ABOUT THE CASE WRITERS

Karl Schmedders is an associate professor in the Department of Managerial Economics and Decision Sciences at the Kellogg Graduate School of Management (Northwestern University), where he teaches quantitative methods for managerial decision making. His research interests include applications of operations research in economic theory, general equilibrium theory with incomplete markets, asset pricing, and computational economics. Dr. Schmedders received his doctorate in operations research from Stanford University, where he taught both undergraduate and graduate classes in operations research. Among the classes taught was a case studies course in operations research, and he subsequently was invited to speak at a conference sponsored by the Institute of Operations Research and the Management Sciences (INFORMS) about his successful experience with this course. He received several teaching awards at Stanford, including the university's prestigious Walter J. Gores Teaching Award. He also has been named the L. G. Lavengood Professor of the Year at the Kellogg School of Management. After teaching a course at WHU Koblenz (a leading German business school) in 2003, he won that school's Best Teacher Award as well for that semester.

Molly Stephens is an associate in the Los Angeles office of Quinn, Emanuel, Urquhart, Oliver & Hedges, LLP. She graduated from Stanford University with a B.S. in Industrial Engineering and an M.S. in Operations Research. Ms. Stephens taught public speaking in Stanford's School of Engineering and served as a teaching assistant for a case studies course in operations research. As a teaching assistant, she analyzed operations research problems encountered in the real world and the transformation of these problems into classroom case studies. Her research was rewarded when she won an undergraduate research grant from Stanford to continue her work and was invited to speak at an INFORMS conference to present her conclusions regarding successful classroom case studies. Following graduation, Ms. Stephens worked at Andersen Consulting as a systems integrator, experiencing real cases from the inside, before resuming her graduate studies to earn a J.D. degree (with honors) from the University of Texas Law School at Austin.

DEDICATION

To the memory of our parents

and

To the memory of my beloved mentor,
Gerald J. Lieberman, who was one of the true
giants of our field

CHAPTER 1

Introduction

1.1 THE ORIGINS OF OPERATIONS RESEARCH

Since the advent of the industrial revolution, the world has seen a remarkable growth in the size and complexity of organizations. The artisans' small shops of an earlier era have evolved into the billion-dollar corporations of today. An integral part of this revolutionary change has been a tremendous increase in the division of labor and segmentation of management responsibilities in these organizations. The results have been spectacular. However, along with its blessings, this increasing specialization has created new problems, problems that are still occurring in many organizations. One problem is a tendency for the many components of an organization to grow into relatively autonomous empires with their own goals and value systems, thereby losing sight of how their activities and objectives mesh with those of the overall organization. What is best for one component frequently is detrimental to another, so the components may end up working at cross purposes. A related problem is that as the complexity and specialization in an organization increase, it becomes more and more difficult to allocate the available resources to the various activities in a way that is most effective for the organization as a whole. These kinds of problems and the need to find a better way to solve them provided the environment for the emergence of **operations research** (commonly referred to as **OR**).

The roots of OR can be traced back many decades, when early attempts were made to use a scientific approach in the management of organizations. However, the beginning of the activity called *operations research* has generally been attributed to the military services early in World War II. Because of the war effort, there was an urgent need to allocate scarce resources to the various military operations and to the activities within each operation in an effective manner. Therefore, the British and then the U.S. military management called upon a large number of scientists to apply a scientific approach to dealing with this and other strategic and tactical problems. In effect, they were asked to do *research on* (military) *operations.* These teams of scientists were the first OR teams. By developing effective methods of using the new tool of radar, these teams were instrumental in winning the Air Battle of Britain. Through their research on how to better manage convoy and antisubmarine operations, they also played a major role in winning the Battle of the North Atlantic. Similar efforts assisted the Island Campaign in the Pacific.

When the war ended, the success of OR in the war effort spurred interest in applying OR outside the military as well. As the industrial boom following the war was running its course, the problems caused by the increasing complexity and specialization in organizations were again coming to the forefront. It was becoming apparent to a growing number of people, including business consultants who had served on or with the OR teams during the war, that these were basically the same problems that had been faced by the military but in a different context. By the early 1950s, these individuals had introduced the use of OR to a variety of organizations in business, industry, and government. The rapid spread of OR soon followed.

At least two other factors that played a key role in the rapid growth of OR during this period can be identified. One was the substantial progress that was made early in improving the techniques of OR. After the war, many of the scientists who had participated on OR teams or who had heard about this work were motivated to pursue research relevant to the field; important advancements in the state of the art resulted. A prime example is the *simplex method* for solving linear programming problems, developed by George Dantzig in 1947. Many of the standard tools of OR, such as linear programming, dynamic programming, queueing theory, and inventory theory, were relatively well developed before the end of the 1950s.

A second factor that gave great impetus to the growth of the field was the onslaught of the *computer revolution.* A large amount of computation is usually required to deal most effectively with the complex problems typically considered by OR. Doing this by hand would often be out of the question. Therefore, the development of electronic digital computers, with their ability to perform arithmetic calculations thousands or even millions of times faster than a human being can, was a tremendous boon to OR. A further boost came in the 1980s with the development of increasingly powerful personal computers accompanied by good software packages for doing OR. This brought the use of OR within the easy reach of much larger numbers of people. Today, literally millions of individuals have ready access to OR software. Consequently, a whole range of computers from mainframes to laptops now are being routinely used to solve OR problems.

1.2 THE NATURE OF OPERATIONS RESEARCH

As its name implies, operations research involves "research on operations." Thus, operations research is applied to problems that concern how to conduct and coordinate the *operations* (i.e., the *activities*) within an organization. The nature of the organization is essentially immaterial, and, in fact, OR has been applied extensively in such diverse areas as manufacturing, transportation, construction, telecommunications, financial planning, health care, the military, and public services, to name just a few. Therefore, the breadth of application is unusually wide.

The *research* part of the name means that operations research uses an approach that resembles the way research is conducted in established scientific fields. To a considerable extent, the *scientific method* is used to investigate the problem of concern. (In fact, the term *management science* sometimes is used as a synonym for operations research.) In particular, the process begins by carefully observing and formulating the problem, including gathering all relevant data. The next step is to construct a scientific (typically mathematical) model that attempts to abstract the essence of the real problem. It is then hypothesized that this model is a sufficiently precise representation of the essential features of the situation that the conclusions (solutions) obtained from the model are also valid for the real problem. Next, suitable experiments are conducted to test this hypothesis, modify it as needed, and eventually verify some form of the hypothesis. (This step is frequently referred to as *model validation.*) Thus, in a certain sense, operations research involves creative scientific research into the fundamental properties of operations. However, there is more to it than this. Specifically, OR is also concerned with the practical management of the organization. Therefore, to be successful, OR must also provide positive, understandable conclusions to the decision maker(s) when they are needed.

Still another characteristic of OR is its broad viewpoint. As implied in the preceding section, OR adopts an organizational point of view. Thus, it attempts to resolve the conflicts of interest among the components of the organization in a way that is best for the organization as a whole. This does not imply that the study of each problem must give explicit consideration to all aspects of the organization; rather, the objectives being sought must be consistent with those of the overall organization.

An additional characteristic is that OR frequently attempts to find a *best* solution (referred to as an *optimal* solution) for the problem under consideration. (We say *a* best instead of *the* best solution because there may be multiple solutions tied as best.) Rather than simply improving the status quo, the goal is to identify a best possible course of action. Although it must be interpreted carefully in terms of the practical needs of management, this "search for optimality" is an important theme in OR.

All these characteristics lead quite naturally to still another one. It is evident that no single individual should be expected to be an expert on all the many aspects of OR work or the problems typically considered; this would require a group of individuals having diverse backgrounds and skills. Therefore, when a full-fledged OR study of a new problem is undertaken, it is usually necessary to use a *team approach.* Such an OR team typically needs to include individuals who collectively are highly trained in mathematics, statistics and probability theory, economics, business administration, computer science, engineering and the physical sciences, the behavioral sciences, and the special techniques of OR. The team also needs to have the necessary experience and variety of skills to give appropriate consideration to the many ramifications of the problem throughout the organization.

1.3 THE IMPACT OF OPERATIONS RESEARCH

Operations research has had an impressive impact on improving the efficiency of numerous organizations around the world. In the process, OR has made a significant contribution to increasing the productivity of the economies of various countries. There now are a few dozen member countries in the International Federation of Operational Research Societies (IFORS), with each country having a national OR society. Both Europe and Asia have federations of OR societies to coordinate holding international conferences and publishing international journals in those continents. In addition, the Institute for Operations Research and the Management Sciences (INFORMS) is an international OR society. Among its various journals is one called *Interfaces* that regularly publishes articles describing major OR studies and the impact they had on their organizations.

To give you a better notion of the wide applicability of OR, we list some actual award-winning applications in Table 1.1. Note the diversity of organizations and applications in the first two columns. The curious reader can find a complete article describing each application in the January–February issue of *Interfaces* for the year cited in the third column of the table. The fourth column lists the chapters in *this* book that describe the kinds of OR techniques that were used in the application. (Note that many of the applications combine a variety of techniques.) The last column indicates that these applications typically resulted in annual savings in the millions (or even tens of millions) of dollars. Furthermore, additional benefits not recorded in the table (e.g., improved service to customers and better

TABLE 1.1 Some applications of operations research

Organization	Nature of Application	Year of Publication*	Related Chapters†	Annual Savings
The Netherlands Rijkswaterstaat	Develop national water management policy, including mix of new facilities, operating procedures, and pricing.	1985	2–8, 12, 20	$15 million
Monsanto Corp.	Optimize production operations in chemical plants to meet production targets with minimum cost.	1985	2, 11	$2 million
United Airlines	Schedule shift work at reservation offices and airports to meet customer needs with minimum cost.	1986	2–9, 11, 17, 26, 27	$6 million
Citgo Petroleum Corp.	Optimize refinery operations and the supply, distribution, and marketing of products.	1987	2–9, 27	$70 million
San Francisco Police Department	Optimally schedule and deploy police patrol officers with a computerized system.	1989	2–4, 11, 27	$11 million
Texaco, Inc.	Optimally blend available ingredients into gasoline products to meet quality and sales requirements.	1989	2, 12	$30 million
IBM	Integrate a national network of spare parts inventories to improve service support.	1990	2, 18, 20	$20 million +$250 million less inventory
Yellow Freight System, Inc.	Optimize the design of a national trucking network and the routing of shipments.	1992	2, 9, 12, 20, 27	$17.3 million
New Haven Health Department	Design an effective needle exchange program to combat the spread of HIV/AIDS.	1993	2	33% less HIV/AIDS
AT&T	Develop a PC-based system to guide business customers in designing their call centers.	1993	17, 20, 26	$750 million
Delta Airlines	Maximize the profit from assigning airplane types to over 2500 domestic flights.	1994	11	$100 million
Digital Equipment Corp.	Restructure the global supply chain of suppliers, plants, distribution centers, potential sites, and market areas.	1995	11	$800 million
China	Optimally select and schedule massive projects for meeting the country's future energy needs.	1995	11	$425 million
South African defense force	Optimally redesign the size and shape of the defense force and its weapons systems.	1997	11	$1.1 billion
Proctor and Gamble	Redesign the North American production and distribution system to reduce costs and improve speed to market.	1997	8	$200 million
Taco Bell	Optimally schedule employees to provide desired customer service at a minimum cost.	1998	11, 20, 27	$13 million
Hewlett-Packard	Redesign the sizes and locations of buffers in a printer production line to meet production goals.	1998	17, 26	$280 million more revenue
Sears, Roebuck	Develop a vehicle-routing and scheduling system for delivery and home service fleets.	1999	11	$42 million
IBM	Reengineer its global supply chain to respond quicker to customers while holding minimal inventory.	2000	18	$750 million in first year
Merrill Lynch	Design asset-based and direct online pricing options for providing financial services.	2002	2, 20	$80 million more revenue
Samsung Electronics	Develop methods of reducing manufacturing times and inventory levels.	2002	2–7	$200 million more revenue
Continental Airlines	Optimize the reassignment of crews to flights when schedule disruptions occur.	2003	2, 11	$40 million

*Pertains to a January–February issue of *Interfaces* in which a complete article can be found describing the application.

†Refers to chapters in this book that describe the kinds of OR techniques used in the application.

managerial control) sometimes were considered to be even more important than these financial benefits. (You will have an opportunity to investigate these less tangible benefits further in Probs. 1.3-1 and 1.3-2.)

Selected Reference 1 at the end of the chapter provides a follow-up on the long-term strategic impact that many of these applications had on their companies. Selected Reference 3 describes some other applications and the key role that operations research plays in increasing the profitability and productivity of numerous companies.

Although most routine OR studies provide considerably more modest benefits than the award-winning applications summarized in Table 1.1, the figures in the rightmost column of this table do accurately reflect the dramatic impact that large, well-designed OR studies occasionally can have.

We will briefly describe some of these applications in Chapter 2, and then we present two in greater detail as case studies in Sec. 3.5.

1.4 ALGORITHMS AND OR COURSEWARE

An important part of this book is the presentation of the major **algorithms** (systematic solution procedures) of OR for solving certain types of problems. Some of these algorithms are amazingly efficient and are routinely used on problems involving hundreds or thousands of variables. You will be introduced to how these algorithms work and what makes them so efficient. You then will use these algorithms to solve a variety of problems on a computer. The **OR Courseware** contained in the CD-ROM that accompanies the book will be a key tool for doing all this.

One special feature in your OR Courseware is a program called **OR Tutor.** This program is intended to be your personal tutor to help you learn the algorithms. It consists of many *demonstration examples* that display and explain the algorithms in action. These "demos" supplement the examples in the book.

In addition, your OR Courseware includes a special software package called **Interactive Operations Research Tutorial,** or **IOR Tutorial** for short. Implemented in Java, this innovative package is designed specifically to enhance the learning experience of students using this book. IOR Tutorial includes many *interactive procedures* for executing the algorithms interactively in a convenient format. The computer does all the routine calculations while you focus on learning and executing the logic of the algorithm. You should find these interactive procedures a very efficient and enlightening way of doing many of your homework problems. IOR Tutorial also includes a number of other helpful procedures, including some *automatic procedures* for executing algorithms automatically and several procedures that provide graphical displays of how the solution provided by an algorithm varies with the data of the problem.

In practice, the algorithms normally are executed by commercial software packages. We feel that it is important to acquaint students with the nature of these packages that they will be using after graduation. Therefore, your IOR Tutorial includes a wealth of material to introduce you to three particularly popular software packages described below. Together, these packages will enable you to solve nearly all the OR models encountered in this book very efficiently. We have added our own *automatic procedures* to IOR Tutorial in a few cases where these packages are not applicable.

A very popular approach now is to use today's premier spreadsheet package, *Microsoft Excel,* to formulate small OR models in a spreadsheet format. The **Excel Solver** (or an enhanced version of this add-in, such as **Premium Solver for Education** included in your OR Courseware) then is used to solve the models. Your OR Courseware includes separate

Excel files for nearly every chapter in this book. Each time a chapter presents an example that can be solved using Excel, the complete spreadsheet formulation and solution is given in that chapter's Excel files. For many of the models in the book, an *Excel template* also is provided that already includes all the equations necessary to solve the model. Some *Excel add-ins* also are included on the CD-ROM.

After many years, **LINDO** (and its companion modeling language **LINGO**) continues to be a popular OR software package. Student versions of LINDO and LINGO now can be downloaded free from the Web. This student version also is provided in your OR Courseware. As for Excel, each time an example can be solved with this package, all the details are given in a LINGO/LINDO file for that chapter in your OR Courseware.

CPLEX is an elite state-of-the-art software package that is widely used for solving large and challenging OR problems. When dealing with such problems, it is common to also use a *modeling system* to efficiently formulate the mathematical model and enter it into the computer. **MPL** is a user-friendly modeling system that uses CPLEX as its main solver. A student version of MPL and CPLEX is available free by downloading it from the Web. For your convenience, we also have included this student version in your OR Courseware. Once again, all the examples that can be solved with this package are detailed in MPL/CPLEX files for the corresponding chapters in your OR Courseware.

We will further describe these three software packages and how to use them later (especially near the end of Chaps. 3 and 4). Appendix 1 also provides documentation for the OR Courseware, including OR Tutor and IOR Tutorial.

To alert you to relevant material in OR Courseware, the end of each chapter from Chap. 3 onward has a list entitled *Learning Aids for This Chapter on the CD-ROM*. As explained at the beginning of the problem section for each of these chapters, symbols also are placed to the left of each problem number or part where any of this material (including demonstration examples and interactive procedures) can be helpful.

Another learning aid provided on the CD-ROM is a set of **Worked Examples** for each chapter (from Chap. 3 onward). These complete examples supplement the examples in the book for your use as needed, but without interrupting the flow of the material on those many occasions when you don't need to see an additional example. You also might find these supplementary examples helpful when preparing for an examination. We always will mention whenever a supplementary example on the current topic is included in the Worked Examples section of the CD-ROM.

The CD-ROM also includes a glossary for each chapter.

SELECTED REFERENCES

1. Bell, P. C., C. K. Anderson, and S. P. Kaiser: "Strategic Operations Research and the Edelman Prize Finalist Applications 1989–1998," *Operations Research,* **51**(1): 17–31, Jan.–Feb. 2003.
2. Gass, S. I., and C. M. Harris (eds.): *Encyclopedia of Operations Research and Management Science,* 2d ed., Kluwer Academic Publishers, Boston, 2001.
3. Horner, P. (ed.): "Special Issue: Executive's Guide to Operations Research," *OR/MS Today,* Institute for Operations Research and the Management Sciences, **27**(3), June 2000.
4. Kirby, M. W.: "Operations Research Trajectories: The Anglo-American Experience from the 1940s to the 1990s," *Operations Research,* **48**(5): 661–670, Sept.–Oct. 2000.
5. Miser, H. J.: "The Easy Chair: What OR/MS Workers Should Know About the Early Formative Years of Their Profession," *Interfaces,* **30**(2): 99–111, March–April 2000.
6. Wein, L. M. (ed.): "50th Anniversary Issue," *Operations Research* (a special issue featuring personalized accounts of some of the key early theoretical and practical developments in the field), **50**(1), Jan.–Feb. 2002.

PROBLEMS

1.3-1. Select one of the applications of operations research listed in Table 1.1. Read the article describing the application in the January–February issue of *Interfaces* for the year indicated in the third column. Write a two-page summary of the application and the benefits (including nonfinancial benefits) it provided.

1.3-2. Select three of the applications of operations research listed in Table 1.1. Read the articles describing the applications in the January–February issue of *Interfaces* for the years indicated in the third column. For each one, write a one-page summary of the application and the benefits (including nonfinancial benefits) it provided.

CHAPTER 2

Overview of the Operations Research Modeling Approach

The bulk of this book is devoted to the mathematical methods of operations research (OR). This is quite appropriate because these quantitative techniques form the main part of what is known about OR. However, it does not imply that practical OR studies are primarily mathematical exercises. As a matter of fact, the mathematical analysis often represents only a relatively small part of the total effort required. The purpose of this chapter is to place things into better perspective by describing all the major phases of a typical OR study.

One way of summarizing the usual (overlapping) phases of an OR study is the following:

1. Define the problem of interest and gather relevant data.
2. Formulate a mathematical model to represent the problem.
3. Develop a computer-based procedure for deriving solutions to the problem from the model.
4. Test the model and refine it as needed.
5. Prepare for the ongoing application of the model as prescribed by management.
6. Implement.

Each of these phases will be discussed in turn in the following sections.

Most of the award-winning OR studies introduced in Table 1.1 provide excellent examples of how to execute these phases well. We will intersperse snippets from these examples throughout the chapter, with references to invite your further reading.

2.1 DEFINING THE PROBLEM AND GATHERING DATA

In contrast to textbook examples, most practical problems encountered by OR teams are initially described to them in a vague, imprecise way. Therefore, the first order of business is to study the relevant system and develop a well-defined statement of the problem to be considered. This includes determining such things as the appropriate objectives, constraints on what can be done, interrelationships between the area to be studied and other areas of the organization, possible alternative courses of action, time limits for making a decision, and so on. This process of problem definition is a crucial one because it greatly affects how relevant the conclusions of the study will be. It is difficult to extract a "right" answer from the "wrong" problem!

The first thing to recognize is that an OR team normally works in an *advisory capacity.* The team members are not just given a problem and told to solve it however they see fit. Instead, they advise management (often one key decision maker). The team performs a detailed technical analysis of the problem and then presents recommendations to management. Frequently, the report to management will identify a number of alternatives that are particularly attractive under different assumptions or over a different range of values of some policy parameter that can be evaluated only by management (e.g., the trade-off between *cost* and *benefits*). Management evaluates the study and its recommendations, takes into account a variety of intangible factors, and makes the final decision based on its best judgment. Consequently, it is vital for the OR team to get on the same wavelength as management, including identifying the "right" problem from management's viewpoint, and to build the support of management for the course that the study is taking.

Ascertaining the *appropriate objectives* is a very important aspect of problem definition. To do this, it is necessary first to identify the member (or members) of management who actually will be making the decisions concerning the system under study and then to probe into this individual's thinking regarding the pertinent objectives. (Involving the decision maker from the outset also is essential to build her or his support for the implementation of the study.)

By its nature, OR is concerned with the welfare of the *entire organization* rather than that of only certain of its components. An OR study seeks solutions that are optimal for the overall organization rather than suboptimal solutions that are best for only one component. Therefore, the objectives that are formulated ideally should be those of the entire organization. However, this is not always convenient. Many problems primarily concern only a portion of the organization, so the analysis would become unwieldy if the stated objectives were too general and if explicit consideration were given to all side effects on the rest of the organization. Instead, the objectives used in the study should be as specific as they can be while still encompassing the main goals of the decision maker and maintaining a reasonable degree of consistency with the higher-level objectives of the organization.

For profit-making organizations, one possible approach to circumventing the problem of suboptimization is to use *long-run profit maximization* (considering the time value of money) as the sole objective. The adjective *long-run* indicates that this objective provides the flexibility to consider activities that do not translate into profits *immediately* (e.g., research and development projects) but need to do so *eventually* in order to be worthwhile. This approach has considerable merit. This objective is specific enough to be used conveniently, and yet it seems to be broad enough to encompass the basic goal of profit-making organizations. In fact, some people believe that all other legitimate objectives can be translated into this one.

However, in actual practice, many profit-making organizations do not use this approach. A number of studies of U.S. corporations have found that management tends to adopt the goal of *satisfactory profits,* combined with *other objectives,* instead of focusing on long-run profit maximization. Typically, some of these *other* objectives might be to maintain stable profits, increase (or maintain) one's share of the market, provide for product diversification, maintain stable prices, improve worker morale, maintain family control of the business, and increase company prestige. Fulfilling these objectives might achieve long-run profit maximization, but the relationship may be sufficiently obscure that it may not be convenient to incorporate them all into this one objective.

Furthermore, there are additional considerations involving social responsibilities that are distinct from the profit motive. The five parties generally affected by a business firm located in a single country are (1) the *owners* (stockholders, etc.), who desire profits (dividends, stock appreciation, and so on); (2) the *employees,* who desire steady employment at reasonable wages; (3) the *customers,* who desire a reliable product at a reasonable price; (4) the *suppliers,* who desire integrity and a reasonable selling price for their goods; and (5) the *government* and hence the *nation,* which desire payment of fair taxes and consideration of the national interest. All five parties make essential contributions to the firm, and the firm should not be viewed as the exclusive servant of any one party for the exploitation of others. By the same token, international corporations acquire additional obligations to follow socially responsible practices. Therefore, while granting that management's prime responsibility is to make profits (which ultimately benefits all five parties), we note that its broader social responsibilities also must be recognized.

OR teams typically spend a surprisingly large amount of time *gathering relevant data* about the problem. Much data usually are needed both to gain an accurate understanding of the problem and to provide the needed input for the mathematical model being formulated in the next phase of study. Frequently, much of the needed data will not be available when the study begins, either because the information never has been kept or because what was kept is outdated or in the wrong form. Therefore, it often is necessary to install a new computer-based *management information system* to collect the necessary data on an ongoing basis and in the needed form. The OR team normally needs to enlist the assistance of various other key individuals in the organization to track down all the vital data. Even with this effort, much of the data may be quite "soft," i.e., rough estimates based only on educated guesses. Typically, an OR team will spend considerable time trying to improve the precision of the data and then will make do with the best that can be obtained.

With the widespread use of databases and the explosive growth in their sizes in recent years, OR teams now frequently find that their biggest data problem is not that too little is available but that there is too much data. There may be thousands of sources of data, and the total amount of data may be measured in gigabytes or even terrabytes. In this environment, locating the particularly relevant data and identifying the interesting patterns in these data can become an overwhelming task. One of the newer tools of OR teams is a technique called **data mining** that addresses this problem. Data mining methods search large databases for interesting patterns that may lead to useful decisions. (Selected Reference 1 at the end of the chapter provides further background about data mining.)

Examples. An OR study done for the **San Francisco Police Department**[1] resulted in the development of a computerized system for optimally scheduling and deploying police patrol officers. The new system provided annual savings of \$11 million, an annual \$3 million increase in traffic citation revenues, and a 20 percent improvement in response times. In assessing the *appropriate objectives* for this study, three fundamental objectives were identified:

1. Maintain a high level of citizen safety.
2. Maintain a high level of officer morale.
3. Minimize the cost of operations.

To satisfy the first objective, the police department and city government jointly established a desired level of protection. The mathematical model then imposed the requirement that

[1]P. E. Taylor and S. J. Huxley, "A Break from Tradition for the San Francisco Police: Patrol Officer Scheduling Using an Optimization-Based Decision Support System," *Interfaces,* **19**(1): 4–24, Jan.–Feb. 1989. See especially pp. 4–11.