国际工商管理百科全书

（第2版）

International Encyclopedia of Business & Management

［英］Malcolm Warner 主编

Operations research 至 Simulation modelling

第6卷

清华大学出版社

国际工商管理百科全书

International Encyclopedia of Business & Management

（第2版）

[英] Malcolm Warner 主编

(第6卷)

Operations research 至 Simulation modelling

清华大学出版社

Using the encyclopedia

The *International Encyclopedia of Business and Management* is designed for ease of use. The following notes outline its organization and editorial approach and explain the ways of locating material. This will help users to make the most of the encyclopedia.

Sequence of entries

The encyclopedia contains 750 entries arranged in a single, alphabetical sequence through seven volumes. Entries are listed in alphabetical order. Note that the sequence follows the order of words rather than that of letter, and that the words *and, in, of* and *the* in entry titles are disregarded. A complete alphabetical list of entries is given in Volume 8 (the Index Volume).

The Index Volume

Volume 8 is devoted to a comprehensive index of the key terms, concepts, countries and names covered in Volumes 1 to 7, allowing users to reap maximum benefit from the encyclopedia. A guide to the index can be found at the beginning of the index. The Index Volume also includes permission acknowledgements, listed in alphabetical entry order.

Cross-references

The encyclopedia has been extensively cross-referenced in order to signpost other entries that are likely to be of interest. There are three types of cross-reference in the encyclopedia:

'See' cross-references

Throughout the alphabetical sequence of entry titles, there are cross-references which direct the user to the entry where a particular topic is discussed either under a different entry title or as part of a larger entry. For example:

Corporate taxation: see TAXATION, CORPORATE

Ethics: see BUSINESS ETHICS; MARKETING ETHICS

'See' cross-references within an entry

Cross-references within an entry direct the user to other entries closely related to the theme under discussion. These other entries will normally give a fuller explanation of the specific theme. These cross-references appear in small capital letters.

'See also' cross-references

At the very end of each entry, 'See also' cross-references guide the user to other entries of related interest, such as more specialized entries, biographical entries and geographical entries, as well as related entries in other disciplines. These cross-references appear in small capital letters in alphabetical order.

Structure of entries

A numbered contents list at the beginning of each entry in the encyclopedia gives the headings of its main sections. The scope and structure of the entry can thus be reviewed and sections of particular interest easily located.

Thematic entries begin with an 'Overview' section that serves as a brief introduction to the topic and a useful summary of the entry's contents. Biographical entries begin with a summary of the significant dates and events in the life of the subject and a list of his or her major works. Every entry is followed by a 'Further reading' section (see below).

Authors

The name of the author or authors is given at the end of each entry. A full list of contributors, showing their affiliation at the time of writing and the titles of the entries they have written, can be found in Volume 8.

Further reading

Each entry has a 'Further reading' section which gives details of all the references cited in the text. Additional suggestions for reading are also provided for those who wish to delve deeper into a particular subject. References cited in the text are preceded with an asterisk (*).

The Further reading list is arranged alphabetically by author/editor and chronologically under the authors'/editors' names. Publications with joint authors are listed under the name of the first author and are listed after any individual publications of that author. Where publications have been issued by an institution, the name of the institution is given as the author. English translations of publications in other languages have been given wherever possible.

Items in the Further reading list have been annotated with a brief description of the level, importance and usefulness of the publication listed.

References and suggestions for further reading are given in the Harvard style. The authors and editors have attempted to provide bibliographic data in the fullest possible detail.

Editorial style

Spelling and punctuation in the encyclopedia have been standardized to follow British English usage. The use of italics has been kept to a minimum and is normally restricted to foreign words and book or journal titles. Abbreviations and acronyms are spelled out in full on their first appearance in an entry. Chinese names have been westernized (i.e. Chen Derong becomes Derong Chen or Chen, D.). In alphabetical lists of names, *Mc* and *Mac* are treated as *Mac* and the next letter in the name determines the position of the entry.

Operations research

Overview

Operations research (OR) devotes its effort to understanding the behaviour of operating systems made up of nature, people and devices, where this last term includes not only machines but the artefacts of modern technology and also laws, common practices, human behaviour and social structures and customs. Such systems are common in business, management and government settings. Operations research work includes both developing theories – called models by operations research workers – and applying their knowledge to improving the operations of these systems. This article offers a brief overview of the history of operations research, general approach, arenas of enquiry and application, methods and professional practice, together with mention of allied fields and a perspective on the future of operations research.

1 History

Operations research is unusual among fields of analysis in that its beginnings – and even its name – can be traced unambiguously to a single event: military operational trials being conducted on England's east coast in the summer of 1938.

During the three preceding years, as part of Great Britain's effort to respond to the growing threat of German air power, scientists had been developing a means of detecting flying aircraft at a distance. By 1937, devices that came to be called radar were in place and experiments were being conducted with them. In the summer of 1938, these experiments incorporated the operations of fighter planes to explore the combined defence effectiveness of long-range detection and active response. The scientists who had come onto the scene to work with the radar then turned their attention to analysing what was happening in these trials; they were able not only to characterize the results but also to show how they could be improved. A British scientist referred to these analyses as 'operational research'. When the term crossed the Atlantic later, it became 'operations research' and the two synonymous usages remain today: 'operational research' in the UK and parts of Europe, 'operations research' elsewhere.

The success of the small British team of analysts that was working directly with military operations in 1938 set a pattern that was soon widely imitated; well before the end of the Second World War most of the military services in Britain, Canada and the USA had established multidisciplinary teams of analysts to work directly with military operations. Their work began by deriving systematic understandings of these operations and helping to improve them. By the end of the war, some 700 analysts from many disciplines had participated in this work (McCloskey 1987).

The basic lesson emerging from the wartime analyses was that complicated operations involving people and machines often exhibit regularities that can be captured in theories, which can then be used to predict the outcomes of future operations. These theories – which operations research workers later came to call models – were also used in planning and addressing strategic issues.

With one exception, the wartime analysts did not invent any significantly new technical tools to support their work; rather, they were able to use methods and techniques that they had brought with them from their various scientific and technical backgrounds. The exception was a theory of search devised to deal with an activity common in warfare: search-

此为试读，需要完整PDF请访问：www.ertongbook.com

ing from a moving platform (such as a ship or an aircraft) for a fixed or a moving target (such as a hostile submarine, ship or aircraft). The goal of analysis based on this theory is to aid in detecting, locating and identifying a target so that appropriate action can be taken. In a peacetime setting, this theory has been applied successfully to finding a sunken treasure ship off the US Atlantic coast (Stone 1992) and for marine rescue operations (Engel and Weisinger 1988).

Generically, the wartime analysts were dealing with operating systems consisting of people, machines and nature in structures that, on close examination, were not as randomly disorganized as the casual observer might expect. After the war, many of the operations research analysts realized that there were many civil operating systems that contained similar generic elements, so that what they had learned in the war could well be applied to managing civil enterprises. Within a decade, many large business and government organizations had added operations research activities to their management procedures. This new initiative, added to the resolve on the part of military services to continue operations research work on their peacetime problems, led to a rapid growth of the operations research community.

The growth of the operations research profession

The first operations research professional society was founded in the UK in 1948; during the next decade professional societies came into existence in seven other countries: USA (1952), France (1956), India (1957), Japan (1957), Belgium (1958), Canada (1958), and the Netherlands (1958). The International Federation of Operational Research Societies (IFORS) came into being in 1959, the UK, US and French societies being the initial members; by 1973, 20 national societies were members, and by 1993 there were 41, with four kindred societies also adhering to the Federation. The worldwide membership in 1993 consisted of over 35,000 persons.

In addition to holding triennial international meetings and publishing a research journal, the Federation publishes *International Abstracts in Operations Research (IAOR)*, a journal that since 1961 has offered conscientious abstract coverage of the world literature in operations research. Three volumes (Batchelor 1959, 1962, 1963) cover the literature before 1961.

Journals

The first two professional journals in the field were the *Operational Research Quarterly* (now the *Journal of the Operational Research Society*), founded in the UK in 1950, and *Operations Research*, founded in the USA in 1952. By 1973 there were 15 primary operations research journals and 22 others of significant interest to operations research workers. By 1997 *IAOR* covered the literature in over 180 journals.

This literature and the professional activities that have prompted its growth since 1938 represent a tremendous development of interest in operations research, and of the extent and power of its theories and their effectiveness in applications to real-world problems.

Books

The first book that described wartime work was Morse and Kimball (1951), which devoted most of its attention to the models that had been used. Although written at the end of the war (but published much later), Waddington (1973) also offers something of the flavour of the day-to-day wartime work.

The first textbook emerging from an educational programme did not appear until over a decade after the end of the war (Churchman *et al.* 1957). Since then, many other textbooks have appeared; Wagner (1969, 2nd edn 1975) did much to standardize introductory operations research courses and is still of interest; more recently, Hillier and Lieberman (2001) has been widely used. Mitchell (1993) and Rivett (1994) focus their primary attention on work in actual operating situations.

There have also been many more specialized treatises as well as two handbooks: one in two volumes (Moder and Elmaghraby 1978) and one in three (Miser and Quade 1985, 1988; Miser 1995); the latter volumes centre their attention on large-scale operations research enquiries. There is also an encyclopedia (Gass and Harris 2001). Much of the material in these volumes is accessible to the general reader; more technical summaries are published in separate volumes.

Educational programmes

In the early post-war years, the only training available was apprenticeship or independent study, but by the mid-1950s some short courses and a few university programmes began to emerge. From then on, the growth of academic interest was substantial throughout the developed world. For example, in the USA and Canada there were over 50 academic curricula in operations research by 1973, and over 100 by 1999. Academic training programmes elsewhere in the world experienced similar growth.

Furthermore, well before the 1990s, graduate training in these operations research curricula had become recognized as an important qualification for entry into professional operations research work, whether teaching, research or practice.

Closely-related specialities

In 1954, a group in the USA founded a professional society called The Institute of Management Sciences and extended its activities and influence internationally. However, its subject, management science (MS), is now recognized as being, for all practical purposes, synonymous with operations research. Indeed, the common subject is often referred to as OR/MS or MS/OR.

In the mid-1950s, interdisciplinary teams involving operations research workers began to investigate large-scale problems generically like those on a smaller scale being explored by operations research workers. The workers on the larger-scale problems, which often involved choices among large-scale systems dominated by technological artefacts, began to call their work 'systems analysis', a term that has continued in this meaning, in spite of its adoption elsewhere, notably in the computer world, for other kinds of work. Thus, here we may regard operations research and systems analysis as being synonymous except for a differentiation in the scale of the work involved. When the work is more dominated by human and political concerns than those of modern technology and administrative structures, systems analysis is often called 'policy analysis'; there is a substantial community that works under this name and a significant literature that reports its work (Quade 1989).

Since the goals of much applied operations research work are to understand the behaviour of operating systems, to offer managers clarifying information and often to support policy and operational decisions, a number of other specialities that were developed in the 1980s and 1990s with related goals have interests that overlap those of operations research; for a discussion of them, see the later section on allied fields.

2 Operations research's scientific aspects

While operations research enquiries take on forms that are as various as the contexts and problems to which they are addressed, it is nevertheless useful to examine an ideal structure that contains the key elements common to much of such work.

It begins with a thorough examination of the context or problem situation. This inspection entails actually viewing its elements, and usually involves capturing evidence by various means, such as recording aspects of chosen phenomena, a process that often includes taking and processing data.

The evidence from this examination is then analysed with a view to discerning its stabilities and embodying them in intellectual constructs called models (the term that operations research workers usually use for their theories). Manipulations with these models can then offer inferences about how the situation might behave under various assumed conditions.

Before these inferences can be considered fully credible, however, they must be tested against the reality of the problem situation. If the comparisons between the actual and predicted behaviour show good agreement, the analysts may view their models and the consequences that can be derived from them with some confidence that they reflect real-world behaviour; if not, then the analysts must return to the context or problem situation to begin the cycle over again.

Depending on the nature of the enquiry, operations research workers may follow this cycle as presented here or may enlist its elements in quite another order. For example, the beginning examination may yield enough evidence to produce immediate understanding without a significant model-building step; an attractive model built on a general overview of a problem situation may seek confirmation in reality; or a concept of how to solve a problem may prompt an exploration of the underlying situation to see if the presumed understanding of the phenomena is confirmed in reality.

This discussion, while cast in terms of problem situations, describes the scientific approach (Kemeny 1959; Miser 1993). The goals of the work may be to solve a practical problem, to fashion a tool that may help solve a class of practical problems, or, more restrictedly, to establish a model that will be useful in a variety of contexts.

3 Arenas of enquiry and application

Operations research devotes its effort to understanding the behaviour of operating systems made up of nature, people and devices, where this last term includes not only machines but the artefacts of modern technology and also laws, customs, common practices, human behaviour and social structures and customs.

The core of this description appears to be a concept of system, but attempts to generalize and apply it based on ideas about systems have not been fruitful. Rather, the concept of an 'action programme' (Boothroyd 1978; Miser and Quade 1988: 489–507) is more useful: a function, operation, activity or response that is related to and given coherence by a human objective, need or problem, together with the collection of people, equipment, portion of nature, organizational elements and management or social structure involved. Operations research can then be described as the analysis of action programmes, its work including understanding their behaviour and using this knowledge to solve problems in them, this latter step being what Boothroyd calls 'articulate intervention'.

One of the important things to note about an action programme is that its elements are usually members of more than one action programme. For example, a person may be involved in an important action programme at work, but may also be involved in others elsewhere, such as membership of a charitable organization, a service club or a church. Similarly, business and management organizations usually contain more than one action programme.

Operations research has entered many arenas in which there are important action programmes. In looking across these arenas, there are common processes. For example, air transportation and retail sales are two important arenas and both have within them queues awaiting service (such as aircraft waiting to land at an airport or people waiting to be served).

Table 1 lists some important arenas in which significant operations research enquiries have taken place; Table 2 lists some common processes that operations research has studied that can be observed in various arenas. Since the terms in these tables have been adapted from *International Abstracts in Operations Research*, examples of work in any of these areas can be identified by consulting recent issues of this journal.

The optimum lot size

There follows an example based on one that Miser and a colleague actually encountered in a furniture factory in 1960. It made, among many other things, a line of open-stock furniture that it associated with its name, and of which it was very proud. Unfortunately, however, the rate of sales had slowed considerably, so that only a few pieces were sold each month. Since manufacturing this line was more complicated than most of the products in the plant, the plant manager wanted to produce a large quantity and put it in the warehouse where it could be drawn on for a long period without the need for setting up the plant for a new lot. On the other hand, the business manager did not want to tie up the firm's capital for so long a period as a large lot would

Table 1 Some important arenas of enquiry for operations research

Advertising	Manufacturing operations
Agriculture and food	Marketing
Commerce	Military and defence operations
Communications	Mineral industries
Community problems	Petroleum-industry operations
Construction and design	Public service
Education operations	Recreation and tourism
Energy supply and distribution	Space operations
Environmental problems	Sports
Financial operations and banking	Transportation
Forestry	Air
Government operations	Rail
Health services	Road
Law enforcement	Water
Library operations	Urban affairs
	Water supply and management

Table 2 Some common processes explored by operations research

Accidents	Management operations and policies
Allocating resources	Marketing
Bidding	Materials
Communication	Measurement
Computer applications	Organizational structures
Control procedures	Performance measures
Decision rules	Personnel policies and manpower planning
Demand	Planning
Design for operating processes	Production
Distribution	Flexible manufacturing
Equipment choice and operation	Just-in-time inventories
Facilities	Material requirements
Financial operations	Project management
Forecasting future conditions	Quality and reliability
Growth factors	Queuing processes
Information generation and use	Research operations
Innovation	Risk
Inspection	Scheduling
Inventory operations	Service operations
Investment	Vehicle routing and scheduling
Location of facilities	Work operations and design
Maintenance, repair and replacement	

Source: Adapted from *International Abstracts in Operations Research* (1993, 44: 14)

involve. The question was whether or not we could shed light on this issue.

Fortunately, the firm had good records on its activities and costs so we were able to gather good data on the situation, which ran something like this: if q is the number of items in the lot, then the average cost of carrying an item in the warehouse is a constant c times q. This perhaps surprising result is exact if the sales rate and the cost of carrying items in storage are both constant (and it can be partially justified by noting that the average storage cost per item in the warehouse obviously increases with the size of the lot). If m is the labour and material cost of producing one item, and s is the set-up cost for the lot (that is, the cost of preparing to manufacture the lot), the cost y of each item in a lot of size q is given by $y = m + cq + s/q$. After differentiating this expression with respect to q and setting it equal to zero, we find that the minimum cost of an item occurs when the lot size is the square root of s/c. It is notable that this expression is independent of the material and labour cost m.

But the criterion of minimizing the unit cost is not the only one that can be considered; another is the maximum profit per lot. Here, if P is the profit for a lot of size q and Y is the selling price for an item, then $P = q(Y - y)$ and the maximum profit is obtained when the lot size is $(Y - m)/2c$. We note that this expression is independent of the set-up cost s.

If we let $c = 10$, $s = 1000$, $m = 40$ and $Y = 500$, then the optimum lot size for a minimum cost per item manufactured is 10, whereas the optimum lot size for maximum profit per lot is 23, considerably larger.

But other criteria can also be considered: maximum return on the investment in the lot and maximum rate of return on the investment, both of which may be of more interest to the firm's treasurer than to the plant manager. The optimum lot size for the first of these criteria turns out to be the same as the one for the minimum cost per item (ten in the example), whereas the optimum lot size for the second is smaller (at a sales rate of two items per month it is six in the example) – a fact not likely to cheer the plant manager. Having results for these four criteria allowed the firm's management to select lot sizes in each case that were suitable compromises among the various financial and operating concerns.

Of course, operating a plant and the business in which it is embedded is considerably more complicated than manufacturing a single item, and operations research workers have developed appropriate theories for such manufacturing and inventory problems and used them to solve practical management problems. For an introductory discussion of the lot-size problem, see Eilon (1959). For a modern overview of where operations research workers have taken this line of thought, see Graves *et al.* (1993).

Some actual examples of successful work

This subsection offers brief descriptions of nine actual cases of successful work, together with references to where they can be explored more fully. All but the sixth are ambitious operations research studies conducted by small teams of analysts; the sixth is a major systems analysis involving an interdisciplinary team of 31 analysts working over a substantial period of time.

Improving the operations and service of a major trucking company

A trucking company in the USA has evolved from a regional operation to one of the largest less-than-truckload motor carriers in the country. At the time analysts entered the picture, it was handling over 15 million shipments annually over a network of 630 terminals. As the company had grown, its terminal managers lacked the information that would allow them to coordinate their activities around the network. An analysis team developed a large-scale interactive information system that allows the company to improve both the design of its network and how shipments are routed. This system, which runs counter to decades of standard practice, has enabled the company to evaluate and implement operating strategies that have yielded substantial cost savings. It has also improved the company's responsiveness to a changing business environment. More important, however, are the company's improvements in transit times and service reliability for its 300,000 regular customers (Braklow *et al.* 1992).

Yield management at an international air carrier

The effective use of its reservations inventory is critical to an airline's operation. Thus, ever since the 1960s analysts at American Airlines have been conducting research on how to manage the revenue from this inventory. They have developed a series of models that effectively reduce this large problem to three much smaller and more manageable subproblems: overbooking, discount allocations and traffic management. The results of the subproblem solutions can then be combined to determine the final inventory levels. The airline has estimated a quantifiable benefit over a three-year period of US$1,400 million and expects an annual revenue contribution of US$500 million from the results of this analysis stream (Smith *et al.* 1992). OR-based yield management models are being applied by hotels and car rental companies as well.

Mortgage valuation at a securities firm

Over a four-year period, a US securities firm built a leading mortgage-backed securities (MBS) department. Owing to the complexity of the securities, standard fixed-income valuation tools are inadequate for MBS. Thus, developing sound and accurate valuation models has been an integral part of the department's growth. These models incorporate a variety of operations research techniques. They allow the firm to value complex MBS quickly and accurately, and hence trade them effectively, hedge the MBS in inventory properly, and structure clients' portfolios to achieve given objectives while staying within specified constraints. Since traders, salespeople and clients use the models to evaluate MBS hundreds of times each day, they form the basic tools necessary for the firm to participate successfully in the mortgage market (Ben-Dov *et al.* 1992).

Improving a city's arrest-to-arraignment system

In 1988, New York City's arrestees were in custody waiting to be arraigned for an average of 40 hours, occasionally for more than 70. Moreover, they were held in crowded, noisy conditions that were emotionally stressful, unhealthy and often physically dangerous. The city engaged an analysis team to address the problems of arraignment delays and the staggering associated costs of operating the arrest-to-arraignment system. In an effort extending over two years, the team developed a comprehensive, easy-to-use model of the system's performance and costs for the four major boroughs of the city and used it to recommend sweeping operational and policy changes. The ones implemented since 1990 have saved the city tens of millions of dollars and have helped lower the average arraignment times to 24 hours or less (Larson *et al.* 1993).

Modernizing a brass company's management practices

In 1990, a brâss company in the USA embarked on a modernization programme. Since customer service is the most important variable in the industry, the goal was to improve it. The company had outgrown its people-intensive systems and needed to implement changes that would lift it to the next level of performance. A new integrated management system relied on four methods: statistical forecasting, cellular manufacturing, planning manufacturing resources and purchasing, and managing finished-goods inventory. The company completed the design and implementation of all four components at modest cost and in record time. The effects proved to be both synergistic and dramatic: the company has made major improvements in customer service while reducing costs (Flowers 1993).

Planning the Netherlands' water resources

In 1976, a severe drought cost the Netherlands more than US$2,500 million in agricultural losses alone, over 4 per cent of the gross domestic product, while worsening water quality problems throughout the country. Thus, the Dutch government undertook a major examination of its water management problems. An interdisciplinary team of analysts from both the Netherlands and the USA conducted the inquiry on which this re-examination was based.

Historically, the Dutch water management problem has been too much water; draining the land and holding back the North Sea. But in recent years, there has been the less dramatic but no less urgent problem of too little fresh water and too much pollution brought on by increased industrialization and a growing population with a high standard of living, with the River Rhine bringing through the Netherlands pollution from upriver countries.

The analysis team undertook to develop a way to assess the multiple consequences of water management policies and to apply it to generating alternative policies and comparing their various possible outcomes.

Since there were a great many tactics that could contribute to these policies – technical, managerial, pricing and regulatory – it was necessary to organize the attractive ones into

此为试读，需要完整PDF请访问：www.ertongbook.com

strategies and then to bring the attractive strategies together into policies. To be able to do this, the analysts had to choose impact measures that were both real and understandable to the wide spectrum of interested parties that would be involved in final decision making.

The team developed a complex set of interactive models to assess the impacts of the tactics, strategies and policies in terms of these measures. The results generated by these models enabled the team to screen the various tactics and strategies so as to eliminate the less desirable ones, so that policies formulated by bringing selected ones together would have favourable impacts. The final step in the analysis was to assess the impacts of these policies and present the findings to the interested parties.

The Dutch government then combined these findings with some additional analyses of its own to formulate a proposal that was eventually adopted by the government and implemented, with beneficent results.

For an overview of this work, together with references to the several volumes that report its details, see Goeller and the PAWN Team (1985).

Improving the efficiency of Chilean forestry

Throughout the 1990s, Chilean forest firms, which have annual sales of approximately US$1,000 million, have been implementing OR models developed jointly with academics. These systems support decisions on daily truck scheduling, short-term harvesting, location of access roads and harvesting machinery and longer-term forest planning. Models used include linear programming with column generation, mixed integer linear programming, heuristic methods and simulation. Annual cost savings and efficiency improvements now exceed US$20 million (Epstein *et al.* 1999).

Shaping the National Defence Forces of the new South Africa

After the Republic of South Africa became a full democracy in 1994, it needed to completely rethink the size and nature of its military. A team involving the South African National Defense Force and external management consultants used OR/MS-based models to support a joint strategic management process. The models included a risk model of various defence contingencies, a growth model of potential force structure elements, a cost model, a mixed integer programming model and a manual model for the user to select different tasks and strategies. The project led to a 22 per cent saving on the approved design and indicated opportunities for additional savings (Gryffenberg *et al.* 1997).

Production-planning and delivery-quotation for semiconductor manufacturing

When a US semiconductor manufacturer found itself disappointing customers and losing money, it cooperated in the development of a customized planning system that takes account of its re-entrant process flows, product substitution possibilities and marketing priorities. The system uses heuristic decomposition to break the overall problem into manageable calcuations including linear programming. It supports production scheduling decisions and product lead time quotation decisions. Its use has raised on-time deliveries from 75 per cent to 95 per cent without increasing inventories, enabling the firm to profitably expand its market share and go from losses to profits (Leachman *et al.* 1996).

Sources for more case descriptions

Each of these cases emerges from a different arena of operations research application. Furthermore, an examination of the details of the analyses behind them as related to other work – although not possible in modest space – would show that many of the tools used had emerged from analyses of common processes.

While most of operations research's primary journals occasionally print articles about successful cases, one of them (*Interfaces*) devotes its primary attention and most of its space to such material. The yearly issue (usually the first of the year) that contains the descriptions of the Edelman Prize finalists is always particularly interesting in this regard (the cases described above were all chosen from such issues). Assad *et al.* (1992) is also a valuable source of good cases related to business and management; they are accompanied by valuable commentary.

Anyone interested in work in a given arena or relating to a chosen common process can locate the relevant literature easily through a search in *International Abstracts in Operations Research.*

4 Methods

In general, the methods of operations research are those of applied science and technology. Some techniques, models and approaches, however, are closely associated with operations research, even though they are also used in other fields.

Since operations research enquiries often entail detailed observations of varying phenomena, the methods of statistics are frequently called into play. From this fact, it follows that the models built to describe such phenomena involve probability elements; a wide variety of such models can be found in operations research work (see MARKOV PROCESSES AND APPLICATIONS; FORECASTING AND STATISTICAL MODELS).

Because many practical problems entail achieving a preferred use of resources under established constraints, there are a variety of models for programming such uses (see LINEAR PROGRAMMING; INTEGER PROGRAMMING; NON-LINEAR PROGRAMMING; INTERACTIVE PROGRAMMING).

Since much operations research work involves seeking preferred procedures or courses of action, analysts have paid a great deal of attention to models that allow optima to be sought (see OPTIMALITY AND OPTIMIZATION; OPTIMAL-DESIGN MODELS; DYNAMIC PROGRAMMING AND THE OPTIMALITY PRINCIPLE).

Because other practical problems entail providing service to various kinds of arrivals which may therefore join queues, queuing systems (see QUEUING SYSTEMS) are studied extensively by operations research workers.

Networks occur in many situations, such as shipping, distribution and planning the succession of tasks in a complex construction project. Thus, operations research workers have devised varied ways of dealing effectively with many aspects of problems involving networks (see NETWORKS AND APPLICATIONS; CRITICAL PATH ANALYSIS).

As the electronic computer came into use, operations research workers were at the forefront of applying this tool to solve business and management problems. This practice is now so widespread as to be regarded as ubiquitous where operations research analysts are at work on practical problems. In addition to providing a convenient means for embodying operations research's models, computers can also be used to simulate business and management operations, either existing or planned, so that how they will behave under changing conditions can be explored. While such simulations can be constructed from first principles (see SIMULATION MODELLING), by the 1990s it was more common for analysts to adapt a commercially available simulation package to the application of concern.

Many operations research problems involve improving transportation operations; transportation models are often used as the basis for solving such problems.

Although classically findings emerging from operations research models support choosing a course of action based on a single criterion of worth, it is also common for choices to involve many criteria and operations research workers have created structures to respond to such situations (see DECISION MAKING, MULTIPLE-CRITERIA). In working towards and reaching a decision, it is also important for the decision maker to have as broad a view of the possibilities as can be brought to bear, based on as comprehensive an understanding of the relevant context as possible.

When game theory emerged on the intellectual scene in the mid-1940s, many operations research workers expected that it would provide an important basis for operations research work. While its overall concept enlightened much thinking about operating and policy problems, the game theory models did not prove to be as useful in practical situations as had been hoped. However, gaming exercises have been able to inform decision makers about the important properties of situations that they face (see GAME THEORY AND GAMING).

References to the many other models that are used in operations research can be found in *International Abstracts in Operations Research.* Introductory explanations of some of the ones mentioned above – and others – may be found in introductory operations research

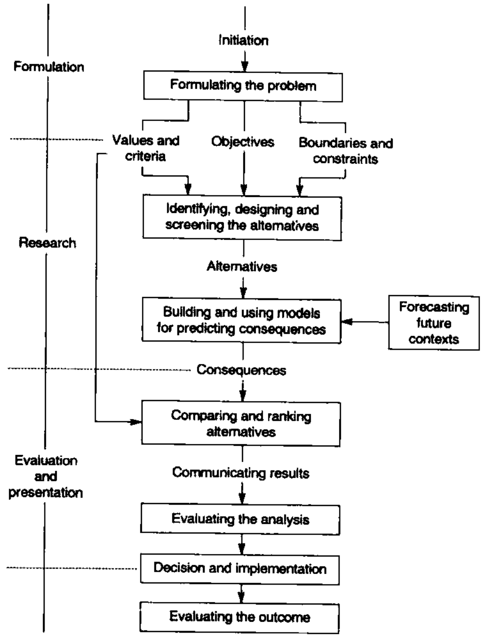

Figure 1 Important elements in an operations research practice engagement that runs from problem formulation through research and implementation to evaluating the outcome
Source: Miser and Quade (1998): 23

textbooks such as Wagner (1975) or Hillier and Lieberman (2001).

5 Practice

An operations research analyst working to solve a practical problem is said to be engaged in practice, an occupation that calls for varied technical and craft skills, as well as an ability to achieve constructive relations with a variety of parties with an interest in the matter being pursued. Practitioners are sometimes attached to the staffs of the organizations that they serve or sometimes they may enter the situation as consultants from an outside organization; each situation has its advantages and disadvantages.

While each situation in practice may properly be regarded as unique, it is nevertheless possible, for the sake of orderly discussion, to describe a central analysis process that contains elements central to most – if not all – operations research practice. Figure 1 offers a synoptic view of these elements. It proceeds from the general recognition that a problem situation exists to the implementation of some policy or course of action and the evaluation of its effects. Since each situation has its own unique properties, few operations research practice engagements follow such a procedure exactly – and sometimes not even approximately. It is common, however, for at least some of these elements to occur in the work.

Formulation

While the person who brought the problem situation to the attention of the operations research analyst may have – along with others – diagnosed the situation and even have a belief about what should be done to improve it, it is commonly the case that the forces actually at the core of the problem situation lie buried deeply enough to make such an early diagnosis questionable and therefore the preconceived fix inappropriate. Thus, typically it is best for the analyst – or the analysis team if the situation is large and complex – to approach it with an open mind and aim to explore it thoroughly before deducing its essential properties and using them to devise a scheme for ameliorating its undesirable aspects.

Therefore, the work begins with a searching and thorough investigation of the problem situation with a view to formulating the problem to be addressed by the later work. It is vital for the client and others in the action programme being explored to cooperate in this work and share in the eventual problem formulation. Once this is done and the client has agreed with the analysis team on the diagnosis and the problem, it is possible for the team to plan the rest of the work to be done.

This early work also identifies the values and criteria that should inform the choice of what eventually will be done to ameliorate the client's concerns, sets up the objectives to be sought by the solution, and agrees with the client on the boundaries and constraints that must be observed in devising it.

Research

This stage extends the information- and data-gathering and analysis that began in the formulation stage, but it is now more sharply focused. The findings that emerge allow the analysis team to identify, design and screen possible alternatives that may help with the problem. Against this background, the analysis team can then build models capable of deducing the consequences of adopting each of the alternatives chosen for further investigation within the framework of possible future conditions.

Evaluation and presentation

With estimates of the consequences in hand, the analysts may compare – and possibly even rank – the alternatives against the criteria chosen earlier in the analysis, plus any new ones that may have emerged during the work. These findings must then be presented to the client and other interested parties in a way that enables them not only to appreciate the results but also to have at least a broad overview of the logic and the factual basis that produced them. These understandings may then enable the client to combine the analysis results with his or her understanding of the situation in which they are embedded in order to adopt a suitable policy or course of action and, when necessary, persuade the various other interested parties to accept it.

Implementation and evaluating the outcome

The client and not the analysts, of course, must decide what to do and how to carry it out effectively, and he or she must see to the implementation and take full responsibility for it. Nevertheless, experience shows that it is very important for the analysis team, or at least some analysts who understand and appreciate what was done, to work cooperatively throughout the implementation stage.

Later, if an evaluation of the efficacy of the results is undertaken, it is also desirable for someone familiar with the input analysis to be involved, although to eliminate a conflict of interest it may be desirable for this person not to have been a member of the original analysis team.

Variations

While it is possible to specify a core diagram of the principal elements of operations research practice, it must be admitted immediately that few, if any, such analyses follow this outline exactly. Rather, since each problem situation is different, the analysis must be adapted to it. Thus, in reviewing a series of cases, one finds variations such as the following.

1. Instead of proceeding linearly from the top to the bottom of the outline in the diagram, the work cycles from intermediate stages back to earlier ones as the research brings new insights and fresh intermediate results that may prompt reconsidering the beginning foundations of the work.
2. Some work may be aimed more at fleshing out the client's understanding of the situation than prompting a significant change, so it may stop at one of the intermediate stages.
3. The relative effort expended in the various stages may vary tremendously from case to case: one case may have to expend its major effort in just the information- and data-gathering stage, after which what needs to be done may be fairly apparent without much further analysis. Another case may proceed fairly expeditiously through the outline of Figure 1 and then have a very long and complicated period of work to achieve what may appear to an outsider to

be the implementation of a fairly simple set of proposals.

4 In some cases, an intermediate stage may dominate the work, owing to such factors as technical difficulty in devising proper models, major uncertainties in forecasting future conditions, complexities of the underlying situation, and so on.

In any case, the procedure outlined here as the basis for discussion must be regarded as one that has stitched together the key elements that may enter operations research practice to varying extents, depending on the peculiarities of the situation being analysed.

We have already noted that there is a branch of operations research devoted to problems of such large scale that they almost always involve major efforts by teams of analysts drawn from a variety of specialities and that in the mid-1950s it came to be called 'systems analysis' (not to be confused with the same term used later in the computer community). While there is a separate literature about such work (Miser and Quade 1985, 1988; Miser 1995), its essential properties are sufficiently close to those of operations research as it is usually understood that inferences back and forth between the two kinds of work are quite appropriate. Thus, in particular, the references just listed may be consulted for further amplification of the material summarized in this section.

The relation between analyst and client

Based on close study in a variety of fields of the ideal relation that should exist between professional practitioner and client, Schön (1983: 286–97) advocates a 'reflective contract' that works in this way:

> in a reflective contract between a practitioner and client, the client does not agree to accept the practitioner's authority but to suspend disbelief in it. He agrees to join the practitioner in inquiring into the situation for which the client seeks help; to try to understand what he is experiencing and to make that understanding accessible to the practitioner; to confront the practitioner when he does not understand or agree; to test the practitioner's competence by observing his effectiveness and to make public his questions over what should be counted as effectiveness; to pay for service rendered and to appreciate competence demonstrated. The practitioner agrees to deliver competent performance to the limits of his capacity; to help the client understand the meaning of the professional's advice and the rationale for his actions, while at the same time he tries to learn the meanings his actions have for the client; and to reflect on his own tacit understanding when he needs to do so in order to play his part in fulfilling the contract.

Not all practice engagements realize this ideal, of course, but successful practitioners achieve a suitable form of it in their work. In any case, the practical arrangements aimed at achieving the ideal of the reflective contract must, of necessity, be evolved in the light of the circumstances and people peculiar to each engagement. For example, for an operations research project of modest size, frequent communication between analyst and client may serve the purpose adequately; for a large project, it may be desirable for one or more members of the client's staff to join the analysis team, not only to facilitate communication but also to facilitate access to information sources. It is usually desirable for the project to make periodic progress reports as the work proceeds. In any case, keeping the client well informed about the work in progress has many virtues, not the least of which is that, as experience shows, it will help to convince all concerned of the value and importance of the work and thus act to assist in convincing the interested parties to accept and act on the findings.

6 Allied fields

There are a number of fields that are allied with operations research, either because they grew out of operations research or because they involve activities that are akin to those of operations research work.

Decision support systems (DSS) (see DECISION SUPPORT SYSTEMS) offer executives computer displays that allow them to manipulate the variables in their operations, thus enabling them not only to view the current outcomes but also to manipulate assumptions about future conditions to see what differences can occur.

The models underlying these displays, however, are ones with which operations research workers are familiar, so that DSS may be viewed, if one wishes to do so, as a special form of operations research application.

Since much operations research work supports business and management decisions, some workers have used this relation as the primary focus for their conception and hence conceive of operations research as a decision science. In addition, there are special decision theories supporting this view. Collectively, then, one can speak of the *decision sciences* (see DECISION MAKING, MULTIPLE-CRITERIA).

Similarly, much work for business and management is aimed at getting suitably digested and discriminating information before responsible executives. Systems for achieving this result are spoken of as *management information systems*, and they are supported by information technology (see INFORMATION TECHNOLOGY).

The field of *operations management* focuses its attention on a selected aspect of the operational context, but one with which most operations research workers would feel at home.

Industrial engineering, which antedates operations research by many decades, began with detailed examination of industrial processes but grew into much more comprehensive enquiries related to industrial operations. Although historically operations research began quite independently of industrial engineering, substantively it can be thought of as an extension of industrial engineering concerns to higher-level management problems and to contexts lying well outside industry. As witness to this fact, many academic departments contain both industrial engineering and operations research living together in harmony, and many practitioners find themselves equally at home under either name.

Some operations research workers are interested in artificial intelligence models (see ARTIFICIAL INTELLIGENCE) and systems analysis and design (see SYSTEMS ANALYSIS AND DESIGN), the latter being concerned with designing operating systems, often involving computers as central elements (this speciality is not to be confused with the large-scale branch of operations research called systems analysis).

7 A future perspective

There is a well-recognized tendency in mature knowledge-based professions for academic theorists to have acquired dominance over the professional view of quality, relevance, and indices of prestige and reward. While this force is strongly present in operations research, the profession retains a healthy balance and combination of interests in both practice and the underlying theories that make the advancement of practice possible. This mix of interests and the communication between them constitute one of the profession's great strengths, on which much of its current success has been built.

Thus, we can expect that this partnership of theory and practice will continue to develop and lead operations research towards deeper and more discriminating models, and its practice towards more comprehensive and effective work, while retaining the core foundations that it has already built for the profession, as summarized in this article and explored more deeply elsewhere in this encyclopedia.

Beyond these professional developments, however, there are many complex problems facing society and its institutions that call for operations research to cooperate with other specialities in working towards improvements. Operations research's record of cooperative work in such contexts gives its workers an important opportunity to lead such efforts and help to give them productive form.

Indeed, operations research and systems analysis can look forward to playing useful roles in helping to ameliorate such problems as these: how to conduct international businesses so as to benefit not only their staffs and stockholders but also the nations and their citizens that are affected; how to meet the world's energy needs while maintaining the quality of the world's atmosphere; how to manage sustainable worldwide economic development while maintaining the quality of the biosphere; how to use the world's waters to support economic development while maintaining their adequate quantity and quality; since the world already produces enough food to feed its population, how to create the economic and operational infrastructures that are needed to supply everyone with adequate nourishment; how to reduce the amount of hazardous waste produced throughout the world

此为试读，需要完整PDF请访问：www.ertongbook.com