the context of knowledge management, Snowden (2002) has argued for a further typology of decisions: the *cynefin*[1] model (figure 1.3). For decision contexts in the *known space*, cause and effect are fully understood. We know everything about the circumstances underpinning the choice, what the alternatives are and what consequences might arise from each. In the *knowable space* cause and effect relationships exist, but there are insufficient data immediately available to make immediate firm forecasts of the consequences of any action. In the *complex space* there are so many interacting causes and effects that predictions of system behaviours are

[1] *Cynefin* is Welsh for 'habitat', although it does not translate quite so simply into English: the word includes the cultural and social as well as the environmental aspects of habitat.

subject to considerable uncertainty. Indeed, the range of actions available may be very unclear. Typically, such complexity arises in social systems. In the *chaotic space* things happen beyond our experience and we cannot perceive any candidates for cause and effect.

Snowden suggests that decision making in the known space tends to consist of recognising patterns in the situation and responding with well-rehearsed actions: recognition-primed decision making. In the knowable space, there is more analysis than recognition, as the DMs learn from the available data about the precise circumstances faced. In statistical terms they need to fit general models to the particular data of the current situation. In the known and knowable spaces, situations are *repeatable*. Essentially identical[2] or extremely similar situations have occurred in the past and the DMs have learnt the underlying cause and effect relationships. Moreover, they have learnt what they would *like* to happen: experience has clarified their preferences and values so that they have clear objectives, often so clear that these are no longer explicitly articulated. Such repeatability is the bedrock of empirical science: see, for instance, our discussion of frequentism in the foundations of probability in section 8.2.

In the complex space the DMs' knowledge is poor: there is much less perceived structure. There are simply too many potential interactions. Situations are so different as to be unique. Analysis is still possible, but its style will be broader, with less emphasis on details. Decisions will be based more on judgement than objective data, and the emphasis will be on developing broad strategies that are flexible enough to accommodate changes as the situation evolves. Before making decisions there may be a need to pause and clarify objectives – i.e. for the DMs to reflect upon how their general preferences and values apply in the current context. Decision making in the chaotic space cannot be analytical because there is no concept of how to break things down into an analysis. The DMs will simply need to take some action and see what happens, probing until they can make some sort of sense of the situation, gradually drawing the context back into one of the other spaces.

Thus, in a sense, the structured/unstructured dimension of decision making curves around from the known to chaotic spaces in the cynefin model (see figure 1.4). Indeed, in many ways the cynefin model adds little to the earlier writings of Simon (1960, 1978) and others. What it does provide, however, is an intuitive representation of the ideas that focuses

[2] No two situations can be entirely identical, by virtue of their different location in space and/or time.

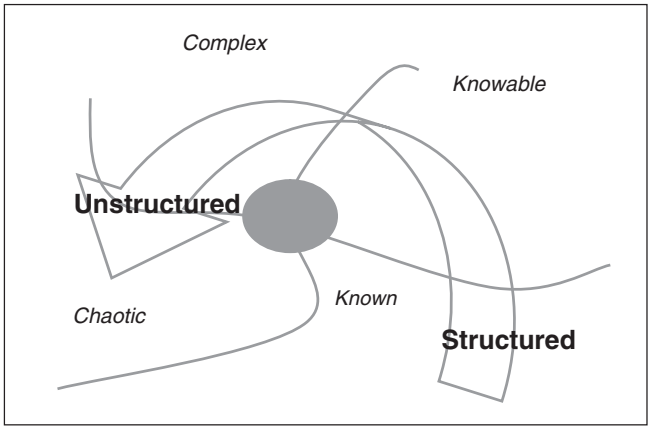

Figure 1.4 The cynefin model and the structured/unstructured dimension of decision making
Note: Compare with figure 1.2.

attention on the knowledge and information available to the DMs. The cynefin model also allows the dynamics of a sequence of decisions in a changing environment to be represented usefully; in other words, as knowledge and understanding of the situation changes, one moves into a different quadrant of the model, suggesting the need for a different form of decision making (see French and Niculae, 2005).

Note that there is much consistency here: the strategy pyramid, Simon's structured/unstructured dimension, Jacques' concept of the time span of discretion and Snowden's cynefin model essentially capture very similar ideas.[3] Each offers a subtle alternative perspective, however, that informs our understanding of the differences between decision contexts.

1.3 Rationalistic versus evolutionary strategic decision making

> Most discussions of decision making assume that only senior executives make decisions or that only senior executives' decisions matter. This is a dangerous mistake. (Peter Drucker)

It is tempting to think that decision making is nicely ordered. First one makes strategic decisions to set context, values and direction; then it is the turn of tactical decisions, to map in the details; operational decisions to allocate resources and manage the work come next; and, finally, the

[3] Indeed, Snowden (2002; see also Snowden and Boone, 2007) also uses the ideas of cynefin to discuss other issues, such as organisational culture and leadership, and in doing so captures many of Jacques' views on these topics.

work itself is driven by almost unnoticed instinctive decision making. This chronologically ordered, logical perspective is, however, more often than not a complete fiction! Decision making is driven by events and developments in the external world as much as by some logical internal rationality. Thus, members of an organisation at any level may change what they do in order to gain some advantage, because the current situation has changed or is not as predicted. They might recognise some inefficiency or, more positively, the opportunities offered by some new technology. They might recognise a change in the behaviour of their customers and respond to that. Slowly a myriad of small changes at the operational or tactical levels can lead to quite significant changes at the strategic level: perceptions of values and stratetegic direction evolve. Such *emergent* strategic development is common in many – all? – organisations (Clarke, 2007; Mintzberg, 1987). Indeed, in our personal lives many major changes come about by events and happenstance rather than conscious decisions.

Our view is that both emergent and rationalistic decision making exist in all organisations. Strategy emerges and evolves in periods of 'calm'. Generally, all changes at different levels in the organisation are roughly, but not perfectly, aligned with an overall planned strategic direction; there is no careful, comprehensive analysis, however. Then some event or a growing awareness that some aspect of the organisation's business is not going as well as it might stimulates a period of reflection and deliberation on some or all parts of the strategy. At such times rationalistic thinking on strategy comes to the fore, decisions are made and a new strategy is adopted. Figure 1.5 illustrates this by 'funnels' of rational thinking

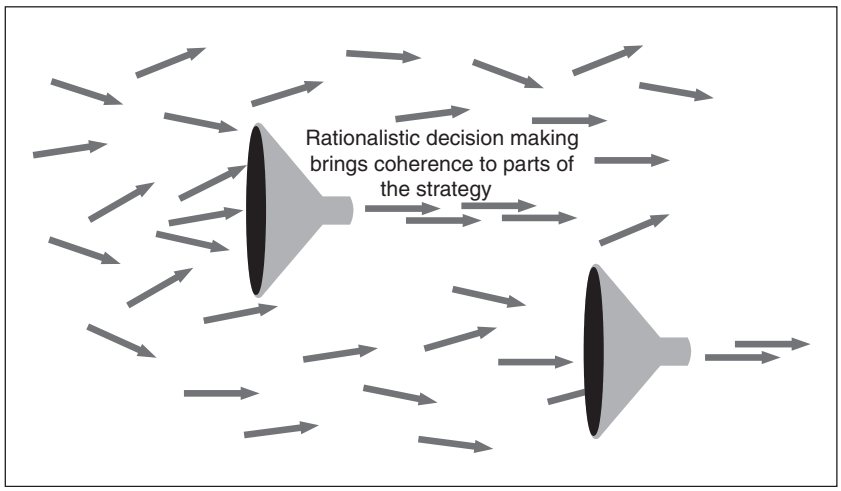

Figure 1.5 The interplay between rationalistic and emergent strategic decision making

that bring coherence, aligning many, but seldom all, of an organisation's activities with a single direction. The funnels may operate at any level of the strategy pyramid and are not necessarily the same size, and some may be much larger, pulling most activities together and reflecting a major strategic deliberation.

In the following chapters our goal is to suggest how best to support these periods of rationalistic decision making. Nevertheless, while that is our emphasis, we are well aware that the context for such support will more often than not have been set by previous emergent strategy.

1.4 Players in a decision

If the people who make the decisions are the people who will also bear the consequences of those decisions, perhaps better decisions will result. (John Abram)

Notionally, the simplest decisions involve just one person: the DM. She[4] provides all the expert knowledge necessary, expresses her own judgements, performs her own analyses and makes her own decisions. In practice, however, this seldom happens. More often decisions are the responsibility of a group of DMs, such as a management board or a government department. They might involve others. They will probably work with accountants, scientists, engineers and other subject experts in order to acquire relevant information. Thus many will contribute to the process that leads to a choice and be a party to the decision making.

The *decision makers* are responsible for making the decision: they 'own the problem'. To be able to take and implement a decision, DMs need to hold the appropriate responsibility, authority and accountability.

- *Responsibility*. Individuals or groups are responsible for a decision if it is their task to see that the choice is made and implemented.
- *Authority*. Individuals or groups have the authority to take a decision if they have power over the resources needed to analyse and implement the choice.
- *Accountability*. Individuals or groups are accountable for a decision if they are the ones who take the credit or blame for the decision process and for the choice that is made, how it is implemented and the final outcome of that choice.

At various points in the decision process, responsibility may pass between different groups of DMs. When this happens, it is very important that

[4] We refer to an individual DM in the feminine and, shortly, the decision analyst (DA) in the masculine, creating a natural contrast in our language.

the appropriate authority and accountability are also passed across. When responsibility, authority and accountability do not pass between groups in a coherent fashion, there is an obvious danger that the decision-making process becomes dysfunctional.

The DMs are accountable to some, but not necessarily all, of the *stakeholders*. Stakeholders share, or perceive that they share, the impacts arising from a decision. They have a claim, therefore, that their perceptions and values should be taken into account – and in many cases they are. The DMs are stakeholders, if only by virtue of their accountabilities; but stakeholders are not necessarily DMs. The obvious stakeholders in a business are its shareholders or partners, but there are many others – e.g. employees, customers, unions, suppliers, local communities. In the public sector, the government and its agencies generally have many stakeholders, such as the public, industry, consumers or political parties; and accountability is correspondingly much broader.

Experts provide economic, marketing, scientific and other professional advice, which is used to formulate and understand the problem and assess the likelihood of the many eventualities that will affect the decision outcome. We often adopt the classical use of the term 'science' and use it to refer to a broad range of human knowledge. The knowledge that experts impart is used in the modelling and forecasting of outcomes of potential decisions. The DMs may have advisers who undoubtedly are experts in this sense, but they are unlikely to be the only experts involved. Other experts may advise some of the stakeholders, informing their perceptions and hence influencing the decision making.

Analysts develop and conduct the analyses, both quantitative and qualitative, that draw together the empirical evidence and expert advice to assess the likelihood of possible outcomes. They work with the DMs to clarify and elicit their uncertainties and values. They will also be concerned with a synthesis of the stakeholders' value judgements. These analyses are used to inform the DMs and guide them towards a balanced decision, reflecting the various expert and stakeholder inputs and the emphases that the DMs wish to give these. Whereas experts support decision making by providing information on the *content* of the decision, such as relevant economic data, the assessment of physical risks or whatever, analysts provide *process* skills, helping to structure the analysis and interpret the conclusions. For this reason, analysts are sometimes referred to as process experts.

Figure 1.6 offers a simplified representation of the interrelationship between experts, stakeholders, DMs and analysts. This separation of roles is, of course, very idealised. Some parties to a decision may take on several