



Saurashtra University

Re – Accredited Grade 'B' by NAAC
(CGPA 2.93)

Rupani, Mehul P., 2010, “*Some contribution in operations research*”, thesis
PhD, Saurashtra University

<http://etheses.saurashtrauniversity.edu/id/eprint/753>

Copyright and moral rights for this thesis are retained by the author

A copy can be downloaded for personal non-commercial research or study,
without prior permission or charge.

This thesis cannot be reproduced or quoted extensively from without first
obtaining permission in writing from the Author.

The content must not be changed in any way or sold commercially in any
format or medium without the formal permission of the Author

When referring to this work, full bibliographic details including the author, title,
awarding institution and date of the thesis must be given.

Saurashtra University Theses Service
<http://etheses.saurashtrauniversity.edu>
repository@sauuni.ernet.in

THESIS FOR AWARDING DEGREE
OF
DOCTOR OF PHILOSOPHY
IN
MATHEMATICS

“ SOME CONTRIBUTION IN OPERATIONS RESEARCH ”

by
MEHUL P. RUPANI
(M.Sc.)

Submitted to
Saurashtra University
Rajkot.

Under the guidance of
Dr. G. C. BHIMANI
Professor
Department of Statistics
Saurashtra University
Rajkot-360005
Gujarat (India)

DECLARATION

I hereby declared that the research work on “ Some Contribution in Operations Research ” is carried out by me and results of this work have not submitted at any other university for the award of the Degree of Doctor of Philosophy in Mathematics

Mehul P. Rupani
Ph. D Student
Saurashtra University
Rajkot - 360005

CERTIFICATE

This is to certify that the Thesis entitled “ Some Contribution in Operations Research ” submitted by Mr. Mehul P. Rupani for the award of the Degree of Doctor of Philosophy in Mathematics is a bonafide research work done independently by him under my guidance.

Dr. G. C. Bhimani
Professor of Statistics
Department of Statistics
Saurashtra University
Rajkot - 360005

ACKNOWLEDGMENT

I would like to thank my guide Dr. G. C. Bhimani for his thoughtful guidance and support. Besides being an excellent guide he has been a great mentor who has helped me develop my self-confidence and always encouraged me to push my boundaries and take on new challenges in life.

This thesis is dedicated to my parents. My parent's integrity, humility, love, and compassion for all people has left an indelible impression in my life. I am eternally grateful for their constant encouragement and for setting the right example in my life.

I express thanks to all my colleagues for their co-operation.

Microsoft Office 2007 was used to prepare this and all calculations were done with the statistical softwares. This thesis contains much of their effort not in terms of paragraphs, table or figures, rather, their understanding and support all the way.

Mehul P. Rupani

CONTENTS

Title	Page no.
Declaration	I
Certificate	II
Acknowledgement	III
Contents	IV-V
1. Introduction	01-09
2. Operation Research the Art and Science of Executive Decision	10-56
2.1 Introduction	10
2.2 Boundaries of Quantative Analysis	19
2.3 Importance of Model-Building	24
2.4 Process of Quantative Analysis	28
2.5 Operations Research, Lilliputian Style	33
2.6 Ending Rainbow	48
3. Computer Simulation of Management Systems	57-114
3.1 Challenge Remaining	59
3.2 Simulation in Perspective	64
3.3 Stock Market Simulation Model	68
3.4 Building a Simulation Model	73
3.5 Generating Random Phenomena	76
3.6 How to Mark Time	90
3.7 Design of Simulation Experiments	98

3.8 Computer Languages	110
3.9 Deus Ex Machine	113
4. Advanced Topics in Network Algorithms	115-134
4.1 Maximum Flow Through a Capacitated Network	115
4.2 Solution of the Assignment Model	120
4.3 Algorithms for other Network Model	134
5. Implementation of Operations Research	135-160
5.1 Introduction	135
5.2 How to put Operations Research to managerial use	137
5.3 How to successfully conductan Operations Research project	146
5.4 How to manage an Operations Research staff	155
5.5 Technical & Technological advances	157
6. Textiles	161-174
6.1 A spin plan of maximum profit	161
6.2 Allotment of drums to winders in cone-winding	164
6.3 Optimum work assignment to siders in spinning	167
6.4 Determination of optimal number of tape stichers in a spinning shed	170
References	175-182

CHAPTER 1

INTRODUCTION

Operations research (also referred to as decision science, or management science) is an interdisciplinary mathematical science that focuses on the effective use of technology by organizations. In contrast, many other science & engineering disciplines focus on technology giving secondary considerations to its use.

Employing techniques from other mathematical sciences — such as mathematical modeling, statistical analysis, and mathematical optimization — operations research arrives at optimal or near-optimal solutions to complex decision-making problems. Because of its emphasis on human-technology interaction and because of its focus on practical applications, operations research has overlap with other disciplines, notably industrial engineering and management science, and draws on psychology and organization science. Operations Research is often concerned with determining the maximum (of profit, performance, or yield) or minimum (of loss, risk, or cost) of some real-world objective. Originating in military efforts before World War II, its techniques have grown to concern problems in a variety of industries.

Operational research encompasses a wide range of problem-solving techniques and methods applied in the pursuit of improved decision-making and efficiency. Some of the tools used by operational researchers are statistics, optimization, probability theory, queuing theory, game theory, graph theory,

decision analysis, mathematical modeling and simulation. Because of the computational nature of these fields, OR also has strong ties to computer science and analytics. Operational researchers faced with a new problem must determine which of these techniques are most appropriate given the nature of the system, the goals for improvement, and constraints on time and computing power.

Work in operational research and management science may be characterized as one of three categories:

- Fundamental or foundational work takes place in three mathematical disciplines: probability, optimization, and dynamical systems theory.
- Modeling work is concerned with the construction of models, analyzing them mathematically, implementing them on computers, solving them using software tools, and assessing their effectiveness with data. This level is mainly instrumental, and driven mainly by statistics and econometrics.
- Application work in operational research, like other engineering and economics' disciplines, attempts to use models to make a practical impact on real-world problems.

2. PRE - WORK

As a formal discipline, operational research originated in the efforts of military planners during world war II. In the decades after the war, the techniques began to be applied more widely to problems in business, industry

and society. Since that time, operational research has expanded into a field widely used in industries ranging from petrochemicals to airlines, finance, logistics, and government, moving to a focus on the development of mathematical models that can be used to analyze and optimize complex systems, and has become an area of active academic and industrial research.

In the world war II era, operational research was defined as “a scientific method of providing executive departments with a quantitative basis for decisions regarding the operations under their control.” Other names for it included operational analysis (UK Ministry of Defence from 1962) and quantitative management.

Prior to the formal start of the field, early work in operational research was carried out by individuals such as Charles Babbage. His research into the cost of transportation and sorting of mail led to England’s universal “Penny Post” in 1840, and studies into the dynamical behaviour of railway vehicles in defence of the GWR’s broad gauge. Percy Bridgman brought operational research to bear on problems in physics in the 1920s and would later attempt to extend these to the social sciences. The modern field of operational research arose during World War II.

Modern operational research originated at the Bawdsey Research Station in the UK in 1937 and was the result of an initiative of the station’s superintendent, A. P. Rowe. Rowe conceived the idea as a means to analyse and improve the working of the UK’s early warning radar system, Chain Home (CH). Initially,

he analyzed the operating of the radar equipment and its communication networks, expanding later to include the operating personnel's behaviour. This revealed unappreciated limitations of the CH network and allowed remedial action to be taken.

Scientists in the United Kingdom including Patrick Blackett later Lord Blackett OM PRS, Cecil Gordon, C. H. Waddington, Owen Wansbrough-Jones, Frank Yates, Jacob Bronowski and Freeman Dyson, and in the United States with George Dantzig looked for ways to make better decisions in such areas as logistics and training schedules. After the war it began to be applied to similar problems in industry.

During the Second World War close to 1,000 men and women in Britain were engaged in operational research. About 200 operational research scientists worked for the British Army.

Patrick Blackett worked for several different organizations during the war. Early in the war while working for the Royal Aircraft Establishment (RAE) he set up a team known as the "Circus" which helped to reduce the number of anti-aircraft artillery rounds needed to shoot down an enemy aircraft from an average of over 20,000 at the start of the Battle of Britain to 4,000 in 1941.

In 1941 Blackett moved from the RAE to the Navy, first to the Royal Navy's Coastal Command, in 1941 and then early in 1942 to the Admiralty. Blackett's team at Coastal Command's Operational Research Section (CC-ORS) included two future Nobel prize winners and many other people who went on

to be preeminent in their fields. They undertook a number of crucial analyses that aided the war effort. Britain introduced the convoy system to reduce shipping losses, but while the principle of using warships to accompany merchant ships was generally accepted, it was unclear whether it was better for convoys to be small or large. Convoys travel at the speed of the slowest member, so small convoys can travel faster. It was also argued that small convoys would be harder for German U-boats to detect. On the other hand, large convoys could deploy more warships against an attacker. Blackett's staff showed that the losses suffered by convoys depended largely on the number of escort vessels present, rather than on the overall size of the convoy. Their conclusion, therefore, was that a few large convoys are more defensible than many small ones.

While performing an analysis of the methods used by RAF Coastal Command to hunt and destroy submarines, one of the analysts asked what colour the aircraft were. As most of them were from Bomber Command they were painted black for nighttime operations. At the suggestion of CC-ORS a test was run to see if that was the best colour to camouflage the aircraft for daytime operations in the grey North Atlantic skies. Tests showed that aircraft painted white were on average not spotted until they were 20% closer than those painted black. This change indicated that 30% more submarines would be attacked and sunk for the same number of sightings.

Other work by the CC-ORS indicated that on average if the trigger depth of aerial delivered depth charges (DCs) was changed from 100 feet to 25 feet, the kill ratios would go up. The reason was that if a U-boat saw an aircraft only shortly before it arrived over the target then at 100 feet the charges would do no damage (because the U-boat wouldn't have time to descend as far as 100 feet), and if it saw the aircraft a long way from the target it had time to alter course under water so the chances of it being within the 20 foot kill zone of the charges was small. It was more efficient to attack those submarines close to the surface when these targets' locations were better known than to attempt their destruction at greater depths when their positions could only be guessed. Before the change of settings from 100 feet to 25 feet, 1% of submerged U-boats were sunk and 14% damaged. After the change, 7% were sunk and 11% damaged. (If submarines were caught on the surface, even if attacked shortly after submerging, the numbers rose to 11% sunk and 15% damaged). Blackett observed "there can be few cases where such a great operational gain had been obtained by such a small and simple change of tactics".

Bomber Command's Operational Research Section (BC-ORS), analysed a report of a survey carried out by RAF Bomber Command. For the survey, Bomber Command inspected all bombers returning from bombing raids over Germany over a particular period. All damage inflicted by German air defences was noted and the recommendation was given that armour be added in the

most heavily damaged areas. Their suggestion to remove some of the crew so that an aircraft loss would result in fewer personnel loss was rejected by RAF command. Blackett's team instead made the surprising and counter-intuitive recommendation that the armour be placed in the areas which were completely untouched by damage in the bombers which returned. They reasoned that the survey was biased, since it only included aircraft that returned to Britain. The untouched areas of returning aircraft were probably vital areas, which, if hit, would result in the loss of the aircraft.

When Germany organised its air defences into the Kammhuber Line, it was realised that if the RAF bombers were to fly in a bomber stream they could overwhelm the night fighters who flew in individual cells directed to their targets by ground controllers. It was then a matter of calculating the statistical loss from collisions against the statistical loss from night fighters to calculate how close the bombers should fly to minimise RAF losses.

The "exchange rate" ratio of output to input was a characteristic feature of operational research. By comparing the number of flying hours put in by Allied aircraft to the number of U-boat sightings in a given area, it was possible to redistribute aircraft to more productive patrol areas. Comparison of exchange rates established "effectiveness ratios" useful in planning. The ratio of 60 mines laid per ship sunk was common to several campaigns: German mines in British ports, British mines on German routes, and United States mines in Japanese routes.

Operational research doubled the on-target bomb rate of B-29s bombing Japan from the Marianas Islands by increasing the training ratio from 4 to 10 percent of flying hours; revealed that wolf-packs of three United States submarines were the most effective number to enable all members of the pack to engage targets discovered on their individual patrol stations; revealed that glossy enamel paint was more effective camouflage for night fighters than traditional dull camouflage paint finish, and the smooth paint finish increased airspeed by reducing skin friction.

On land, the operational research sections of the Army Operational Research Group (AORG) of the Ministry of Supply (MoS) were landed in Normandy in 1944, and they followed British forces in the advance across Europe. They analysed, among other topics, the effectiveness of artillery, aerial bombing, and anti-tank shooting.

With expanded techniques and growing awareness of the field at the close of the war, operational research was no longer limited to only operational, but was extended to encompass equipment procurement, training, logistics and infrastructure.

Academic Denis Bouyssou describes the historical development of operational research from the 1940s to the 1970s as follows. “The historical development of Operational Research (OR) is traditionally seen as the succession of several phases: the ‘heroic times’ of the Second World War, the ‘Golden Age’ between the fifties and the sixties during which major theoretical achievements were

accompanied by a widespread diffusion of OR techniques in private and public organisations, a 'crisis' followed by a 'decline' starting with the late sixties, a phase during which OR groups in firms progressively disappeared while academia became less and less

concerned with the applicability of the techniques developed”.

CHAPTER 2

Operation Research

The Art and Science of Executive Decision

2.1 Introduction :

You may well be curious to know how a subject with so abstruse a name as operations research could beget a several hundred page introductory text purportedly dealing only with principles. The ambiguous term operations research was coined during World War II. At that time, it was an apt description of the subject matter. Unfortunately, the name stuck, even though present-day applications of operations research are considerably more diverse than they used to be.

Now there is a worldwide confederation of professional societies named Operations Research. The staffs of many industrial organizations bear the title. So do departments in leading universities, which have gone on to sanctify the term by granting advanced degrees bearing its name. These vested interests are so well entrenched that the name operations research is unlikely to be supplanted in our lifetime.

Disgruntled though we may be, saddled as we are with a title that is undisruptive if not downright misleading, we nevertheless must show our respect. After all, the scientists who originated the term were on the winning side of the war. (Who knows what might have happened if the other side had invented the

approach first?) Numerous synonyms for operations research are in common use. The British like operational research. A frequent American substitute is management science. (The popularity of this name is fostered by yet another international professional society called the Institute of Management Science. The Operations Research Society and the Institute of Management Science regularly hold joint meetings, and their membership overlaps to a large extent.) As a beginning student, fortunately, you can afford to assume a lofty indifference to the whole matter, leaving this semantic bone of contention for your seniors to wrangle over.

For convenience, and with reasonable accuracy, you can simply define operations research as a scientific approach to problem-solving for executive management. An application of operations research involves:

- Constructing mathematical, economic, and statistical descriptions or models of decision and control problems to treat situations of complexity and uncertainty.
- Analyzing the relationships that determine the probable future consequences of decision choices, and devising appropriate measures of effectiveness in order to evaluate the relative merit of alternative actions.

It is sometimes believed that operations research refers to the constant monitoring of an organization's ongoing activities—and, in fact, decision and control problems often do concern certain daily "operations" of the organization. Examples of this sort include production scheduling and

inventory control, facility maintenance and repair, and staffing of service facilities, to name a few applications.

But many operations research studies treat other kinds of decisions that bear on daily operations only indirectly. These studies usually have a planning orientation. Illustrations include determining the breadth of a firm's product line, developing a long-term program for plant expansion, designing a network of warehouses for a wholesale distribution system, and entering a new business by merger or acquisition.

It is bad enough that the word "operations" inadequately describes the diversity of present-day applications. To make matters worse, the word "research" creates the false impression that the method is a "blue-sky approach." On the contrary, in the past decade operations research has proved time and again to be a powerful and effective approach for solving critically real management problems. You will learn most of the reasons in this chapter, and you will know the full story after reading the main chapters of this book.

Of course, fundamental research in the methods of operations research continues, mainly at universities and at governmental and industrial research laboratories. Unlike the situation with basic research in other sciences, however, relatively little time elapses between an important discovery in operations research and its implementation by experienced practitioners in industrial groups.

Better decisions in a complex and uncertain environment.

A preferable term to describe the subject of this book is decision

analysis. An emphasis on making decisions or taking actions is central to all operations research applications.

Decision analysis separates a large-scale problem into its subparts, each of which is simpler to manipulate and diagnose. After the separate elements are carefully examined, the results are synthesized to give insights into the original problem. You may wonder why such complex decision-making problems arise in the first place.

One reason is that in today's economy, technological, environmental, and competitive factors typically interact in a complicated fashion. For example, a factory production schedule has to take account of customer demand (tempered by the likelihood of a price-cut by competitors), requirements for raw materials and intermediate inventories, the capacities of equipment, the possibility of equipment failures, and manufacturing process restrictions. It is not easy to make up a schedule that is both realistic and economical.

Other reasons for complexity in real decision-making situations are that the organization (perhaps only half-knowingly) may be pursuing inconsistent goals, the responsibility and authority for making the required decisions may be greatly diffused within the organization, and the economic environment in which the company operates may be uncertain.

To be successful, an operations research approach must improve the managerial decision-making process—the improvement being measured by the net cost of obtaining it. You should keep in mind the distinction between

improved decision-making and improved performance, or more succinctly, between a good decision and a good outcome. For example, by all prior analysis, your betting on the Irish Sweepstakes may not appear to be a good decision (economically or morally); but after betting, the outcome will be good if you win, Improving decision analysis is important because the only thing you control is your decision prior to the uncertain outcome.

Distinguishing characteristics. There are many ways to approach management problems, and most of these ways are related. Certainly, there is no clear boundary line isolating the solutions derived by professional operations researchers from those derived by such people as industrial engineers, or economists specializing in economic planning, or accountants or financial analysts oriented toward management information systems. But most operations research applications possess certain distinguishing characteristics. Specifically, a suggested approach to a particular problem must contain all the following qualities before we would call it an operations research approach:

- i. A Primary Focus on Decision-Making,** The principal results of the analysis must have direct and unambiguous implications for executive action.
- ii. An Appraisal Resting on Economic Effectiveness Criteria.** A comparison of the various feasible actions must be based on measurable values that unequivocally reflect the future well-being of the organization.

In a commercial firm, these measured quantities typically include variable costs, revenues, cash flow, and rate of return on incremental investment. A recommended solution must have evaluated the tradeoffs and have struck an optimum balance among these sometimes conflicting factors.

iii. Reliance on a Formal Mathematical Model. The procedures for manipulating the data should be so explicit that they can be described to another analyst, who in turn would derive the same results from the same data.

iv. Dependence on an Electronic Computer. This characteristic is not really a desideratum but rather a requirement necessitated by either the complexity of the mathematical model, the volume of data to be manipulated, or the magnitude of computations needed to implement the associated management operating and control systems.

In science we trust. To embrace operations research, a company must believe that applying the scientific method is relevant to the analysis of managerial decisions. This statement is not the platitude it may seem to be at first reading. The adoption of operations research calls for an act of faith in the benefits of a systematic approach to decision-making, and not all corporation executives are ready to make that act as yet.

It may sound strange, at this late date, to hear a plea for faith in science—and operations research is a science. After all, the legitimacy of the scientific

method in the study of other subjects, such as physical phenomena, is hardly open to question. After hundreds of years of experience, chemists and physicists have developed efficacious laboratory techniques. But the virtue of applying scientific procedures to decision-making problems of significance is not so well-established; its recognition still calls for what the poet Coleridge described, in another context, as “the willing suspension of disbelief.” Here is why.

Rarely, if ever, can a company perform what most people would regard as a bona fide “scientific” experiment to test the merit of an operations research solution. Consider a company that is contemplating using a mathematical model to arrive at its annual operating plan. Since the company’s economic environment differs from year to year, it can never exactly repeat history, and therefore can never prove indisputably that the model solution will produce a realized improvement over the company’s current planning approach.

Consider a second illustration. Suppose that an operations research model has been suggested for controlling a company’s inventories. Again, testing whether the new system will definitely yield an improvement over the present approach is inherently limited. Although you could use historical data to compare how the suggested rule would have operated in the past, the comparison does not represent a truly scientific experiment with controlled variables. For one thing, you can only assume that historical data are indicative of what will happen in the future. For another, if the suggested rules improve service and customers recognize the improvement, then there may be an

increase in customer demand. In other words, the very operation of the suggested policy can alter the environment.

Thus, the historical data may not be typical of the future. And because the decision system itself influences the environment, it is not really possible to operate both the present and new systems “in parallel” (Occasionally, you can run part of the system under the new set of rules and the other part under the old set of rules. Explain why this test, also, is not a truly controlled experiment.)

Of course, before a manager accepts a specific operations research solution, he should perform various tests of reasonableness, including historical comparisons. But at some point after making such tests, even in an ideal situation, the manager will have to accept as axiomatic that a scientific approach has intrinsic merit. We make three amplifying observations before leaving this conclusion.

First, even though a company may be convinced about the worthiness of the scientific method to aid decision-making, it need not accept the results of a particular operations research study as being valid. After all, the specific project may have been ill conceived or poorly executed.

Second, a trust in science does not imply the abandonment of hunch and intuition. On the contrary, the history of science itself is studded with cases of important discoveries made through chance, hunch, serendipity—even dreams. Behavioral scientists have not yet developed ways to induce such flashes of brilliance consistently. But most executives who use their

hunches well also seem, to possess a high level of knowledge and understanding about their activities. So the question is not when to apply science and when to rely on intuition, but rather how to combine the two effectively.

Third, the inherent difficulty of demonstrating that a suggested solution is a sure-fire improvement is not unique to operations research. Because of the inability to duplicate history, an act of faith is also required to accept any other proposed solution—including maintaining the status quo.

Past, present, and future. Although the term “operations research” was coined during World War II, the scientific origins of the subject date much further back. Primitive mathematical programming models were advanced by economists Quesnay in 1759 and Walras in 1874; more sophisticated economic models of a similar genre were proposed by Von Neumann in 1937 and Kantorovich in 1939. The mathematical underpinnings of linear models were established near the turn of the 19th century by Jordan in 1873, Minkowski in 1896, and Farkas in 1903. Another example of early development is the seminal work on dynamic models accomplished by Markov, who lived from 1856 to 1922. Two further illustrations are the innovative suggestions for economical inventory control, published in business and industrial engineering journals during the 1920’s, and the pioneering studies of waiting line phenomena completed by Erlang, who lived from 1878 to 1929.

Even though these early starts received recognition and acclaim, only recently have mathematical models for decision analysis taken hold in business. Why? At least two factors are important. First, the competitive pressures of doing business have increased tremendously since World War II. Executives of large corporations now find it essential in maintaining profits to improve on the traditional ways of collecting and analyzing data. Second, the fantastic development and widespread adoption of high-speed electronic computers have fostered the growth of more sophisticated means for assessing decision alternatives.

There are many reasons to believe that the process of implementing operations-research-oriented systems will quicken. For example, new technological developments in what is called time-shared computing bring the power of an electronic computer literally into an executive's office. It is a pipedream to suppose that, in the next few years, most corporation presidents will have computer consoles on their desk tops for querying at a moment's notice. But already, financial vice presidents in several industrial companies do have such consoles to evaluate major investment alternatives. The future is getting closer all the time.

2.2 BOUNDARIES OF QUANTITATIVE ANALYSIS

As should be obvious, quantitative analysis can never provide the entire basis for all strategic decisions. It is inconceivable that the selection of a

corporation president by a company's board of directors, for example, could (and should) ever rest solely on the manipulation of quantitative data, although some numerical information may be relevant.

It is probably less apparent that even when quantitative analysis is of central importance for a managerial decision process, an operations-research-oriented system never supplies all the information required for action, no matter how sophisticated the system's design. Furthermore, a truly successful implementation of an operations research system must apply behavioral as well as mathematical science, because the resultant system must interact with human beings. And finally, the very process of constructing an operations research system involves the exercise of judgment in addition to the logical manipulation of symbols and data. We discuss below each of these boundaries on quantitative analysis.

Problems solved and unsolved. As you read the chapters of this book, you will learn the ways an executive can be aided by the different operations research models that are treated. We therefore limit our comments in this section to more generally applicable remarks.

We have already mentioned that at the very inception of implementing an operations research system it is necessary for experienced executives to discern the relevance of the model. This alone is not enough, of course. Since the corporate owners hold these men responsible for wisely managing the firm, executives must continue to exercise their judgmental duty well

beyond initial acceptance of the model. In one way or another, they must monitor the system to ensure that the underlying model remains valid, and in particular that it continues to be used properly to provide insights into the real decision-making problems of the company. (Managers must guard against thinking of the model as being reality, and hence of the accompanying answers as being sacrosanct.)

A newly implemented operations research system may well bring about a restructuring or an amplification of information. As a result, executives may act differently from how they might have acted without such information. There is no getting away from the fact, however, that an executive, not the model, takes the action.

In short, an operations research model is never sufficient unto itself; it cannot become entirely independent of judgment supplied by knowledgeable managers. This boundary on quantitative analysis is always manifest, because the number of questions that managers can pose is boundless, whereas the kinds of answers that a single model can provide are inherently limited.

Systems are for people. The above discussion also suggests that there is more to a successful implementation of an operations research system than the mere design of a mathematically correct model. Clearly, the system must operate in the larger context of managerial activity. The model must take account of the data sources, with respect both to quality of the data and to the goals and expertise of the people responsible for collecting the data.

The system must also reflect the information requirements of the managers who review the analytic results, especially the needs for descriptive and interpretive commentary.

Most experienced practitioners of operations research know how to solve these so-called problems of communication. But there is a more fundamental limitation on quantitative analysis: rarely, if ever, is a suggested operations research system in perfect harmony with previously existing managerial attitudes and predilections. To ignore this fact is to invite internal conflict, subterfuge, and sometimes downright sabotage of a new system.

For example, a corporate planning model may call for the development of realistic sales forecasts. You would expect that marketing executives should ordinarily be entrusted to provide these figures; but the traditional orientation of the marketing department may make it impossible for these personnel to articulate anything other than sales goals. If the motivational drive of the sales organization is to set up targets and then try to meet them, and if it is then called upon to enunciate both targets and realistic forecasts, severe organizational conflicts may break out.

What can be done about this kind of limitation on quantitative analysis? Presently, considerable research is under way by behavioral scientists to discover successful means of instituting effective organizational change. Such developments in administrative science will certainly have a fundamental effect on the actual degree of success and speed of implementation of operations research systems.

The art of management science. The problem-solving ingenuity of professional operations researchers is still a limiting factor in the spread of quantitative analysis. Despite the enormous growth in the acceptance of management science models, there are precious few “standard” applications. Even in areas of decision-making where the relevance of mathematical models has become well established, designing particular applications in specific companies requires significant skill on the part of the management scientist. Model formulations remain tailor-made to a large degree.

Conceivably, in the next decade some of the well-developed applications of operations research will have become so widely adopted that procedures for building these models can be codified, as many of the techniques in industrial engineering and managerial accounting have been. The unabated expansion of quantitative analysis into previously untouched areas of decision-making is so enormous, however, that the need for imaginative and talented problem-solving will remain undiminished for some time.

In other words, a considerable amount of “art” is still required for the successful practice of management science. This in turn means that whether you are a managerial user or a practitioner of operations research, you must have some facility with both the artistic and the scientific ingredients of the subject. A text-book, such as this one, can teach you many of the scientific aspects, and give you a modicum of practice in the art through the study of toy examples and the formulation and solution of small-scale problems.

Unfortunately, however, it can do no more than make you aware of the artistic elements.

To help you understand this interplay between the art and the science of apply-ing operations research, we offer an analogy with the fine arts. A knowledge of scientific principles, such as the chemistry of paint, the physiology of the eye, the physics of light, the psychology of color, and the laws of perspective, helps the artist master fully the craft of painting. Likewise, such knowledge also distin-guishes the true connoisseur from the casual, albeit appreciative, Sunday museum-goer. By the same token, an understanding of the fundamentals of operations research is essential not only for the practitioner, but for the manager who wants to make truly effective use of the approach. If today's business world continues to become more complex, an executive will not be able to compete successfully in the role of a casual onlooker, or he himself may end up as a museum exhibit.

2.3 IMPORTANCE OF MODEL-BUILDING

If you study this text diligently, you will learn a considerable amount of mathematical technique. But a benefit that far transcends the mastering of specific algorithms is the facility you will gain in formulating, manipulating, and analyz-ing mathematical models. Model-building is the essence of the operations research approach. It is the counterpart to laboratory experimentation in the physical sciences.

Constructing a model helps you put the complexities and possible

uncertainties attending a decision-making problem into a logical framework amenable to comprehensive analysis. Such a model clarifies the decision alternatives and their anticipated effects, indicates the data that are relevant for analyzing the alternatives, and leads to informative conclusions. In short, the model is a vehicle for arriving at a well-structured view of reality.

A mixed bag. The word “model” has several shades of meaning, all of which are relevant to operations research. First, a “model” may be a substitute representation of reality, such as a small-scale model airplane or locomotive. Second, “model” may imply some sort of idealization, often embodying a simplification of details, such as a model plan for urban redevelopment. Finally, “model” may be used as a verb, meaning to exhibit the consequential characteristics of the idealized representation. This notion conjures up in the mind those television commercials dramatizing how love and happiness will result after a single application of the sponsor’s product.

In operations research, a model is almost always a mathematical, and necessarily an approximate, representation of reality. It must be formulated to capture the crux of the decision-making problem. At the same time, it must be sufficiently free of burdensome minor detail to lend itself to finding an improved solution that is capable of implementation. Striking a proper balance between reality and manageability is no mean trick in most applications, and for this reason model-building can be arduous.

You will find three pervasive and interrelated themes in operations

research model-building. The first is an emphasis on optimization. Concentrating on decisions that are optimal according to one or more specified criteria has been the forcing wedge for attaining improved decision-making. Typically, the optimization is constrained, in that the values of the decision variables maximizing the stated objective function are restricted so as to satisfy certain technological restraints. Often, the model includes restrictions that mirror the impact of dynamic phenomena.

The second theme is derivation of the analytic properties of a mathematical model, including the sensitivity of an optimal solution to the model's parameters, the structural form of an optimal solution, and the operating characteristics of the solution. To illustrate, if you have a mathematical model leading to an inventory replenishment policy, you will want to know how the rule depends on forecasts of customer demand, the specification of the rule (such as, "when down to n , order again"), and the long-run frequency of stockouts and the average inventory level.

The third theme is explicit recognition of system interactions. One of the difficult tasks in writing an elementary text is to convey how, in real applications, the model-building effort is oriented toward management system considerations. The result of an operations research analysis must be integrated into the management information, decision-making, and control systems fabric of the organization. Operations research applications cannot be undertaken in isolation from the surrounding managerial environment. For these reasons, an operations research project should be regarded, at least in part, as a systems

effort.

In one easy lesson. Many of the scientists who pioneered present-day applications of operations research are still alive and carry forward their individual banners of progress. One cannot help but be struck by the way each of these men treats nearly all of the significant decision-making problems he encounters by using his own specialty, such as linear programming, dynamic programming, inventory theory, or simulation, etc. This ability to apply a single solution technique or mathematical construct to a diverse range of problems—and to do so effectively—attests not only to the sheer genius of these innovators but to the flexibility of their approaches.

The experience of these scientists notwithstanding, most operations research analysts, when faced with a difficult managerial decision problem, usually do not find it self-evident that a single solution technique or model is patently most appropriate. For example, an analysis of what markets a company should serve, what products it should manufacture, what investments it should undertake, or where it should locate its plants and warehouses rarely leads to an immediate selection of a linear programming, or a dynamic programming, or a simulation approach. This being the case, you may well wonder how you will go about building or selecting a model when faced with a particular decision problem.

We know that the notion of model-building, as described in a textbook, carries with it an aura of mystery. Regrettably, it is virtually impossible to

provide you with a checklist for infallibly selecting and developing a model. But rest assured, there is considerable evidence that most students who have been trained in either the sciences, engineering, mathematics, business administration, or economics have little trouble building models in practice, provided they are inclined to do so. And nowadays, rarely, if ever, will you be faced with applying operations research unaided by an experienced practitioner. Therefore, you can count on being tutored at least the first time you use operations research.

2.4 PROCESS OF QUANTITATIVE ANALYSIS

We outline below the stages that are standard in applying quantitative analysis. An experienced practitioner takes these steps almost instinctively, and frequently does not attach formal labels to them. Actually, the components are not entirely distinct, and at any point in time, several of the phases proceed in concert. As a beginner, however, you will find it helpful to look over the entire process seriatim, so that you can plan ahead accordingly.

A prelude to a quantitative analysis of a decision problem should be a thorough qualitative analysis. This initial diagnostic phase aims at identifying what seem to be the critical factors—of course, subsequent analysis may demonstrate that some of these factors are not actually so significant as they first appear. In particular, it is important to attain a preliminary notion of what the principal decisions are, what the measures of effectiveness are among these choices, and what sorts of tradeoffs among these measures are likely

to ensue in a comparison of the alternatives. There will be trouble ahead unless you get a good “feel” for the way the problem is viewed by the responsible decision-makers. Without this appreciation, you may encounter considerable difficulty in gaining acceptance and implementing your findings. What is worse, your results could very well be erroneous or beside the point.

Formulating the problem. The preceding diagnostic should yield a statement of the problem’s elements. These include the controllable or decision variables, the uncontrollable variables, the restrictions or constraints on the variables, and the objectives for defining a good or improved solution.

In the formulation process, you must establish the confines of the analysis. Managerial decision-making problems typically have multifold impacts, some of them immediate and others remote (although perhaps equally significant). Determining the limits of a particular analysis is mostly a matter of judgment.

Building the model. Here is where you get down to the fine detail. You must decide on the proper data inputs and design the appropriate information outputs. You have to identify both the static and dynamic structural elements, and devise mathematical formulas to represent the interrelationships among these elements. Some of these interdependencies may be posed in terms of constraints or restrictions on the variables. Some may take the form of a probabilistic evolutionary system.

You also must choose a time horizon (possibly the “never-ending

future”) to evaluate the selected measures of effectiveness for the various decisions. The choice of this horizon in turn influences the nature of the constraints imposed, since, with a long enough horizon, it is usually possible to remove any short-run restrictions by an expenditure of resources.

Performing the analyses. Given the initial model, along with its parameters as specified by historical, technological, and judgmental data, you next calculate a mathematical solution. Frequently, a solution means values for the decision variables that optimize one of the objectives and give permissible levels of performance on any other of the objectives. The various mathematical techniques for arriving at solutions comprise much of the contents of this text.

As pointed out previously, if the formulation of the model is too complex and too detailed, then the computational task may surpass the capabilities of presentday computers. If the formulation is too simple, the solution may be patently unrealistic. Therefore, you can expect to redo some of the steps in the formulation, model-building, and analysis phases, until you obtain satisfactory results.

A major part of the analysis consists of determining the sensitivity of the solution to the model specifications, and in particular to the accuracy of the input data and structural assumptions. Because sensitivity testing is so essential a part of the validation process, you must be careful to build your model in such a way as to make this process computationally tractable.

Implementing the findings and updating the model. Unfortunately, most tyro management scientists fail to realize that implementation begins on the very first day of an operations research project. There is no “moment of truth” when the analyst states, “Here are my results,” and the manager replies, “Aha! Now I fully understand. Thanks for giving me complete assurance about the correct decision.”

We consider the entire process of implementation in Chap. 22. But we mention here the importance of having those executives who must act on the findings participate on the team that analyzes the problem. Otherwise, the odds are heavy that the project will be judged only as a provocative, but inconclusive, exercise.

It is common for an operations research model to be used repeatedly in the analysis of decision problems. Each time, the model must be revised to take account of both the specifics of the problem and current data. A good practitioner of operations research realizes that his model may have a long life, and so documents its details as well as plans for its updating.

What’s it all about ? Having learned the basic components of the quantitative analysis process, you should step back to see what the entire approach accomplishes.

The major effort is constructing a mathematical representation of a complicated situation, along with gathering the required data. The model is

essentially approximate—elaborate enough to capture the essentials, yet gross enough to yield computable solutions. The balance between detail and tractability is found by a trial and error process, involving considerable examination of preliminary findings and extensive sensitivity analysis.

When operations research is applied in a planning context, the solution usually consists of a most favorable set of values for the decision variables, with some information as to the cost of deviating from these values. When management science is used for developing an operating system, such as a means for controlling inventories, then the solution consists of a set of decision rules. Often, these rules are embodied in a computer program. For an inventory system, the computer routines analyze historical demand data, permit judgmental adjustments if specified, signal when replenishment is to take place, and calculate the reorder amount.

Only rarely does an operations research solution represent a precise forecast of what will happen in the future. Such an accurate prediction would be of interest; but the crux of the decision problem is to select among alternatives, not to forecast. A well-built model makes a valid comparison among the alternatives. In case this distinction between accurately predicting an outcome and legitimately comparing alternatives is puzzling, consider the following illustration.

A company is about to decide whether or not to open a new plant in Europe. An operations research model is constructed that contains forecasts of sales, costs, and revenues; the resultant solution probably indicates

anticipated production levels. If the economic advantage of opening this new plant is relatively insensitive to a range of reasonable values for the forecasted figures, then the company can make the correct expansion decision. It does not need to commit itself, at the same time, to production levels; they would be determined subsequently when more accurate demand forecasts are available.

2.5 OPERATIONS RESEARCH, LILLIPUTIAN STYLE

Before explaining how and why quantitative analyses have been valuable in aiding executive decision-making, we will examine a few highly simplified illustrations of operations research models. Since our only purpose is to show what mathematical decision models look like, we make no pretense about the realism of these formulations; you will find more practical versions in the subsequent chapters.

One-Potato, Two-Potato Problem. A frozen-food company processes potatoes into packages of French fries, hash browns, and flakes (for mashed potatoes). At the beginning of the manufacturing process, the raw potatoes are sorted by length and quality, and then allocated to the separate product lines.

The company can purchase its potatoes from two sources, which differ in their yields of various sizes and quality. These yield characteristics are displayed in Fig. 1.1. Observe that from Source 1, there is a 20% yield of

French fries, a 20% yield of hash browns, and a 30% yield off flakes; the remaining 30% is unrecoverable waste. The figures for flakes and waste are also 30% for potatoes from Source 2. but the yield of French fries is relatively higher.

Potato Yields
for a Unit of Weight.

Product	Source 1	Source 2	Purchase Limitations
French fries	.2	.3	1.8
Hash browns	.2	.1	1.2
Flakes	.3	.3	2.4
Relative Profit	5	6	

How many pounds of potatoes should the company purchase from each source? The answer depends, in part, on the relative profit contributions of the source. These relative figures are calculated by adding the sales revenues associated with the yields for the separate products, and subtracting the costs of purchasing the potatoes, which may differ between the two sources. (We have used the term relative profit contribution because we are ignoring other variable expenses, such as sales and distribution costs. These depend only on the products and not the sources of the raw potatoes, and so do not affect the purchase allocation decision.) Suppose the relative profit contribution is 5 for Source 1 and 6 for Source 2. Even though Source 2 is more profitable, it does not follow that the company should purchase all of the its potatoes from Source 2.


At least two other factors are relevant to the purchase decision: the

maximum amount of each product that the company can sell, and the maximum amount that the company can manufacture—given its production facilities. To keep the Exposition simple, suppose that the two factors, in concert, imply that total production cannot exceed 1.8 for French fries, 1.2 for hash browns, and 2.4 for flakes, where these constants are measured in terms of an appropriate unit of weight (such as millions of pounds). These restriction can be expressed mathematically as follows.

Let p_1 denote the amount (in weight) of potatoes that will be purchased from Source 1, and P_2 the amount from Source 2. Then the values for P_1 and P_2 are con-strained by the linear inequalities

$$.2P_1 + .3P_2 \leq 1.8 \quad \text{for French fries}$$

$$.2P_1 + .1P_2 \leq 1.2 \quad \text{for hash browns}$$



Feasible Purchasing Policies

$$(1) \quad \begin{aligned} .3P_1 + .3P_2 &\leq 2.4 \quad \text{for flakes} \\ P_1 &\geq 0 \quad \text{and} \quad P_2 \geq 0. \end{aligned}$$


The nonnegativity restrictions $P_1 \geq 0$ and $P_2 \geq 0$ are imposed because a value such as $P_1 = -4$ would have no physical significance.

All the values for P_1 and P_2 satisfying (1) are shown in the shaded region. Notice that each line in the diagram is represented by a restriction in (1) expressed as an equality. The arrow associated with each line shows the direction indicated by the inequality signs in (1). Explain why a pair of values for P_1 and P_2 that satisfies both the French fries and hash brown constraints will also satisfy the flakes constraint.

Optimal values for P_2 and P_1 are found by making the relative profit contribution as large as possible, consistent with the constraints. Therefore, the optimization problem is to

$$(2) \quad \text{maximize } (5P_1 + 6P_2)$$

subject to (1). In this simple problem, the solution can be exhibited graphically, as



Each of the parallel straight-line segments represents different combinations of P_1 and P_2 that give the same value for the linear objective function $5P_1 + 6P_2$. The highest segment still having a point in the feasible

constraint region is the optimal value of the objective function, and such a point is an optimal solution. You can see that there is only one optimal solution in this example; it occurs at the intersection of the French fries and hash brown constraints. Consequently, you can calculate the optimal values by solving the associated simultaneous linear equations

$$\begin{aligned} &.2P_1 + .3P_2 = 1.8 \quad \text{for French fries} \\ (3) \quad &.2P_1 + .1P_2 = 1.2 \quad \text{for hash browns.} \end{aligned}$$

Verify that the optimal answers are $P_1 = 4.5$ and $P_2 = 3$, as shown in Fig. 1.3, giving an objective-function value of 40.5.

This problem illustrates what is termed a linear programming model. Real applications of linear programming usually involve hundreds of constraints and thousands of variables. You will learn how to formulate and solve such models in Chaps. 2 through 7.

Secretary Problem. An executive wishes to hire a new secretary, and is about to ask a placement service to send qualified girls for him to interview. He has found from past experience that he can determine from an interview whether a girl, if hired, will turn out to be terrific, good, or just fair. He assigns a relative value of 3 to a terrific secretary, 2 to a good one, and 1 to a fair one. His previous experience also leads him to believe that there is a .2 chance of interviewing a girl who will be a terrific secretary, a .5 chance that she will be a good one, and a .3 chance that she will be a fair one.

He wishes to see only three girls at most. Unfortunately, if he does not

hire a girl immediately after an interview, she will take another job: hence, he has to decide right away.

If the first girl he sees is terrific, he will hire her immediately, of course. And if she is fair, he has nothing to lose by interviewing a second girl. But if the first girl looks good, then he is not sure what to do. If he passes her by, he may end up with only a fair secretary. Yet if he hires her, he surrenders the chance of finding a terrific girl. Similarly, if he chooses to see a second girl, he will again face a difficult decision in the event that she turns out to be good.


The selection problem can be displayed conveniently by a so-called decision tree, shown The circled nodes represent the interviewed girls, and the branches from these nodes show the chance events and their probabilities. The boxes indicate where a decision must be made, and the number at the end of a branch gives the relative value of stopping the decision process at that point.

The problem of finding an optimal decision strategy can be solved by what is termed dynamic programming, and in particular, by a process known as backward induction. You will study dynamic programming models and solution techniques in Chaps. 8 through 12, and Chaps. 17 and 18. The solution process is so simple in this example that you can compute the optimal hiring strategy very easily, as we have done below.

Suppose that the executive does end up interviewing, and hence hiring, a third girl. Then the expected value associated with the uncertain event is

$$(4) \quad 3(.2) + 2(.5) + 1(.3) = 1.9.$$

In other words, the average value of a girl selected at random for an interview is



Legend:
 T Terrific
 F Fair
 G Good

Decision Tree for Secretary Problem.

1.9. Assume that this expectation legitimately represents the executive's evaluation of the chancey event. Mark the number 1.9 above the circled Node 3

Next consider what happens if the executive does interview a second girl, and she turns out to be good. If he decides to stop, then he obtains a value 2. But if he continues,, then he can expect to receive only the value 1.9. So he should stop when the second girl looks good. Put an X on the branch indicating "Continue" when the second secretary is good; this signifies not to take that action.

Now you are ready to determine the correct decision if the first girl looks good. By stopping, the executive would obtain the value 2. But if he

continues, then the expected value associated with the chancey outcome of the second interview, and possibly the third, is

$$(5) \quad 3(.2) + 2(.5) + 1.9(.3) = 2.17.$$

The first term in (5) is for the event, of seeing a terrific secretary, whom he hires; the second term is for the event of seeing a good secretary, whom, he hires, as you already determined in the preceding paragraph; and the third term is for the event of seeing a fair secretary, and consequently continuing to the third chancey event that has a value of 1.9, given in (4). Since 2.17 is larger than 2, the executive should pass up the first girl if she turns out to be good. Mark 2.17 above the circled Node 2 and put an X on the branch indicating “Stop” when the first secretary is good.

To summarize, the optimal policy is to stop after the first interview only if the girl is terrific, and to continue after the second interview only if the girl is fair.

The overall expected value of the interviewing process, given that the executive acts optimally, is

$$(6) \quad 3(.2) + 2.17(.5) + 2.17(.3) = 2.336.$$

Explain why. Mark this number above the circled Node 1 Since the quantity 1.9, calculated in (4), also represents the expected value if the executive interviews only a single secretary and hires her, the difference $(2.336 - 1.9 = .436)$ is the incremental value from interviewing as many as two more girls.

Where-or-When Production Problem. The name of this problem arises from the observation that the associated mathematical model has several interpretations. One is in terms of deciding optimal production levels at each of several plants in a single time period; another is in terms of choosing optimal production levels at a single plant in each of several time periods. (The model also can be interpreted as a combination of the two problems, that is, as a where-and-when problem.)

Starting with the multiplant version, suppose a company has N plants, and must manufacture a total of D units of a single item during a stated time period. Hence, letting x_t denote the amount of production at Plant t , the levels $x_1, x_2,$

\dots, x_N must satisfy the constraints

$$(7) \quad x_1 + x_2 + \dots + x_N = D \text{ and all } x_t \geq 0.$$

Assume that the cost of producing x_t at Plant t is given by $(1/c_t) x_t$, where $c_t > 0$ is known from historical accounting information. Consequently, optimal values for the x_t are those that

$$(8) \quad \text{minimize} \left(\begin{array}{ccc} x_1^2 & + & x_2^2 & + & \dots & + & x_N^2 \\ c_1 & & c_2 & & & & c_N \end{array} \right)$$

subject to (7).

This optimization problem can be solved by dynamic programming methods as well as by some simple nonlinear programming techniques, which are discussed in Chaps. 14 and 15. The numerical answers can be easily

computed from the insightful formula

$$(9) \quad \text{optimal } x_t = \frac{c_t \cdot D}{c_1 + c_2 + \dots + c_N} \quad \text{for } t = 1, 2, \dots, N,$$

which yields the associated minimum cost

$$(10) \quad \frac{D^2}{c_1 + c_2 + \dots + c_N} \quad (\text{optimal policy}).$$

Turning to the multiperiod version, suppose you interpret x_t as being the level of production in a single plant during Period t . Notice that in this version, all the costs are due to production, and no storage costs are incurred while the units are inventoried from Period 1 to the end of Period N , when the demand requirement D must be met. Given this view, you can state what would be the optimal value for x_t if the preceding levels x_1, x_2, \dots, x_{t-1} were already determined [not necessarily by (9)], namely,

$$(11) \quad \text{optimal conditional } x_t = \frac{c_t \cdot (D - x_1 - x_2 - \dots - x_{t-1})}{c_1 + c_2 + \dots + c_N}$$

(x_1, x_2, \dots, x_{t-1} are specified).


As you can verify, calculating x_1, x_2, \dots, x_N recursively (that is, successively, one by one, starting with x^* from (11) yields the same values as computing each of them from (9).

Economic Order Quantity Problem. The formulation below is perhaps the simplest such model. Its precise assumptions are only rarely satisfied in real life. Nevertheless, the resultant solution turns out to be

sufficiently close to optimal for many practical situations as to make it a very useful approximation.

Consider a company that consumes (or sells) an item at the rate of, say, M units per week. For simplicity, suppose there is no uncertainty about this consumption. Hence, if the inventory level is kM units, then this stock is depleted in exactly k weeks. Further, suppose the rate M is unchanging over time, so that the company must regularly place a replenishment order. The decision problem is to determine the most economical order quantity. (Assuming that the delivery time for an order is also known exactly, each replenishment action is initiated early enough so that the order arrives just when the inventory level falls to zero.)

Let the order quantity be denoted by Q . Then the level of inventory can be pictured by the sawtooth pattern shown. Observe that each time a replenishment arrives, the inventory level shoots up by the order quantity Q . Then the level diminishes, as shown by the downward slope of the sawtooth, which equals $-M$.



Legend:
 M demand per week
 Q order quantity

Pattern of Inventory Levels.

An optimal order quantity strikes a balance between the costs associated with replenishing and with holding inventory. Specifically, assume that a fixed setup cost K is incurred each time an order is placed, that a purchase cost c is paid for each item ordered, and that a holding cost h is assessed for each unit of inventory held per week. The setup cost is related to the effort expended in placing and receiving the order. The holding cost is associated with storage, insurance, and the capital tied up in inventory.

Let the economic criterion of effectiveness be measured as average cost per week. Then the contribution due to setup costs is $K (M/Q)$, since there are M/Q , setups per week. The contribution due to purchase costs is cM , since M items are consumed per week. And the contribution due to holding costs is $h(Q/2)$, since $Q/2$ is the average level of inventory, as you can see Adding the components, you have

$$(12) \quad \text{average cost per week} == AC = \frac{KM}{Q} + cM + \frac{hQ}{2}$$

The economic order quantity that minimizes AC is

$$(13) \quad \text{optimal } Q = \frac{\sqrt{2KM}}{h}$$

which can be found by setting the derivative of AC with respect to Q equal to 0, and solving for Q . It follows from (13) that the optimal order quantity only doubles when the demand rate quadruples. Also note that the optimal quantity is determined by the ratio of setup to holding costs.

OR Airline Problem. The One-Ride Airline Company is opening a reservation service to be located in a suburban shopping center. A passenger making reservations will be able to telephone the office and state his request. The OR Airline Company wants to decide how many telephone lines to install for answering reservation calls. It can easily compute the telephone and personnel expenses that vary with the number of lines. But it also wishes to compare the level of service for several different numbers of lines. In particular, suppose the company seeks to determine the percentage of time all the lines will be busy and the average length of such busy periods.



The Distribution of Telephone Intearrival Times.


This sort of analysis is classified as queuing or waiting line theory, We could construct an explicit model and subject it to rigorous mathematical analysis, but instead, we explain here how the method of simulation can be used to deter-mine the service figures. To keep the expla-nation easy,

Suppose the company obtains data show-ing the statistical frequencies

of minutes between successive incoming telephone calls. As a first approximation, assume that these successive interarrival times are completely independent (such independence does not hold precisely if, for example, a passenger calls up, finds the lines busy, and immediately redials the number). A convenient way to summarize this distribution is to use a pie diagram, such as the one where we assume for simplicity that the time between incoming calls never exceeds 5 minutes.

Imagine a pointer, or spinner, affixed to the center of the pie diagram—the mechanism would look something like a wheel-of chance at carnivals, or a device that is often included in a child's game to determine how many advances a player's piece may take at each turn. You can simulate the traffic of incoming calls by giving the pointer a succession of sharp spins, and jotting down the resultant sequence of interarrival times.


In addition, suppose the company has a frequency distribution of the number of minutes that incoming calls require. Assume that these service times are independent of each other and of the interarrival times. Then another pie diagram and spinner mechanism can be constructed for generating the service times.



Simulated History for One Telephone Line.

You are now ready to simulate the system. To begin, suppose there is only a single telephone line. Then a simulated history may look like that. The instants of incoming calls are recorded with X's on the time axis, and are determined by successive spins of the pointer mechanism for interarrival times. The telephone line becomes busy as soon as the first call arrives. The length of the busy period is determined by a spin of the pointer mechanism for service times. Notice that the calls arriving at the instants circled are not answered because the single telephone line is busy. You can obtain a good estimate of the percentage of time the line will be busy and the average length of a busy period by calculating the corresponding statistics for a fairly long simulated run.

Suppose next that there are two telephone lines. Then the same sequence of incoming calls can lead to a service history like that. Observe that more incoming calls are answered. Here a busy period is defined to be a length of time during which both telephone lines are tied up. Note how these periods are shorter than for the particular history. As you can see, the same approach can be used to estimate service for any number of lines.



Simulated History for Two Telephone Lines.

Commentary. As we stated at the beginning of this section, the above examples illustrate how mathematical models are constructed to analyze decision problems. We have made no attempt to be “realistic” in these examples. But as you continue reading this book, you will discover how to build practical models, and learn ways to solve problems with them.

2.6 ENDING RAINBOW

Today, virtually every major corporation employs personnel who are responsible for applying operations research. Usually, these people constitute a staff group at headquarters level. The group often reports to the controller, the chief financial officer, or the head of corporate planning; but with growing frequency, companies are also assigning operations researchers to report directly to line managers. In a parallel fashion, operations research activity has enjoyed a widespread growth within federal and local governments, as well as in other nonprofit organizations. This section tells why management science has succeeded so well.

Merits of a rational process. Obviously, executives must and do make decisions all the time. For a particular situation, an operations research model may yield the very same conclusion that an experienced manager would arrive at solely on intuitive grounds. Therefore, the benefits of using operations research have to be evaluated in terms of its long-run impact on the entire managerial process.

The proper comparison is well represented by the question, “If a company does not use operations research to guide the decision process, then what will it use, and will the answers be consistently as good?” Corporations that apply operations research—even when the approach does not meet all of the company’s initial expectations—find that analyzing complicated managerial problems this way is a sounder method than traditional means. This assertion is amply borne out by the steadily increasing support given to operations research in both the private and public sectors.

Qualities inherent in this particular rational approach make it a valuable method, regardless of whether a company’s operations research personnel attain the highest level of accomplishment for a specific project. These benefits include:

- Emphasis on assessing the system-wide interactions and ramifications of decision alternatives. Intrinsic in an operations research approach is the construction of a model that synthesizes the segments of an enterprise that are affected by a decision choice. Each individual part is constructed by personnel who are the most knowledgeable about the relevant data.

- Impetus to developing a full range of decision alternatives. The number of action possibilities that can be analysed increases tremendously by the application of mathematics and computers.
- Focus on resolving the critical issues. The approach proceeds in the fashion of establishing implications of the form, “if Hypothesis H is true and Action A is taken, then Result R will occur.” The method fosters interdepartmental communication. As a consequence, clashes of opinion within an organization can be sorted into disagreements over the probable truth of different hypotheses, and over the assumptions used in deriving the implications of different actions.

Important ancillary benefits emerge from the direction provided by the model for gathering data, quantifying the value of additional information, and documenting factual knowledge that may be required in subsequent decision analyses.

Having listed several of the important merits stemming from the rational analytic process of operations research, we hasten to point out that the advantages will occur to a greater or lesser degree, depending on the skill used in carrying out the study.

Managerial cutting edge. The preceding paragraphs dealt with why operations research is helpful in analyzing decision problems. Now we discuss how the approach is beneficial. We have classified the ways into four somewhat arbitrary and partially overlapping categories:

i. Better Decisions. Frequently, operations research models yield actions that do Improve on intuitive decision-making. A situation may be so complex (because of Intricate interrelationships among decisions, voluminous data pertinent to operations, and uncertainties of market activity) that the human mind can never hope to assimilate all the significant factors without the aid of operations-research-guided computer analysis.

Of course, in the past managers have made decisions in these situations without the aid of operations research. They had to. But the depth of their understanding and the quality of their decisions improve with the application of such models, as considerable experience has shown. Particular decisions may ultimately turn out to be wrong, but the improved decision-making process reduces the risk of making such erroneous decisions.

ii. Better Coordination. Sometimes operations research has been instrumental in bringing order out of chaos. The following example, drawn from an actual application, illustrates what can happen.

During special campaigns, a food manufacturer runs advertising that significantly increases sales volume. But manufacturing production facilities are limited, the supply of the foodstuff is limited, and the sales response is often erratic. In the past, consequently, the marketing and manufacturing divisions have been at opposite poles in terms of cooperative actions. An operations-research-oriented planning model becomes a vehicle for coordinating marketing decisions within the limitations imposed on

manufacturing capabilities.

iii. Better Control. The managements of large organizations recognize that it is extremely costly to require continuous executive supervision over routine decisions. Operations research approaches combining historical experience with the scientific method have resulted in standardized and reliable procedures for handling everyday activities and for signaling dangerous trends. Executives have thereby gained new freedom to devote their attention to more pressing matters, except for those unusual circumstances which, when they arise, necessitate reviewing the the course of everyday action. The most frequently adopted applications in this category deal with production scheduling and inventory replenishment.

iv. Better Systems. Often, an operations research study is initiated to analyze a particular decision problem, such as whether to open a new warehouse. After-wards, the approach is further developed into a system to be employed repeatedly. Thus the cost of undertaking the first application may produce benefits that are longer lasting than originally envisioned.

Where the action is. By this time, applications of operations research are so common in industry and government, and so diverse in the functional areas of decison-making, that we cannot hope to provide a complete survey. To give you some idea, however, we mention that there are numerous applications in industries such as aircraft, apparel, chemicals, cement, glass, computers, electronics, farm and industrial machinery, food, metal

manufacturing and products, mining, motor vehicles, paper and wood products, petroleum refining, and pharmaceuticals, as well as in commercial banks, insurance companies, merchandising firms, public utilities, and transportation companies.

Depending on the industry, the applications pertain to extraction of natural resources, manufacturing, transportation and warehousing, plant size and location, inventory management, scheduling of men and machines, forecasting, new product development, marketing, advertising, cash management and finance, portfolio management, mergers, and both short- and long-range corporate planning.

Most companies' early operations research projects deal with monthly or quarterly scheduling, annual planning, inventory control, and other fairly well-defined areas of decision-making. After the operations research group demonstrates its capability in these areas, a company then applies its operations research talents to the study of high-level strategic problems, such as selecting new plant sites, entering new markets, acquiring overseas affiliates, and so forth.

Operations research has provided a significant advance in the techniques of long-range strategic planning. Even senior executives have difficulty piecing together all the important considerations involved in a well-designed long-range plan. What is more, the operations research approach lends itself to the formulation of contingency plans, that is, a complete strategy indicating which courses of action are appropriate for various future events. In addition,

the findings may include directions for obtaining and then utilizing critical information about such future events. In this way, the operations research model suggests the actions to be taken immediately and the ones to be postponed, and when to undertake a reassessment. For these reasons, more and more frequently, boards of directors of large corporations find strategic proposals being justified on the basis of extensive operations research studies.

Profitable applications in nonprofit organizations. The growth of operations research in government and nonprofit corporations has been phenomenal. A long succession of military applications began during World War II. Now governmental applications involve health, education, and welfare; air and highway traffic control; air and water pollution; police and fire protection; voter and school redistricting; and annual planning and budgeting by program, to name only a few.

Certainly, much of the credit for the great impetus in the adoption of operations research in the public sector goes to the RAND Corporation (located in Santa Monica). Many fundamental concepts and techniques in operations research can be traced to the innovative ideas of RAND scientists. At present, there are several other research organizations similar to RAND serving the Federal Government ; a number of these (RAND among them) work on state and city management problems as well as on military projects.

How different is operations research applied in the public sector from that used in the private sector? The answer depends on the aspects of

operations research you are considering. Specifically, the problem-solving characteristics, the ensuing benefits, the emphasis on model-building, the limitations of the scientific method, and the analytic process discussed above, all hold with equal validity in applications to governmental and nonprofit organizations. There does seem to be a noticeable difference, however, in the sorts of solution techniques used. For example, linear programming has gained widespread acceptance in industry, but is employed only occasionally in governmental planning. The opposite has been true of simulation models. Sometimes it is claimed that the lack of clear-cut objective functions to be optimized in nonprofit organizations raises a significant difference between applications in the private and public sectors. Industry most often measures improvement in terms of contribution to profit, but this criterion is by no means the only relevant one for decision-making. Businesses are always compromising among different objectives. Therefore, we feel that the absence of a profit measure is less important than it might seem at first glance.

Probably the most important difference between the public and private sectors concerns the exercise of decision-making responsibilities. The organizational structures of big corporations are complex, and the authority for taking actions may not always be precisely defined. But these structures are simple indeed in comparison with most governmental structures. The difference can be stated this way. In a commercial company, there is no one left to pass the buck to, once the necessity for making a decision reaches top management (or the board of directors). In an organization such as the Federal

Government, even the President's decisions are subject to the review of— and thus become partly the responsibility of— Congressmen, who, along with the President, are publicly elected. Understand that diffuse responsibilities and authority only make it difficult to apply operations research outside of industry. As the record shows, plenty of applications are being made in governmental and nonprofit organizations.

CHAPTER 3

Computer Simulation of Management Systems

If you have patiently proceeded from, one chapter to the next, you have studied a perhaps bewildering variety of operations research models and techniques. Students often ask, in effect, “Is this arsenal of tools powerful enough to encompass all the important managerial decision problems requiring data analysis?” The answer is no, not by a long shot. To see why, reflect on the kinds of problems that you know can be effectively analyzed by the operations research tools presented thus far. As you become aware of gaps, you will see more clearly why so many significant types of decision-analysis problems are still not solvable by these approaches, and therefore must be attacked in other ways. In the next few paragraphs we summarize the limitations as well as the strengths of operations research tools including linear and dynamic programming, inventory and queuing theory.

You have already learned that linear programming models are most successful in aiding the planning efforts of corporate enterprises. If the planning horizon is 10 years or longer, a corresponding multiperiod linear programming model typically deals only with annualized data. The effects of the resultant plan on week-to-week and month-to-month operations are left implicit. Analogously, if the planning horizon is much shorter, say three months to a year, the corresponding model usually ignores the day-to-day and week-to-week variations. Thus, for the most part, linear programming analysis falls short of prescribing rules that translate a recommended plan into operating

procedures for time spans shorter than the periods in the model.

A second limitation of linear programming analysis relates to uncertainty about the future. Imprecise forecasts to some degree exist in all planning studies.

Frequently, this uncertainty is not really the essence of the planning problem, or it reflects a lack of knowledge about only a few parameters in the model. In such cases, sensitivity analysis, as discussed in Chap. 5, suffices to determine the impact of uncertainty. But on other occasions uncertainty pervades the entire model, and standard sensitivity analysis is too clumsy and computationally burdensome for analyzing the impact of uncertainty.

To illustrate, consider a chemical manufacturing company that seeks a long-range strategy for the development and marketing of new products. Substantial research and investment costs are associated with each product, and the actual size of the product's market is uncertain. Furthermore, most of the profits that are generated from a successful product will be used to finance the research and development of new products. A linear programming model that manages to capture the dynamic elements of this situation, but treats the uncertainty aspects by simply using average values, is not likely to yield a good strategy.

In contrast, dynamic programming models can analyze multiperiod planning problems containing uncertainty, and so can be used to determine optimal strategies. But, as compared with linear programming applications, these dynamic programming models in practice can treat only drastically

simplified systems. As you learned in Chaps. 10 and 17, unless the underlying system is characterized by only a few state variables, the computational task of solving a dynamic programming model is horrendous.

A similar limitation holds for those dynamic probabilistic models that are amenable to mathematical analysis, such as the inventory and queuing phenomena you studied in Chaps. 19 and 20. To solve these models, you not only must restrict yourself to a small-scale system, but you also must simplify the way the system can operate. To illustrate, a realistic analysis of waiting lines in a job-shop is intractable using mathematical queuing theory like that presented in Chap. 20 and Appendix III. Those models serve only as rough approximations to realistic queuing phenomena.

Thus, despite the great diversity of applications of mathematical programming and probabilistic models, many important managerial decision-making problems must be analyzed by other kinds of techniques.

3.1 CHALLENGE REMAINING. The expanding scientific literature on operations research bears witness that there is steady progress in finding techniques to overcome the above-mentioned limitations. But for now and the foreseeable future, the approaches given in the preceding chapters cannot be relied on to provide a complete analysis of managerial decision-making problems pertaining to:

(i) Choice of Investment Policies for Strategic Planning. A major corporation's investment policy, to be comprehensive, should include

provisions relating to research and development of new products, expansion into new markets, choice of selection criteria for major projects, measurement and evaluation of risk, means of financing by debt and equity, reinvestment of earnings, disposition of liquid assets, evaluation of mergers and acquisitions, and divestment of assets. A full-fledged operations research model for the analysis of alternative policies must recognize the impact of the uncertain and dynamic nature of investments, as well as provide a means for screening the enormous variety of investment decisions that face an organization.

(ii) Selection of Facilities in Operations Planning. Several examples in this category were already discussed in Sec. 20.1. They included the determination of the number of checkout stands in a supermarket, the number of gasoline pumps at a service station, and the number of elevators in a new office building. There are numerous other examples dealing with personnel staffing, plant layout, and machine capacity decisions. Typical facilities selection questions are of the form: “How many?” “How large?” “Where located?”

(iii) Design of Information-Feedback Scheduling and Operations Rules. Illustrations of decision problems in this category are equally numerous, although you may not think of them right away, unless you have had some previous work experience. An important example is the design of scheduling rules for a job-shop manufacturing plant, or an equipment repair facility, or a computer center. Such rules for a manufacturing plant take account of promised due dates to customers, the requirements for, and the availabilities

of, machine capacities, the deployment of skilled labor, and the provisioning for raw materials. As information on new orders arrives, and as completed orders leave the system, the shop schedule has to be updated and revised.

Another example of an information-feed back system is a scheduling procedure for routing transport facilities. To illustrate a freighter shipping company in making a schedule of its ocean going equipment for several months ahead, must take into account cargo demands at various ports, ship capacities and speeds, uncertainties in sailing times due to vagaries in the weather, and delays due to port congestion. Many shipping lines that own a large fleet of vessels must reschedule daily as they receive more accurate information about uncertain events. Similar problems arise in the scheduling of patients into a hospital, and the timing of traffic lights on a major thoroughfare.

What makes the three types of problems described above so difficult to analyze? It is the combined effect of uncertainty, the dynamic interactions between decisions and subsequent events, the complex interdependencies among the variables in the system, and, in some instances, the need to use finely divided time intervals. Such total systems problems are too big and too intricate to handle with linear and dynamic programming models, or standard probabilistic models.

Frequently, actual decisions arising from these three types of problems involve spending at least several hundred thousand dollars, and vitally affect the future operating costs and efficiencies of a company. Thus, management

is highly motivated to employ a systematic approach to improve on intuitive, or “seat-of-the-pants,” analysis. So far, the best operations research approach available is digital computer simulation, Simulation approach. Our main concern in this chapter will be to describe simulation and the kinds of problems you encounter in employing this technique. We do not show you in detail how to design and run simulations. Such instructions are in texts devoted to simulation and in manuals distributed by computer manufacturers to explain special simulation programming languages.

In brief, the simulation approach starts by building an experimental model of a system. Then various specific alternatives are evaluated with reference to how well they fare in test runs of the model.

If you think about it, you will recall occasions when you have been involved in a simulated environment. For example, an amusement park, like Disneyland, offers you many attractions, such as the jungle boat-ride and the Matterhorn bobsled, that try to simulate actual experience. Less frivolous examples are planetarium shows and the environments in a museum of natural history. You may have learned how to drive an automobile in a mock-up mechanism with a steering wheel and gas and brake pedals. And if you have been in the armed services, you will remember that boot camp or basic training consists mainly of simulated exercises.

It is usually too inconvenient and expensive to solve managerial decision problems by environmental analogue simulations, such as the field combat war games that are used in boot camp and basic training. Rather, it is preferable

to represent a complex system by a computerized mathematical model. In a computer, the only thing that can be shot is an electronic circuit.

The uncertainties, dynamic interactions, and complex interdependences are all characterized by formulas stored in the memory of the high-speed digital electronic computer. The system simulation begins at a specified starting state. The combined effects of decisions, and of controllable and uncontrollable events, some of which may be random, cause the system to move to another state at a future instant in time. The evolutionary process continues in this fashion until the end of the horizon. Frequently, the time intervals are finely divided and extend over a fairly long horizon. As a consequence, the simulation experiments involve a vast number of calculations, rapidly performed by the computer. This feature of years of history evolving in a few minutes on a computer is termed time compression.

The only game in town. Most operations research analysts look upon digital computer simulation as a “method of last resort”—hence the title of this section, “When All Else Fails. . . .” There are two reasons for this gloomy attitude.

The first reason is the nature of most simulation results. When the model includes uncertain events, the answers stemming from a particular simulation must be viewed only as estimates subject to statistical error. For example, a simulated queuing model yields only an estimate of a waiting line’s average length or the associated probability of a delay. Therefore, when

you draw conclusions about the relative merit of different specific trial policies as tested by a simulation model, you must be careful to assess the accompanying random variations.

The second reason for diffidence about simulation involves the nature of the applications themselves. If a system is so complicated that it is beyond the reach, of such operations research tools as linear and dynamic programming or standard probability analysis, then the required model-building effort and the subsequent analysis of the simulated results are likely to be difficult. Many an unwary analyst has found, to his chagrin, that the simulated world is as unfathomable as the real world he hoped to approximate—he allowed so much to go on in the model that it hampered his finding any insightful information.

The above two reasons also suggest why electronic computers are indispensable in performing simulations. To obtain sufficient statistical accuracy for reliable decisions, a considerable number of simulation runs are usually necessary. Each experiment is so complicated that it would be virtually impossible to perform the simulation manually in a reasonable period of time. It is not surprising, then, that computer simulation is often an expensive way to study a complex system.

3.2 SIMULATION IN PERSPECTIVE

As you read in the preceding section, many important managerial decision problems are too complex and too large to be solved by mathematical programming and standard probability analysis. In such cases,

real-life experimentation, even if feasible, is usually too costly a way to analyze the alternatives. These observations establish the need for other problem-solving approaches, but do not by themselves justify computer simulation. Here we discuss why computer simulation is a useful technique, as well as what its limitations are.

Unlike the situation with mathematical programming, there are as yet no underlying principles guiding the formulation of simulation models. Each application is ad hoc to a large extent. Computer simulation languages come the closest to providing any general guidelines. [SIMSCRIPT and the General Purpose Systems Simulator (GPSS) are the two best-known languages; we say more about these programs in Sec. 21.8.]

The absence of a unifying theory of digital simulation is both a boon and a bane. On the positive side, you can build a simulation model containing arbitrarily high-order complexities and a huge number of dynamic interdependencies, as well as nonstationarities and correlated random phenomena. On the negative side, the more complicated the model, the more you will have to rely on embryonically developed statistical theory to perform the data analyses. As mentioned above, the very intricacy of the model can make it difficult to assess the model's validity. If the model is very complicated, you may have to expend a great deal of computer time on replication to obtain trustworthy answers and nearly optimal policies. Given the considerable research interest in simulation techniques, however, many of the current deficiencies in the theory and design of simulation experiments are bound to

be eliminated in the years ahead.

Objectives. You would construct a simulation model to assist in analyzing managerial decision problems with one or more of the following purposes in mind:

(i) To Describe a Current System. Consider a manufacturing firm that recently has witnessed an increase in its customer orders, and has noticed a consequent marked deterioration in meeting due-dates promised to its customers. This company may want to build a simulation model to study how its current procedures for estimating due-dates, scheduling production, and ordering raw material are giving rise to the observed delays.

(ii) To Explore a Hypothetical System. Consider a hospital that is contemplating the installation of a new computerized inventory replenishment system for its medical supplies. It may want to build a simulation model using historical data to test what the average level of inventory investment would be, and how often there would be shortages of various supplies under the proposed plan.

(iii) To Design an Improved System. Consider a job shop in which machine capacities are allocated by priorities assigned to each job. The company may want to build a simulation model in order to find an effective way to assign such priorities so that the jobs are completed without long delays and, at the same time, so that equipment utilization is acceptably high. We turn next to the steps in constructing and applying a simulation model.

So you want to build a simulation. The outline to follow describes the way

you would go about constructing a simulation:

Step 1. Formulate the Model. This step is much the same as that for other operations research models. There is an ever-present danger, however, of including too much detail in a simulation model and, as a result, consuming excessive amounts of computer time to perform the experiments. The best guard against this tendency is to keep your specific purpose constantly in mind. For example, if a model is to aid in the choice between two different locations for a new warehouse, it is probably not necessary to simulate activities on a hour-to-hour, or even day-to-day basis; weekly aggregates ought to suffice. If, on the other hand, a model is to aid in the choice between one or two loading docks at a new warehouse, then it may be necessary to simulate activities occurring in intervals as small as 5 to 15 minutes.

Step 2. Design the Experiment. You will reduce the chance of making mistakes and wasting time if you work out the details of the experimental procedures before running the model. This means that you need to think out carefully what operating characteristics of the simulated system you plan to measure. Further, you must consider the statistical tools you intend to apply to take account of the experimental fluctuations in the measurements.

Step 3, Develop the Computer Program. The simulation experiments will be performed entirely by a high-speed electronic calculator. That is, each historical evolution of the model, including the generation of random events, will take place within the computer. If the simulated model has a very simple structure, you may find it easiest to use a standard programming

language, such as FORTRAN, PL/1, or ALGOL, to develop the computerized version. More likely, you will find it preferable to employ one of the several simulation languages, such as SIMSCRIPT or GPSS, that are available on many large-scale electronic computers.

When you undertake an actual application, you will find that the above steps are not completely separate and sequential. For example, if you have already become familiar with, say, the GPSS simulation language, then you may want to formulate the model, initially, in terms of this language. We give more detail on each of these steps in the sections below.

3.3 STOCK MARKET SIMULATION EXAMPLE

An investor, Wynn Doe, wants to evaluate a particular strategy for buying and selling common stocks. To keep the exposition straightforward, suppose he does all of his trading in a single stock. At present, he holds 100 shares of the stock, which currently has a price of 810 a share. Again for the sake of simplicity, assume that the stock price can change each day by only ± 1 , so that some of the possible stock prices are $P, P-1, P, P+1, P+2, \dots$. The investor makes, at most, one transaction each day, and pays a commission of 2% of the transaction value whenever he buys or sells; of course, he need not make a transaction every day.

Wynn Doe wants to test the profitability of the following rule for buying and selling that has been suggested by his broker Benton Churning;

(i) If you own the stock, then sell it whenever the price falls.

(ii) If you do not own the stock, then buy it whenever the price rises.

According to this rule, if Wynn Doe owns the stock he will hold on to it while the price stays the same or rises; if he does not own the stock, he will refrain from buying it as long as the price stays the same or falls.

In order to evaluate this strategy, Wynn Doe must also postulate how he believes the stock price will fluctuate from day to day. After analyzing historical data, he formulates the price-movement model shown. To illustrate, if the share prices on Monday and Tuesday are both \$10, then he believes that the price on Wednesday will be \$11 with probability i , \$10 with probability $\frac{1}{2}$, and \$9 with probability $\frac{1}{2}$, as can be seen in the second row. If, instead, Tuesday's price is \$9, then he believes that the share price on Wednesday will be \$10 with probability $\frac{1}{2}$, \$9 with probability i , and \$8 with probability $\frac{1}{2}$, as can be seen in the third row. Notice that as the stock price increases, the investor thinks there is probability that it will increase again, and analogous statements hold if the share price remains the same or decreases.

To begin testing Churning's rule by manual simulation generate a specific history of price movements according to the probabilities given. A simple mechanism for doing this is to toss a pair of unbiased coins, using the correspondences shown. Verify that the assignments of the outcomes of a toss of two unbiased coins yield the postulated probabilities.

Suppose you simulate 20 days of activity, starting on Day 1 and ending on Day 20. Then you must toss the two coins 20 times; a particular sequence of tosses is recorded To determine the associated sequence of stock prices, you have to specify the initial conditions, namely, the stock price on Day 0 and whether it represents a fluctuation from the preceding day. the price on Day 0 is \$10, which presents no change from the preceding day. Given these initial conditions and a toss having a head and a tail on Day 1, the stock price for Day 1 is \$10, according to the second row of Then on Day 2, since yesterday's price remained the same, the toss of two tails implies that the share price falls to \$9, again according to the second row Proceeding to Day 3, since yesterday's price decreased, the toss of two heads causes the share price to be \$10, according to the third row Check the

Day	Coin Toss	Yesterday's Price	Today's Stock MoToss
0	—	—	10*
1	H/T	Same*	10
2	ZT	Same	9
3	2H	Decreased	10
4	ZH	Increased	10
5	ZH	Same	11
6	H/T	Increased	12
7	ZH	Increased	12
8	ZT	Same	11
9	ZH	Decreased	12
10	H/T	Increased	13
11	2T	Increased	12
12	H/T	Decreased	11
13	ZT	Decreased	11
14	ZH	Same	12
15	H/T	Increased	13
16	H/T	Increased	14
17	ZT	Increased	13
18	H/T	Decreased	12
19	H/T	Decreased	11
20	ZT	Decreased	10

Legend; H/T A Head and A Tail

2H Two Heads

2T Two Tails

Simulated Price Movements

entries for Days 5, 10, 15, and 20.

You can now determine how well Churning's suggested rule for buying and selling has performed on this particular simulated 20-day history of price movements. The details are shown notice that the history of prices from has been copied for easy reference. The entries in the column labeled "Decision" are a direct consequence of the price history and the suggested rule. The entries in the last three columns are determined after some auxiliary calculations.

To illustrate, on Day 2, the investor sells his 100 shares at a price \$9; but he must pay a 2% commission, which amounts to $(.02 \times \$9 \times 100 = \$18)$; thus he receives only $\$882 (= \$900 - \$18)$ from the sale. On Day 3, he repurchases the stock. Once again he must pay a 2% commission, so effectively the stock price is \$10.20 a share. Since he has \$882 cash, he can purchase only 86 shares, leaving him $\$4.30 (= \$882 - 86 \times \$10.20)$ cash. Notice that at the end of the 20th day, the investor's cash position—\$931.90—is worse following the rule than it would have been if he had sold his 100 shares on Day 0 and thereby received 8980 cash, after paying the commission.

Given all the model's assumptions, is Churning's rule profitable? Probably your immediate reaction is, "No." But wait a minute. Suppose instead of arbitrarily selecting 20 days as the length of the simulation, you had picked either 6 or 16 days instead. What would your answer have been then? Or suppose you rerun the simulation with a new history of 20 tosses.

Will the rule still look poor at termination? The issue of whether the rule is any good really depends in part on the statistical variability in the result obtained on Day 20, and on the significance of looking at a horizon of 20 versus 200, versus 2000, versus any other number of simulated days.

As you think further about the model, you will realize that the evaluation issue is complicated by the fact that as the horizon lengthens, there is an increase in the possible range of variability in the investor's wealth position at the end of the horizon. Further, even ! the rule implies an upward drift in the expected wealth position as the horizon lengthens, there is at least an initial increase in the probability that the investor may go broke along the way.

Day	Stock Price	Decision	Shares Held	Value of Stock	Cash
0	10		100	1000	
1	10		100	1000	
2	9	Sell	0	0	882.00
3	10	Buy	86	860	4.80
4	10		86	860	4.80
5	11		86	946	4.80
6	12		86	1032	4.80
7	12		86	1032	4.80
8	11	Sell	0	0	931.88
9	12	Buy	76	912	1.64
10	13		76	988	1.64
11	12	Sell	0	0	895.40
12	11		0	0	895.40
13	11		0	0	895.40
14	12	Buy	73	876	1.88
15	13		73	949	1.88
16	14		73	1022	1.88
17	13	Sell	0	0	931.90
18	12		0	0	931.90
19	11		0	0	931.90
20	11		0	0	931.90

Twenty-Day Test of Wynn Doe's Trading Rule.

So as you can see, even this simple-minded simulation gives rise to some difficult questions concerning what to measure and how to design a scientific experiment to test the effectiveness of the rule. What is more, if you take the trouble to run the mode) by hand for another 20 periods, you will quickly appreciate the desirability of letting an electronic computer do all the coin tossing and arithmetic.

3.4 BUILDING A SIMULATION MODEL

We now return to a more general discussion of the steps involved in using computer simulation. In this section we examine three aspects of model building: specifying the model's components; testing its validity and reliability; determining its parameters and measuring its performance,

Model components. The structure of most simulation models is conveniently described in terms of its **dynamic phenomena** and its entities. The dynamic phenomena in the stock market simulation of the preceding section include the investor's activity of buying or selling the stock, according to the stated **decision rule**, and the factors governing the movement of stock prices. The entities on any day include the amount of stock the investor holds, his cash position, and wealth. Typically, the entities in a model have **attributes**. To illustrate, the amount of stock the investor holds has a monetary value, given the associated price of the stock. Further, there are **membership relationships** providing connections between the entities. For example, the investor's wealth on any day includes both his cash and stock positions.

At any instant of a simulation, the model is in a particular **state**. The description of the state not only embodies the current status of the entities but frequently includes some historical information. For example, the state of the system at the beginning of a day in the stock market simulation is described by yesterday's price, how yesterday's price differed from the price on the day before, the number of shares held, and the cash position.

A model also can encompass **exogenous events**, that is, changes that are not brought about by the previous history of the simulation. To illustrate, the investor in the stock market simulation may have decided to add \$1000 more cash from his savings on Day 21, regardless of how well he has done using the tested strategy.

Knowing the state of the system and the dynamic phenomena, you can then go on to determine the subsequent activities and states. Frequently, simulation models having this evolutionary structure are called **recursive or causal**.

Note that in building a causal model, you must resolve the way activities occur within a period. For example, on each day of the stock market simulation, first the price is determined, then the decision to buy or sell is exercised. Actually, the price of a stock may change several times during a day, so the model we constructed is only a rough approximation to reality. The model also assumes that if the investor sells the stock, he receives the cash at the end of the day; and analogously, if he purchases the stock, he pays the cash at the end of the day. Such financial transactions do not always occur so

rapidly in practice.

Model validity and reliability. After building a simulation model, you are bound to be asked, “How realistic is it?” The more pertinent question is, “Does the model yield valid insights and reliable conclusions?” After all, since the model can only approximate reality, it must be evaluated by its power to analyze the particular managerial decisions you are studying.

Once the purpose of the simulation experiment is defined, you construct each piece of the model with a commensurate amount of detail and accuracy. A caveat is in order here. As simulation experts can attest, it is easy for a novice to build a model that, component by component, resembles reality; yet when the pieces are hooked together, the model may not behave like reality. So beware not to assume blindly that the entire simulated system is sufficiently accurate, merely because each of the component parts seems adequate when considered in isolation. This warning is especially important, because usually the objective of a simulation model is to fathom the behavior of a total system, and not that of the separate parts.

Model parameters and performance measures. It is one thing to describe the pieces of a simulation model abstractly, and it is another to collect sufficient data for a trustworthy representation of these pieces. Limited availability of data may very well influence the way you build a simulation.

You must be particularly cautious when you are dealing with extrapolated

data and nonstationary performance measures. (Remember the story of the cracker barrel manufacturer who, not so very long ago, forecasted that he would be selling millions of barrels today. He assumed, unquestioningly, that his sales trend would continue as it had in the past.)

You also must watch out for cyclical or periodic phenomena. When these are present, you must be judicious in selecting the variables to measure in the experiments. If you look only at “ending values,” for example, then your conclusions may be very sensitive to the exact length of the horizon that you simulated.

3.5 GENERATING RANDOM PHENOMENA

Most applications of simulation models encompass random phenomena. For example, in simulated waiting line models, the random variables include arrival and service times; in inventory models, the variables include customer demand and delivery times; and in research and development models, the variables include events of new product discoveries. Frequently, such simulations require thousands, and sometimes hundreds of thousands, of draws from the probability distributions contained in the model. How an electronic computer makes these draws is the subject of this section.

Uniform random numbers. As you will see, the basic building block for simulating complex random phenomena is the generation of random digits. The following experimental situation is an illuminating description of what we

mean by generating a sequence of uniform random numbers,

Suppose you take ten squares of paper, number (them 0, 1, 2, . . . , 9, and place them in a hat. Shake the hat and thoroughly mix the slips of paper. Without looking, select a slip; then record the number that is on it. Replace the square and, over and over, repeat this procedure. The resultant record of digits is a particular realized sequence of uniform random numbers. Assuming the squares of paper do not become creased or frayed, and that you thoroughly mix the slips before every draw, the n th digit of the sequence has an equal, or uniform, chance of being any of the digits 0, 1, 2, . . . , 9, irrespective of all the preceding digits in the recorded sequence.

In a simulation, you typically use random numbers that are pure decimals. So, for example, if you need such numbers with four decimal places, then you can take four at a time from the recorded sequence of random digits, and place a decimal point in front of each group of four. To illustrate, if the sequence of digits is 3, 5, 8, 0, 8, 3, 4, 2, 9, 2, 6, 1, . . . , then the four-decimal-place random numbers are .3580, .8342, .9261, Suppose you have to devise a way for making available inside a computer a sequence of several hundred thousand random numbers. You would probably first suggest this idea; perform something like the “slips-in-a-hat experiment” described above, and then store the recorded sequence in the computer’s memory. This is a good suggestion, and it is sometimes employed. The RAND Corporation, using specially designed electronic equipment to perform the experiment, actually did generate a table of a million random digits. The

table can be obtained on magnetic tape, so that blocks of the numbers can be read into the high-speed memory of a computer as they are needed. Several years ago, this tabular approach looked disadvantageous, because considerable computer time was expended in the delays of reading numbers into memory from a tape drive. But with recent advances in computer technology and programming skill, these delays have been virtually eliminated.

Experts in computer science have devised mathematical processes for generating digits that yield sequences satisfying many of the statistical properties of a truly random process. To illustrate, if you examine a long sequence of digits produced by these deterministic formulas, each digit will occur with nearly the same frequency, odd numbers will be followed by even numbers about as often as by odd numbers, different pairs of numbers occur with nearly the same frequency, etc. Since such a process is not really random, it is dubbed a **pseudo-random number generator**.

Computer simulation languages, like those discussed in Sec. 21.8, invariably have a built-in pseudo-random number generator. Hence, you will rarely, if ever, need to know specific formulas for these generators. But if you want to strengthen your confidence in the process of obtaining the numbers, then you can study the example of a pseudo-random number generator given below. If not, go on to the discussion of how to generate random variables.

Congruential method. To begin, we need to review the idea of **modulus arithmetic**. We say that two numbers x and y are **congruent, modulo** m , if the quantity $(x - y)$ is an integral multiple of m . For example,

letting $m = 10$, we can write

$$\begin{array}{ll}
 3 = 3 \pmod{10} & 4 = 4 \pmod{10} \\
 13 = 3 \pmod{10} & 84 = 4 \pmod{10} \\
 \text{(1)} \quad 313 = 3 \pmod{10} & 124 = 4 \pmod{10} \\
 48,653 = 3 \pmod{10} & 1,000,004 = 4 \pmod{10}.
 \end{array}$$

To find the value of, say, $857 \pmod{10}$, you calculate the integer remainder of 857 divided by 10, which is 7.

One popular approach for generating pseudo-random numbers is the so-called **Multiplicative Congruential Method**. The general formula for producing the random numbers is

$$\text{(2)} \quad r_n = ar_{n-1} \pmod{m},$$

where the parameters a and m , and the seed r_0 are specified to give desirable statistical properties of the resultant sequence. Note that because of the modulus arithmetic, each r_n must be one of the numbers $0, 1, 2, 3, \dots, m - 1$.

Clearly, you must be careful about the choice of a and r_0 . For example, if $a = 1$, then $T_n = r_a$, for all n . Or if $r_0 = 0$, then $r_n = 0$, for all n . The values of a and r_c should be chosen to yield the largest cycle or period, that is, to give the largest value for n at which $r_n = r_0$ for the first time.

To illustrate the technique, suppose you want to generate ten-decimal-place numbers u_1, u_2, u_3, \dots . It can be shown that if you use

$u_a = r_a \times 10^{-10}$, where

$$(3) \quad r_n = 100,003 r_{n-1}, \text{ (modulo } 10^{10})$$

$r_0 = \text{any odd number not divisible by } 5,$

then the period of the sequence will be 5×10^9 ; that is, $r_a = r_0$ for the first time at $n = 5 \times 10^9$, and the cycle subsequently repeats itself. Given that you want ten-decimal-place numbers, this is the maximum possible length of period using (2). (There are other values for a that also give this maximum period.)

Verify that the selection of r_a in (3) eliminates the possibility that $r_n = 0$; so u_n satisfies $0 < u_n < 1$.

Let us look at an example of (3). Suppose $r_0 = 123,456,789$. Then

$$(4) \quad r_1 = (100,003)(123,456,789) = 12,346,049,270,367 \\ = 6,049,270,367 \text{ (modulo } 10^{10}),$$

so that $u_1 = .6049270367$, and

$$(5) \quad r_2 = (100,003) - (6,049,270,367) = 604,945,184,511,101 \\ = 5,184,511,101 \text{ (modulo } 10^{10}),$$

n	U_n
1	.60492 70367
2	.51845 1 1 101
3	.66636 33303
4	.332 1 1 99909
5	.99544 99727
6	.9836 1 99181
7	.94266 97543
8	.80343 92629
9	.33660 77887
10	.78869 33661
11	.70269 00983
12	.1 1790 02949
13	.383 1 9 08847
14	.23804 26541
15	.97953 79623
16	.73484 38869
17	.59322 16607
18	.94573 49821
19	.33541 49463
20	.50087 46389

The Multi-plicative

Congruential Method: $a =$

100,003; $\tau_0 = 123,456,789.$

so that $u_2 = .5184511101$. The decimals u_m , $m = 1, 2, \dots, 20$, are shown in Fig. 21.5. Notice that the rightmost digits in this sequence form a short cycle 7, 1, 3, 9, 7, 1, 3, 9, Thus the statistical properties of the digits near the end of the number are far from random.

While (2) works reasonably well for some types of simulation models, it has poor serial correlation properties that make it dangerous to use for dynamic systems. A simple device for rectifying this deficiency is to intermix several sequences, each being generated with a different value for the seed τ_0 , and possibly a different value for a . For example, you can sequentially rotate among, say, 10 of these generators.

The advantage of using a pseudo-random number generator in lieu of a recorded table of random numbers is that only a few simple computer instructions are required to generate the sequence. Therefore, the approach uses only a small amount of memory space and does not require reading magnetic tape.

Generating random variables. We turn next to an explanation of how to employ a sequence of uniform random numbers to generate complex probabilistic phenomena. The treatment below suggests several techniques that can be used; but it is by no means exhaustive. Further, the examples that illustrate the techniques are chosen more for expository ease than computational efficiency. In an actual situation, you should seek the advice of a computer science specialist to determine the appropriate technique for your model.

Inverse Transform. Method. The following is the simplest and most fundamental technique for simulating random draws from an arbitrary single-variable probability distribution. Let the distribution function for the random variable be denoted by

$F(x)$ = probability that the random variable has a value less than or equal to x .

For example, suppose the random phenomenon has an exponential density function

$$(6) \quad f(t) = \lambda e^{-\lambda t}, \quad t \geq 0,$$

then

$$(7) \quad F(x) = \int_0^x \lambda e^{-\lambda t} dt = 1 - e^{-\lambda x}$$

Now $0 < F(x) < 1$, and suppose $F(x)$ is continuous and strictly increasing. Then given a value u , where $0 < u < 1$, there is a unique value for x such that $F(x) = u$. Symbolically, this value of x is denoted by the inverse

function $F^{-1}(u)$. The technique is to generate a sequence of uniform random decimal numbers U_n , $n = 1, 2, \dots$, and from these, determine the associated values as $x_n = F^{-1}(u_n)$.

The correctness of this approach can be seen as follows. Consider any two numbers u_a and u_b , where $0 < u_a < u_b < 1$. Then the probability that a uniform random decimal number u lies in the interval $u_a < u < u_b$ is $u_b - u_a$. Since $F(x)$ is continuous and strictly increasing, there is a number x_a such that $F(x_a) = u_a$, and a number x_b such that $F(x_b) = u_b$, where $x_a < x_b$. The Inverse Transform Method is valid provided that the true probability of the random Variable having a value between x_a and x_b equals the generated probability $u_b - u_a$. This true probability is $F(x_b) - F(x_a) = u_b - u_a$ by construction, so that the method is indeed valid.


To see how this method works, return to the exponential distribution (6) and (7). Let v_n , denote a uniform random decimal number. Set

$$(8) \quad v_n = 1 - e^{-\lambda x_n},$$

so that

$$(9) \quad x_n = \frac{-\log_e(1 - v_n)}{\lambda} = \frac{-\log_e u_n}{\lambda}$$

where $u_n = 1 - v_n$ and hence is itself a uniform random decimal number. Thus, you generate a sequence of uniform random decimal numbers u_1, u_2, u_3, \dots , and by (9) compute x_1, x_2, x_3, \dots , to obtain a random exponentially distributed variable. A diagrammatic representation of the technique is shown



The idea can also be applied to a probability mass function $P(j)$, Suppose $j = 0, 1, 2, 3, \dots$, so that

$$(10) \quad F(x) = \sum_{j=0}^x p(j)$$

then the inverse function can be written as

$$(11) \quad x^n = j \text{ for } F(j-1) < u^n \leq F(j),$$

where we let $F(-1) = 0$. For example, suppose the probability mass function is the binomial distribution:

$$(12) \quad p(j) = \binom{k}{j} p^j (1-p)^{k-j} \text{ for } j = 0, 1, \dots, k,$$


where $0 < p < 1$, and k is a positive integer. In particular, assume $k = 2$ and $p = .5$; then $p(0) = 1/4, p(1) = 1/2$ and $p(2) = 1/4$, so that by (11) you have

$$0 \text{ for } 0 < u^n < .25$$

$$1 \text{ for } 0 < u^n \leq .25$$

$$(13) \quad x^n = \begin{cases} 1 & \text{for } .25 < u^n \leq .75 \\ 2 & \text{for } .75 < u^n \leq 1. \end{cases}$$

Since u_n is a uniform random decimal number, there is a 1/4 probability that u_n lies between 0 and .25, a 1/4 probability that it lies between .25 and .75) and a 1/2 probability that it lies between .75 and 1. A diagrammatic representation of the technique is shown. Of course, many continuous distribution functions $F(x)$ do not have analytic inverse functions as does the exponential distribution. The Inverse Transform Method can still be applied in these instances by employing a discrete approximation to the continuous function, that is, by storing the values of $F(x)$ for only a finite set of x . The accuracy of the approximation can be improved by interpolating between the stored values. In several computer simulation languages (such as GPSS), you need only specify this discrete approximation, and the corresponding random phenomenon will be automatically generated.



Inverse Transform Method for a Probability Mass Function.
 Example: $p(0) = \frac{1}{4}$, $p_1 = \frac{1}{2}$, $p_2 = \frac{1}{4}$.

Tabular Method. The rule in (II) is easily implemented for an electronic computer by means of a few standard programming instructions. But if the range of possible values for j is large, then an excessive amount of time may be consumed in searching for they that satisfies the inequalities in (11). A faster version of the Inverse Transform Method can be employed at the expense of using part of the computer's internal memory for storing a long table. We illustrate the idea with the binomial example (13).

You can store the inverse function in computer memory in the form

$$\begin{aligned} & 0 \text{ for } s = 1, 2, \dots, 25, \\ \text{(14) } G(s) &= 1 \text{ for } s = 26, 27, \dots, 75, \\ & 2 \text{ for } s = 76, 77, \dots, 100. \end{aligned}$$

Given a value of u_n , let d_a be the number formed from the first two digits of u_n . Set $s_n = d_a + I$, and let $x_n = G(f_n)$. For example if $u_a = .52896. . .$, then $s_n = 52 + 1 = 53$, and so $x_n = G(53) = I$ in (14). Once the function $G(s)$ has been stored in memory, only a few calculations must be performed by the computer to produce x_n .

Method of Convolutions. Sometimes you can view a random variable as the sum of other independently distributed random variables. When this is so, the probability distribution of the random variable is a convolution of probability distributions, which may be easy to generate. (Occasionally, you can obtain a workable approximation to a complex probability distribution by using a weighted sum of independently distributed random variables. For

this reason, the approach also has been called the **Method of Composition.**)

To illustrate, consider a random variable having a gamma density function

$$(15) \quad g(y) = \frac{\lambda(\lambda y)^{k-1} e^{-\lambda y}}{(k-1)!}, \quad y \geq 0, \quad k \text{ a positive integer,}$$

Such a variable can be considered as the sum of k independent random variables, each drawn from the same exponential density specified in (6). Consequently, adding k independent values of x , as given by (9), yields a random variable with the distribution in (15).

Similarly, a binomial random variable, as specified in (12), can be viewed as the sum of k draws of a variable described by

$$(16) \quad i = \begin{cases} 1 & \text{with probability } p \\ 0 & \text{with probability } 1 - p. \end{cases}$$

You can therefore obtain a binomially distributed variable by adding k values of i . Each of these values for i is determined by the rule

$$(17) \quad i = \begin{cases} 1 & \text{for } 0 < u \leq p \\ 0 & \text{for } p < u \leq 1, \end{cases}$$

where u is a uniform random decimal number.

Method of Equivalent Transformations. Sometimes you can generate a random variable by exploiting a correspondence between its probability distribution and that of a related random variable.

For example, consider the Poisson distribution written in the form

$$(18) \quad p(j) = \frac{(\lambda T)^j e^{-\lambda T}}{j!} \text{ for } 0, 1, 2, \dots,$$

which has mean λT . In terms of the waiting line models you can interpret j as the number of customers arriving during a period of length T , where the interarrival times for the customers are independently and identically distributed exponential random variables with the density function specified in (6). Consequently, you can generate a Poisson distributed random variable by making successive independent draws of an exponentially distributed variable — using (9) to obtain such values. You stop making draws as soon as the sum of $j + 1$ of these variables exceeds T . The distribution of the resultant j is

(18) Explain why.

Normally distributed random variables. Unfortunately, the distribution function for the Normal density with mean 0 and variance 1,

$$(19) \quad F(x) = \int_{-\infty}^x \frac{1}{\sqrt{2\pi}} e^{-t^2/2} dt,$$

does not yield an analytic formula for the inverse function $F^{-1}(a)$. Of course the Inverse Transform Method can be used by employing a discrete approximation, and interpolating between values. But there are other methods for generating a Normally distributed random variable. Only a few are presented here.

One technique requires generating a pair of independent uniform random decimal numbers u and v , and in turn yields a pair of independently distributed

Normal random variables x and y having the distribution function in (19).

Specifically, compute

$$(20) \quad x = (-2 \log^e u)^{1/2} \cos 2\pi v$$

$$y = (-2 \log^e u)^{1/2} \sin 2\pi v$$

Alternatively, you can apply the Method of Convolutions and invoke the Central Limit Theorem. This technique employs the sum of k independently and identically distributed uniform random variables. Specifically, let u_i for $i = 1, 2, \dots, 12$, be independent draws of a uniform random decimal number; then compute

$$12$$

$$(21) \quad x = \sum_{i=1}^{12} u_i - 6$$

$$i-1$$

The distribution of x will have mean 0 and variance 1, and will be approximately Normal. The approximation is poor for values beyond three standard deviations from the mean.

A third approach is to compute

$$(22) \quad x = \left[\frac{(1-u)^{-1/6.158} - 1}{-1} \right]^{-1/4.874} - \left[\frac{u^{-1/6.158} - 1}{-1} \right]^{-1/4.874}$$

$$.323968$$

where u is a uniform random decimal number.

Correlated random variables. There are straightforward ways to generate variables from a multivariate Normal distribution, and from other joint probability distributions, as well as random variables having serial

correlation. The techniques go beyond the scope of this book, but can be found in most texts on computer simulation.

3.6 HOW TO MARK TIME

A dynamic systems simulation model can be structured in either of two ways. One approach, which is the more obvious, views simulated time as elapsing period by period. The computer routine performs all the transactions taking place in Period t , and then proceeds to Period $t - j - 1$. If the events in Period t imply that certain other transactions are to occur in future periods, then the computer stores this information in memory, and recovers it when the future periods arrive. You already saw an illustration of this approach using fixed-time increments in the stock market simulation of Sec. 21.4. Another example is given below.

In some simulations, the periods have to be relatively short. But there may be many of these periods in which no transactions occur. For such models, there is a second approach that lets the simulation advance by variable-time increments. This idea is illustrated in the second example below.

Time-step incrementation — inventory model. Suppose you wish to evaluate the operating characteristics of a proposed inventory replenishment rule. Assume that you can specify the probability distribution for each day's demand, and that daily demand is identically and independently distributed. If demand exceeds the amount of inventory on hand, the excess represents

lost sales. Let us postulate that, during a daily time period of the simulation model, the sequence of events is: first, any replenishment order due in arrives; then demand occurs; and finally, the inventory position is reviewed, and a reorder is placed if the replenishment rule indicates it should be. An order placed at the end of Period t arrives at the start of Period $t + L$, where L is fixed and $L > 1$.

To keep the exposition simple, assume that the replenishment rule is to order Q units whenever the amount of inventory on hand plus inventory due in is less than or equal to s , where $Q > s$. Verify that the inequality $Q > s$ implies there is never more than one replenishment order outstanding. (Since our focus here is simulation and not inventory theory, we do not comment further on the reasonableness of the replenishment rule; we do point out, however, that the model is an approximation to that in Sec. 19,6.)

A simulation model of this inventory system is easily constructed by stepping time forward in the fixed increment of a day, beginning with Day 1 ($t = 1$). To start the simulation, you must specify the initial conditions of the level of inventory on hand, the amount due in, and the associated time due in. You must also designate the number of periods that the simulation is to run; let the symbol "HORIZON" denote this value.

A flow chart of the simulation is shown in Fig. 21.8. The initializing is done in Block 1. For example, you can let the amount INVENTORY ON-HAND = Q , the AMOUNT DUE-IN = 0, the TIME DUE-IN = 0, and $t = 1$. Verify that when Block 2 is reached, the answer is "No," and you proceed at

Block 4 to generate a value of demand q for Day 1. Here is where you use an approach from the preceding section.

At the end of Day 1, INVENTORY ON-HAND is diminished by q , unless q exceeds the amount available, in which case the amount of-INVENTORY ON-HAND becomes 0. This calculation is performed at Block 5.


At Block 6, a test is made to determine whether a replenishment order is to be placed. If so, the AMOUNT DUE-IN becomes Q ., and the TIME DUE-IN becomes $I + L$ (since at the start $t = 1$), as indicated in Block 7. If a replenish-ment order is not placed, you continue directly to Block 8, where the time step is incremented by 1; that is, the simulation clock is advanced to Day 2.

If Day 2 goes beyond the HORIZON you specified, the simulation terminates. Assuming that you set the HORIZON > 1 , the simulation returns from Block 9 to Block 2.

At some day, TIME DUE-IN will equal (, and then the simulation branches from Block 2 to Block 3, where the amount of INVENTORY ON-HAND is augmented by Q , and the AMOUNT DUE-IN is reset to 0.

The flow chart does not indicate where you would collect statistical data on the operating characteristics of the system. In programming the model, you would keep a tally at Block 5 of the level of INVENTORY ON-HAND at the end of a day, as well as of the amount of lost sales and the number of days when a stockout occurs. You would tabulate at Block 6 the number of

days an order was placed.



Inventory Model Simulation Flow Chart.

Then, before terminating the simulation at Block 10, you would summarize these tallies into frequency distributions, along with their means, standard deviations, and other statistical quantities of interest.


Suppose the item is a “slow mover,” that is, there is a high probability that demand $q = 0$ on any day. Then the time-step method may be inefficient, because there will be many consecutive days when the computations in Blocks 2, 5, and 6 will be identical. Such redundancies can be eliminated by using

the technique illustrated below.

Event-step incrementation—waiting line model. Suppose you want to examine the operating characteristics of the following queuing system, which is simple to describe but proves difficult to analyze mathematically. Customers arrive at the system according to a specified probability distribution for interarrival times. The system has two clerks, A and B. When both servers are busy, arriving customers wait in a single line and are processed by a first come, first served discipline. The service times for each clerk can be viewed as independent draws from a specified probability distribution; but each clerk has a different service time distribution. Neither the interarrival nor the service time distributions are exponential.

After thinking about the way this system evolves over time, you will discover that the dynamics can be characterized by three types of events: a customer's arrival, a customer's service begins, and a customer's service ends. Each event gives rise to a subroutine in the computerized version of the system,

A simulation model using variable-time increments also contains a **master program** having an **event list**, which is repeatedly updated as the master program switches from one **event subroutine** to another. At the start of a simulation run, the event list is usually empty; but at a later instant in the run, it indicates when some of the future events are to occur. The role this event list plays will be clearer as you examine the event flow charts



Waiting Line Model Chart of Events.

Assume that you specify the initial conditions of the simulation as: a customer arrives, say, at Time 0, there are no customers in line, and both the clerks are free. The master program starts with the event subroutine CUSTOMER ARRIVES, shown as Block 1 in Fig. 1. Interarrival time at Block 2. The information that this next arrival event occurs at the implied future time is entered into the event list. A determination of whether both clerks are busy is made in Block 3. Since the answer is “No” at the start, the master program switches to the event subroutine SERVICE BEGINS, as indicated in Block 4. To keep the flow diagram uncluttered, we have suppressed the details that would specify that the computer must keep track of which clerk serves the customer, an item of information that is needed when the master program switches to the subroutine SERVICE BEGINS in Block 7.

The first instruction in SERVICE BEGINS is at Block 8, which records that the selected clerk is now busy. Then the service time of the customer is determined in Block 9, using the appropriate service time probability distribution for the selected clerk. The information that a service-ends event occurs at the implied future time is entered into the event list. The subroutine then transfers back to the master program with the instruction in Block 10 to find the **next imminent** event in the event list. So far, this can be either the arrival of the next customer or the completion of service of the first customer. Suppose it is the latter, so that the master program switches to the subroutine SERVICE ENDS in Block 11.

The first instruction in SERVICE ENDS is at Block 12, which records that the server is now free again. Then a test is made at Block 13 to see whether the waiting line designated by the symbol “LINE” is empty. Verify that the answer is “Yes.” Also check that when control switches back in Block 14 to the master program, the next imminent event will be the arrival of the second customer.

Later in the run, the LINE will contain customers, and then the answer is “No” at Block 13. As a result, the length of the LINE is decreased by 1 in Block 15, and control transfers in Block 16 to the subroutine SERVICE BEGINS. Explain what can happen subsequently if Block 10 leads to the CUSTOMER ARRIVES subroutine.

A simulation run progresses as each event subroutine either switches to another event subroutine or instructs the master program to increment

time to the next imminent event. As you can imagine, considerable skill is required to write a simulation program that uses event-step incrementation. In particular, expertise is needed to program the updating of the event list efficiently as future events are generated by the subroutines. Many simulation languages of the type discussed in Sec. 21.8 already include a master program that maintains an updated list of events; to employ these languages, you only have to specify the separate event subroutines.

We have glossed over a number of details in describing the queuing simulation model. We briefly mention a few of these before going on to the next section, which considers the design of simulation experiments. First, note that the charts do not show a test for terminating a simulation run. Of course, you must include such a calculation; you might state it by means of a time horizon or limit on the number of customers arriving. Second, observe that tabulating statistics on the operating characteristics is not an easy process because of the variable-time increments between successive events. Care must be taken to measure, for example, not only the frequency with which the waiting line has n customers, but also the associated fraction of the simulated horizon. Finally, recall that the initial conditions were chosen arbitrarily. If the queuing system in fact tends to be congested, then the effect of letting $LINE = 0$ at the start will take a while to wear off. Specifying appropriate initial conditions is part of the tactics of designing a simulation experiment.

3.7 DESIGN OF SIMULATION EXPERIMENTS

After constructing a simulation model, you face the difficult task of designing a set of runs of the model and analyzing data that emanate from these runs. For example, you must decide the

- Starting conditions of the model.
- Parameter settings to expose different system responses.
- Length of each run (the number of simulated time periods and the amount of elapsed computer time).
- Number of runs with the same parameter settings.
- Variables to measure and how to measure them.

If you are not careful, you can expend an enormous amount of computer time in validating the model to see whether it behaves like a real system, in estimating the system responses of the model to different parameter settings, and in discovering the response relationships among these parameters. Even then, and even after collecting a vast amount of data, you still may not have sufficiently accurate information to guide a managerial decision.

Surprising as it may seem, there has been relatively little development of statistical techniques aimed at constructing efficient designs of simulation experiments. By and large, professional management scientists have tried to “make do” with standard statistical tools to analyze experimental data from simulations. These techniques at best are only moderately successful, because most of them are not constructed for the analysis of multidimensional time series data. In particular, many (but not all) of the commonly used statistical

tools assume that separate observations of the variables being measured are uncorrelated and drawn from a Normal distribution with the same parameters.

We cannot possibly summarize all of the standard statistical techniques that can be applied to analyze simulation data. Instead, we discuss certain design procedures that enable you to employ many techniques ordinarily found in a this section requires a. knowledge of statistical methods modern text on experimental statistics. We also give a brief a statistical approaches that are particularly well suited to the simulations.

In search of Normality, Suppose you have constructed a queuing model to test two different service disciplines. For example, your application may-be a model of a job-shop production system, and the two disciplines lor processing orders are “first come, first served” and a particular priority scheme. Assume further that the difference in the two disciplines is to be measured solely in terms of the average waiting lime (exclusive of service) for orders. How might you ascertain what this difference is ?

This question is more difficult than it may appear at first glance. Since your measurements will be random variables, you must consider their statistical variability and be on the watch for certain kinds of complications. In any single simulation run, the waiting times of successive orders will be **serially cor-related** (sometimes called **autocorrelated**) ; that is, there is a greater like-lihood thai ihe $(n + 1)$ st order will be delayed if the n th order waits, than if the n th order commences service immediately. The extent of variation in

waiting times may itself be affected by the two different disciplines. The model may be unstable and the trend of waiting times may be ever upward. Even if the system does approach equilibrium, which may require a considerably long run, waiting times need not be Normally distributed. To ignore all these considerations and simply compare the average waiting times from a simulated run of each discipline is to court disaster.

Suppose you can demonstrate, on theoretical grounds, that the queuing model is stable, and that the effects of the starting conditions eventually fade away. Then it can be proved that even though the waiting times of successive orders are autocorrelated, the expected value of the sample average of these waiting times, taken over a sufficiently long run, is approximately that implied by the equilibrium distribution.

More precisely, let x_i , for $i = 1, 2, \dots, q$, represent q successive data observations of the random variable in a given simulation run, and define the **time-integrated average** as

$$(1) \quad x = \sum_{t=1}^q x_t .$$

Let μ represent the so-called ensemble mean of this random variable, as calculated from the equilibrium distribution. Then for q sufficiently large, we have the approximation

$$(2) \quad E[x] \approx \mu$$

Furthermore, it can be shown that the sampling distribution of x is

approximately Normal. You can calculate an estimate of the variance of this distribution as follows. Assuming that the process is **covariance-stationary** (the covariance between x_t and x_{t+k} depends only on k and not on t), and that the associated autocorrelations tend to 0 as k grows large, you first estimate these autocorrelations by :

$$(3) \quad r_k = \frac{1}{q-k} \sum_{t=1}^{q-k} (x_{t+k} - \bar{x})(x_t - \bar{x}) \text{ for } k = 0, 1, 2, \dots, M,$$

where M is chosen to be much smaller than q . (Unfortunately, a discussion of how much smaller M should be is too complicated to be given here, but can be found in the statistics literature under the subject title autocorrelation and **spectral analysis**.) The appropriate estimate of the variance of x is

$$(4) \quad V_x = \frac{1}{q} \left(r_0 + 2 \sum_{k=1}^M (1 - k/M) r_k \right)$$

Note that if, in fact, the time series is known to be free of autocorrelation, then the terms r_k , for $k = 1, 2, \dots, M$, would be eliminated from (4). The presence of positive autocorrelation, however, implies greater statistical variability in x as compared with the case of uncorrelated observations.

We now can look at two commonly employed approaches to statistical analyses. For the first method, consider making one very long run of each service discipline; specifically, take T consecutive observations in each run. Then you can apply (1) through (4) with $q = T$. If T is sufficiently large, the

statistic $(\bar{x} - \mu) \sqrt{N}/\sigma$ is approximately Normally distributed with mean 0 and variance 1. This fact allows you “to perform standard statistical procedures for hypothesis testing and constructing confidence intervals for μ , as well as to use modern Bayesian analysis. To compare the effect of the two service disciplines on average waiting time, you can apply standard statistical theory for discerning the difference between the means of two Normally distributed variables that have possibly unequal and estimated variances.

For the second method, consider making n independent replications, that is, n different runs. Suppose you want to have T observations in total from the n replications and that you take T/n observations from each run (assume T/n is an integer). Then for each replication p , calculate a time-integrated average \bar{x}_p , for $p = 1, 2, \dots, n$, using (1) with $q = T/n$. Afterwards compute the grand average

$$(5) \quad \bar{x} = \frac{1}{n} \sum_{p=1}^n \bar{x}_p$$

For any T/n , if n is large enough, the sampling distribution of \bar{x} is approximately Normal due to the Central Limit Theorem for the mean of independently and identically distributed random variables (namely, the \bar{x}_p). If you let T/n be large enough, the approximation is improved because of the near-Normality of the sampling distribution for each \bar{x}_p . What is more, it follows from (2) that when T/n is sufficiently large,

$$(6) \quad E[x] \approx \mu.$$

To determine the accuracy of x , you can estimate the variance of the sampling distribution of x from the variation in x_p , using

$$(7) \quad V_x = \frac{1}{N} \frac{\sum_{p=1}^n (x_p - x)^2}{n - 1}$$

Once again, if n and T are large, the quantity $(x - \mu)/\sqrt{V_x}$ is approximately Normally distributed with mean 0 and variance 1, and so the same sorts of statistical analyses can be performed as in the one-long-replication procedure.

Although the preceding discussion has related to a comparison of two different service disciplines in a job-shop model, these statistical approaches are generally applicable. In summary, assuming that the simulated system does approach an equilibrium, then under widely applicable conditions, you can legitimately average the successive observations of a simulated time series. As the number of observations grows large, this time-integrated average, in a probabilistic sense, converges to the desired ensemble mean implied by the equilibrium, distribution. (You can find the subject of probabilistic convergence treated in detail in texts on stochastic processes under the heading of ergodic theorems.) And furthermore, under widely applicable conditions, the time-integrated average is approximately Normally distributed. (You can look up the topic of Normal approximations in advanced statistics texts under the heading of the Central Limit Theorem for correlated random variables.)

Therefore, in many situations you can apply Normal-distribution theory if you either replicate simulation runs and then take a grand average of the individual time-integrated averages, or if you take a single time-integrated average from a very long run. A comparison of the relative merits of these two approaches as well as of other methods goes beyond the scope of this text. (The issues involved concern the amount of bias introduced by the starting conditions of the simulation and the stability properties of V_x .)

Sample size. Assume that you take a sufficient number of replications or let the simulation run long enough to justify using the Normal distribution to approximate the sampling distribution of the calculated averages. You still may need even more replications or a longer run to obtain the accuracy you require for decision analysis. The determination of an appropriate sample size for a simulation is no different from sample-size determination in ordinary statistical problems. Therefore, you can find the question discussed in detail in every text on statistical analysis.

We do emphasize, however, the influence of the number of observations on the accuracy of the statistical estimates. Whether you use the single-long-run approach, depicted in (1) through (4), with $q = T$, or the n -replication approach, depicted in (5) through (7), with $q = T/n$, the true variance of the sampling distribution for the calculated mean equals the reciprocal of the total number of observations T multiplied by another factor that is independent of T . Therefore, to reduce the standard deviation of the sampling distribution

of either x or 5 from a value of s , say, to $(.1)s$, you must increase the total number of observations to $100T$. More generally, to reduce the standard deviation by a factor of $1/k$, you have to take k^2 as many observations.

Usually, you cannot know how many observations to take at the start of a simulation, because you do not know the factors that multiply $1/T$ in the expressions for the true variances of the sampling distributions of x and x . For this reason, a commonly used procedure is to sample in two stages. In the first stage, you take a relatively small number of observations, and thereby calculate an estimate of the factor that multiplies $1/T$. With this estimate, you determine the remaining number of observations to take in the second stage to give the required accuracy. In actual applications, you may be surprised to find how many observations are needed to yield reasonable accuracy in the estimates. As pointed out above, the root of the difficulty is often the presence of positive autocorrelation. We discuss below a few approaches for coping with inherently large variation in the statistical estimates.

Variance-reduction techniques. There are a number of ways to improve the accuracy of the estimate of the ensemble average for a given number of data observations. These techniques are explained in texts on simulation under the heading of **Monte Carlo** or **variance-reduction methods**. Their use in management-oriented simulations is not yet widespread. We give only a couple of illustrations to suggest what is involved.

To assist in the exposition, we return to the example above of simulating a job-shop production system. Suppose, for the sake of definiteness, that you are simulating under the “first come, first served” discipline, and that you want to estimate the average waiting time of an order.

The first device we examine is sometimes called the Method of Control Variates, or alternatively, the Method of Concomitant Information. We present a highly simplified example of the idea. By elementary considerations you know that the interarrival times and the waiting times of each order are negatively correlated - roughly put, the longer the time since the previous order arrived, the shorter the waiting time of the latest order. State why. Therefore, suppose in a particular simulation run that the observed average of interarrival times is greater than the true average. Then you can use this information to add a positive correction to the observed average value of the waiting times. Similarly, suppose the observed average of interarrival times is smaller than the true average. Then you can make a negative correction to the observed average value of the waiting times. The technique explained below calculates either a positive or negative correction, whichever is appropriate.

Specifically, from the input data for the simulation you have the value of the true mean interarrival time, say, $1/\lambda$. Then let x_t represent the waiting time of Order t , and y_t , the interarrival time between Orders $t - 1$ and t . Consider the measurement

$$(8) \quad z_t = x_t + y_t - \frac{1}{\lambda}$$

and its time-integrated average

$$(9) \quad z = \frac{1}{T} \sum_{t=1}^T z_t = \frac{1}{T} \sum_{t=1}^T x_t + y_t - \frac{1}{\lambda} = \bar{x} + \bar{y} - \frac{1}{\lambda}$$

Note that the expectation of z is the same as that of x , since y is an unbiased estimator of $1/\lambda$. So you can use z as a consistent estimate of the average waiting time. But if x_t and y_t are sufficiently negatively correlated, then the variance of z will be less than the variance of x . A sample estimate of the variance of z can be calculated by assessing the variation in z from several replications, or by substituting z_t for x_t in (1), (3), and (4) above.

A more sophisticated method than (8) is to calculate $z_t = x_t + a(y_t - 1/\lambda)$ where now the value of a is specifically chosen to make the variance of z small. Under ideal conditions, a can be set such that $\text{Var}(z) = \text{Var}(x)(1 - p^2)$, where p is the correlation between x and y .

Before going on, we caution that the preceding example is meant only to be illustrative of the control variate idea. If you actually apply the technique to a queuing model like a job-shop production system, you should select a control variate that would absorb more of the sampling variation that would be accounted for by the interarrival times of orders. In fact, you probably should use several control variates instead of only one.

The second variance-reducing device we examine is called the **Method of Antithetic Variates**. The aim here is to introduce negative correlation between two separate replications of the simulation, so that the variance of the combined averages is less than if the replications were independent. {The idea can also be extended to more than two replications.)

Suppose in the job-shop production simulation that the interarrival times are determined by the Inverse Transform Method of Sec.21.5. Let u_t , for $t = 1, 2, \dots, TV2$, be the corresponding uniform random decimal numbers for generating the interarrival times y_t in the first simulation run of the model. Then in the second simulation, by using the values $1 - u_t$, which are also uniform random decimal numbers, the two time-integrated sample averages will be negatively correlated. State why.

Notice that the two simulations involve a total of T observations. Whether the mean of the two separate negatively correlated averages has less statistical variation than does the average of T autocorrelated observations from a single run depends on the extent to which the antithetic variates u_t and $1 - u_t$ induce negative correlation. Thus, the answer depends on the particular model being simulated, and the specific values of the model's parameters.

The crucial factor in deciding when to use variance-reduction techniques is whether, in fact, a given approach diminishes the variance of the estimates, and if so, whether the reduction is sufficient to warrant the extra computations required.

Multivariate analysis. The discussion so far has been partly misleading in that we have discussed examples involving the measurement of only a single operating characteristic for a system, such as average waiting time, and the comparison of only two alternatives, such as two different service disciplines. In real applications of simulation models, there are usually several operating characteristics of relevance and a multitude of alternatives to evaluate.

Multivariate analysis is by no means a new subject in statistics literature, but techniques for the analysis and design of experiments involving multivariate time series are just emerging. The reason for this relatively late development is that only recently has the availability of electronic computers made it practical to perform such data analyses.

By employing the approaches previously described to yield measurements that are Normally distributed, you have at least partially opened the storehouse of standard multivariate statistics. But still it is no simple matter to design a simulation experiment that can legitimately apply, say, latin squares, factor analysis, or multivariate regression.

Progress in devising helpful tools for multivariate analysis and complex experimentation is being made on two fronts. One important development, known as spectral analysis, aims at exploring the nature of serial correlation and periodicities in time series. The other front seeks methods for finding optimal levels of the decision variables; two such developments are response surface and stochastic approximation techniques. You can find these

developments explained in the technical statistics literature.

3.8 COMPUTER LANGUAGES

Unless you become both an operations research specialist and a computer-programmer, you personally will not have to translate your simulation model into a workable computer program. You should, however, know the major steps involved in this translation.

If your model is fairly simple and is a common application of simulation, then a so-called **canned program** may be available in which all you need do is specify a modest amount of input information. The best examples of this type of program are inventory control simulators. There are a number of canned programs that test the effectiveness of inventory replenishment rules. To employ these routines, you must supply the specific rules, such as “when down to 4, order 10 more/” or a formula to calculate the rules, given demand data. You also supply as part of the input either actual historical data on customer demand or a probability distribution for demand. The computer program then simulates the system for whatever number of time periods you designate, and calculates statistics such as the frequency of stockouts, the average inventory level, the number of orders placed, etc.

More typically, your model will require some special computer programming. If the simulation is only moderately complex, is to be used infrequently, and is to be programmed by personnel inexperienced in simulation techniques, then using a general purpose language, such as FORTRAN, PL/

1, or ALGOL, is probably the easiest way to accomplish the task. This type of computer language is familiar to all programmers of scientific problems; a programmer requires only the details of your model to translate it into computer language.

There is an important drawback to employing languages like FORTRAN, PL/1, and ALGOL. The programmer has to write, from scratch, subroutines for certain kinds of calculations that are included in almost all simulations. In the vernacular, the-programmer has to “reinvent the wheel/” For example, most simulations require generating random variables, and so a subroutine is needed for each such variable in the model. In addition, since you want to collect statistics on the system’s operating characteristics, subroutines have to be written to calculate these statistics, and a fair-sized associated programming effort must be accomplished to format the output of the simulation runs. Even a moderately complex model requires careful attention in organizing the data within the computer memory, writing a master routine for sequencing events in their proper order, and keeping track of simulated time within the computer.

Several computer languages have been developed for the specific purpose of easing the programming task of building a simulation model. These programs require that you specify only the probability distribution functions, and they automatically generate random events according to the distributions you indicate. Several of the languages collect statistics on whatever operating characteristics you want to examine, and report the results on predesigned

output forms. These languages also properly sequence events and keep track of time as it elapses in the model.

With such advantages, you may wonder why all simulations are not programmed in one of these languages. At present, there are several good reasons. One is that the languages differ to some extent from FORTRAN, PL/I or ALGOL, and hence require a programmer to become familiar with a new system.

One of the most powerful simulation languages is SIMSGRIPT; it requires a knowledge of FORTRAN and is fairly complex because of its considerable flexibility. At the other extreme of complexity is the General Purpose Systems Simulator (GPSS). It is a self-contained language that is easy to learn by beginners, but, accordingly, is restricted in its scope.

A second reason for not employing a simulation language is that it may not be available on the computer you want to use. This is rarely the determining factor today because SIMSGRIPT and GPSS programs are available for many computers, and there is widespread access to computer service bureaus that have these programs.

A third reason becomes important if the simulation is complex and is to be run frequently. A price you pay in using a simulation language is that it often runs slowly and consumes large amounts of a computer's high-speed memory. As a result, you may find it costly to perform many experiments, and your model may literally not fit into the available memory capacity of the computer.

As further technical improvements in simulation languages continue, and as management scientists gain more experience in employing computer simulation, it seems likely that such languages will be the common mode of solution.

3.9 DEUS EX MACHINA

So far we have discussed only simulation models that to some degree represent approximations to real situations. Their orientation has been to provide a simulated environment in which to test the effects of different managerial policies. A related class of simulation models tries to encompass goal-seeking or purposeful behavior. These models display what is termed **artificial intelligence**.

Some of the popular examples of artificial intelligence programs include computer routines for playing such games as chess and checkers. There also have been a few applications to managerial problems. One group of applications focuses on the behavioral patterns of individual decision-makers. A measure of such a model's success is how well it yields decisions agreeing with those of (the individual whose behavior is allegedly represented).

Another group of applications deals with complex combinatorial problems, like those discussed. 13. They are sometimes referred to as **heuristic programming** methods. For example, several of these models have been designed to derive good schedules for intricate sequencing problems. The following illustration suggests how they work.

Suppose the goal of the model is to schedule orders through a job-shop with maximum equipment efficiency. The computer starts by tentatively scheduling a few orders. It then selects another order to schedule, and examines various feasibility restrictions, due dates, and equipment efficiency. As a consequence, the computer may have to reschedule some of the previous orders. In brief, the computer model uses a number of “look-back” and “look-ahead” rules, and proceeds by educated trial-and-error toward a feasible schedule. If the rules are sufficiently sophisticated, then usually the schedule is good. Frequently, the schedule is nearly optimal according to the specified efficiency criterion, assuming the heuristic rules are promulgated with reference to this criterion.

Management scientists have also employed computer models for **operational gaming**. Some of the early applications, known as **management games**, involved several teams of players, each representing a business firm. A team made decisions about pricing, production quantities, advertising, etc. The computer served the two-fold purpose of keeping the accounting records, and of calculating the net impact of the decisions made by the several teams. More recently, such applications have been used to train personnel in administrative procedures, and to explore the system dynamics of an industry in which the competing firms are employing specified strategies.

CHAPTER 4

Advanced Topics in Network Algorithms

4.1 MAXIMUM FLOW THROUGH A CAPACITATED NETWORK

In the development of sophisticated techniques to solve difficult network models, analysis of the following problem is of central importance: given a network with arc capacities, where Node 0 is the source of all flow and Node $p + 1$ is the sink, what is the maximum amount of flow that can be routed from source to sink? Formally, the model is described as

$$(1) \quad \text{maximized } F$$

subject to

$$(2) \quad \sum_{(0,j) \text{ in network}} x^{0j} = F \quad \text{for } k = 0$$

$$(3) \quad \sum_{(k,j) \text{ in network}} x^{kj} - \sum_{(I,k) \text{ in network}} x_{ik} = 0 \text{ for } k = 1, 2, \dots, p$$

$$(4) \quad - \sum_{(i, p+1) \text{ in network}} x_{i, p+1} = -F \quad \text{for } k = p + 1$$

$$(5) \quad 0 \leq x_{ij} \leq u_{ij} \text{ for all } (i,j) \text{ in network,}$$

where u_{ij} are nonnegative integers.

A simple method can be used to solve the problem. To make the basic idea of the

algorithm transparent, assume

$$(6) \quad u_{ij} = 1 \text{ for all } (i,j) \text{ in network.}$$

Once you see how to find a solution with (6), you will have no difficulty in understanding the minor modification required to solve the general case.


Begin with any feasible flow. The steps in the technique either determine that the flow is maximal, or discover another solution with increased flow:

Step 1. Starting at Node 0, put a (+) on each arc (0,j) without flow and **label** Node j with a check mark (\surd). Put the mark (\surd) on Node 0.

Step 2. Consider any Node j that is labeled (\surd) Put a (+) on every flowless outward arc (j, k) if Node k is not labeled, and label Node k with (\surd). Then put a (-) on every inward arc (k,j) with flow if Node k is not labeled, and label Node k with (\surd). Finally, cross the check (\surd) on Node j to indicate that the node also has been **spanned**.



Step 3. Continue with the operation in *Step 2* until Node p + 1 is labeled or all labeled nodes have been spanned. A **breakthrough** occurs as soon as Node p + 1 is labeled, because a **flow-augmenting path** has been discovered from Node 0 to Node p + 1. The path can be found by tracing back from Node p + 1 the arcs that have been marked with a (+) or a (-). Add a unit of flow on each arc with a (-) and remove the flow from each arc with a (+). Return to Step 1. If, however, Node p + 1 remains unlabeled at the

termination of Step 2, then the current solution is maximal.



There are several ways to proceed. Follow the one below by making light pencil marks on

- (i) Scan Node 0: put + on arcs (0, 4) and (0, 5) and label Nodes 4 and 5 with (\surd) . Mark Node 0 as (\surd) .
- (ii) Scan Node 5: put — on arc (3, 5) and label Node 3 with (\surd) . Mark Node 5 with (\surd) .
- (iii) Scan Node 3: put — on arc (1, 3) and label Node 1 with (\surd) . Mark Node 3 with (\surd) .
- (iv) Scan Node 1: put + on arc (1, 2) and label Node 2 with (\surd) . Mark Node 1 with (\surd) .
- (v) Scan Node 2: put + on arc (2, 6) and label Node 6 with (\surd) . A flow augmenting path has been found.




Check your work with the result in Fig. 1.2. The solution therefore improves if you

- (i) Add a unit of flow on the arcs (2, 6) and (1,2).
- (ii) Remove a unit of flow from arcs (1, 3) and (3, 5).
- (iii) Add a unit of flow on arc (0, 5).

The revised solution is given in Fig. 1.3 on the next page.

Verify that the flow is now maximal by repeating the steps of-the algorithm. The sequence of nodes scanned will be Nodes 0, 4, and 5. It will not be possible to label any other node.

To remove the arc capacity restriction (6), the algorithm is modified in two respects. Put a (+) on every outward arc with less than capacity flow in the scanning process. Then, when a flow-augmenting path has been found, route as much flow as possible on the path, taking into account the amount of unused capacity on each (+) arc and the current level of flow on each (–) arc. Thus for the maximum flow problem (1) through (5) the algorithm now terminates according to the following stopping rule.



Flow = 3

One Unit: ① → ② → ③ → ⑥

One Unit: ① → ③ → ⑥

One Unit: ① → ⑤ → ⑥

Maximum Flow Solution.

STOPPING RULE. Every arc from a spanned node to an unlabeled node is at its full capacity and every arc from an unlabeled node to a spanned node is at zero flow.

At each iteration, *Step 3* results either in one or more units of increased flow or in termination; therefore the algorithm is finite, since the maximum possible flow is bounded. It remains only to show that when the algorithm terminates, the solution in fact is optimal. To do this, partition all the nodes into two classes, say, C_0 and C_{p+1} . Put Node 0 in C_0 and Node $p + 1$ in C_{p+1} . Such a separation is called a **cut**, and we define the **cut capacity** as the sum of all the u_{ij} such that Node i is contained in C_0 and Node j in C_{p+1} .

The cut capacity for any partition provides a limit on the maximal value of flow possible. (If the cut capacity equals 0, then Node 0 literally is cut off from Node $p + 1$ and no flow is possible between source and sink.) Consequently, if a feasible flow equals any cut capacity, the flow must be

optimal. Furthermore, because of the conservation of flow restrictions (2), (3), and (4), the value of F in any feasible routing must equal the sum of the flows along all arcs (i,j) minus the sum of the flows along all arcs (j,i) , where Node i is contained in C_0 and Node j in C_{p+l} .

With the above observations, the optimality proof is at hand. Consider the feasible solution when the algorithm terminates. Define C_0 to be all the spanned nodes, and C_{p+l} the remaining nodes. According to the Stopping Rule, there is no flow from any node in C_{p+l} to any node in C_0 , so the total flow equals the sum of the flows on all arcs from nodes in C_0 to nodes in C_{p+l} . All of these arcs contain flows at capacity levels; therefore the total flow in the solution equals the cut capacity, and no further improvement is possible. The preceding argument is summarized by the fundamental result below. MAX FLOW/MIN CUT THEOREM : The maximum flow F in the network structure (2) through (5) is equal to the minimal cut capacity relative to the source and sink. A corollary result is that the algorithm results in integral values for all x_{ij} .

4.2 SOLUTION OF THE ASSIGNMENT MODEL

Recall the assignment model, discussed in detail in Sec. 6.4:

$$(1) \quad \text{minimize } \sum_{i=1}^n \sum_{j=1}^n c_{ij}x_{ij}$$

subject to

$$(2) \quad \sum_{j=1}^n x_{ij} = 1 \quad \text{for } i = 1, 2, \dots, n$$

$$(3) \quad \sum_{i=1}^n x_{ij} \quad \text{for } i = 1, 2, \dots, n$$

$$(4) \quad x_{ij} = 0 \text{ or } 1 \quad \text{for all } i \text{ and } j.$$

By the nature of the problem, a feasible solution contains only n variables equal to 1, whereas a basic solution includes $n + n - 1$ variables. Consequently, when the simplex algorithm for networks is applied to the assignment model (1) through (4), each basis contains $n - 1$ routes at zero level. This observation suggests that the special structure of the assignment model is not fully exploited by the simplex algorithm. In this section, three other approaches will be explained.

The first makes use of the maximum flow problem of the preceding section. The second combines the principles of the maximum flow and shortest-route algorithms. The third demonstrates a further connection between the assignment problem and the shortest-route problem. The methods will be illustrated by an example, which is first solved by the simplex algorithm to provide a basis for comparison. As you will learn in Sec. 1.3, the three approaches are of additional significance because they generalize to other network optimization models.

Simplex algorithm. Consider the assignment problem shown. If you apply the procedure explained in the advanced material of Sec. 7.4 for calculating relative costs to assist in picking an initial basic solution, you obtain the amounts. A starting basis is exhibited. Notice that three routes (in

the second row) are at zero level, indicating degeneracy. You can follow the iterations by examining which contain the succession of improvement potentials and trial solutions. Observe that the value of the objective function does not change until the final solution.

2	10	15	0	
1				1
10	18	20	9	
0	1	0	0	1
15	24	26	10	
		1		1
12	25	27	8	
			1	1
	1	1	1	1

Total Cost = 54

Assignment Model Initial Basic Solution.

	$c_{ij} - v_i - w_j$				v_i
	1	1	4	0	0
	0	0	0	0	9
	4	5	5	0	10
	3	8	8	0	8
w_j	1	9	11	0	

Assignment Model Relative Costs:
 $c_{ij} - v_i - w_j$.

2	10	15	0		v_i
Q	0	-3	1		-8
10	18	20	9		0
Q	Q	Q	Q		0
15	24	26	10		6
1	0	Q	5		6
12	25	27	8		-1
-3	-8	-8	Q		-1
w_j	10	18	20	9	

Assignment Model Improvement Potentials for Initial Solution.

1				1
0	1	0		1
		1	0	1
			1	1
	1	1	1	1

Total Cost = 54 - 5(0) = 54

Assignment Model Second Basic Solution.

2	10	15	0		v_i
Q	0	-3	-4		-8
10	18	20	9		0
Q	Q	Q	-5		0
15	24	26	10		6
1	0	Q	Q		6
12	25	27	8		4
2	-3	-3	Q		4
w_j	10	18	20	4	

Assignment Model Improvement Potentials for Second Solution.

1				1
	1	0		1
		1	0	1
0			1	1
	1	1	1	1

Total Cost = 54 - 2(0) = 54

Assignment Model Third Basic Solution.

				v_i
	2	10	15	0
	\emptyset	2	-1	-2
	10	18	20	9
	-2	\emptyset	\emptyset	-5
	15	24	26	10
	-1	0	\emptyset	\emptyset
	12	25	27	8
	\emptyset	-3	-3	\emptyset
w_j	8	18	20	4

Assignment Model Improvement Potentials for Third Solution.

	1			1
	0	1		1
		0	1	1
1			0	1
	1	1	1	1

Total Cost = $54 - 2(1) = 52$

Assignment Model Optimal Basic Solution.

				v_i
	2	10	15	0
	-2	\emptyset	-3	-4
	10	18	20	9
	-2	\emptyset	\emptyset	-5
	15	24	26	10
	-1	0	\emptyset	\emptyset
	12	25	27	8
	\emptyset	-3	-3	\emptyset
w_j	8	18	20	4

Assignment Model Improvement Potentials for Optimal Solution.

Maximum flow approach. An alternative method is based on the maximum flow algorithm in the preceding section. for finding a good initial solution. The technique, in bare outline, consists of two steps:

Step 1. Given row constants u_i , for $i = 1, 2, \dots, n$, and column constants w_j , for $j = 1, 2, \dots, n$, yielding nonnegative relative costs $(c_{ij} - u_i - w_j) \geq 0$, determine whether a feasible solution exists using only routes with relative costs equal to 0. If so. stop, since the solution is optimal; otherwise go to Step 2.

Step 2. Revise v_i and w_j such that at least one new route has relative cost equal to 0. Return to Step 1.

The details of each step are explained with reference to the previous

example. You can always begin Step 1 with the constants used by the method for selecting an initial solution. However, in order to compare this approach with the simplex algorithm, it is convenient to start with the values indicated

The maximum flow algorithm is employed in Step 1 to find whether there exists a feasible solution using only routes with 0 entries Tabular short-cuts

				v_i	
	2	10	15	0	0
	0	0	3	0	
	10	18	20	9	8
	0	0	0	1	
	15	24	26	10	10
	3	4	4	0	
	12	25	27	8	8
	2	7	7	0	
w_j	2	10	12	0	


Assignment Model Relative
Costs: $c_{ij} - v_i - w_j$.

				v_i	
	2	10	15	0	0
	0	0	3	2	
	10	18	20	9	8
	0	0	0	3	
	15	24	26	10	12
	1	2	2	0	
	12	25	27	8	10
	0	5	5	0	
w_j	2	10	12	-2	

Maximum Flow Algorithm
Revised Relative Costs.

are available to carry out the procedure; but since the main purpose of this discussion is to make clear how the maximum flow model is of value, the exposition will not streamline the format to aid the calculations.

A network flow diagram comprised only of the routes with relative cost equal to 0 is constructed The node designation r_i , corresponds



to Row i and similarly k_j to Column j . All arcs out of the source node and into the sink node have capacity 1, corresponding to the row and column constraints of the problem.

If the maximum flow for the network were four, the corresponding routing would be an optimal assignment, since a unit of flow on arc (r_i, k_j) would imply $x_{ij} = 1$. A trial solution with three units of flow is exhibited in Fig. 1.15 The maximum flow algorithm is carried out in Fig. 1.16 to demonstrate that flow cannot be augmented in this network. Therefore you must go to Step 2.

The following rationale makes plausible a way to revise the u_i and w_j . Since the flow in




Figure is maximal but less than four units, it is necessary to introduce at least one new arc (r_i, k_j) . Given the nature of the flow algorithm, it is reasonable to restrict attention to those arcs such that Node r_i is labeled but Node k_j is not. If such an arc is added, then the steps of the flow algorithm will continue and permit at least one more Node k_j to be spanned.


In altering the u_i and w_j however, you should be careful not to destroy the equality $c_{ij} - u_i - w_j = 0$ for routes now having flow, as well as for those marked with (+) in the spanning process. Otherwise, you may not be able to continue the scanning where you left off in Step 1. Also, the relation $c_{ij} - u_i - w_j \geq 0$ must be preserved for all i and j . A rule that achieves all these conditions is

- (5) (i) Add c to u_i if Node r_i is labeled,
(ii) Subtract c from w_j if Node k_j is labeled, where
- (6) $c =$ smallest relative cost for arcs between every labeled Node r_i and unlabeled Node k_j .


Nodes r_3 and r_4 are labeled, whereas Nodes k_1 , k_2 , and k_3 are not; therefore examine the entries at the intersection of Rows 3 and 4 and Columns 1, 2, and 3 to obtain

(7) $c = \text{minimum}(3, 4, 4, 2, 7, 7) = 2 = c_{41}$.

The revised u_i and w_j are shown in Fig. 1.14 along with the new relative costs. Observe that arc (r_4, k_1) now has relative cost equal to 0, but arc (r_1, k_4) has a positive relative cost. The associated network appears The maximum



flow algorithm continues from and, in the process, labels Nodes $k_1, r_1, k_2, r_2, k_3,$ and finally the sink. The (+) and (-) signs on the arcs indicate the flow-augmenting path that produces the optimal solution in



We now recapitulate the algorithm. Starting with trial values for u_i and w_j in Step 1, apply the maximum flow algorithm to the associated network. If the resultant total flow equals n , the solution is optimal and you stop. Otherwise, proceed to Step 2, where you revise the u_i and w_j according to the rule in (5) and (6). Return to Step 1 with the resultant, somewhat altered, network: at least one new arc will have been added (where it is assumed the corresponding

c_{ij} equals the value c), and some unused arcs may have been dropped. Restart the maximum flow algorithm with the previous flow pattern. Again check for a total flow of n , and so forth.

Proof that the solution at termination is optimal is exactly the same as that given for the simplex method applied to the transportation model. Convergence in a finite number of iterations is established by noting that

- (i) Each time Step 2 occurs, at least one more node is spanned in the subsequent Step 1. Since the number of nodes is finite, Step 2 eventually results in a breakthrough at the succeeding Step 1.
- (ii) Only a finite number of breakthroughs can occur, because each results in increased flow, and total possible flow is bounded.

An implication of the convergence proof is that no more than $.5(n^2 + 3n - 2)$ applications of Step 2 are required. This bound is typically far in excess of what actually occurs; however, you should note that the bound is considerably smaller than $\binom{2n-1}{n}$, which is the simplex algorithm bound calculated in terms of the number of basic solutions possible.

The network flow approach has a certain resemblance to the dual simplex method in that the dual constraints are satisfied at every iteration, but a feasible solution is not obtained until termination. In sharp contrast to both the standard and dual simplex methods, the network approach does not maintain a basic solution.

Minimal cost/maximum flow approach. To initiate Step 1 of the previous approach, you had to find constants u_i and w_j such that all the relative costs $(c_{ij} - u_i - w_j)$ were nonnegative. You can assume without loss of generality that all $c_{ij} \geq 0$, since you can always add a positive constant c^* to every cost element. Then you can begin Step 1 in the previous approach by letting $u_i = w_j = 0$. If you do so, the sequence of solutions turns out to be a minimal cost routing among all routings with the corresponding amount of flow. This point of view leads to another statement of the algorithm that provides a helpful insight in generalizing the method to more complex network problems. The idea is to increase total flow in such a way that the routing for each higher level of flow incurs minimal total cost. Thus, if F units of flow have been so routed, the method seeks a flow-augmenting path with least cost, and increases flow on this path. The technique for finding the path employs the shortest-route algorithm.

The approach is summarized below:

Step 1. Construct a new network' based on the current solution as follows. Include each arc in the original network that currently is flowless, and let c_{ij} be the arc's path length. If flow occurs between Nodes r_i and k_j , add arc (k_j, r_i) and let its path length be $-c_{ij}$. It is possible to augment flow in the original network on any path between source and sink of the new network.

Step 2. Find a shortest path from the source to the sink in the new network. In the usual manner, increase flow on this path in the original network. If all the assignment model's constraints are satisfied, stop; otherwise, return to

Step 1.





FIGURE I.19. Minimal Cost/Maximal Flow Network.



Shortest Route.

The approach is illustrated for the example Assume you have applied the algorithm to obtain the minimal cost routing for three units of flow. The solution is the same as that in Fig. 1.15, and has a cost of 28 ($=2 + 18 + 4 - 8$).

The new network constructed according to Step 1 appears Notice the source is connected only to Node r_3 and the sink to Node k_3 , since the other such arcs are at capacity flow. The network contains arcs (k_1, r_1) , (k_2, r_2) , and (k_4, r_4) because the current solution has flow in the opposite direction. The remaining arcs have zero flow in the current solution.

The shortest-route algorithm of Sec. 7.6 is applied to the network and yields the result in Fig. 1.20. Sending a unit of flow along the indicated path produces the optimal solution in agreement with Fig. 1.18. Notice that a unit routed on arc (k_4, r_4) implies that flow is to be removed from arc (r_4, k_4) in the original network. Since the shortest path has length 24, the final solution has total cost of 52 ($= 28 - 24$). The length of the shortest routes from each Node r_i is an optimal value for $-u_i$ and similarly, the length from each Node k_j is an optimal value for w_j .

Shortest-route approach. you discovered how to convert a shortest-route problem into the format of an assignment model. In this chapter, you learned an efficient method for solving the shortest-route problem without making such a conversion. Now you will see how to solve the assignment model by putting it into the form of a shortest-route problem and then applying the shortest-route algorithm.

The general idea is to solve a 1 x 1 assignment model, which is a trivial task. Then use the answer to solve a 2 x 2 model, and continue in the same fashion, adding one more row and column at each iteration until the n x n solution is obtained. Given the solution for a p x p problem, any remaining row and any remaining column can be chosen to form the (p + 1) x (p + 1) problem.

As before, the technique is explained in terms of the example in Fig. 1.4 Suppose you have solved the 3 x 3 problem consisting of Rows 1, 2, and

4 and Columns 1, 2, and 4 of Fig. 1.4. For convenience, the array is displayed and the circled c_{ij} correspond to the optimal solution for that table. Notice that the solution is the same as that

	(Col. 1)	(Col. 2)	(Col. 4)	
(Row 1)	②	10	0	1
(Row 2)	10	⑬	9	1
(Row 4)	12	25	⑧	1
	1	1	1	

3 x 3 Assignment Problem.

	(Col. 1)	(Col. 2)	(Col. 4)	(Col. 3)	
(Row 3)	15	24	10	26	1
(Row 1)	2	10	0	15	1
(Row 2)	10	18	9	20	1
(Row 4)	12	25	8	27	1
	1	1	1	1	


4 x 4 Assignment Problem.

To form the 4 x 4 problem, add the remaining third row and column of Fig. 1.4, which gives Fig. 1.22. Then calculate relative costs $(c_{ij} - u_i)$, where $u_1 = 0$ for the new row and the other u_i are the values c_{ij} for the optimal routes in the 3 x 3 problem. The result appears. Notice the table has the

Node	1	2	3	4	v_i
0	15	24	10	26	0
1	0	8	-2	13	2
2	-8	0	-9	2	18
3	4	17	0	19	8


Relative Costs:
 $c_{ij} - v_i$.

same appearance as that for a shortest-route problem, such as Fig. 6.9 on p. 179. The network is drawn using the node designations appearing



Shortest Route Network for 4×4 Assignment Problem.

Apply the shortest-route algorithm to find a best path from Node 0 to Node 4 corresponds to the optimal solution to the original assignment problem.



Shortest Route from Node 0 to Node 4.

The only subtlety in demonstrating the validity of the approach is to establish that when the shortest-route network is drawn, the sum of the path lengths around every loop is nonnegative. Recall that such a condition is required for the shortest-route algorithm. Observe that only loops possible involve Nodes 1, 2, and 3, which correspond to the rows columns considered in the 3×3 problem. As usual, the optimal solution to the 3×3 problem must remain optimal in terms of the relative costs. A loop among the Nodes 1, 2, and 3 is tantamount to a particular solution to the 3×3 problem; therefore no loop can have a negative cost, since the optimal 3×3 solution has zero cost

4.3 ALGORITHMS FOR OTHER NETWORK MODELS

The techniques in the preceding section generalize readily to other network optimization problems.

The only alteration needed to adapt the maximum flow approach to handle the standard transportation problem in Sec. 6.2 is to let the arc capacity from the source to Node r be S_r , and similarly the arc capacity from Node k to the sink be D_k . For this problem, the number of times Step 2 is executed cannot exceed

$$(1) \quad \sum_{i=1}^n x_{ij} = 1 \quad \text{for } j = 1, 2, \dots, n$$

Analogously, the only essential change required in the minimal cost maximum flow approach to cover the general network optimization model (1) through (4) in Sec. 6.8 is to revise Step 1 so as to include each arc in the original network that at present has flow below capacity. (The statement concerning Nodes i and k applies to any Nodes i and j in the network.) As the advanced material at the end of Sec. 6.8 indicated, it may be necessary to transform the model initially so that there is only a single source and sink.

The shortest-route approach also can be extended to treat the standard transportation problem, which is not surprising since such a problem is conceptually equivalent to an assignment model.

CHAPTER 5

Implementation of Operations Research

5.1 Introduction

In the past few years, many significant technical breakthroughs have been made in operations research. Even greater progress occurred, however, in implementing operations research in commercial and governmental enterprises.

In the early 1960's, a practicing operations researcher had to be both a scientific expert and a master of the "art of persuasion." Ethical, but convincing, salesmanship was needed then because relatively few companies firmly believed that operations research was a profit-yielding activity. Most executives classified the effort as blue-sky research and development, and, in fact, several major corporations placed their operations research group in an R & D department.

Since then, the picture has changed dramatically. Only rarely now is an operations researcher called upon to defend his *raison d'être*. Today, executives show pride in employing computer models that have been designed to assist them in analyzing complex decision problems. (Many managers guard their computer models as a part of their territorial imperative.) In short, very few executives in leading corporations still ask, "Why do we need operations research?" They know why.

The questions that managers do raise are, "What areas of application

are the most profitable?” “Is our company spending too little or too much on operations research?” “How can I best use operations research?” In other words, the present interest of businessmen is learning how to get the maximum benefit from operations research.

The sections below explore the practical implications of this managerial attitude and offer some insights into the implementation process. Although the sections are entitled “How to . . .,” the chapter is not really a comprehensive nuts-and-bolts manual of procedures for guaranteeing profitable implementation. Rather, the chapter provides a few guidelines for making operations research work effectually. The orientation of the discussion is toward people, not mathematical techniques.

As you will soon discover, the key to the successful conduct of operations research is the joint exercise of good judgment by executives and professional operations researchers. In particular, the managers and technicians must decide in concert what projects to pursue, what goals to sight, what level of effort to expend, and what timetable to follow. These are the subjects analyzed below.

You may be curious to know why implementation is so much easier now. The reasons may not be obvious, as many of the most useful mathematical techniques have existed for more than 15 years. Part of the answer has already been given in Sec. 1.7, which explained how executives have come to appreciate the powerful analytic assistance of operations research. But this recognition is by no means the only factor contributing to a

widespread acceptance of the approach.

An equally important reason is that computing equipment has become vastly more accessible, and that the number of programming specialists has increased concomitantly. The improvement in so-called software computer programs that assist in solving large-scale optimization and simulation problems as well as the development of time-shared computing systems have diminished the task of systems design and lessened the difficulties of obtaining numerical solutions. Another reason is that the population of well-trained and experienced operations researchers has expanded sufficiently to provide the required professional manpower.

5.2 HOW TO PUT OPERATIONS RESEARCH TO MANAGERIAL USE

What commonly distinguishes an executive familiar with employing operations research from a first-time user is his recognition of the need to exercise responsibility vis-a-vis the conduct of the project. For quite understandable reasons, tyro managers are usually “stand-offish” in their involvement. Such a posture is ill advised and can be expensive to the company, even when the operations research application ultimately succeeds.

In essence, line managers must take responsibility to see that the right problem is analyzed and that adequate controls are exercised to monitor the progress of the application. Experience has shown repeatedly that ignoring this responsibility is detrimental to all and may easily be the root cause of

failure, despite the expertise and sincerity of the operations research technical staff. This section suggests some ways for an executive to ensure that an operations research effort is well directed and aimed at bettering the entire organization.

What benefits should a manager expect ? Operations research can be employed to mount a massive analysis, when warranted, of an important and intricate decision-making problem. As you will readily observe in applying operations research, the approach inherently requires adhering to systematic procedures and paying careful attention to details. (No other approach for solving complex management problems even comes close in demanding so much discipline in analysis.) The combined utilization of advanced mathematical techniques and enormous computing power permits a thorough exploration of relevant alternatives. A good operations research study will leave no doubt in an executive's mind that all reasonable courses of action have been investigated, and will make crystal clear the relative merits of specific alternative actions and their possible consequences.

A central ingredient to a sound operations research investigation is extensive sensitivity testing. Careful managerial scrutiny of comparative case studies provides the principal means by which an executive can confirm his understanding of the underlying model, its assumptions, and its data. Furthermore, the benefits a manager receives from a planning-oriented model stem largely from such insightful sensitivity testing. Rarely, if ever, does an

executive seek “numbers” as answers; rather, most decision-makers want a quantitative assessment of what risks are at stake with different actions, what changes in direction are likely to yield profit improvements, and what avenues are promising for further investigation (such as the development of new products, expansion into new markets, location of new plant sites, etc.). Often sensitivity testing reveals that the uncertainty of an allegedly critical factor is actually not very important. In making a good decision, whereas another factor, previously thought insignificant, is truly pivotal.

To determine whether an operations research project is meeting acceptable standards of quality, some easy items to check are ‘the ready availability of input data and model assumptions in a form understandable to nontechnicians, the summarization of results and the backup detail printed in the format of managerial reports, and reasonable turnaround times for running additional analyses having slightly modified input data or assumptions. The best way for a manager to make these checks is to ask questions and probe the answers. A competently designed model should provide a manager comprehensible answers to his spontaneous “why does. . .” and “what if. . .” questions without requiring a mammoth crash effort. (We hasten to add, however, that it is unfair to expect such rapid service at the initial stages of a study. A line executive should continue to ask questions throughout the duration of the project, and monitor whether the effort required to answer these questions eventually becomes routine and commensurate with the value of the analytic assistance provided.)

Another indicator of project quality is the extent to which the analysis results in a recommended strategy, as distinct from a suggested single decision. To illustrate, the output from a long-range capacity-expansion study should not be merely a string of recommended equipment purchases and forecasted production levels. Rather, the output should indicate the decisions to take immediately, should include recommendations of when to make the next set of decisions, given the present data, and should establish the circumstances for reviewing, and possibly revising, these future decisions. Even, the immediate decision recommendations should be qualified to the point of ascertaining what other alternatives are appropriate if the data are varied within a plausible range of values and any restrictive assumptions are relaxed.

What limitations should a manager recognize? This question was partially answered in Sec. 1.3, and you may want to review that material. Three more cautions are added here.

First, when an operations research model is used to reduce costs, the percentage savings may be relatively small. But if this percentage is applied to a large cost base, the absolute savings can pay for the operations research study many times over. Occasionally, a planning model will uncover a costly error in current operating procedures; in such an instance, the savings may be large. Most often, profit improvements stem from executives possessing a deeper understanding of the problem area, and hence developing a keener sense for taking correct actions and maintaining control in an uncertain and

competitive environment. It is impossible to assign a precise dollar improvement figure to this type of impact; nevertheless, the benefits are real and are perceived and valued by company management. In a preponderance of successful applications, the beneficial effects are truly manifest in the altered decision behavior of executives and managers at several levels of the corporation.

Second, although an operations research model often uses the mathematics of optimization, the resultant solution should not be viewed as necessarily yielding an optimal answer to the real problem. After all, as the text has stressed through-out, a model is inherently an approximation to reality, and therefore an optimal solution to this approximation need not be the “final” answer to the actual decision problem. The important issue, however, is not whether a proposed solution is optimal, but whether the solution yields a significant enough improvement over the alternatives to make it worthy of acceptance.

Third, while providing a solution to one set of problems, the operations research model may create, in turn, another set of problems. For example, the analysis may demonstrate the need for an improved information gathering system, or for a restructuring of operating policies. And, ensuring the continued maintenance of an up-to-date model does, itself, pose new managerial problems.

When should, a manager initiate an operations research project? It is helpful to distinguish between so-called one-shot or infrequent decision

problems and recurring decision analyses (like devising an annual plan, scheduling men and equipment, and replenishing inventories).

In special studies, a decision to apply operations research depends on the economic and strategic importance of the decision, the time span available for performing the analysis, and the relevance and availability of data. It is difficult and hazardous to apply operations research under “time pressure.” Consequently, a manager should consider employing the approach when the stakes are sizeable, the decision does not have to be made next Monday morning, data are available for the analysis, and the choice is not so governed by political and personality considerations within the company that economic analysis is of only minor import.

In planning situations, the decision to apply operations research also depends on the economic and strategic stakes of the problem and the available data. But planning applications differ from special studies most critically in the longer time horizon over which the model can be developed and tested. As we point out in later sections, controlling progress in the conduct of an operations research study is important; nevertheless, the corporation will not grind to a halt if a couple of weeks’ delay postpones the completion of an operations research planning model. (And there always is a couple of weeks’ delay!)

The decision to develop a computerized model for daily operations usually is more involved. Numerous companies have successfully constructed such models for as diverse applications as inventory control, tanker-fleet routing,

and job-shop scheduling. Often the economic benefits are small percentage wise, the systems design effort is staggering, and the implementation process is painful. Hence, this type of application is usually justified in terms of producing economic benefits that will extend over a relatively long term.

Sometimes executives misjudge whether the available data are sufficiently accurate as to warrant using an operations research approach. Applications of statistical techniques to the design of industrial research experiments, to the monitoring of continuous production processes and machinery, and to the auditing of voluminous accounting transactions demonstrate that mathematical techniques can be effective in analyzing sparse data that are subject to variability and measurement errors. Inaccurate or limited data do not per se negate the application of a mathematical technique. Even if there are no historical data at all, managers may be able to impart their experience-based knowledge by means of probabilistic statements. Hence, it is inappropriate for an executive to reject using operations research solely on the grounds of less than perfect factual information.

Sometimes executives shy away from operations research because they feel that their company personnel are not sophisticated enough to use mathematical techniques. This fear may be well founded, but the apprehension also may be based on a limited or even erroneous understanding of the degree of sophistication that is actually required. And all too often, senior managers underestimate the capability of their experienced personnel to learn how to apply operations research.

Many successful applications have been made by personnel who are trained in accounting, engineering, economics, or business, and who have been away from school for years. Their first-hand knowledge of the company more than compensates for their initial unfamiliarity with technicalities of operations research. In addition, the widespread availability of easy-to-use canned computer programs has removed much of the burden in going from a model formulation and actual data to a numerical solution and sensitivity analysis. And, finally, although the mathematical methods employed to obtain a numerical solution may be advanced, the solution itself may be easy to interpret and to implement. (A good example of this type of application is inventory control. The computations of a reorder point and replenishment quantity can sometimes be intricate; nevertheless, the resultant ordering policy may simply be of the form “when down to 4, order more” and thus may be easily understood.)

How can a manager get what he pays for? Perhaps the most difficult responsibility that an executive faces in controlling the progress of an operations research application is to strike the right balance between conducting the effort as a “research project” and as a “task-force assignment.”

Estimating how profitable or beneficial an application will be in a particular company is central to the research aspect. For example, many companies are able to reduce inventory investment by at least 25% by adopting scientific inventory control, but the level of reduction in a specific company can only be estimated after the operations research project is begun and

some trial tests are completed. Similarly, most medium-sized oil refineries are able to cut costs by \$1000 a day when using a linear programming model to make a weekly operating schedule, but an estimate of savings at any particular refinery can only be made after a preliminary model is built and run on a trial basis. Thus, an executive should view the initial phases of an operations research effort as exploratory.

It is erroneous, however, for management to view the entire project as research. Companies with the best record of implementing operations research plan each project from the very beginning as an effort to improve current procedures. The line managers who are involved share a sense of urgency about completing the effort and remain vigilant in keeping the study practical and pertinent to the actual decision problems.

Standard control techniques for managing include formulating a statement of goals, assigning task responsibilities, developing and updating a time schedule for completing various tasks, and planning for managerial reviews. It is the nature of operations research studies to encounter delays and unforeseen difficulties. Hence, expect that the unexpected will occur. The inevitability of these contingencies is the very reason why an operations research project needs careful managerial control.

Most operations research efforts require two to three man-years of effort and extend over a period of three to nine months. Naturally, if the project is important and complicated, these figures will be exceeded. The economic benefits of a well-conceived and controlled application, should far

outweigh the expense of developing and operating the system.

5.3 HOW TO SUCCESSFULLY CONDUCT AN OPERATIONS RESEARCH PROJECT

This section outlines the components of a successful operations research application and expands on several of the factors already discussed above; the context here is the conduct of a selected project.

Managerial guidance and participation. Both top management and operating management must recognize their respective roles in the evolution of a project. Since an operations research application typically cuts across different departments, the effort must have the sincere sponsorship of top management and the needed entrees into line activities. Furthermore, top management must watch that the corporation's best interests are held paramount and that the study is not diverted so as to serve the interests of individual groups at the expense of the company.

Operating management must actively participate in the project's goal formulation, administration; and evaluation. It is both difficult and foolish to impose an operations research system on an operating management that has not been a party to the system's design. Anyone with only a modicum of experience knows that the best of plans can be so cleverly sabotaged by a group of unwilling personnel [hat the promulgator looks like a fool. But more is at issue than just personality conflicts. When operating management has not been actively engaged in the study, there is substantial likelihood that the

proposed methods of the system will not be sufficiently comprehensive and flexible to handle the inevitable exigencies. Thus, if operating management has not participated in the evaluation (and, as a result, has little confidence in the worth of the endeavor), trouble looms ahead, even with the -most insistent encouragement of top management.

Project planning and control. The need for monitoring the progress of a project has been underscored. Now we highlight several factors in this process that are critical to success.

- The project team should realize at the outset where managerial judgment will be required. Specific plans should be made to obtain this counsel, and these provisions may in turn require some preparatory educational effort. People, not computers, make managerial decisions.
- The technical phase should be executed carefully, because if it is poorly done, the outcome can be disastrous. The team should recognize, however, that the mathematical side of the study will represent probably only a minor part of the total effort of developing and implementing the application.
- The data requirements should be ascertained early, and the information collection indicated soon enough to avoid long delays in the project. Often, this phase is poorly executed in an operations research study, even when the project is led by an experienced practitioner.

- Managers and operating personnel should be alerted to any transitional difficulties that may arise in testing and installing a new system. For example, when scientific inventory replenishment rules are implemented, total inventory investment usually rises for the first few months. (Can you explain why?) Top management is likely to express consternation unless properly fore-warned.
- The team should be careful to document the model's components and assumptions, and to record the input data and sources. In a large-scale effort, assumptions made several months earlier are easily forgotten. Furthermore, as test results and new data are examined, the model is inevitably altered. So it is essential that the team systematically catalogues each revision.

Credibility. Just like pregnancy, there is no such thing as a little credibility. Either an executive believes that the operations research representation of his problem is valid or he dismisses the results as worthless. The following paragraphs discuss how to develop a model that legitimately earns the trust of managers.

The project team should realize from the very beginning that the economic benefits of an operations research application never prove themselves and are never self-evident. And to make matters worse, a reliable “before and after” comparison is always extremely difficult to perform. There are two reasons why.

First, sufficient data about past operations may not be available, at least not in a form convenient for tabulation and analysis with acceptable accuracy. Hence, in its enthusiasm to design and implement a new approach the operations research team should not slight the job of installing a data-gathering system to reflect the true economic impact of a change. And when past data are insufficient, the team should start collecting current data long before it institutes new procedures. The team also should recognize the need to design a controlled experiment that focuses the effects to be evaluated. Unless the team heeds these cautions, it will, itself, be unable to prove factually that an improvement has occurred.

Second, only in exceptional circumstances can a team make a completely parallel comparison between two systems operating under different sets of procedures. There is no guarantee that an approach that looks attractive in terms of last year's operations will be just as attractive during this year's activities (or vice versa). Further, because managerial decisions at one point, in time may have a specific effect on business conditions later, it may be futile to attempt to show with great precision how anything but an actually operating system behaves over an extensive period of time.

Thus, it is hard to prove precisely how well an operations research approach, would have performed historically, or how much better an implemented operations research approach is faring as compared to what the previous system would have done. Management and the professionals must realize at the outset that they are limited in providing irrefutable evidence

that improvements actually result from an operations research approach. But it is important to remember that the same limitations exist in measuring the impact of any competing problem solution.

The above observations mean that by and large credibility ought to be established during the course of the project and not relegated to the end. Most executives express the following doubts about an operations research model: “How do I know that it uses the right data? . . . makes realistic assumptions? . . . computes the economic consequences correctly? . . . and encompasses the enormous number of relevant detailed considerations?” If you pause to think, it does stagger the imagination that the essence of a complex decision-making problem can be transferred to the “brain” of an inanimate electronic device. The following analogy may help to explain the psychology of establishing credibility and suggests some ways of allaying those doubts expressed by managers.

Suppose you are handed a telephone book for the first time and told that the volume contains the correct telephone numbers of everyone in the city. In a moment, you surely would realize that the claim is an overstatement. After all, telephones are installed and removed every day, so the telephone book is only an approximate representation of all the telephone numbers in the city. (In this sense, the volume of listings is a “model.”) What really concerns you is whether the approximation is worth using. How would you find out ?

Probably you would start by looking up a telephone number you already know (perhaps, your own). If you find that the listed number is correct, you then might select a person whose number you do not know, look up the number, and place a call to see if the book in fact gives the right listing. After several more tries of this kind, assuming you are successful each time, you would be willing to start using the book. And most likely you would continue to use it until you observed an increased frequency of wrong numbers. Then you would complain to the telephone company, or go back to relying upon the Information Operator.

Now consider the telephone company's objectives. It wants to provide a model or system that gives you the right numbers. There are many possible systems (or models, if you like) for providing this service. The telephone company has discovered that the most economical solution is to publish one book containing every listed number and to distribute the volume to you and all other subscribers. The company knows full well that you will use only a miniscule fraction of all the numbers; even so, you will judge the system's merit on the validity of this small fraction.

The preceding analogy is relevant to the design of an operations research system in several ways. An executive first tests the validity of an operations research model by asking questions about data and conclusions; he knows the right answers to some of the questions and has some intuition about others. His confidence builds if the forthcoming answers are straightforward, comprehensible, and correct. He will start to rely on the

model until his confidence is shaken by some “obvious” mistakes.

The operations research project team should try to anticipate what questions managers may ask and what data yield answers. This task is helped by discussing the detailed designs of the data reports and numerical summaries with the executives involved. The computer analysis should include not only summary reports similar to standard management information reports but also detailed backup analyses that clearly show the “how” and “why” of the summary figures. Much of [he output may be examined infrequently; but it is there “just in case.”

The telephone book analogy should not be pushed too far, because it is impractical and impossible to provide every number that an executive might possibly request. But novice operations researchers invariably make the mistake of providing far too little backup information, documentation, and analysis. As a consequence, they are frequently put in a position that is embarrassing to them and annoying to a manager, namely, having to go “back to the drawing board” to obtain the information that executives want in order to understand the model’s results.

The above discussion stresses the output requirements of a well-conducted operations research analysis. Of course, the team also must employ other means of effective communication. These are familiar to professional task-force leaders and amount to maintaining an open dialogue between the managers, who ultimately have to judge the merit of the results, and the team members. To repeat, managerial guidance and participation is a sine qua non

for establishing credibility.

Responsive and responsible implementation. Truly effecting change within a corporate organization, whether the change be installing a new computer system or reassigning managerial responsibilities, is usually a difficult job. Aside from any special aversion that personnel may have to computer-based systems analysis, there are few, if any, problems of implementing change that are peculiar to an operations research project. As is true for effecting most significant changes within a company, the support of top management is vital, adequate educational training of operating personnel is necessary, a carefully worked out plan for introducing the changes is essential, and the implementation process must be controlled and monitored to sense and then correct difficulties that may arise. Unfortunately, there is no substitute for experience in knowing how to implement change skillfully.

One problem does deserve additional discussion. The difficulty is reminiscent of that encountered in factory mechanization many years ago. Certain operations research applications, especially those involving daily operations, may drastically change the character of the decision-maker's job. For example, developing a computer scheduling model for the processing of orders in a factory, or the routing of ships between ports, or the purchasing of materials from vendors, may transform a job requiring long familiarity with the decision problem into one of routinely supplying raw data to a computer. An operations research approach may remove the fun, challenge, exercise of

judgment, sense of contribution, and mystique in a job. Rarely is top management willing to forego the resultant economic benefits for these reasons. But the project team must face up to the likely reaction of individuals who will be so affected. The team should recognize that the implementation process will arouse hostility; accordingly, they should provide post-implementation procedures to control a situation that might easily deteriorate because of a hostile environment.

Systems design. If the application is to be used again after the initial testing and analysis, then the ultimate success of the project depends upon the model's long-term viability. In the early years of commercial applications of operations research, many companies achieved noteworthy success for a while; later they discovered that their efforts had dissipated with the changing of business conditions and the promotion or resignation of operations research personnel. Now experienced firms realize the necessity for building systems support to main-tain and update a continuing operations research application.

This point would not merit special mention except for a commonly observed phenomenon that most executives still find paradoxical. The typical operations researcher, although having expertise in model building and analyzing complex problems, is usually ill equipped and frequently disinterested in the above-mentioned systems requirements. Consequently, experienced companies include systems-oriented personnel in an operations research project team to devise procedures for maintaining the model in good working

order.

5.4 HOW TO MANAGE AN OPERATIONS RESEARCH

STAFF

In keeping with the tenor of the chapter, this section highlights only a few issues that pertain directly to the profit-making impact of a corporate operations research activity.

Location and size. The proper placement of an operations research group within a large corporation is no longer a subject of much debate among professionals. No standard pattern has evolved, even within an industry, and these technical staffs have successfully operated under the guidance of controllers, chief planners, vice-presidents of manufacturing, as well as chiefs of research and development departments. Today, pragmatic considerations dominate the location decision. And divisionalized companies operating under a policy of decentralized management frequently have operations research activities at both the corporate and division levels.

The size of an operations research staff is an unreliable indicator of the group's productivity; a small staff of six talented professionals may have a much greater profit impact on a corporation than a group of 20 that contains only two or three top-notch scientists. In operations research, quantity is a very poor substitute for quality.

Corporate responsibilities. Top management expects the operations research staff to exercise a high degree of intellectual integrity. This means not only that the group must meet demanding professional standards, but also that the staff must seek truthful conclusions and refrain from organizational partisanship.

The operations research group manager must be careful not to overcommit his staff. In an effort to please, many groups undertake more projects than they can accomplish in a reasonable period of time. As a result, all the users become dissatisfied. An operations research group should have a systematic way to decide what projects to accept and how to allocate its own scarce professional resources to best serve the needs of the entire company.

Cooperation with users. The preceding sections emphasized the importance of working with line managers in the conduct of operations research projects. Here this subject is treated from the technical staff's point of view.

The group should always keep in sight the way an operations research model typically assists managers. In most applications, the model-building effort provides insights into the quantitative implications of specified data and assumptions. Ultimately, it is the managers who make the decisions and are held responsible for the outcomes. Hence these executives must assess the relative likelihood of various assumptions and weigh the risks associated with different courses of action. An operations researcher should avoid the

trap of believing that his model is true reality.

When a corporate operations research staff is first establishing its reputation, the group will work under less than ideal cooperative arrangements with its users.-The requesting organization may be pleased to see the project completed successfully, and even may pay for the project. Nevertheless, the user organization may not readily provide other necessary kinds of help, which often include the collection of data and careful managerial review of intermediate results; consequently, the effort may get bogged down waiting for essential assistance in the line organization. But when the operations research staff has progressed to where it can pick and choose from among several worthwhile projects, then a major selection criterion should be the willingness of the user organization to commit its personnel time to the project team. A good index of user interest and involvement is the extent to which it will allocate the lime of its people to assist in the application.

5.5 TECHNICAL & TECHNOLOGICAL ADVANCES

Operations research has advanced so rapidly that speculating about even the near future is risky. But several developments are now underway that certainly will have a major impact over the years immediately ahead. We somewhat arbitrarily separate these into technical and technological advances.

Technical progress. Two avenues of research are particularly important for new applications of operations research to executive decision

problems. The first is the development of efficient techniques for analyzing very large-scale problems, such as linear programming models containing thousands of constraints. The second is the development of practical methods for solving realistic combinatorial problems (as discussed in Chap. 13).

Contributions will continue to be made, of course, in the theory of operations research, including the development of improved nonlinear programming algorithms, multi-item inventory replenishment models, and statistical procedures for simulation experiments. Business applications including decision models in areas such as marketing, advertising, pricing, purchasing, personnel development, and finance also will receive more attention from management scientists. In addition, tremendous progress will be made in the application of operations research in other than profit-making enterprises.

Technological progress. Although the technical breakthroughs often occur independently of advances in computer hardware, the applications of the techniques to actual problems do depend significantly on the state of computer technology. (Linear programming models, for example, did not have widespread acceptance until comprehensive and easy-to-use computer programs were generally available. The same is true for simulation models.) The development of time-shared computing systems seems to be the most significant technological advance with regard to future operations research applications.

Time-sharing has already reduced the amount of effort required to

build and test a model; and this mode of computing has markedly cut the turnaround time between successive trial runs. But the more important impact is the stimulus to building new types of models and attacking different types of managerial problems. Corporations in the avant-garde are constructing very detailed financial planning models and are continually running economic analyses of new action alternatives as they arise in managerial deliberations. Time-sharing now allows the decision-maker to perform sensitivity analysis on-line. Being able to execute hundreds of thousands of calculations in a matter of seconds and getting the results immediately is an impressive and important step forward in the science of executive decision making.

An inevitable future impact of time-sharing will be the development of operating systems for decision problems that relate to the scheduling and daily allocating of scarce resources, and to the supervision of production processes. The rate of progress and acceptance in these areas depends on the growth in the general availability of time-shared computing equipment and software systems that are appropriate to the needs of commercial installations.

Administrative science. The time has past when an operations researcher can build a mathematical model and remain impervious to the behavioral characteristics of the individuals affected and the organizational milieu. Visionaries among operations research professionals are fully aware that new developments such as those described above exert tremendous

strains on the managerial fabric of a corporate organization. To enhance the adoption of these technical and technological advances by industry and government, management and behavioral scientists together will have to find ways by which executives can deal effectively with computerized systems as beneficial change agents.

CHAPTER 6

TEXTILES

6.1 A Spin Plan for Maximum Profit

A project was undertaken in a cotton-spinning mill processing the counts 2fls, 30s, 31s, 40s and 60s. It was desired to determine the quantity to be spun in each count subject to the availability of resources so as to obtain maximum profit.

It was not possible to allot all the frames to that count which gave the maximum profit per frame shift, since capacity of the machinery at back process, the availability of cotton reserve, and difficulty in marketing the entire production of a single count might not permit such a step. Knowing the profit margin available in the various counts considered for production, the problem was to determine the counts to be spun and the quantity to be produced in each, so as to obtain the maximum possible profit subject to. of course, the restriction imposed by the capacity of machinery available to various departments and cotton reserve on hand for each count.

Further there were certain number of doubling frames in the mills, with the help of these frames, two single threads of any particular count, could be twisted together (doubled) and sold in the market. The profit margin for doubled yarn was again different from that if the yarn was sold as single yarn. Taking into consideration the margin in doubled yarn and the capacity available in doubling, it was found necessary to work out how much quantity in each count should be doubled, so that the resultant profit would be maximum.

There are 85 spinning frames including 27 new Texmaco (N.T.) frames. From the past performance, it has been found that production N.T. frames was invariably higher by 10% as compared with the other frames of different make. For simplifying the problem slightly, it was better to take the machinery capacity in the spinning department in terms of certain standard units. Accordingly the 27 N.T. frames had been taken to be equivalent to 30 standard units, i.e., 27x1.1, thus giving weight age for the increased production of these frames. The total number of frames available in the mill could, therefore, be taken as 88 in terms of standard units. A similar procedure could be adopted to standardise the units in the other sections also.

In the sections Drawing, Fly Frames. Spinning and Doubling, the capacity was limited and there would be difficulty in feeding material at the subsequent processes in the coarse counts such as 20s. The availability of machines in these Departments should therefore be taken as a restriction in finding a solution to the problem for maximising the profit.

For standard units, production figures should be known in each of these sections to relate the requirements with the availability of machinery. From the records of the mills, the production figures in each section were analysed and standards of production of the standard units assumed were evolved for each count.

Since planning period was one month, the cotton stock for different counts both on hand and the expected arrival in the course of the month-had been taken into account. Again for the purpose of relating the figures of

availability and requirements of the raw materials, we need to know the quantities of materials required at each stage to produce one kg of yarn. For this purpose, waste figures in each count at various departments were analysed and the required ratios were estimated.

The cost of yarn consisted of three parts-raw material cost, processing cost and overhead cost. Since the overall overhead cost for the mill is constant for each month, the margin between the sale rate and the cost excluding the overhead was considered for maximisation instead of the direct profit figures.

For the present project the latest sale rates as obtained, from the management and the average cost figures for the previous month had been considered.

The Linear Programming formulation of the problem was considered for solution. The problem consisted of determining the quantity to be manufactured in single yarn and double yarn in each count, in such a way that the actual requirements as regards the machinery in each department and raw material did not exceed the availability and at the same time the resultant profit would be maximum.

The solution to the problem is arrived at by employing simplex technique.

The management of the mill was really impressed by the solution given. The solution offered could give additional profit of about Rs. 32,500 per month over the programme they had in force. Naturally the management was eager to implement the results and reallocate the frames to different counts in the

manner suggested by the solution obtained. However, they wanted to modify the programme so as to take into consideration the additional quantity of cotton they had purchased in 30s and 40a since the start of working out the problem. Taking them into account the problem was reformulated and solved. This modified programme would enable the mill to earn a profit of more than Rs. 60,000/- per month over the profit under the spinning plan followed by the mill earlier. The management had already taken steps to reschedule the allotment of frames to the various counts in accordance with the solution.

The solution arrived at cannot hold good for ever. It is to be modified every time there are changes in selling prices or cost of production in the various counts. However, due to practical considerations, since it is not possible to go on changing the plan every now and then, it is suggested that the plan be revised once in a month, taking into account, all possible changes that might have occurred in the factors affecting the solution.

6.2 Allotment of Drums to Winders in Cone-Winding

In a textile mill, production efficiency in the Cone-winding department was found to be around 60% in the case of 20s count. The winding capacity in the mill was a limiting factor and the management wanted that the production efficiency in this department must be stepped up. Observations made on the workers in the winding department revealed that they were tightly engaged all the time and operator idleness was not the reason for the lower efficiency

obtained in the section. This gave an indication that the workload might have to be reviewed with a view to improve the machine efficiency.

It is common experience that when the allotment of drums to winders is high, the machine efficiency is low and the operator efficiency is high; on the other hand, when the number of drums allotted to winders is low, the machine

efficiency may be better but the operator efficiency will be reduced. It is, therefore, necessary to fix the work assignment at a level which will minimise the total of the two cost components. The assignment for which the cost is minimum or the profit is maximum is known as optimum workload.

For any given number of drums allotted to a winder, the machine efficiency will depend upon:

- (1) average number of breaks per 10G cop changes ;
- (2) average number of cop changes per cone;
- (3) average time taken to attend to end breaks ;
- (4) average time taken for a cop change;
- (5) drum speed;
- (6) average time taken for doffing;
- (7) average yarn content per cop.

Extensive data were collected to obtain the values for the factors listed as above. Using finite queuing theory the machine utilisations for different allotments of drums were calculated. From the knowledge of drum speed

and the corresponding theoretical production, the expected production for various allotments of drums was arrived at. From cost considerations, the marginal profit per machine shift in winding was worked out with respect to each allotment of drums. The results are summarised in the following table.

**MACHINE UTILISATION AND MARGINAL PROFIT FOR
DIFFERENT ALLOTMENT OF DRUMS**

Number of drums	Machine utilisation	Expected production per spindle (in lbs)	Marginal profit per machine shift (Rs.)
18	61.6	8.53	93.34
17	65.0	9.00	99.75
16	68.5	9.48	105.01
15	72.0	9.96	109.09
14	75.4	10.44	113.67
13	78.4	10.85	116.92
12	81.0	11.21	118.87
11	83.2	11.51	118.81
10	85.0	11.76	117.46

It is obvious from the above table that the margin of profit per machine shift is maximum corresponding to an allotment of 12 drums per operator. This allotment increases the marginal profit per machine shift by

Rs. 19.12 over the existing allocation of 17 drums per operator. Based on the findings, the management immediately reduced the workload to 12 drums per operator. Consequently, the machine utilisation increased to over 80% and the net increase in the marginal profit per annum was found to be Rs. 30,000 for the two machines working on this count.

6.3 Optimum Work Assignment to Siders in Spinning

Proper fixation of workloads to the workers in the various departments of a textile mill offers good scope for cost reduction. In any spinning mill, the siders in the spinning section form a major proportion of the total number of workers employed and determination of the work assignment for this category of workers is very important from the cost, point of view. Generally, a group of spindles will be assigned to a sider and he will be required to attend to end breaks, creel breaks and creel changes occurring within that group.

Arbitrary fixation of the number of spindles in any such group may lead to an increase in cost. For, when the number of spindles allotted to a sider is large, the chance of finding more number of ends un pieced at any time will be greater and this will result in not only loss of production in spinning, but also increased waste. Consequently the component of cost due to loss of production and additional waste may go up. If, however, the allotment of spindles per sider is small, the number of siders required to

cover all the ring frames in the shed will be more and naturally the labour wages component of cost will increase. Hence the problem is one of finding the stage (optimum workload) at which the total cost of labour wages and the loss due to idle spindle and bonda waste is minimum.

A study was taken up in a spinning mill for arriving at an optimum workload for the siders in spinning for 60s count. The number of spindles to be allotted will naturally depend upon the following factors.

- (1) end breakage rate in ring frame;
- (2) rate of creel breakages;
- (3) frequency of roving bobbin changes;
- (4) time taken for piecing an end break;
- (5) time taken for piecing a creel break
- (6) time taken for making r. bobbin change;
- (7) time taken for patrolling and cleaning the sider.

Extensive data were collected to determine the values for the above factors with a reasonable degree of accuracy.

from the consideration of the statistical distribution of the end breakage and that of creel breakages in general, a model had been developed using queueing theory, to estimate the proportion of time that a an end remains un pieced and also the proportion of time that a spindle is rendered idle due to creel breaks, corresponding to any particular spindle allotment.

From an analysis of cost figures available with the mills, the loss per shift on account of end remaining un pieced was estimated as 30 paise.

Similarly the cost associated with an idle spindle was worked out as 15 paise per shift. The total wages paid to a sider including the fringe benefits like P.F., E.S.I., leave with wage, etc. accounted to Rs. 13.50 per day.

knowing the the proportion of time that an end remains un pieced and a spindle is idle for a given allotment of spindles to a sider, the sum of the siders' wages and the loss due to bonds waste and idle spindles per frame per shift was evaluated. The results are presented ir the table.

Number of spindles allotted	estimated proportion of time thai an end remains unpieced	expected proportion of idle time due to creel break	total cost component/ frame shift (Rs.)
600	0.70%	0.10%	0.90
700	0.92%	0.14%	8.90
800	1.23%	0.18%	8.34
900	1.64%	0.25%	8.14
1000	2,23%	0.34%	8.28
1100	3.14%	0.50%	8.97
1200	3.14%	0.50%	8.97
1300	7.30%	1.26%	13.66

6.4 Determination of Optimal Number of Tape Stitchers in a Spinning Shed

The mill had an installed capacity of 39,960 spindles (in 94 spinning frames). A system of snap survey in the spinning shed every shift was in vogue in the mills. An observer goes round the spinning section noting down the idle spindles in each frame and the causes for them. Among the causes for idle spindles were the apron cut, feed cut, tape cut, etc. Efforts were being made to eliminate the idle spindles due to each of these causes. Having tackled the other causes, the present study was directed towards investigation of the problem of tape cut and the optimum number of tape stitchers required per shift. There were 3 tape stitchers engaged in a shift which was considered high.

Snap survey results of the past six months revealed that on an average 40 spindles were idle at any time due to the tape cut. This incidence of tape cut was considered very high in view of the fact that three tape stitches were engaged in every shift. It was felt that an elaborate study would bring out all the relevant factors. The factors studied included the following:

- (i) The average number of tape cuts occurring in a shift;
- (ii) average time taken by tape stitchers for replacing and stitching the tapes.

The actual occurrence of tape cut was studied for a fortnight in the department. It has been observed that on an average 48 tapes are cut per shift in spinning section.

As soon as the tape cut is noticed, the tape stitcher walks to the frame and then to the exact place of tape cut with the stitching machine. The required lengths of tapes are cut and kept ready. The stitcher inserts the tape through the swindles, takes it round the drum, joins the ends, stitches the ends and fixes the tape on the spindles. The stitcher has to go round the spinning shed searching for the tape cuts and mending them wherever observed.

The stitching time, varied from 2.5 to 5.5 minutes and the bulk of the stitching time was centered at 3 minutes.

One important aspect to be considered in determining the optimum number of stitchers is the interference delay. The occurrence of tape cuts is purely random. In that case assuming only one tape stitcher is there, when one tape is being stitched, another tape has to wait for stitching till the tape stitcher completes his job on the tape he is attending. One tape idle for a minute results in a loss of 4 spindle minutes. The optimum number of tape stitchers has to be decided based on the cost due to idleness of spindles and the wages to the stitchers.

From the actual observation of the number of tape cuts and time of stitching during- the shift, instead of estimating the interference delay based on any rigid statistical distribution, the actual distribution has been simulated with the help of random number tables.

Since the frequency distribution of the servicing time is also known from the collected data, the stitching time in each case also has been simulated. The simulated time of occurrence of tape cuts and stitching time were recorded

at different periods of the day.

The interference delay when the number of tape stitchers is varied from 1 to 3 is shown in the table below:

INTERFERENCE DELAY FOR VARIOUS ALLOTMENTS OF STITCHERS

No. Of tape stitchers	No. of tape cuts per shift	Interference minutes	Delay spindle minutes
1	48	36	144
2	43	1	4
3	48	0	0

In addition to considering the interference due to employment of 1,2 or 3 tape stitchers, we should also consider the cost associated with each and arrive at an optimum number of tape stitchers required based on the minimum total cost.

The cost components associated with the breakage and stitching of s are as follows:

- (i) Cost of idle spindle due to interference delay;
- (ii) cost associated with the wage for the stitcher.

It has been estimated that idleness of one spindle per shift will cost Rs 0 25. Daily wage of each tape stitcher has been estimated as Rs. 15 The cost of idle spindle per day and total cost taking into consideration the wages of the stitchers and other costs is presented in the table boew:

TOTAL COST FOR DIFFERENT ALLOCATIONS

No. Of. Ape stitchers per day	Wage of the tape stitchers per day (Rs.)	Cost of idle Idle spindle minutes per day	Spindle Cost (Rs.)	Total cost per day (Rs.)
3 (1 per shift)	45.00	432	0.25	45.25
6 (2 per shift)	90.00	12	0.01	90.01
9 (3 Per shift existing)	135.00	0	0	135.00

It is seen that the total cost is minimum when one tape stitcher is employed. By increasing the tape stitchers from one to three, loss of Rs. 89.75 per day, i.e., Rs. 31,405,50 will be incurred per annum. Hence one tape stitcher is considered to be optimum for stitching tapes.

From the present set up, it is seen that even after employing three tape stitchers, the number of idle spindles due to tape cut per sbifi came to 40 which shows clearly that either the tape stitchers do not properly carry out their duties or they are not being utilized properly. This maybe due to the fact that

- (i) the tape stitchers are not going round the frames observing the tape cuts and stitching them immediately, and
- (ii) the tape stitchers are engaged in other activities. In order to enable the stitcher to stitch the tapes without delay as soon as tape cut is observed, it is

suggested that an indicator may be fixed on each frame.

The arrow of the indicator should be raised by the siders attending on that frame. This will enable the stitcher to stitch the tapes as soon as it is cut.

If a stronger thread is used for stitching, it may last longer.

Due to friction, wear is caused and tape gels worn out throughout the length and breaks. Hence spindles should be properly checked and alignment of spindles should be adjusted.

All these recommendations were immediately put to practice by the management.

REFERENCES

Kirby, M. W. (Operational Research Society (Great Britain)). Operational Research in War and Peace: The British Experience from the 1930s to 1970, Imperial College Press, 2003.

Nelson, Rosser T, "Labor and Machine Limited Production Systems," Management Science, Vol. 13 (1967), pp. 648-671.

Shubik, Martin, "Bibliography on Simulation, Gaming Artificial Intelligence and Allied Topics," Journal of the American Statistical Association, Vol. 55 (1960), pp. 736-751.

Teichroew, Daniel and John F. Lubin, "Computer simulation- Discussion of the Technique and Comparison of Languages," Communications of the ACM, Vol. 9 (1966), pp. 723-741.

Dantzig, G. B. Linear Programming and Extensions, Princeton University Press, Princeton, N.J., 1963.

Dreyfus, S.E., "An Appraisal of Some Shortest-Path Algorithms," Operations Research, 17(3), (1969).

Ford, L. R. and D. R. Fulkerson, Flows in Networks, Princeton University Press, Princeton, N.J., 1962.

Gale, D., The Theory of Linear Economic Models, McGraw-Hill New York, 1960.

Gass, S. I., Linear Programming, Fourth Ed., McGraw-Hill, New York, 1975.

Geoffrion, A. M., ed., Perspectives on Optimization: A Collection of Expository Articles, Addison-Wesley, Reading, Mass., 1972.

Goyal, S. K., "Improving VAM for Unbalanced Transportation Problems," J. Operational Res. Society, 35(12):1113-1114(1984).

Hillier, F. S., and Gerald J. Lieberman, Introduction to Operations Research, Third Edition, Holden-Day, San Francisco, Calif., 1980.

Hu, T. C., Integer Programming and Network Flows, Addition-Wesley, Reading, Mass., 1969.

Ravindran, A., "On Compact Book Storage in Libraries," Opsearch, 8(4),(1971).

Taha, H. A., Operations Research An Introduction, 2nd Ed., Macmillan, New York, 1976.

Wagner, H. M., Principles of Operations Research, 2nd Ed., Prentice-Hall, Englewood Cliffs, N. J., 1975.

Wiest, J. D., and F. K. Levy, A Management Guide to PERT/CPM, 2nd Ed., Prentice-Hall, Englewood Cliffs, N. J., 1977.

Conway, R. W., W. L. Maxwell, and L. W. Miller, Theory of Scheduling, Addison-Wesley, Reading, Mass., 1967.

Cooper, R. B., Introduction to Queueing Theory, 2nd ed., Elsevier North-Holland, New York, 1981.

Cox, D. R., and W. L. Smith, Queues, Wiley, New York, 1961.

Giffin, W. C., Queueing: Basic Theory and Applications, Grid, Columbus, Ohio, 1978.

Gross, D., and C. M. Hariss, Fundamentals of Queueing Theory, Wiley, New York, 1974.

Khintchine, A. Y., *Mathematical Methods in the Theory of Queueing*, Hafner Publishing Co., New York, 1960.

Kleinrock, L., *Queueing Systems, Vol. I*, Wiley, New York, 1975.

Kleinrock, L., *Queueing Systems, Vol. II*, Wiley, New York, 1976.

Lee, A. M., *Applied Queueing Theory*, St. Martin's Press, New York, 1966.

Newell, G. F., *Applications of Queueing Theory*, Chapman and Hall, London, 1971.

Saaty, T. L., *Elements of Queueing Theory*, McGraw-Hill, New York, 1961.

Stidham, S., Jr., "A Last Word on $L = W$," *Operations Research*, 22(1974), 417-421.

Takacs, L., "On Erlang's Formula," *Am. Math. Stat.*, 40(1969), 71-78.

White, J. A., J. W. Schmidt, and G. K. Bennett, *Analysis of Queueing Systems*, Academic Press, New York, 1975.

Banks, J. and J. S. Carson III, *Discrete Event System Simulation*, Prentice-Hall, Englewood Cliffs, N. J. 1984.

Bratley, P., B. L. Fox, and L. E. Schrage, *A Guide to Simulation*, Springer-Verlag, New York, 1984.

Fishman, G. S., *Principles of Discrete Event Simulation*, Wiley, New York, 1978.

Law, A. M., and W. D. Kelton, *Simulation Modeling and Analysis*, McGraw-Hill, New York, 1982.

Pegden, C., *SIMAN*, Systems Modeling Corporation, P. O. Box 10074, State College, Pa., 1984.

Pritsker, A. A. B., Introduction to Simulation and SLAM II, 2nd ed., Halstead Press, New York, 1984.

Taha, H.A., Operations Research-An Introduction, Prentice-Hall Inc., New Jersey, 1997.

Ackoff, Russell L., Progress in Operations Research, Vol. I, John Wiley & Sons, Inc., 1961.

Ackoff, Russell L. and Patrik Rivett, A Manager's Guide to Operations Research, John Wiley & sons, Inc., 1963.

Ackoff, Russell L. and Maurice W. Sasieni, Fundamentals of Operations Research, John Wiley & Sons, Inc., 1968

Baumol, W. J., Economic Theory and Operations Analysis, Prentice-Hall, 1963.

Boot, J. C. G., Mathematical Reasoning in Economics and Management Science: Twelve Topics, Prentice-Hall, 1967.

Bowman, Edward H. and Robert B. Fetter, Analysis for Production and Operations Management, R. D. Irwin, 1967.

Carr, Charles R. and Chaeles W. Howe, Quatitative Decision Procedure in Management and Economics-Deterministic Theory and Applications, McGraw-Hill, 1964.

Churchman, C. West, The System Approach, Delacorte Press, 1968.

Churchman, C. West, Russell L. Ackoff, and E. Leonard Arnoff, Introduction to Operations Research, John Wiley & Sons, 1954.

Fishburn, Peter C., Decision and Value Theory, John Wiley & Sons, 1964.

Hertz, David B. and Roger T. Eddison, Progress in Operations Research, Vol. II, John Wiley & Sons, 1964.

Hertz, David B., New Power for Management-Computer Systems and Management Science, McGraw-Hill, 1969.

Hillier, Frederick S. and Gerald J. Lieberman, Introduction to Operations Research, Holden-Day, Inc., 1967.

Horowitz, Ira, An Introduction to Quantitative Business Analysis, McGraw-Hill, 1965.

Miller, David W. and Martin K. Starr, Executive Decisions and Operations Research, Prentice-Hill, 1960.

Teichroew, Daniel, An Introduction to Management Science Deterministic Models, John Wiley & Sons, 1964.

Theil, Henri, Optimal Decision Rules for Government and Industry, North Holland, 1964.

Cyert, Richard M. and James G. March, A Behavioral Theory of the Firm, Prentice-Hall, 1963.

Greenlaw, Paul S., Lowell W. Herron, and Richard H. Rawdon, Business Simulation in Industrial and University Education, Prentice-Hall, 1962.

Hammersley, J. M. and D. C. Handscomb, Monte Carlo Methods, John Wiley & Sons, 1964.

Markowitz, H. M., Bernard Hausner, and H. W. Karr, SIMSCRIPT: A Simulation Programming Language, Prentice-Hall, 1962.

Meier, Robert C., William T. Newell, and Harold L. Pazer, Simulation in Business and Economics, Prentice-Hall, 1969.

Mize, Joe H. and J. Grady Cox, Essentials of Simulation, Prentice Hall, 1968.

Naylor, Thomas H., Joseph L. Balintfy, Donald S. Burdick, and Kong Chu, Computer Simulation Techniques, John Wiley & Sons, 1966.

RAND Corporation, A Million Random Digits with 100,000 Normal Deviates,”
The Free Press, 1955.

Blake, K. and G. Gordon, “System Simulation with Digital Computers,” IBM
Systems Journal, Vol. 3 (1964)

Burr, Irving, W., “A Useful Approximation to the Normal Distribution
Function, with Application to Simulation,” Technometrics, Vol. 9 (1967)

Clark, Charles E., “Importance Sampling of Monte Carlo Analyses.”
Operations Research, Vol. 9 (1961)

Cohen, Kalman J. and Edwin J. Elton, “Inter-Temporal Portfolio Analysis
Based on Simulation of Joint Returns,” Management Science, Vol. 14 (1967),

Efron, R. and G. Gordon, “A General Purpose Digital Simulator and Examples
of its Application: Part I — Description of the Simulator,” IBM Systems
Journal, Vol. 3 (1944)

Ehrenfeld, S. and S. Ben-Tuvia, “The Efficiency of Statistical Simulation
Procedures,” Technometrics, Vol. 4 (1962)

Fetter, R. B. and J. D. Thompson, “The Simulation of Hospital Systems,”
Operations Research, Vol. 13(1 1965)

Fishman, George S., “Problems in the Statistical Analysis of Simulation
Experiments : The Comparison of Means and the Length of Sample Records,”
Communications of the ACM, Vol. 10 (1967)

Fishman, George S., “The Allocation of Computer Time in Comparing
Simulation Experiments,” Operations Research, Vol. 16 (1968)

Fishman, George S. and Philip J. Kiviat, "The Analysis of Simulation-Generated Time Series," *Management Science*, Vol. 13 (1967)

Gafarian, A. V. and C. J. Ancker, "Mean Value Estimation from Digital Com-puter Simulation," *Operations Research*, Vol. 14 (1966)

Ghare, P. M., "Multichannel Queuing System with Bulk Service," *Operations Research*, Vol. 16 (1968) [23]

Gordon, G., "A General Purpose Systems Simulator," *IBM Systems Journal*, (1962)

Greenberger, Martin, "Method in Randomness," *Communications of the ACM*, Vol. 8(1965)

Hertz, David B., "Risk Analysis in Capital Investment," *Harvard Business Review*, Vol.42 (1964)

Hertz, David B., "Investment Policies that Pay Off," *Harvard Business Review*, Vol. 46 (1968)

Hull, R. E. and A. R. Dobell, "Random Number Generators," *SIAM Review*, Vol. 4 (1962)

Kabak, Irwin W., "Stopping Rules for Queuing Simulations," *Operations Research*, Vol. 16 (1968)

Krasnow, Howard S. and Reino A. Merikallio, "The Past, Present, and Future of General Simulation Languages," *Management Science*, Vol. 11 (1964)

MacLaren, M. Donald and George Marsaglia, "Uniform Random Number Generators," *Journal of the Computing Machinery*, Vol. 12 (1965)

Markowitz, H. M., "Simulating with Simscript," *Management Science*, Vol. 12 (1966)

Maxwell, W. L. and M. Mehra, "Multiple-Factor Rules for Sequencing with Assembly Constraints," *Naval Research Logistics Quarterly*, Vol. 15 (1968)

McKenney, James L., "Simultaneous Processing of Jobs on an Electronk puter," *Management Science*, Vol. 8 (1962)

Nelson, Rosser T., "Queuing Network Experiments with Varying Arrival Service Processes," *Naval Research Logistics Quarterly*, Vol. 13 (1966)

Nelson, Rosser T., "Labor and Machine Limited Production Systems," *Management Science*, Vol. 13 (1967)

Nelson, Rosser T., "Dual-Resource Constrained Series Service System," *Research*, Vol. 16 (1968)

Scheuer, Ernest and David S. Stoller, "On the Generation of Normal Random Vectors," *Technometrics*, Vol. 4 (1962)

Shubik, Martin, "Bibliography on Simulation, Gaming Artificial Intelligence and Allied Topics," *Journal of the American Statistical Association*, Vol. 55 (1960)

Taft, Martin Israel and Arnold Reisman, "Toward Better Curricula Through Computer Selected Sequencing of Subject Matter," *Management Science*, Vol. 13(1967)

Teichroew, Daniel and John F. Lubin, "Computer Simulation—Discussion of the Technique and Comparison of Languages," *Communications of the ACM*, Vol. 9 (1966)

Tocher, K. D., "Review of Simulation Languages," *Operational Research Quarterly*, Vol. 16 (1965)

Westlake, W. J., "A Uniform Random Number Generator Based on bination of Two Congruential Generators," *J. of ACM*, Vol. 14 (1967)