OPTIMIZATION MODELS FOR PLANNING AND ALLOCATION: TEXT AND CASES IN MATHEMATICAL PROGRAMMING

Roy D. Shapiro

OPTIMIZATION MODELS FOR PLANNING AND ALLOCATION: TEXT AND CASES IN MATHEMATICAL PROGRAMMING

Roy D. Shapiro

Graduate School of Business Administration
Harvard University

JOHN WILEY & SONS
New York Chichester Brisbane Toronto Singapore

PREFACE

During the past decade, the teaching of management science has begun to shift away from a pedagogy based entirely on the theory toward one that emphasizes the application of models in actual practice. Rather than focus primarily on the theoretical development of algorithms, the latter approach stresses the practical considerations involved in the recognition and formulation of appropriate models and the subsequent analysis of the economic and managerial implications of the solution results.

Over the past three decades, the development of management science techniques has far outpaced our understanding of how and when and, in fact, whether to use them. The successful practice of management science requires the recognition that these concepts and techniques are applied to problems involving people—people who interact with each other and their environment within the confines of an oftentimes complex organizational structure. Thus, the study of management science is incomplete without exposure to actual decision problems faced by actual managers—problems fraught with the ambiguity and uncertainty that often accompany decisions made in the world of real organizations. The cases in this book provide this exposure.

For many students who are accustomed to textbook exercises, an approach that emphasizes problem identification as well as problem solution may seem disconcerting. For most students, the entire educational process has attuned them to finding answers (using well-defined methods) to precisely structured questions. In practice, decision problems can rarely be solved that easily. First, there is almost never a single right answer; second, determining what the problem really is often requires far more insight and ability than its analysis once it is precisely defined. Thus, providing the student with real situations that lacks rigid structure is an essential element in a management education.

Mathematical programming is one of the most widely used and certainly one of the most successful of management science methodologies. Over the nearly four decades since its inception, it has attracted the attention of managers, engineers, social scientists, and public policy makers as one of the most powerful and flexible tools available to support decision making. Hence an understanding of these tools is useful to all levels of management; yet due to their mathematical complexity, these are typically not well understood by the nonexpert. In view of this need the book is structured not to teach mathematics but to demystify the often elusive concepts underlying these optimization models.

Optimization Models for Planning and Allocation emphasizes the formulation of mathematical programming models, the use of computer output to gain insight as well as answers, and the discussion of the difficult implementation problems that often accompany the development and use of quantitative models in the corporate decision-making process. While the purpose of the text is not to train technicians, it is important that both line and staff management understand the technology that can be brought to bear in making decisions for planning and allocation. Thus, the material in the book has been structured to be suitable for a wide audience from top management who will use this technology to middle management who will initiate and/or manage the development of appropriate models to staff who will participate in the actual model design and development. Observations of the uses of quantitative models in corporate settings has reinforced the conviction that problems often result from the gap between line management and technical staff. Thus, the aim of this book is to educate managers who might effectively act as a liaison between these two groups, understanding *and* being able to communicate the needs of both. Both students with a great deal of technical training and experience and those with limited backgrounds can use the material and benefit from this orientation.

This focus, rather than the more traditional focus on theory and algorithms, is intensified by a sequence of *cases* that highlights both the formulation of optimization models and the subsequent analysis of the economic and managerial implications of the solution results. Typically, a case is a description of a real-life management situation that the student is expected to analyze, discuss in class, and suggest some recommendations for action. Rather than having students memorize facts and techniques, this pedagogical approach emphasizes learning by "doing," through the process of exposure to a series of cases covering a variety of different problems in different industries.

There are two reasons that the cases in the book were chosen. First, I want students to follow the progress of management science projects. Exposure to the choices real analysts face in designing models for actual situations is invaluable; it illustrates the benefits and pitfalls in the use of mathematical programming practice. Second, I think that students should have the opportunity to do their own model development in an environment as close to the real one as is possible in the classroom—an environment of people and thus personalities, of organizational hierarchy and thus political complexity. Developing management science approaches to a case situation forces the student to come to grips with these issues and to grapple with problem definition, implemention, and communication with different levels of the organization—as well as with methodology. Perhaps more importantly, these cases provide "hands-on" management science experience. Charles Gragg, one of the more outspoken and best remembered advocates of the use of cases in business education, wrote:

> *It can be said flatly that the mere act of listening to wise statements and sound advice does little for anyone . . . no amount of information, whether theory or fact, in itself improves insight or judgment or increases ability to act wisely under conditions of responsibility. . . . In the process of learning, the learner's dynamic cooperation is required. Such cooperation from students does not arise automatically, however. It has to be provided for and continually encouraged.*[1]

The cases herein provide this encouragement.

Optimization Models for Planning and Allocation has been designed for a graduate course in applied mathematical programming, especially for students who may soon be in a position to develop and/or implement such models. The book has *not* been designed for the "technical course in optimization theory" of the sort offered by many operations research or mathematics departments. However, enough underlying theoretical and conceptual development has been provided to allow future managers to understand the techniques. Nor is the text appropriate for those who have had no prior exposure to management science, although the text does, for the sake of completeness, review the initial conceptual underpinnings that would be presented in a first course such as "Introduction to Management Science" or "Quantitative Methods in Management." The cases, however, clearly the *raison d'etre* for this book, can be taught at several different levels. Many of these cases have been used successfully with audiences ranging from public policy makers and practicing business executives to students in MBA programs and in graduate operations research courses.

Apart from the desirability of a prior course in management science, which covers, among other topics, the basics of linear programming and the presumption of some familiarity with basic mathematical concepts (e.g., graphing functions, using equations), there are few prerequisites for using this book. I have tried to deal as little as possible with complex mathematics and have

[1]Charles Gragg, *Because Wisdom Can't Be Told*, distributed by the Intercollegiate Case Clearing House, ICCH No 9-451-005

structured the material so that neither experience with linear algebra, familiarity with vector and matrix notation, nor knowledge of calculus is required. I have assumed that students can handle equations containing summations, although I have used that notation as sparingly as possible. When uncertainty is explicitly brought into the picture, I presume at least a fleeting prior exposure to expected values and cumulative probability distributions, and use decision diagrams as occasional illustrations.

Chapter 1 briefly discusses a taxonomy of problem solving approaches and indicates how optimization models fit. Chapters 2 and 3 provide an in-depth discussion of linear programming, the topic that not only is most important in its own right, given the vast majority of actual applications, but that introduces the major features of all mathematical programming methodologies. Chapter 4 gives some insight as to how multiple criteria may be incorporated into the optimization process.

Chapters 5 through 9 offer a range of special mathematical programming structures which have found or are beginning to find "their place in the sun." Chapters 5 and 6 introduce combinatorial optimization, covering networks and integer programming. Chapters 7, 8, and 9 deal with more advanced topics: incorporating uncertainty into the optimization process, nonlinear optimization, and an approach for sequential decision problems—dynamic programming.

February 1983 *Roy Shapiro*

CONTENTS

Chapter 1

MODELS AND OPTIMIZATION

In today's world, it would be hard to find many successful managers, engineers, or social scientists who had not benefited from meaningful insights into the "systems" with which they deal which were made possible by the creation and use of *models*—those physical or mathematical abstractions that, although simplified, reflect the key interactions of the system variables. Indeed, one of the key truths of "modern management" involves the recognition that today's problem solving requires talents beyond those inherent in the combination of intuition and experience. Central to the material in this book is the concept of the model as an insight-producing aid to managerial decision making. In this vein, this prefatory chapter is meant to serve as a brief, but important, initial discussion of modeling in general and optimization modeling in particular, prior to the detailed study that begins in Chapter 2.

Since the dawn of civilization, we have constructed simplified models of our environment so as to understand that environment better. Modeling as a means of analysis and, ultimately, insight, has been with us since our earliest beginnings. The first model was, without a doubt, a physical model, constructed crudely. This early experimentation is a far cry from the analysis made possible today to the scientist, engineer, or manager armed with a mathematical model, an optimization algorithm, and a computer. The advances in management science and the phenomenal developments in computer hardware and software have created an environment where an understanding of the construction and use of models is an indispensible part of the education of a modern manager. There is scarcely an industry today that has not experienced the significant impact of model-based planning and analysis. Unfortunately, this impact has not always been positive. One key reason for this often referred to "failure of management science models" is the refusal of management to become involved in the modeling process. Models are too important to be left to analysts alone. The manager, if he or she wishes the model to be useful and *used*, must be part of the development. This requires some grounding in the "art" of modeling and the "science" of solution techniques. It is to this end that this casebook is devoted.

We begin by suggesting a taxonomy of decision-making or problem-solving approaches so as to see where the mathematical model fits in, and then how optimization fits in the continuum of mathematical models.

A TAXONOMY OF PROBLEM-SOLVING APPROACHES

A team of aircraft engineers needs to determine appropriate wing dimensions and shape for a new aircraft which they have designed. Perhaps the conceptually simplest approach to solving this problem is to build a variety of aircraft, utilizing a range of alternative wing configurations. These aircraft could then be subjected to extensive *in-flight* monitoring to determine the performance characteristics of each. This approach—*direct experimentation*—has the advantage of representational accuracy, that is, each wing being tested is the actual alternative under consideration. On the other hand, the time, expense, and difficulty that would be experienced in implementing this series of experiments would severely limit the number of alternatives that could be tested. As such, this approach would likely be deferred until all but a few alternative designs had been screened out.

A second approach, one utilized extensively in the engineering sciences, would be the construction of a variety of *physical models*. These models could then be subjected to a series of tests in simulated environments (e.g., a wind tunnel) to predict how the real wing might perform in flight. Note that this alternative offers a reduced level of representational accuracy. That is, in certain ways, the actual wing would no doubt act differently in flight than the model in a wind tunnel. On the other hand, this approach allows less time consuming and less costly evaluation of far more alternatives than would be possible with direct experimentation.

A third approach would be the formulation of a *mathematical model*. The mathematical model would replace with equations those characteristics of performance the physical model would reveal in the wind tunnel. For example, the mathematical model of wing performance would incorporate laws of motion, momentum, and air flow—in essence, equations that capture precisely how the *system variables* interact. Of the three approaches discussed, this would be the easiest, quickest, and least expensive to implement, especially with modern day computers. Unfortunately it would also require the greatest degree of simplification—it would provide the least representational accuracy.

Indeed, the major difficulties in using a mathematical model are (1) assuring that the model captures the key features of the system under study as well as the interrelationships among these features that determine causes and effects, and (2) interpreting the "answer" to a simplified picture in a way that will provide useful insight in dealing with the actual environment, which is far more complex than the model reflects.

The major benefit of the mathematical model is its ability to evaluate a vast array of alternatives in a fraction of the time and at a fraction of the cost necessary for the construction and testing of the same number of physical alternatives. Figure 1-1 illustrates these features.

While all three approaches might be feasible for the above example, the mathematical model may be the only feasible choice for the analysis of a variety of managerial or social problems. Consider, for example, the plight of a manager trying to determine an economically viable location for the firm's new factory. While it might be possible to actually construct a plant at a site under consideration and then observe the economic impacts of that choice, such a direct experimentation approach would be ludicrous. A physical model would be of no help, since the performance to be monitored is economic, not subject to physical laws. Instead, as illustrated by Figure 1-2, the manager might describe the interaction of the system variables via a set of equations describing—for example, transport of raw materials to the site, the production process, and the distribution of finished goods to the firm's warehouses, capturing the economics of each set of activities within those equations. Similarly, competitive and market forces might be modeled with a series of relationships that describe how customers or competitors might react to specified actions by the firm. This mathematical model might then be applied to a set of locations under consideration and, for each, would predict its economic performance over time.

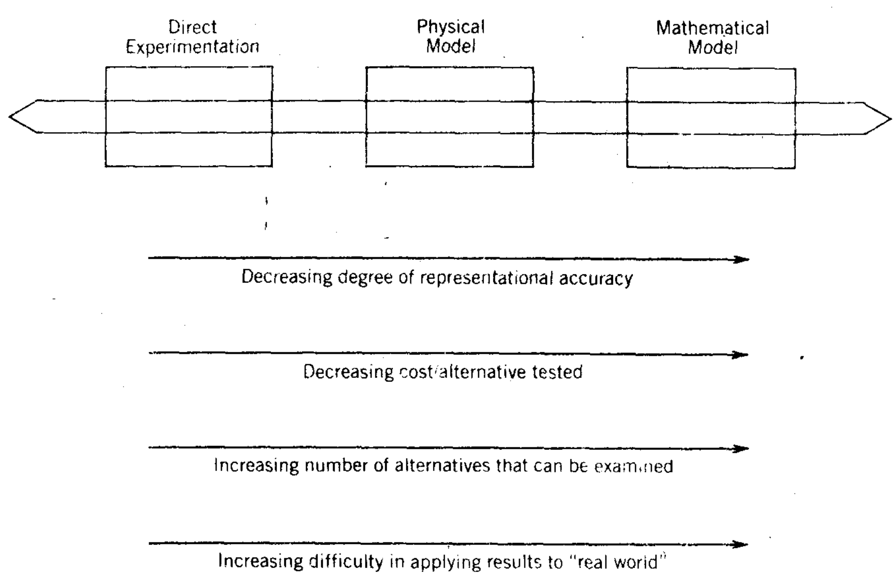

FIGURE 1-1. A continuum of problem-solving approaches.

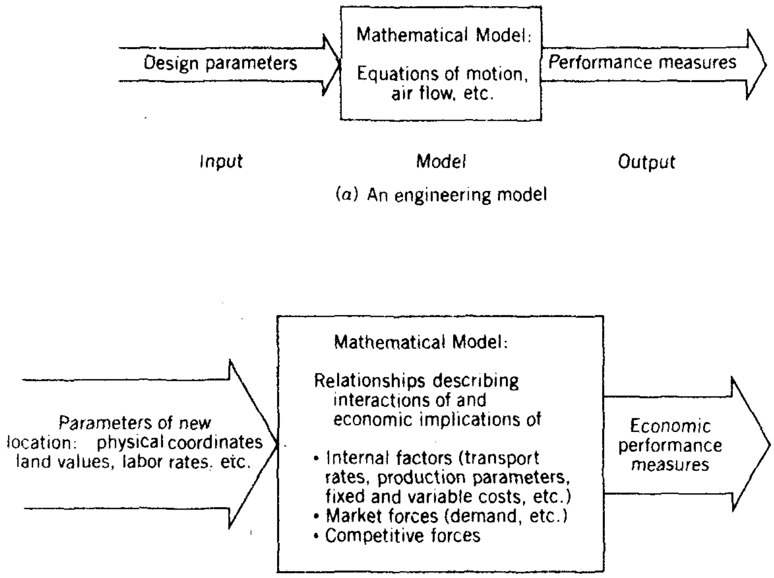

(*a*) An engineering model

(*b*) A managerial model

FIGURE 1-2. Input-output schematics for two models.

EVALUATION VERSUS OPTIMIZATION

A further categorization of problem solving approaches is helpful in understanding mathematical models. The preceding examples have been examples of *evaluative* models. That is, they operate by evaluating alternatives one at a time—they neither generate alternatives nor do they choose the "best." Those activities are undertaken by the human problem solver or decision maker. Simulations are perhaps the best known and most commonly used examples of evaluative models. They permit a great deal of flexibility, and a creative modeler with a good understanding of the underlying system can guarantee a relatively high level of representational accuracy. On the other hand, there is a practical limit to the number of alternatives that can be evaluated. Typically, evaluative models are used in a *satisficing* mode; that is, the analyst or manager will evaluate candidate plans until he or she is satisfied with the performance of one or more, rather than in an *optimizing* mode, where the best possible plan is determined. With an evaluative model, there can be no guarantee that the chosen alternative—the best alternative evaluated—is the best *possible* alternative. That guarantee requires the use of an optimization model.

Most simply, an optimization model is an evaluative model with added features. Those added capabilities are:

1. The ability to automatically generate new alternatives.

2. The ability to test whether a given alternative that has just been evaluated is the "best."

The first step in constructing an optimization model is the same as that necessary in constructing an evaluative model-the description of the system: the specification of the set of equations (or inequalities) that characterizes the interrelationship of key system variables. Most of the effort in any analysis will, in practice, be devoted to understanding the system and translating that understanding into quantitative form. This is as it should be; a model without understanding has little value.

The second step is the adoption of a measure of system effectiveness, a measure by which we may judge, "Plan A is *better* than Plan B." Sometimes, in a physical or simplified economic system, the criterion to be chosen is easy—efficiency, speed, cost, or profit. Often, however, in a social, political, or more realistic economic system, conflicting objectives make the choice of any one criterion difficult: Decreased inflation or decreased unemployment? A healthy economy or a healthy natural environment? Even for the individual firm, there are issues to be resolved—for example, short-term gain or long-term stability? Chapter 4 will suggest optimization methods to help ease this criterion-selection difficulty.

Step three is the choice of an algorithm that will optimize; that is, consistent with any restrictions which may be imposed by or on the system, as embodied by the mathematical description of that system (step one), the algorithm will choose that plan that, when evaluated along the dimensions suggested by the chosen measure of effectiveness (step two), achieves the best score of *all* possible feasible plans. We might think of optimization algorithms as operating as shown by the schematic flow chart of Figure 1-3. The process starts with an initial candidate plan, either user-generated or generated by the algorithm itself. The parameters (e.g., operating levels for the various activities) of the plan are fed to an evaluative model that, through the mathematics which describe the interrelationships among the system variables, "scores" the

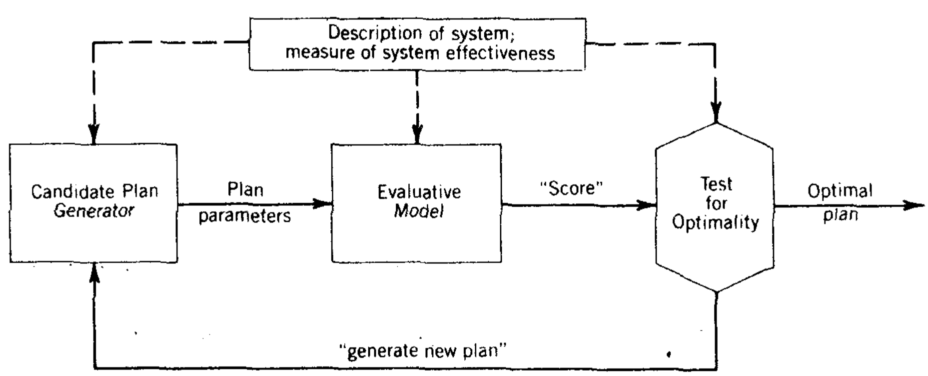

FIGURE 1-3. An optimization schematic.

plan with respect to the measure of system effectiveness specified. Next is a key to the optimization algorithm's success—a test for optimality, a test that determines whether any further improvement (with respect to the chosen criterion) is possible. If not, then the candidate plan under study is optimal. If improvement is possible, the algorithm cycles back and generates a new candidate plan, one that is created within the restrictions or guidelines embodied in the system description. Especially successful algorithms guarantee that the new plan will provide better performance (i.e., will score better) than the previous plan. Another key success factor is the (possibly empirical) guarantee that this cycling will end quickly.

THE MODELING PROCESS: A FRAMEWORK FOR ANALYSIS

This book focuses less on the theory behind the specific details of various optimization algorithms than on the practical considerations involved in the recognition and formulation of appropriate models and the subsequent analysis of the economic and managerial implications of the solution results. The schematic "bow tie" in Figure 1-4 illustrates our view of the modeling process: A necessary precursor to the formulation or use of a mathematical model is *understanding*: What drives the system under consideration? What are the key factors in predicting system behavior? What is the managerially relevant problem? What are the key measures of success? Only after these questions

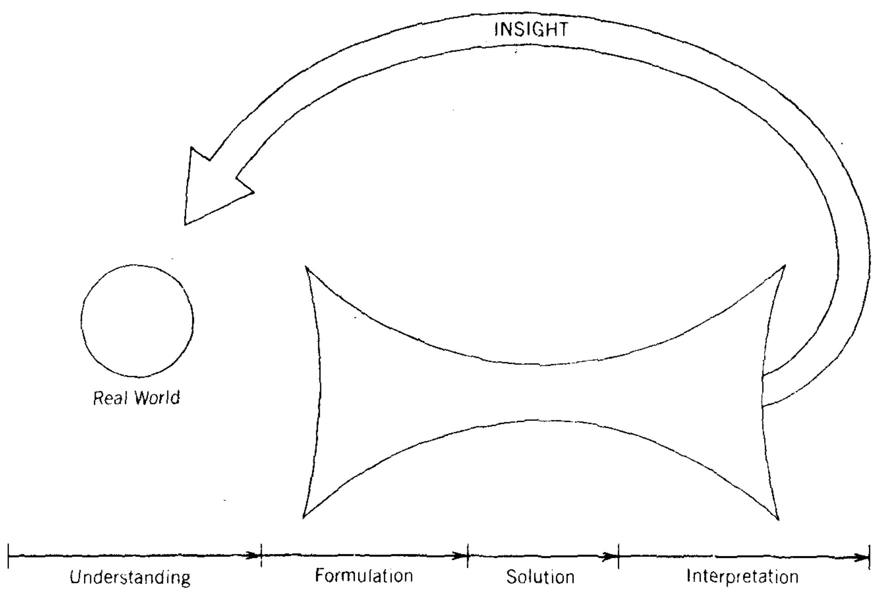

FIGURE 1-4. "Bow-tie" schematic of the modeling process.

have been answered can modeling begin. *Formulating* the model consists conceptually of extracting from complex, real-world situation those variables and relationships essential to the operation of the system of interest, and of narrowing down the multitude of variables and issues to those crucial to solving the problem at hand. This simplification process is key: not enough simplification and, at best, solution will be overly time-consuming and costly; at worst, solution techniques will prove to be unable to handle the model's complexity or size. Too much simplification and the ultimate users of the model results will be unwilling to use them due to the lack of realism and accuracy in the model specification.

Although, for ease of illustration, we often discuss formulation as separate from underlying relationships that drive the system, in practice the formulation exercise often greatly increases understanding. There are countless examples of model development in corporate settings where the true value of the model comes from the discipline imposed by formulating that model. Involved managers who must express their understanding in quantitative terms often find that their understanding of the system has increased. Questioning of underlying assumptions, quantifying perceptions, and specifying relationships increases insight, in and of itself.

Moving forward as shown in Figure 1-4, formulation is followed by solution. First, the selection of the appropriate solution methodology given the characteristics of the formulated system relationships; second, the collection of data—in practice, the most time-consuming effort; and, third, the actual running of the appropriate algorithm to produce a numerical "solution."

Altogether too many courses and textbooks stop at this point, having given short shrift to formulation and focussing on solution algorithms. In a managerial context this is disastrous. It is perhaps the source of the often heard complaint about the analyst who, after having disappeared for months to formulate a complex optimization model, drops a two-inch thick stack of computer output on the manager's desk, cheerfully exclaiming, "Here's the answer." The "bow-tie" in Figure 1-4 indicates, however, that at this point we are only halfway through the process—that *we have no answers, only solution results*. These solution results require interpretation. This interpretive task is, in some sense, a reversal of the formulation task. Whereas formulation requires that we narrow down the broad expanse of the real-world system, interpretation suggests broadening the narrow meaning of the solution results to apply them to the actual situation.[1] Interpretation involves recognizing the simplifications made and checking the solution results against them, verifying that nothing of import was lost in the translation. Interpretation is truly an art. Certainly the most difficult of the three modeling tasks, its goal is to take solution results—a set of numbers—and create insight into the management problem, insight that will be key in the ultimate management decision.

[1]This simplification or *narrowing* of the vast array of real world factors, followed by the interpretation or *expansion* of the numerical solution results, gives rise to the "bow-tie" shape of Figure 1-4.

MATHEMATICAL PROGRAMMING

The material in this book covers the class of optimization models that has had the widest application to management problems—mathematical programming models. These models, probably the best developed and most used of management science techniques—certainly the most successful—are concerned with allocating scarce resources to competitive activities in a way that is "best" for the enterprise as a whole. The resource allocation problem is one of the most important decision problems managers face—how to allocate productive capacity or scarce materials to a variety of possible products, how to allocate available finances among various capital investment opportunities, how to allocate manpower among possible R&D projects, and so forth. In view of the difficulties frequently encountered in making these decisions as well as their importance, managers are grateful for tools to bring to bear for aid in the decision process. This, perhaps, explains the success mathematical programming has had.

Note that the example applications in the paragraph above all have several characteristics in common:

1. All require a quantitative specification of *how much* of the available resources are to be allocated to each of the competing activities.
2. For each, there is implicitly some objective(s) to be optimized. One focus of this book is to help the manager make the objective or objectives *explicit*.
3. For each, there are constraints limiting the amount of resources which can be "consumed" by the competing activities as well as, possibly, additional constraints restricting the makeup of the portfolio of activities the manager can choose.

Formally, mathematical programming seeks to determine the values of certain *decision variables*[2] (alternatively, the levels of specified *activities*) subject to various restrictions *(constraints)* that are expressed as equalities or inequalities in terms of the decision variables. The constraints, in practice, reflect financial, availability, technological, organizational, or political considerations as well as a multitude of others. The values of the decision variables are determined so as to optimize (minimize or maximize) some chosen criterion that, when expressed in terms of the variables, is called the *objective function* (or simply the objective). The form of the objective, the constraints, and the values the decision variable may take determine the type of mathematical programming algorithm that need be used.

[2]The optimal values of the decision variables represent the optimal plan or *program*; thus the origin of the term "mathematical programming."

For example, if all mathematical relationships may be expressed by linear functions, the model is a linear programming model. Linear programming is the most strikingly successful of all management science applications. Its ease of use and the wealth of sensitivity information provided along with the optimal solution contribute to its success. Chapters 2 and 3 deal with these models by discussing, in some detail, the key features of linear programming. The format of these two chapters mirror the three tasks discussed in the preceding section—formulation (Chapter 2), solution, and interpretation (both covered in Chapter 3). Chapter 4 introduces multicriteria issues and ways of dealing with them.

In the remainder of the book we systematically relax the assumptions that underlie continuous linear programming. Chapters 5 and 6 focus on combinatorial optimization—situations in which we wish to optimize over discrete classes or sets of items, where the solution sought specifies the best combination or grouping of these discrete entities. In this case, "yes-no" variables are key. We make the transition between linear programming and combinatorial optimization by considering a class of optimization models that is a special case of both model categories—network models. These structures, in use since the late 1940s, are discussed in Chapter 5. Chapter 6 introduces integer programming, a class of optimization models of growing importance. Hitherto unsolvable except for relatively small models, the rapid increase in speed and decrease in cost which characterizes computer technology of the 1970s and 1980s has made the solution of larger and larger models possible.

Chapters 7, 8, and 9 deal, quite briefly, with a set of methodologies that, as yet, have had only a minor impact on business practice, but, as computer capabilities continue to increase and computer costs continue their descent, these now-advanced topics will find their way more and more into managerial application. All are tools of which the manager of tomorrow should be cognizant. Chapter 7 relaxes the assumption of deterministic coefficients and introduces uncertainty explicitly into the linear programming framework. Chapter 8 relaxes the assumption of linearity and briefly surveys the broad expanse of nonlinear programming. Chapter 9 considers sequential decision problems, both deterministic and stochastic, introducing dynamic programming.

As is said above, this book focuses on formulation and interpretation, not on algorithmic details. Nevertheless, the serious student of optimization models will want some background knowledge of the Simplex method, George Dantzig's 1947 discovery that forms the core of all mathematical programming; perhaps the most successful algorithm devised in the last half-century. Appendix A gives a brief treatment of this algorithm, at a level that requires little mathematical sophistication. Appendix B illustrates, for the interested reader, another important solution methodology—the transportation method. Like the Simplex method, this algorithm for solving the transportation problem—probably the most often used and successful special linear programming structure—has historical as well as practical significance. Appendix C discusses a