Part I

Management Science

1 What Is Management Science?

1.1 Introduction to Management Science

What is Management Science (or a closely linked term, as we said at the start, "Operational Research" – or "Operations Research" outside Europe)? Well, that is a difficult question, and one that we'll spend this chapter investigating. As we said in our introduction, we can start with the beginning of the definition that appears both on the website of the Operational Research Society in the UK (www.orsoc.org.uk) and on the website set up by the American Operational Research/Management Science organization INFORMS, which is devoted to looking at what OR is and how it can help (www.scienceofbetter.org./). This says: "In a nutshell, operations research (O.R.) is the discipline of applying advanced analytical methods to help make better decisions" (INFORMS, 2006). So, Management Science/Operational Research is aimed at those who have to make decisions, or understand situations better, in industry, commerce, government, defence – in a whole multitude of arenas. And it uses a panoply of analytical methods to better understand those decisions or situations in order to help those decision-makers.

Operational Research, as it was originally known, has been around under that name for over 60 years. An excellent history of the field by Kirby and

Capey (later extended into a major work (Kirby, 2002)) begins: "At the end of the Second World War, operational researchers could congratulate themselves on their substantial and, on occasion, decisive contributions to the allied war effort in a number of theatres. In the North Atlantic, for example, they had assisted in the defeat of the U-boat weapon by devising, inter alia, optimal convoy tactics, the most effective settings for depth charges, and efficient servicing schedules for long-range aircraft . . ." (Kirby & Capey, 1998). These successes led to OR (as it became abbreviated) being taken up enthusiastically by many organizations, most notably in the UK by the National Coal Board and organizations within the iron and steel industry, where it scored major successes in its ability to analyse complicated problems and both facilitate rational decision-making and enable increases in efficiency to be made. This era saw the birth of the subject as a postgraduate degree in various universities, and a gradual move thereafter into undergraduate education.

As management in the 1970s looked to put their organizations on a more "scientific" footing (Locke, 1981) and the abilities of scientific computing (which Operational Research could build upon to substantially increase its power) increased exponentially, the scope for Operational Research could have been expected to have increased dramatically. But the subject faced a dilemma, and to some extent a division, the effects of which can still be seen today. The field of Operational Research, as originally interpreted – using mathematical and analytical techniques to solve well-defined problems, based on concrete measurable entities – continued, and is alive and well, particularly in the US (as "Operations Research"). But many practitioners saw that the real problems that they wanted to solve were less well-defined, and included variables and entities that could not be easily defined and measured unambiguously, and issues that needed wider exploration. That is not to say the subject split, but there were clearly different directions for the subject. (Some would call the latter practitioners "Management Scientists", considering that "Operational Research" tends to have a more mathematical flavour than "Management Science". However, as is often the case in these types of issues, there is no agreement on these labels, and many commentators would not recognize this distinction. INFORMS, to which we have referred previously, was actually formed from a merger between the Operations Research Society of America (ORSA) and The Institute for Management Sciences (TIMS) (also based in the US), but the motivation for setting up TIMS, although partly founded on the subject areas that might not flourish within ORSA, was also based on the perceived predominance of military work and membership policies of ORSA (Lathrop, 1957).)

The problems with the Operational Research approach were most famously espoused in two papers by Russ Ackoff. In the first (Ackoff, 1979a), he declared that the underlying paradigm of current Operational Research, which he termed "predict and prepare", was not suitable for modern organizations, which do not simply passively respond to the environment but actively engage with it – and this is a topic we'll return to over the first two chapters of this book. But the attempt by mathematicians to model these messy situations was "equated by managers to mathematical masturbation and to the absence of any substantive knowledge or understanding of organizations, institutions or their management". In his second paper (Ackoff, 1979b), he gave his proposals for the future of OR. He talked about making Operational Research participative and "based on planning with people, not for them" (taken from Pidd's (2001) summary, where he explains how this echoes Mintzberg's later analysis of big corporate planning exercises that are left unimplemented, and also a major critique of what came to be called in the mid-twentieth century "scientific management" or "Taylorism", that of separating planning from action). Again, we'll come back to this, particularly in Chapter 2. But Ackoff also explained how this would require a change in Operational Research practice and education, from applying mathematical techniques in a vacuum to a practice in which "understanding how people worked in organizations was fundamental. Thus mathematics becomes the servant rather than the master."

This led to a reawakening of Operational Research, particularly in the UK, as it came to grapple with the nature of the situations within which it was now being asked to intervene. One of the important influences in the UK was Eden, who a few years after Ackoff's paper explained that "although our heritage is the 'application of science', it is a narrow vision of science which has dominated the profession . . . it has largely been the application of applied mathematics, statistics, and computer sciences. Thus, while the definition of OR in the UK is not narrow, the practice has proven to be so" (Eden, 1982). Repeating Ackoff's quote above, he continues: "The emphasis on this particular sort of science has meant that the profession has recruited predominantly from disciplines that characteristically attract convergent, rather than divergent, thinkers. A reinforcement of the profession is thus made up of those who believe in a form of objectivity that is a poor match for the realities of organizational life" (Eden, 1982). A defining moment of this reawakening came as the UK Operational Research Society set up a Commission to look into the practice and future of Operational Research (reported in Mitchell, 1986), which described OR practice with significant amounts of structuring messy problems, and "little explicit use of those mathematical techniques which are most commonly associated with OR (for example, mathematical programming and queuing theory). These and the insights they offer, along with many other technical devices, help constitute a tool-kit from which the practitioner may draw as need dictates. . . . Better tools may give better results, allow more jobs to be done and save on tedious work, but they seem to affect underlying methodology only slightly. Methods are renewed and extended but the essential methodology persists. The main methodological drive, as inferred by the Commission, is pragmatism. . . . [The means used] usually entail working closely with the client or his client, and almost continuous negotiation, stage by stage, of how the work should be moving."

Here I must, as author, declare my bias, and the bias in this book. I was trained as a mathematician and then as a traditional Operational Researcher in the 1970s in one of the first Operational Research Master's courses in the UK. And I firmly believe that mathematics, separate from the analysis of "real" social systems, has an important role in certain isolated areas. But for the main role of Management Science, that of taking real-world practical situations and using our analytical and modelling skills to bring greater understanding and better decision-making, mathematics divorced from reality is not helpful. The best exposition of the path that I have been on I think is summed up in one of my favourite books, *The Glass Bead Game* by Hermann Hesse, which was first published in 1943 (see Hesse, 2000). This book described an almost magical game: "All the insights, noble thoughts, and works of art that the human race has produced in its creative eras, all that subsequent periods of scholarly study have reduced to concepts and converted into intellectual property – on all this immense body of intellectual values the Glass Bead Game player plays like the organist on an organ." This wonderful game was attractive and satisfying, just as the mathematics of Operational Research can be to the experienced exponent. Indeed, in the Glass Bead Game, "Mathematicians in particular played it with a virtuosity and formal strictness at once athletic and ascetic. It afforded them a pleasure which somewhat compensated for their renunciation of worldly pleasures and ambitions." The game is almost weirdly like the attempt by Operational Researchers/Management Scientists to capture phenomena in analytical form, as Hesse describes: "Men like Abelard, Leibniz, and Hegel unquestionably were familiar with the dream of capturing the university of the intellect in concentric systems, and pairing the living beauty of thought and art with the magical expressiveness of the exact sciences." The people who played this game mostly "lived in a state of political innocence and naïveté such as had been quite common among the professors of earlier ages". They did not let the issues of "real-life" get in the way of their game, much like some modern-day Operational Research professors. But the hero of *The Glass Bead Game*, who rises to become the *Magister Ludi* – the most expert and most important player of the game – by the latter stages of the book comes to see that the whole point of such learning is to apply it in the real world and to bring useful benefit to the world. In *The Glass Bead Game*, this meant leaving the mountain where the game-players lived and moving into the world to become a humble teacher. For Management Scientists, the same understanding means that we must sacrifice beautiful, neat mathematical treatises for work that is practical and makes a difference to the world. (A colleague who read through the draft of this book pointed out that the hero of *The Glass Bead Game* dies after his first intervention in the real world – being useful in the world might not be the most comfortable path!)

When a group from a major business school in Europe, INSEAD, looked at the practice and literature of OR in 1993 (Corbett & van Wassenhove 1993), they concluded that the OR community and the management community didn't take much notice of each other's literature; OR tools were being used, but there was an expanding gap between OR publications, which were mainly theoretically orientated, and management publications addressing real practical needs (something I as an editor of a leading OR journal have been trying to change!). To take just one instance, the field of project management is one where there is a whole literature of complex mathematical works purporting to model "real" situations, but whose analysis is virtually never used or applied in actual real cases (Williams, 2003). So, Operational Research/Management Science is truly "the beautiful game" but it needs to bring benefit and use to the world, not just indulge in self-gratification (to take Ackoff's distinctive metaphor). And it is to avoid being tarred with this latter brush that this book takes the stance of the Management Scientist rather than that of the Operational Researchers – albeit these actual terms are fairly interchangeable nowadays.

> This book is about "real world" Management Science: modelling situations we find in practice and trying to bring rationality and analysis rather than simply "nice" intellectually satisfying models.

1.2 The nature of problems

So, what are the problems that the Management Scientist investigates? The word "problem" here is surprisingly difficult to define. Mitchell (1983) has a go at this by saying that "The word 'problem' is used to describe many situations in which an individual finds himself. These situations have three common threads.

(i) The individual is dissatisfied, or surprised, by what is happening. He might believe that his circumstances must be capable of improvement, or his understanding is at fault, or rules or beliefs which he has held inviolate are changing.
(ii) The individual believes he can and/or should respond, by action or by revising some or all of his beliefs.
(iii) The individual does not find it obvious what action he should take, or even if any is available, or how he should revise his beliefs.

These features of problems are subjective. They suppose that an individual perceives some normal state for himself which satisfies him and offers no surprises. . . ."

This certainly does define a problem situation, to which we can relate in our everyday life: we wish to change our job, or we need to work out the best airline routes to book, or the best ways to manage our finances. But (as Mitchell does go on to explain) the organizational problems with which Management Scientists deal are not as simple as that. And, indeed, the very word "problem" (with its associated assumptions that we can "solve" a problem and then we've finished) does not fully describe what it is that Management Scientists face in their work. The classic description of the various types of concepts in the area of problems was given, for example, in Pidd's (2003) illustration shown in Figure 1.1, which leans heavily on Ackoff's (1974, 1979a, 1979b) pioneering work. In this, the top axis is taken to be three points on a spectrum.

At the left-hand end of the spectrum are "puzzles". In these, there is no ambiguity about the formulation of what needs to be solved; the issues and options are clear, and the answer is unarguably the correct one. These require logical thought, and sometimes aren't easy to solve, but are not the domain in which we want to model. We come across such "puzzles" frequently: the popular

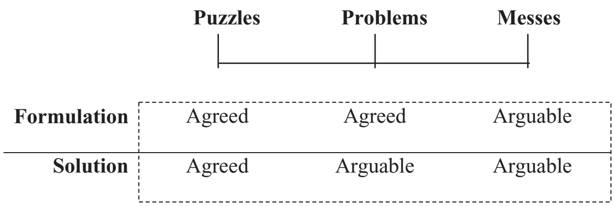

Figure 1.1 Puzzles, problems and messes. From Pidd. Copyright 2004 John Wiley & Sons Limited. Reproduced with permission.

"Sudoku" game would be an example, but so would many examination questions given to undergraduate students.

The nature of the real world is such that puzzles can only be hypothetical – in reality, issues and options are never perfectly clear. But as clarity decreases while we move along to the right-hand end of Pidd's spectrum, there is a range of real-life situations. In the mid-point of the spectrum comes what Pidd simply calls "problems", epitomized by the well-defined OR problems that were tackled in the 1960s. Here the formulation of the problem is usually pretty well agreed, but there are a variety of approaches to solving it. Pidd takes the seemingly straightforward question "how many depots do we need in order to provide daily replenishment of stock in all our supermarkets?", and then goes on to show how even this seemingly innocuous and unambiguous question can itself be questioned, and various aspects explored. Experience shows that two analysts will rarely have the same approach, although many approaches *may* tend towards similar answers. But, essentially, "problems" were what the Operational Research world started tackling, and where some mathematicians have got stuck. And many of these are important: "How do we schedule our railways?"; "how do we optimize loading of ship-containers so that we use the minimum number of containers?"; "how much stock should we order?" But such well-defined questions are fairly rare; and frequently, when they do arise, they arise only as the mid-point after the analyst has undertaken considerable problem structuring.

As OR proved adept at tackling a range of problems, it was faced with the requirement to tackle the many situations in which there is a lot of ambiguity, no agreement about the issues, or about concept relationships, or about what is going on, or whether a solution exists at all. These situations are termed "messes" by Ackoff. These are the sort of situations into which Management Scientists are often called, where the definition of the problem itself is not clear or agreed upon, let alone how to approach it. In such situations, an analyst cannot move straight in with his/her toolbox of mathematical tools and start modelling – the situation must be defined, structured, agreed and made amenable to analysis (Pidd talks about "taming" the messes). This is the "mess" – and the competencies needed to tackle such messes are quite different from the divergent problem-solving competencies needed to tackle well-defined traditional Operational Research problems. But these are the situations Management Scientists generally face: "it's all a mess – why?", or, "it's all going wrong – what should we do?", or "what are our options, and how do we choose?", or "I think I know what the problem is but the CEO thinks it's something entirely different".

This last phrase captures a further reason why our "messes" are so "messy" – the sort of situations into which Management Scientists are called usually involve groups of participants, indeed sometimes groups of decision-makers. If a problem for an individual, as defined by Mitchell above, is centred on an individual's beliefs, or rules or perceptions, then in a group – which is likely to have different objectives, beliefs, rules, or maybe even perceptions – then clearly the issues are going to be compounded. This was at the heart of Operational Research's identity problems in the late 1970s/early 1980s we discussed earlier: as Eden and Sims (1979) put it: "There is a major difference in emphasis between discussion within the profession and the behaviour of the OR consultant working within a social organization. In the first instance we see, within textbooks, journals, project reports, proceedings of conferences, an attention to the problem – its characteristics, structure, content – that is to say our attention is directed to an objective reality, a system of interacting variables that as a consequence of manipulation could be made to behave differently. The description of the problem is implicitly contained by the form of the model which is used for its solution; a solution is discovered (usually through mathematical/numerical manipulation) which will enable the system to operate in a preferred manner. . . . However, when we study what is going on in the process of behaving as a consultant we see a part of OR practice which we apparently feel unable to discuss, reflect upon, or theorize about. We see a complicated drama unfold which involves power, influence, negotiation, game playing, organization politics, complex social relationships with real people not merely office holders. In this environment problems are not self-evident at all; under the guise of the same problem title each actor sees a reality which is unique to him, that which comes to be known as the *real* problem for the consultant depends then upon his *own* reality and that belonging to those actors to whom he chooses to listen."

Because of this effect of the organization on our approach, some have taken up the simple classification given by Roth and Senge (1996). They take firstly the underlying complexity of the problem situation itself, which they call "dynamic complexity". More specifically, they say that dynamic complexity characterizes the extent to which the relationship between cause and the resulting effects are distant in time and space. This mirrors Simon (1982), who says that a complex system is essentially "one made up of a large number of parts that interact in a non-simple way. In such systems the whole is more than the sum of the parts, not in an ultimate, metaphysical sense but in the important pragmatic sense that, given the properties of the parts and the laws of interaction, it is not a trivial matter to infer the properties of the whole." But as well as this, Roth and Senge (1996) also take the complexity of the group effect, which they call "behavioural complexity": "Behavioural complexity characterizes the extent to which there is diversity in the aspirations, mental models, and values of decision makers." In political situations with high behavioural complexity they suggest as an example Jews and Palestinians in Gaza; as an example of low behavioural complexity they suggest a group of financial analysts solving a technical problem. They then propose a simple matrix as shown in Figure 1.2.

As we have seen, "tame problems" can be treated by ordinary, traditional Operational Research methods. "Wicked problems" are those where "complex underlying social realities are inescapable, and different groups of key decision-makers hold different assumptions, values and beliefs which are in opposition to each other." Messes we have already defined. Roth and Senge (1996) go on to say: "The research traditions that deal with behavioural complexity . . . and dynamic complexity . . . have remained largely separate. What

		Dynamic complexity	
		Low	High
Behavioural complexity	Low	Tame problems	Messes
	High	Wicked problems	Wicked Messes

Figure 1.2 Types of problem. From Roth and Senge. *Journal of Organizational Change Management* 9, Issue 1, page 92-106, Copyright 1996 Emerald Group Publishing Limited all rights reserved. Reproduced with permission

befuddles organizational decision making is that the two coexist and interact in what we have termed 'wicked messes' – needing new research methods and new synthesizing of the old methods." This is the area in which Management Scientists often find themselves, and the area we shall consider as we go through this book.

As we have moved away from treating tamer problems to dealing with wicked messes, Operational Research / Management Science has had to change also. One of the key texts in describing this movement is Rosenhead and Mingers (2001). They describe "traditional Operational Research" as dealing with the following:

- A single problem, which can be formulated with a single objective.
- A single person who has to weight the situation and make a decision about the problem: the "decision-maker". That decision-maker has objectives that could be defined. Furthermore, we could regard concrete decisions coming from that decision-maker through a chain of command to the workforce, who would carry out these actions.
- The problem is clear and unarguable, and we could assume consensus about what the problem was.
- Decisions are to be made based on these abstract objectives and then actioned; people in the analysis can be regarded simply as passive objects.
- Finally, they say that the aim of the analysis was to minimize or even abolish uncertainty about the future and as far as possible pre-take future decisions, to map out the future path.

The modern "alternative paradigm", with which they say we need to approach problems, recognizes:

- that there is a search for alternative solutions acceptable to a number of participants on separate dimensions, rather than a single all-purpose optimization;
- that the analysis and models are simple and transparent in order to be comprehended and therefore bought-into by the various participants;
- that the analysis needs to (both pragmatically and ethically) conceptualize people as active subjects;
- that the analysis needs to facilitate planning from the bottom up as well as, or instead of, from the top down; and, finally
- that we need to accept uncertainty, the aim of the analysis being to keep options open.

Now this argument covers a whole gamut of issues, but we can see the issues of multiple views of the problem and reality, and the need to deal with groups of decision-makers, with models that are transparent so as to keep them “on side” or “bought in”. More of these issues will arise during the course of this book.

This book will cover “messy” situations, in which not only is it unclear how to solve “the problem” but also it is not clear what “the problem” is. Since these problems are situated in real human situations, they will often be “wicked messes”, in which perceptions, values, assumptions, even the underlying understanding or “reality” vary between the stakeholders.

1.3 The Management Science approach

If ‘What is Management Science?’ is a question that has not been answered satisfactorily, even more so does the question “What is the Management Science approach?” not have a uniformly agreed answer. Different authors give different answers (again with Operational Research and Management Science put together).

- Some answers are technique-based: “Management Science uses mathematical and computer techniques” or “Management Science uses mathematical techniques such as linear programming”. And it is true – some Management Science does use such techniques, and we’ll have a look at many of these techniques in Chapter 5. But this does not define the heart of Management Science.
- Some answers define the Management Science approach in terms of other approaches: “Management Science uses the System Analysis approach” or even “Management Science uses the scientific approach” – the former being too restrictive, and neither being helpful to tell us what this approach is.
- Some define Management Science in terms of quantification: “Management Science measures things”; “Management Science looks to optimize systems”. Again, this is often the case, and Chapter 7 will look at issues of measurement and quantification. But there is Management Science that structures and does not quantify; and there is Management Science

that explains but does not offer solutions, let alone "optimum" solutions.

Mitchell (1983) comes down to the most inclusive definition that he can: Operational Research/Management Science "may be broadly defined as an activity which seeks: (a) to help groups (or individuals) to solve their problems (b) by using methods that would enjoy consensus support among scientists."

Underlying all of our work is the idea that we aim to apply the "Scientific method". This is an ill-defined term, but it is generally based on seeking an objective process (to avoid bias) by which empirical measurement data are collected to support, refute or help develop hypotheses about how the world operates. So, in many ways, this is our starting-point. Mitchell (1983) points to methods used by scientists that are clearly applicable to Management Scientists: using the rules of logic or mathematics to draw conclusions; checking beliefs or results for mutual validity; and so on. So the metaphor he comes out with is of Management Science as the intersection of two cultures: the culture of the problem domain (i.e. Management) and the culture of Science from which the analyst draws his/her methods – and s/he needs to understand both cultures.

But, following this definition, there are elements that seem to be agreed to come within the Management Science process:

- the need to structure the problem (or situation, or area);
- building some analytical framework of the issues, which is generally termed a model (this idea seems to be key, so we'll discuss it next);
- the need to collect data and measure, bringing the model into line with the "real world";
- a dialogue throughout with the "client" or "problem owner" (which very occasionally might be the same person as the analyst, but these are generally expected to be distinct).

We'll look at what a "model" is first, then we'll move onto where it fits into the Management Science process.

Pidd (1995) identifies six pictures or metaphors of Operational Research – the differences being to some extent the assumptions about the world upon which they are based. Two, *decision mathematics* and *optimization in social systems*

– similar to our third bullet point above – regard the underlying nature of our situations as non-problematic; his next two, *problem solving* and *management science* (defined in a more narrow sense) look further at the organization but still largely see this as a well-defined machine, and so fit into the picture above. But the final two perspectives, *systems perspectives* and *intervention and change*, start to look at issues of alternative perspectives of reality, the roles of the participants, and the nature of Management Science interventions. Therefore, we'll go on to look at these issues as we move through the final section of this chapter and Chapter 2.

1.4 What is a "model"?

At an intuitive level, modellers know what a model is – but it is surprisingly difficult to come up with an all-encompassing definition that tells us what is the essence of "a model", suitable for all types of Management Science models. A good start is always the dictionary. The Collins English Dictionary (1986) tells us that a model is "a simplified representation or description of a system or complex entity, especially one designed to facilitate calculations and predictions". This definition tells us that:

- a model *represents* or *describes* something real;
- a model *simplifies* that real entity;
- the production of a model has a *purpose*, generally to make some sort of calculations or predict how the entity will behave.

The first two aspects are our starting point: our models take something in the "real world", simplify it and attempt to represent or describe it. Now the term "the real world" needs some further consideration (as even this idea is not unproblematic), as does the word "simplify". But let's look first at the idea of "representing" or "describing" the world. Can anything that represents or describes the "real" world be termed a model in the Management Science sense of the word? We are all familiar with one popular form of representing the real world – that is, a painting or photograph. Does this constitute a model? Intuitively, management scientists would think that it doesn't, but why not? The answer to this question actually lies in the third point – the purpose of the model. A painting represents a single, static, representation of reality, which, having been created, is not changeable. As a contrast, we want to *manipulate* a model to tell us something useful, such as to explore alternative realities or to explain why the differences between these realities occur. This is because a model not only defines parts or conceptual elements of the whole, but must also define the *relationships* between the concepts.

Since we wish to manipulate these definitions, they must be *formal*, *theoretically based* definitions of reality that can be manipulated. This means that the "language" of the model will need to be as consistent, unambiguous and precise as possible. This often means using some form of mathematics rather than the English language, which tends to be inconsistent, ambiguous and imprecise. But we are not limiting ourselves to mathematical models, and indeed many of our quantitative mathematical models will be developed from qualitative models expressed in "English" terms – but in formalized formats so that the concepts are made as consistent, unambiguous and precise as we can make them and the relationship structures (for example, causal relationships) between these concepts are expressed.

Figure 1.3 shows an example of the most straightforward type of modelling, where we take something in the "real" world and express it in mathematical terms, and manipulate that model to gain some extra understanding. This is one of the oldest "Operational Research" type of models, dating back to Harris in 1913, although its development is normally credited to Wilson (1934). The example begs a lot of questions and makes unspoken assumptions – it is definitely *not* planned to give an idea of what real-life Management Science is like – but it is provided here to give a flavour of what a "mathematical model" means to any reader who may not have come across the idea before. It aims to illustrate how to approach a structured situation, take the concepts we have defined and express them as mathematical terms.

But we must have a vision for the scope of modelling. Those who see simple equations as the only sort of "model" will have a very jaundiced view of what modelling can do for them. In contrast, Management Scientists would say that modelling can be used – indeed must be used – to represent the whole breadth of reality as we see it (subject to various caveats below and in Chapter 8), and although one technique may be more useful than another, modelling must be available to model any aspect of our project.

So, let us say that we have an explicit representation of "the real world" – what is this "real world"? Traditional Operational Research at its most

Economic Order Quantity

Shop X sells 2 400 boxes per year of a product. Customers come along continuously through the year and buy from the stock of the product kept in Shop X's warehouse. Every so often, therefore, Shop X has to replenish its stock from the supplier. It is expensive putting in an order: Shop X estimates that it costs a total of £320 a year in clerical costs for processing the order and shipping costs to obtain the stock. So it seems sensible to put in only a few, large orders a year. But it's also expensive keeping the product in stock: Shop X estimates that it costs around 24% of the cost of product to keep it in stock for a year – in interest charges on the capital tied up (the product costs the Shop £40 per box) and storage costs (as well as breakages, pilfering, insurance etc.) That suggests it is sensible to put in lots of small orders a year. So, how many times per year should the Shop put in an order?

Let's build a model. Let's first look at all of the aspects in this description that we might need in our model, and represent them by a mathematical variable. It doesn't matter what you call them, but we'll use the some of letters traditionally used in this problems. Let's denote:

the size of the batch we order each time by	Q	(boxes)
the cost to Shop X for ordering each year	B	(£/year)
the cost to Shop X for holding stock each year	H	(£/year)
the total cost to Shop X each year	T	(£/year)

So, how do we calculate these costs? Well, the cost to the shop for ordering each year we know is £320 for each order. The shop is going to order 2 400 boxes a year in batches of size Q, so there must be 2 400/Q orders each year. The cost of ordering each year will therefore be:

$$B = 320 \times (2\,400/Q)$$

If we wanted to calculate the total cost of holding stock, we must think about how much stock is held during the year – and here we must start to make some assumptions. We don't actually know how

Figure 1.3 A mathematical modelling assignment: the "economic order quantity" variable the demand is (nor for that matter how variable delivery times are), so we don't know how much so-called "safety stock" the Shop will hold. If they're very risk averse, they might still be holding a lot of stock when a new batch comes in. For the purposes of this example, we'll put that question to one side and *assume* that demand is regular and continuous, and that new batches come in just as the stock empties out. Thus the stock will start at size Q and gradually decrease to zero, then be replenished up to Q, gradually decrease, and so on. The *average* stock held is therefore $Q/2$. The cost per box per year of holding stock is 24% of £24. And so the cost of holding stock:

$$H = \text{(average amount of stock held)} \times \text{(cost per box per year of holding stock)}$$
$$= (Q/2) \times (0.24 \times 40)$$

The total cost is then simply:

$$T = B + H = 320 \times (2\,400/Q) + (Q/2) \times (0.24 \times 40) = 768\,000/Q + 4.8 \times Q$$

We can now plot a graph of how the total cost varies as the batch size varies:

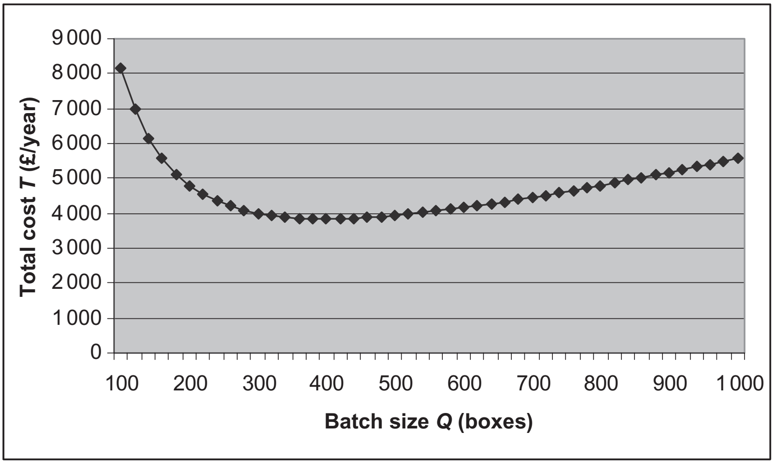

And this – or using calculus – will give the least cost when the batch size is $Q = 400$ boxes.

Figure 1.3 *Continued*

simplistic viewed the world as an absolute reality, which the modeller sought to represent. However, in practice the modeller gains much of his/her knowledge about the reality s/he is seeking to model through human actors who will each have their own world-views (or *Weltanschauung*). So first we may find ourselves modelling an individual subject's *perceptions* of the real world rather than the reality itself, and then we will be working with groups with perhaps inconsistent or even incompatible views of "reality" – and we shall need to discuss what we mean by "real" at all. We'll return to this area in a few pages, but it is worth noting that some would argue that models should not even try to be representations of reality, but simply inventions for debate.

The second bullet point on page 15 said that a model is a *simplified* representation of reality. This is a necessary fact of life – we cannot reproduce reality exactly in a finite period of time. This is not a disadvantage of modelling but indeed one of its most powerful advantages: that we seek to abstract the key elements of reality to provide us with the information we need. This enables us to analyse the model and come to some simplified conclusions about the real world which would be impossible to come to if we had to deal with all the richness, complexity and detail of the real world. Although perhaps not a model in our sense of the word (since it cannot really be manipulated), think of a metro map (an example taken from Pidd, 2003): this is a simplified view of reality – as in reality stations are not spread neatly around, nor are the lines straight, nor are the railway lines single lines in gaudy colours – but the map provides its readers with the information they need to travel on the railway. And, indeed, a precise map of the railway lines with all of the detail of crossing-points, branch-lines and differing depths, with all the lines coloured accurately (i.e. all a similar metallic colour!) would be totally useless for the purpose. So we're back again with our third bullet point: the degree of simplification that we impose on reality to produce our representation depends crucially on the purpose for which we are building the model. Again, this will recur later in the book, particularly in Chapter 8.

So, if the *purpose* of the modelling is so key, for what sort of "purposes" might we be modelling? This takes us back to the first few sections of the book – it might be to help a decision-maker to make a better decision, or help a manager control a system more effectively, or increase understanding in a system or a situation; any of the purposes we have discussed above.

Let's summarize all of the above in a definition that might not be perfect, but will be a start:

> A model represents or describes perceptions of a real situation, simplified, using a formal, theoretically based language of concepts and their relationships (that enables manipulation of these entities), in order to facilitate management, control, understanding or some other manipulation of that situation.

We've left some considerations of what "reality" is, which we'll come to in a few pages. And even on this simple view, there are many other important things that constitute a *good* model – this just looks at what a model is – but we'll leave those to Chapter 8, where we'll look at a much more structured definition of different types of model using various dimensions of model characteristics.

1.5 The Management Science process

So, modelling enables the analyst to build a representation of the real world that s/he can manipulate in place of manipulating the real situation. Indeed, some take this idea so far as to "allow modellers to create computer-based learning environments (or microworlds)" – in that case "for policymakers to 'play with' their knowledge of business and social systems and to debate policy and strategy change" (Sterman, 1988). A traditional view of Management Science looked to modelling a real situation and using that to come up with a "solution" or "answer" or "new understanding", which could then be fed back to the "real world" – that is, something like Figure 1.4.

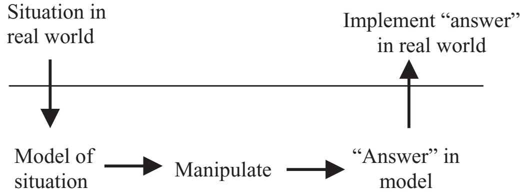

Figure 1.4 Traditional use of modelling

However, as we've been discussing, the types of situation in which Management Scientists have become involved have changed. Furthermore, the old social ideas that "problems" are best left in the hands of the "experts" who will come up with the "best solution" have declined, making the type of relationship where the analyst takes away a problem and comes back with a report now unsustainable. Problems are not objective entities we can point to, and give to the "experts" to solve: problems are subjective constructs, that are only problems as participants and observers perceive them. With our greater understanding of the client–consultant relationship, we shall be seeing in Chapter 2 that we have continuously to be discussing the situation with the client group, trying to model their understanding and negotiating what the "problem" actually is. Thus what we will be describing in this book is much more like Figure 1.5.

Thus we will describe continual iterations of the client-interaction process, in which we have to negotiate the definition of the problem with the client, and continue to analyse and build models in constant dialogue with the client. This modelling relationship in which we need to structure the problem, build models and collect data can be expanded in Figure 1.6.

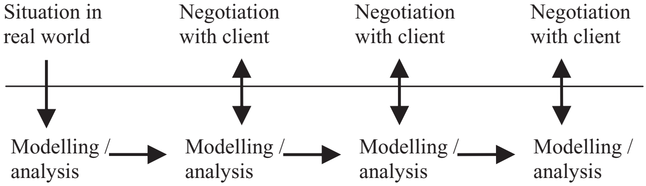

Figure 1.5 More current view of modelling

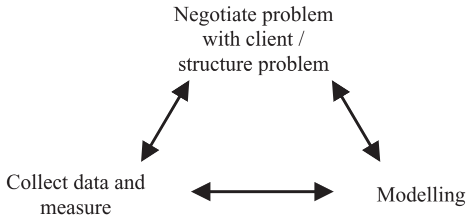

Figure 1.6 Modelling in the Management Science process

In Chapter 2 we will discuss in more detail what a Management Science "intervention" into a situation really means. But it is worth noting two points here first.

Firstly, the end of the process in Figure 1.5 is not nice and clean as it is in Figure 1.4. Indeed, as we'll see in Chapter 2, the whole idea that a "problem" is "solved" and then the Management Science work is complete is unhelpful to the analyst and confuses the process. Eden (1987), for example, says that the idea of "problem-finishing" is much more helpful than "problem-solving" – a problem can be finished in various ways, such as on reaching a "satisficing" solution (where the client "feels that it is obvious what must be done"), or the client might begin to feel equipped to manage the problem on their own, or the problem might be dissolved in that it has been redefined and no longer exists in its original form. This will be important in Chapter 2 as we consider what a "successful" Management Science intervention is.

Secondly, we will be considering an intervention into a complex decision situation to help decision-makers. We can differentiate this from a lot of Operational Research work, which can bring a lot of benefits by solving problems that are important in that they are repeated many times. For example, if you were to look at what are considered to be the best pieces of Operational Research work in the US, such as those that win the US IFORMS Edelmann prize (an annual prize, described for 2004 in Spencer and Graves (2005)), you will find that all of them (except one) describe an ongoing process, with regular incremental decisions, for which the task of the Operational Research group was to supply a "decision-support system" to enable these decisions to be taken, including a user-friendly computer program. Thus, typical of this work is scheduling containers in a terminal, or planning production, or scheduling rolling-stock. These are really "problems", in the sense we discussed above, and the skill of this work is in developing good, effective, fast algorithms. In virtually none of this work are the wider views of decision-makers described, as each individual decision is very small – the impact on the company comes as many thousands of such decisions are automated and optimized. This is not to deny the great worth of this work in solving problems – it is simply a different situation to the position in which Management Scientists generally find themselves, and which we will address in this book. (The one exception, incidentally, was described in Butler *et al.* (2005) and was, again, building a user-friendly decision-support system, but this time in support of looking at alternatives to reduce plutonium – effectively a one-off decision.

Here, the view of all stakeholders is considered and the paper looks at the rationality of their decision-making and at the surrounding decision issues and political factors – a paper following the sort of Management Science we will describe.)

1.6 What is "true"?

Before we explore in more detail how Management Scientists intervene into a situation, we do need to consider some underlying philosophy. Now, you might be surprised that we need to include philosophy in a book about "science", but experience shows that we do come across issues within modern Management Science practice (which, of course, deals with social and management structures and phenomena) that will lead us astray if we haven't thought through these questions beforehand.

There isn't scope within this book to give a full exposition of all the philosophical stances that have been taken towards research and consultancy within management. That would take a whole book – and, indeed, an excellent book has already been published (Johnson and Duberley, 2000) which does analyse the main management research philosophical standpoints; a quick introduction to some of the issues can be found in Chapter 3 of Easterby-Smith, Thorpe and Lowe (1991). Besides, the point of this section is not to lay out a completely philosophical foundation – it is to identify those areas where we as Management Science practitioners can come unstuck, and to give some of the background to what people have thought of these issues.

Essentially, the questions that arise come in two interrelated areas. The first area is ontology, or what is "real". Johnson and Duberley (2000) describe a spectrum of beliefs, from

- a "realist" or "objectivist" ontology, which assumes that a reality (including a physical, natural and social reality) exists independently of us, the observers – a reality exists whether or not human beings can actually cognitively perceive it, to
- a "subjectivist" ontology, which assumes that what we perceive (and thus assume is an external reality, again physical, natural and social) is merely a creation of our consciousness or our cognition.

The second area is epistemology: how we know what is "true". Again, Johnson and Duberley (2000) describe a spectrum of standpoints from the objectivist epistemology, which assumes that we can study the external world objectively, to a subjectivist epistemology, which assumes that this is impossible. (This doesn't mean there is a division between two opposed viewpoints; rather that there is a spectrum of emphases or views (Reed, 1985).)

Following on from these two fundamental questions is the issue of how we view interventions into organizations. However, we'll begin by looking at these first two issues – what views have been taken, and why does it matter to us as practitioners?

After describing all of the standpoints and how they developed, Johnson and Duberley (2000) give a diagram (which is reproduced in Figure 1.7) showing how the standpoints are positioned on these two fundamental dimensions.

Positivism underlies the scientific tradition, and can be traced through many major figures in the history of science: through Descartes (whose championed rationalism, and said that true knowledge could only be gained through an individual rationally looking at an external reality), Locke (who championed empiricism) and Hume (who championed empiricism and analysis of causality). But the key exponent of the idea of positivism is clearly Auguste Comte (1853), who saw true knowledge as only coming from that which is "positively

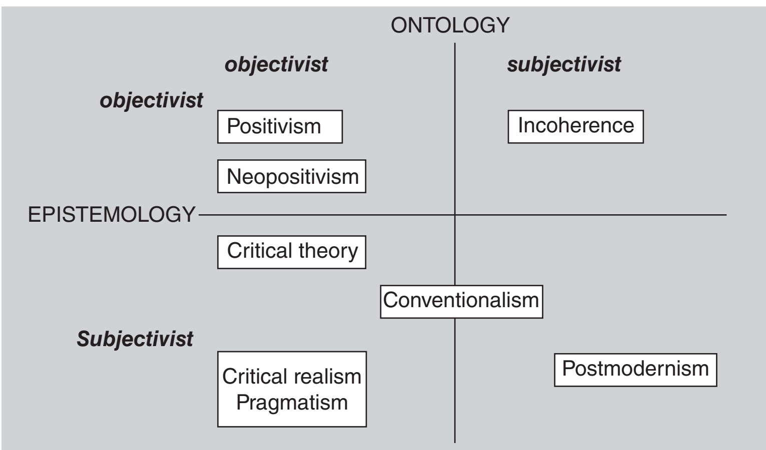

Figure 1.7 Main philosophical positions. From Johnson and Duberley. Copyright 2000 SAGE Publications. Reproduced with permission

given" – that is, which is directly available from the senses, making the assumption that the world outside the observer really exists in an objective sense, and that to find out about the world you should use objective observation and measurement methods.

A major move forward in positivist thinking came particularly from a group of intellectuals in Vienna in the 1930s, sometimes called the "Vienna Circle", who developed positivism into what we now would call "logical positivism". Taking Comte's view stated above, along with the claim that observation of the real world can be neutral and value-free, they would want to test empirically any theoretical statement or postulated theoretical mechanism – and any explanation for the world that cannot be empirically tested, cannot be sustained.

Positivist thinking took a further step forward with the well-known work of Popper (1959), who disagreed with inductive means of generating knowledge (i.e. collecting data, and inferring theories and laws from the data, particularly the causal relationships within that data) and laid the basis for deduction and falsification of theories (i.e. the idea that while the scientific method cannot prove a theory, by showing empirical data which disagrees with the theory, it can disprove it). Popper's work has been very influential in the social sciences, but we can now perhaps see two strains of positivist thinking: one inductive and one deductive.

What does this mean for our practice as Management Scientists? Positivism forms the basis for what most natural scientists see as "the scientific method"; in many senses it gives the foundations from whence Operational Research sprang. From positivist thought springs a number of implications about the way we carry out our analysis:

- observers who are independent from what they observe, and carry out their observation objectively and without influence from internal values or interests;
- the search for causality in the behaviour of the systems we study;
- the aim to operationalize the concepts we study so that they can be conceptualized and then measured quantitatively;
- the need to look at large samples, or repeat experiments many times, to be able to generalize about what we observe;
- we derive knowledge using the "hypothetico-deductive" methods – that is, by hypothesizing general laws and then deducing the types of empirical

experiment that will falsify (or demonstrate the truth of, if using the inductive method) these laws.

Positivism gives the basis for Management Science work, and is a standard against which we will often be measured. Thus our starting-point should be as objective, independent observers, seeking operationalizable concepts and looking for causality to demonstrate "scientifically" and objectively results that can be relied upon.

So, why do we need go any further than positivism and the "scientific method". Well, as we look very briefly through some of the other philosophical traditions in Figure 1.7, we shall see that there are issues that arise where we have to modify our practice – life isn't always as simple as the positivist world would imply.

The first issue, which we can only look at briefly, focuses on the idealistic view of the independent, value-free observer, looking at a system with a completely scientific open-minded view (the first implication of the positivist viewpoint above). The ideas of conventionalism stem from the work of Immanuel Kant (2003), particularly the *Critique of Pure Reason* first published in 1781. His argument was that "pure reason" claimed access to knowledge of the real world beyond what it is possible for human observers to have: we cannot know the world as it is, only as we perceive it through our cognitive or mental structures; we come to view the world not with empty minds, but with minds that shape the way we perceive the world. Conventionalists will say that we cannot escape from the socialization to which we are all subject, and the pre-understandings that come from that socialization, so we cannot know the extent to which our mental structures have shaped the reality we observe (indeed, some would ask whether there is such a reality, or whether what we take to be an external social world is simply a creation of our consciousness – in other words, some conventionalists would take a subjectivist ontology, which we'll return to below).

But again, what does this mean for our practice as Management Scientists? Whatever your views about conventionalism as it comes to scientific research about the natural world, as we come to study the social world of management, it is clear that we come with pre-suppositions and assumptions that colour our observation.

More than this, as we come to collect data, which is key to our process in Figure 1.6, we need to consider what those data are: even when we believe we are collecting data on a "real" phenomenon (i.e. we have an objectivist ontology), often we are not collecting "real" data but someone's perception of that data, so our observation of the world is coloured not only by our mental processes but also those of the subject from whom we're collecting data. For example, if we were modelling the progress of a project (using the "Earned Value" method), one requirement would be to identify the progress on each sub-task of the project, which could involve measuring some physical progress (if it was, say, laying down a path, which could be measured), but might involve asking the participants how far through an activity they were (if, say, it was a systems engineering design task), which will elicit an entirely subjective response ("Oh, I think I'm around 75% of the way through").

And, of course, if this is true for apparently "real" data, there is a whole spectrum of data from the "real" to the clearly subjective. For example, suppose we are to investigate getting the best new production plan for a large factory ready. Some of our data will be physical: how big is the factory? how fast can our materials be moved? As we move to questions such as how many of each product will we sell, we can approach that using quantitative "hard" calculations of past performance, but that must be coloured by subjective beliefs about the future; if we ask what is the probability that new technology will overtake the products that we sell and make them obsolete, that is clearly a piece of "data" that is positioned firmly within the subjective and the social. If we move further to considering what is important to the future of our business – what does "best" mean in our assignment – then we are considering values that have no real physical realization but are socially constructed.

As we use such data, we need to be aware of the status of the data (i.e. how subjective it is) and the colouring that might have occurred both from any human data-provider as well as within us as data-collector. So, in our practice, the following are two clear modifications we must make to our positivist stance.

> We are observing human systems, and so many of the data we collect are filtered through human perceptions. We need to be clear about the extent to which these data are objectively reliable and the effects those perceptions might have had.

We ourselves are not value-free, independent observers: we must identify wherever possible our own *a priori* understandings and bear them in mind both as we analyse what we observe and as we make claims to the knowledge we generate.

Like Kant, Habermas (1972, as well as many other publications) rejected the idea of a neutral observer, which ignores the effect of the knower upon that which the knower is seeking to know, but he went further in placing the contamination of the "ideal" neutral observer within the social and cultural framework within which the observer sits. His work thus tries to take account of the power and dominating relationships that there are within all human relationships, and clearly within management situations. One particular focus is on communication: in his theory of communicative rationality (Habermas, 1984), he claims that it is the communicative structure which shows us where decision-making will be found, pointing out that where communication is dominated by sincere communicative speech acts (discourse unfettered by the coercive use of power) there will be good exchange of rationality (while, on the other hand, where power is being used to limit free communication, there might be failures in decision-making). From his work and other similar writers grew the ideas of "Critical Theory", which is particularly concerned with the effect of socially and historically constituted power relations on thought, and is closely related to considerations of privilege and oppression (more details are given in Mingers, 1992).

While this clearly opens up huge avenues of important epistemological and ethical debate (the latter returning in Chapter 12), we will limit ourselves here to a simple statement of an implication for our Management Science practice – although we will explore this more in the coming chapters. All study of management involves study of groups of humans, who will be placed within historical and current social structures.

As we study management issues, we must be aware of the social and power relationships within the groups, and the effect that they will have on our perceptions and our ability to gather "true" data. This is true of all data, but most obviously for data that are a representation of social factors such as values and group views. This also affects the relationship between the analyst and those being analysed, and the effect of their power relationship must also be examined.

Critical Theory, or Habermas at least, still retains a realist or objectivist ontology. And surely, as Management Scientists, wouldn't we hold onto that ontology even if we have epistemological doubts? Well, if we turn to postmodernism, and the work of one particular author, Foucault, we do see some elements that are important to our analysis. Postmodernism is notoriously difficult to pin down, so we shall here restrict ourselves to a brief summary of some of Foucault (1980). The key idea within much of Foucault's work is the idea of a "discourse", which has become a widely used concept in the social sciences. Discourses are a "set of ideas and practices which condition our ways of relating to and acting upon particular phenomena: a discourse will be expressed in all that can be thought, written or said about a particular topic, which by constituting the phenomenon in a particular way influences behaviour" (Johnson and Duberley, 2000).

We can observe such effects in the management world in particular: by defining a particular way of looking at management issues, we therefore define how that issue is thought about and thus limit what is done or even thought. One example might be project management: as soon as an enterprise is entitled a "project", this defines how people within the project will be able to think about the enterprise; according to Hodgson (2002), although the proponents of project management claim its toolkit to be "universal and politically neutral", enforcement of project management terminology leads to an imposed ontology and specific way of thinking in a company, and he uses Foucauldian analysis to suggest that these claims "serve to establish significant power effects within organizations". In Hodgson (2004), he says: "The key effect of the application of project management models and techniques is enhanced control over the conduct of employees, based on close surveillance and the limited delegation of discretion to those subjects involved in project work. In particular, the quantification and detailed planning involved in project management serves to 'enhance the "calculability" of individuals through developing measures of routine predictability and control' (quoting Metcalfe)" (also incidentally showing a Habermas-ian interest in the power relationships involved). "Critical authors . . . [focus] on who is included in and who is excluded from the decision-making process, analysing what determines the position, agendas and *power* of different participants, and how these different agendas are combined and resolved in the process by which *decisions* are arrived at" (Cicmil and Hodgson, 2006). Fournier and Grey (2000) give three tenets of critical work, one of which is that it "aims to prevent oppression/exploitation" (typically exploitation of employees, women, ethnic minorities or the environment) (quoted in Cicmil and Hodgson, 2006). Some authors leaning towards a Marxist view have also argued against the false objectivity of Operational Research/Management Science and claimed that (traditional) Operational Research suppresses dialogue and works in favour of those in power, denying agreement by consensus (see e.g. Rosenhead and Thunhurst, 1982).

Again, what are the implications for our Management Science practice? Firstly, there is the point of power.

> As we study decision-making processes, we need to be aware of the underlying power structures and the effect on the situation under study.

We have also said that when we try to analyse a "real" situation, the ways we define our problems and approaches can themselves change that "reality" and, to that extent, our objective reality is only real insofar as we have defined it – hence, at least a nod towards a subjectivist ontology.

> As we study management, our formalizations themselves can change or even define the "reality" we are seeking to analyse, and such effects need to be recognized and action taken.

There is a further implication of these points for our modelling. We have already said that models represent or describe perceptions of a real situation – but some phenomonologists, such as Checkland, whose work will look at in Chapter 3, would argue that models are not actually representations of reality but simply inventions for discussion, and Figure 3.2 shows how "models" form part of the debate about how to proceed. Certainly, we shouldn't give our models a status of reality beyond that which they actually possess – as we will discuss in Chapter 12 when we look at the so-called Barabba's (1994) law ("Never say the model says"), this can be a real problem in how our work is used and viewed. The role of models in our practice – about which Checkland has been very influential – needs careful consideration, and this will come up throughout this book.

> Models are constructions that we use in our work – they are not the reality!

This has been an extremely brief gallop through centuries of philosophical thought, but has already made some points that will affect our practice. A pragmatic approach to praxis will need to be based on a certain degree of realism and objectivism in ontology and a certain degree of objectivism in epistemology; however, the expert Management Scientist will be wise enough to understand the limitations of an objectivist approach and take cognizance of the issues we have highlighted above, and their implications, which we will explore in the coming chapters.

Our views of these two fundamental questions – how we regard ontology and epistemology – will influence how we view interventions into organizations and how we regard the knowledge we can or should gather within our interventions. This will be covered in the following chapters, particularly Chapter 2.

If the reader is looking for the explicit methodology underlying this book, and has expected that to be expounded in this chapter, s/he will have been disappointed. The stance taken in this book is that there usually *cannot* be a single methodology – or, indeed, a single paradigm – to approach the work. Successful Management Science often needs to take a multi-methodological and multi-paradigmatic approach, viewing the world "through different lenses". Some might feel that we can't mix up apparently incompatible paradigms, but a pragmatic view is that this is often beneficial and sometimes even inescapable. This is a matter we will return to in Chapter 4.

1.7 Conclusion and going forward

In this chapter, we've sketched out the basis for exploring how to carry out Management Science. We've described the sort of problems with which today's Management Scientists get involved – messy, having regard to social and "soft" issues as well has "hard" physical issues. We've described the idea of intervening in management situations, with that intervention centred on an analytical approach using modelling as the distinctive feature of the Management Science approach. We've talked about some difficult philosophical issues, which are necessary as we take a critical view of management situations and need to have confidence in what we conclude.

But how do we carry out this work? The remaining chapter in Part I will discuss what a Management Science intervention is, and effective ways of interacting with the problem situation. Then, Part II will look at the toolbox of techniques we might use. Part III will describe some practical skills necessary as part of the Management Scientist practice. And finally, Part IV will continue this discussion with a wider look at what it means to practise Management Science.