Optimisation in Economic Analysis

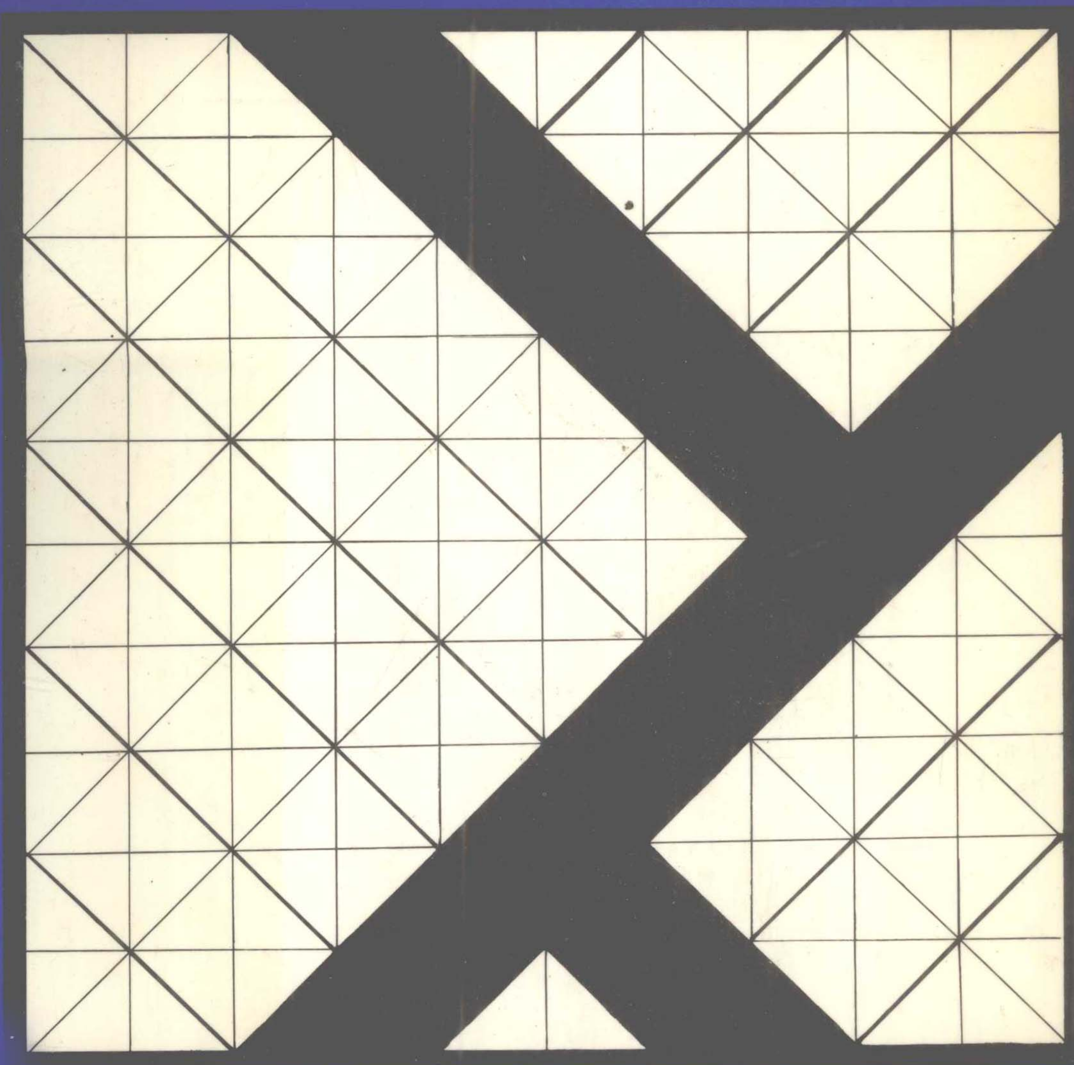

Gordon Mills

Optimisation in Economic Analysis

GORDON MILLS

Professor of Economics, University of Sydney

London
GEORGE ALLEN & UNWIN
Boston Sydney

George Allen & Unwin (Publishers) Ltd,
40 Museum Street, London WC1A 1LU, UK

George Allen & Unwin (Publishers) Ltd,
Park Lane, Hemel Hempstead, Herts HP2 4TE, UK

Allen & Unwin Inc.,
9 Winchester Terrace, Winchester, Mass 01890, USA

George Allen & Unwin Australia Pty Ltd,
8 Napier Street, North Sydney, NSW 2060, Australia

First published in 1984

British Library Cataloguing in Publication Data

Mills, Gordon
Optimisation in economic analysis.
1. Economics, Mathematical 2. Mathematical optimization
I. Title
330.01'515 HB135
ISBN 0–04–311001–0
ISBN 0–04–311002–9 Pbk

Library of Congress Cataloging in Publication Data

Mills, Gordon.
Optimisation in economic analysis.
Includes bibliographies and index.
1. Mathematical optimization. 2. Economics, Mathematical. I. Title.
HB143.7.M54 1984 330'.01'51 84–359
ISBN 0–04–311001–0
ISBN 0–04–311002–9 (pbk.)

Set in 10 on 12 point Times by Mathematical Composition Setters Ltd., 7 Ivy St. Salisbury, Wilts.
and printed in Great Britain by
Mackays of Chatham

Optimisation in Economic Analysis

Preface

This book is well travelled. My teaching of optimisation techniques began at the University of Bristol in 1962, when Ronald Tress (then head of the Department of Economics) and my other colleagues there were so kind (rash?) as to agree to my proposal for an undergraduate course in optimisation. In consequence, a first draft of Chapters 1–4 was written in 1970.

In 1971–2, my appointment as visiting professor in the Department of Economics of the University of Virginia resulted in my teaching a mathematical economics course in which I included non-linear optimisation, with some emphasis on the Kuhn–Tucker conditions; this yielded a first draft of Chapter 7. From 1977 onwards, a course assignment in the Department of Economics of the University of Sydney renewed my interest in the teaching of optimisation, and eventually led to rough drafts of most of the remaining chapters.

The first draft of the complete text was prepared in Sydney in 1980 and 1981, with the able help of Swee Kuen Pan, who has outstanding skill in the typing of mathematical work.

Much of the final revision was done during the later part of the academic year 1982–3, while I was a visiting scholar at the University of Virginia, on sabbatical leave from the University of Sydney. The work was completed during the (northern) summer of 1983, while I was a visitor at the London Business School. The generous and friendly nature of my two hosts is greatly appreciated.

Most of the final draft was typed by Swee Kuen Pan. Some later chapters were typed by Sarah Harrod at the London Business School.

The publisher's (anonymous) readers provided a helpful blend of encouraging praise and trenchant criticism; and Nicholas Brealey of Allen & Unwin insisted on careful and extensive revision.

My wife Pauline suffered the usual pains associated with having a husband working on a book; she also proof read the typescript, prepared the index and claims to have improved the grammatical quality of the text.

To all these people go my thanks. Of course, I take sole responsibility for the book's remaining imperfections.

Sydney
September 1983 GORDON MILLS

Contents

Introduction

THE ROLES FOR OPTIMISATION IN ECONOMIC ANALYSIS

Because the fundamental economic problem is one of making the best use of limited resources, mathematical optimisation methods have an important position in the economist's tool-kit. Indeed, various different kinds of economic study employ models of optimising behaviour, and hence use such methods.

One major field in which optimisation methods are widely used is neoclassical economic theory, which perceives all economic agents as optimisers. Or, to be more specific, the theory postulates that an economic agent chooses values for decision variables (e.g. consumption quantities, prices charged by manufacturers) so as to optimise the value taken by the target or objective function (e.g. profit, utility). The choice of values for the decision variables is, of course, circumscribed by various limitations inherent in the situation being modelled (e.g. the restricted availability of some resources), and these limitations are expressed as constraints in the mathematical model. The optimising analysis is then used to explain or predict behaviour. In recent decades, economists have attempted to come to grips more effectively with the complexities of the real world by building more detailed models, and this in turn has led to the deployment (and, on occasion, even the invention and development) of more powerful mathematical methods.

Of course, the applicability of optimisation techniques is not restricted to the particular view of the world that is implicit in such neoclassical theory. This proposition is most amply demonstrated by the use of these techniques in economic planning. Here the aim of the study is explicitly *prescriptive:* given the target of the economic agent, what investments should be made, what output levels should be set, what purchases should be made, so as to meet this target to the greatest degree possible? The use of such planning techniques transcends geographical and political boundaries, as may be seen readily by considering some examples.

The transnational oil company pursues its capitalistic goals with the help of planning tools such as mathematical programming procedures which help it to decide (*inter alia*) what pattern of purchases of crude oils of differing specifications will enable it to meet stated market demands at minimum cost. Equally, similar optimisation techniques are used (albeit with different goals) at a variety of levels in the planning bureaucracies of the centrally planned economies of Eastern Europe.

The universality of interest in the use of optimisation techniques to guide planning decisions may be illustrated also from the history of what has come to be called 'linear programming', one of the most widely used planning techniques. When consulted in 1938 about a practical problem concerning the most economical utilisation of resources in a local enterprise a Russian mathematician, L. V. Kantorovich, came to realise that there was an entire class of such problems that could be formulated in terms of a linear optimising model; he also realised that the development of an effective mathematical procedure to find the optimal solution to such a problem would be a far from trivial task. In 1947, an American mathematician, G. B. Dantzig, independently came to the same understanding (while considering some planning problems occurring in large military enterprises), and discovered the first effective general-purpose mathematical procedure for computing optimal solutions to such problems.

Since then, development of linear programming and related mathematical optimisation techniques has gone on apace in these two and many other countries. Some of the purely mathematical research, and most of the development of skills in the practical use of these planning techniques, has taken place in planning agencies and in the planning sections of productive enterprises, government departments and so forth. This work has contributed to the growth of an entire profession of people who specialise in model-building for planning decisions, who practise what is now often called management science or operational research, and for whom mathematical optimisation techniques are an important tool.

Much the same set of mathematical optimisation methods is used in the two areas of economic theory and applied planning (prescriptive) analysis. However the style of use differs. Traditionally the economic theorist has been concerned primarily with so-called *comparative statics* analysis: this postulates a change in some exogenous factor (e.g. an increase in taxation, a discovery of a new oil field, an increase in pay in an industry) and studies the behaviour of the relevant economic agents before and after the change; each of the 'before' and 'after' situations is analysed in a static way, i.e. as if it had prevailed for a long time and all agents had adjusted fully to it; the analysis is not concerned with the process of transition, in which the agents move from the old to the new situation. Thus the study requires only a double application of a static, single-period model; and hence the analysis may not be at all complex.

In contrast, a proportion of the more recent work in economic theory deals with *dynamic* situations, in which a principal focus is the way in which economic agents adjust their behaviour over time. These studies address questions such as: what is the optimal consumption pattern for an individual over a lifetime, given that the individual's income rate first increases, and then (later in life) decreases? and how and at what rates will society use up stocks of an exhaustible natural resource? In such dynamic analysis, the distinctive feature is not merely that time passes; rather it is that some underlying features necessarily change with the passage of time. For example, the individual grows older; the resource stock is diminished. And in such analyses, the transitions matter: the rates of change (per unit of time) of the various economic magnitudes are an essential feature of the situation being studied.

For such studies in *economic dynamics,* the optimising models are necessarily more complex. Because the values of the economic variables can change over time, all variables must be identified by time – i.e. for each variable, a distinction must be made between the value at time t and the value at each other time. If time is divided into discrete periods (as is often the practice), then a dynamic model necessarily deals with a significant number of such periods, and this makes the model larger than that for a (corresponding) analysis in comparative statics, where the model represents only a single time-period.

Most prescriptive studies for policy decisions also require dynamic formulations, and hence here too multi-period models are commonly used. In many such contexts, the stylised representations often used in economic theory are inappropriate; instead, the specific history of the situation must be included in the model. Where the entity is not newly begun (as, for example, where the model represents production opportunities open to a long-established firm), it is usually necessary to include the history of past investment decisions, at least where these lead to a heterogeneous collection of productive capacities. Such studies are termed exercises in *historical dynamics.* Because of the need to include historical detail, the models tend to be even more complex than those used in dynamic studies in economic theory.

A further reason for the additional complexity of applied prescriptive models is the empirical need to define commodities very narrowly, a practice that can lead to very large dimensionality. Models with thousands of commodity variables are by no means uncommon, and the number of resource and other constraints can also be very considerable. In much economic theory, on the other hand, it is often sufficient to postulate only two commodities; where the model contains more than two, the stylisation may permit all or many products to be regarded as having similar general properties, and hence not requiring individual treatment.

Thus developments in the use of optimisation methods in economic analysis (a term used here to embrace both theory and prescriptive policy studies) have depended not only on the research of mathematicians (and economists, and others) who have developed the new mathematical techniques, but also on the astounding growth in computational power afforded by the advances in micro-electronics, and needed to handle the great dimensionality required in many models.

Nevertheless the dazzling results stemming from these conceptual and engineering advances must not be allowed to hide the fact that the building of the mathematical model is still an art that requires good judgement and sense. No matter how great the intellectual prowess in developing optimising procedures, and no matter how striking our ability to make enormous optimising calculations in remarkably small time-spans, if the model is a foolish, or merely a poor, representation of the problem in hand, then the (nominally) 'optimal' results will not bring benefits; indeed, if they are misleading, they can be positively dangerous in the hands of a susceptible and enthusiastic innocent.

Although there can be significant problems in modelling the technology and in ascertaining the resource or other constraints, the greatest difficulty is often found in establishing an acceptable maximand (or minimand, as the case may be). Does the firm really want to go the whole way in maximising profits? If the individual is to be regarded as maximising utility, how is utility to be conceptualised and measured? It seems that it is doubts on this aspect, above all, that lead some to criticise and reject formal optimisation in economic analysis.

Although such wholesale rejection is usually no more than an act of nihilism, those who favour and practise formal optimisation must never forget that model-formulation is usually the weakest link in the whole business. In empirical prescriptive work (though not in economic theory), some consolation is to be had from the fact that, with the great computational power now at our disposal, it is often possible to use a maximand (which is admitted to be an unsatisfactory formulation) simply to lead the calculations to an initial solution, which (it is hoped) is at least 'sensible', and perhaps even 'good'. Further calculations can then be used to explore other good solutions in the neighbourhood of the first. The choice among these alternative policies is then made directly by the decision-makers, who first use the formal optimising calculations to guide them to the more promising outcomes in the realm of the possible, and then use their judgements about what is preferable to select among alternatives, even when those preferences can not be made sufficiently explicit for formalisation in a maximand.

THE ORGANISATION OF THE BOOK

The style of the book is informal, and I have attempted to present arguments in a detailed step-by-step manner, often with intuitive explanations, in order to facilitate understanding by readers who do not consider themselves specialists in mathematical work. In general, computational procedures are not considered; mathematical proofs are rarely given, but I do try to convince the reader of the plausibility of the results. At the same time, such informality is no excuse for lack of rigour. My aim is to state theorems with precision, and to demonstrate the great care that is needed for safe and successful application of the mathematical tools.

Although the greater part of the book is devoted to explanation of mathematical tools, there are some parts dealing with the art of model formulation; and, throughout the book, there is discussion of economic uses and interpretations.

The first four chapters deal with linear models, that is to say models where the maximand is a linear function of the decision variables, and where the constraints are also linear in those variables. Chapter 1 introduces in some detail the components of the linear model; the discussion also serves to illustrate the elements of the general problem in constrained optimisation. Chapter 2 identifies the nature of the optimal solution in a linear model and gives a very brief introduction to linear programming, i.e. to the basic ideas underlying the computational procedures used to calculate the optimum in a linear model. This mathematical analysis is continued in Chapter 3, which defines and interprets the so-called 'duality' properties of the optimum solution, properties that are of crucial significance in many economic applications. Chapter 4 provides more examples of the formulation of linear models; it considers the possibility of alternative formulations, and discusses the equivalence between optimisation and competitive equilibrium.

These linear models may be used to represent static or single-period situations, and the illustrations in the first four chapters are all of this kind. However, a great strength of such modelling is that it may be extended to multi-period situations; it is used in this way both in economic theory and in prescriptive policy studies. Chapter 6 treats such optimisation over time: the first few sections deploy linear models in this new context; the later parts deal with the concept of the planning horizon, and consider the question of consistency in decision-making over time, i.e. where the agent can change later parts of a plan before they have been implemented.

Neoclassical economic theory supposes that, at the outset of the optimisation, a list of the efficient solutions is available; the economist's task is then to apply the relevant economic criterion in order to select the most desirable of these efficient solutions. This view supposes (in effect) that some engineers or other generous souls have carried out some prior optimisation in order to produce the list of efficient points. For work in economic theory, this approach *may* be entirely satisfactory, since the objective is to establish the *qualitative* nature of the optimum. For *empirical* work in prescriptive policy studies (or elsewhere), no such short-cut is ever possible. Instead, the analysis must begin with the fundamental building blocks. This is what the linear (programming) model does, and thus it is no surprise to find programming models much used in empirical work. This theme, in which traditional neoclassical theory is contrasted with programming analyses, is explored in detail in Chapter 5, in the context of production theory.

A further comparison between neoclassical theory and programming models relates to the nature of the optimal solution. The traditional theory uses calculus-based models, in which it is supposed that non-negativity (and similar) requirements are met automatically. Programming methods in contrast (and more realistically) give explicit recognition to such requirements, and the optimal solution is not uncommonly a 'boundary' solution, i.e. at a point where the non-negativity (or other) requirement is an effective constraint. Indeed, it is one of the particular strengths of linear programming that it recognises this point, and uses the property to guide efficient methods of calculation.

For non-linear problems, the traditional calculus approach has been extended in recent years. The crucial innovation is the development of the Kuhn–Tucker conditions; these characterise an optimum when the problem has *inequality* constraints, which is of course the norm in economic models. These developments are treated in Chapter 7.

The Kuhn–Tucker conditions may be used in (qualitative) inquiries in economic theory, and also to guide explicit numerical calculations for small-scale models. But the conditions themselves do not comprise a computational procedure and hence are not applicable, as they stand, to large-scale numerical cases. The first section of Chapter 8 considers the nature of the non-linear optimum, to show why computation is difficult in the general case, and briefly describes some general computational procedures. The next two sections deal with relatively powerful computational procedures for *separable* and *quadratic programming*, two special cases that are important in economic analysis. The related topics of *indivisibilities* and *diseconomies of scale* (and the problems they create for competitive markets) are considered in the remaining parts of the chapter.

Multi-period situations can be regarded as merely one (important) class of multi-stage decision processes, in which a number of decisions are taken sequentially (or may be so regarded). Often, there are a number of alternative computational procedures that may be used for such a problem. However, the conceptual approach known as *dynamic programming* commonly leads to the most effective computational procedure; this is the subject of Chapter 9.

Chapter 10 presents detailed accounts of some further economic applications using formal optimising methods. The aim of the chapter is to convey to the student something of the practicalities of using such methods in economic analysis. In the same spirit, the chapter ends with some advice on reading such research in the journals and research monographs.

Finally it must be noted that the book does not discuss uncertainty. Recent years have seen very important developments in the modelling of uncertainty (mainly by using probabilistic concepts), and in consequence significant methods and results have been established, both in economic theory and in prescriptive work. But these problems and methods generally require models that are even more complex than those noted earlier in this introduction, and it seems best not to include them in a book intended for the student's first approach to the field of optimisation.

To help the student with limited mathematical background, Appendix A identifies the mathematical prerequisites, and offers advice on reading for any student who wishes to undertake some revision of these topics. Some students (especially those with a business orientation) may come to the book with a limited knowledge of the fundamental concepts of economic theory. Although some such concepts are sketched briefly when first used in the text, fuller and more formal definitions of some of the terms are given in a glossary (Appendix B).

Traditionally the student of economic theory has not been led to pay sufficient attention to the practicalities of applying the analytical models. As should be clear already, this book seeks to avoid that shortcoming, and it is hoped that such a student will wish to read the whole book. However, any student of theoretical inclination who insists on a narrow specialisation may omit, without significant loss of continuity, a few 'very applied' sections, such as: 1.10, 2.5, 4.1, 4.2, 5.10, 8.2, 8.3 and 8.8.

Similarly the business-oriented student should read the whole book to obtain the widest perspective. But lack of economic background may encourage such a student to omit, again without significant loss of continuity, sections 4.4, 4.5, 4.6 and 8.7.

Some chapters need not be studied in the order in which they are presented. The first three chapters must be read at the outset. Thereafter, for an initial reading, the student may pass over any or all of Chapters 4, 5 and 6 if it is desired to come quickly to the non-linear optimisation of Chapter 7. On the other hand, the student with business and management interests might read the first three sections of Chapter 4, include or pass over Chapter 5 according to taste, and then read Chapters 6 and 9 before turning to the non-linear and related topics discussed in Chapters 7 and 8. But all students are urged to return later to the omitted chapters, in order to get a balanced and full picture of economic optimisation.

1

The Formulation of Linear Models

1.1 PROGRAMMING PROBLEMS

Many decision problems facing economic agents can be represented by *programming* models. The word 'programming' seems to have been used initially to denote that the aim of the economic calculation is to find a programme, or set of decisions. The mathematical techniques for performing such calculations have become known as 'mathematical programming'. A distinctive feature of these techniques is that they are directed towards *finding* numerical solutions (as distinct from merely describing the general properties of solutions); hence the emphasis is on computational efficiency.

The programme or solution that is sought is the one that optimises some function; depending on the context, the aim is to maximise something desirable (e.g. profits, utility, output of steel, balance-of-payments surplus) or to minimise something undesirable (e.g. costs, the time it takes to complete a set of tasks). As in all optimisation problems in economics, resources are scarce, and hence the programming model has constraints to reflect this. The model may also need to include constraints to represent other requirements that the decision variables must meet.

This and the following three chapters deal with *linear* programming problems; in other words, the models optimise linear functions, subject to linear constraints. This might turn out to be a serious restriction on our field of interest. But as will be seen, a wide variety of problems can be represented satisfactorily by linear models. In many cases, the problem naturally takes a linear form; in *some* of the cases where this is not so, the problem may be represented approximately by a linear model, with results that are a sufficiently good approximation to be worth using. However, there are many cases where a linear model will not serve, and some of these are considered in later chapters.

Until further notice, it is supposed that the commodities involved in the models are continuously divisible, i.e. they may be measured in fractions of a unit if necessary. On many occasions, this will be conceptually precise (e.g. tons of steel); on other occasions it will be a good approximation (e.g. if weekly production is of the order of 2,000 cars, it will be permissible to round to the nearest car any fractional result that comes out of the programming model). In yet other cases, the divisibility assumption is not tenable. For example, if a model solution proposes that 1.42 aircraft be flown from A to B, the answer clearly will not do as it stands. Furthermore, it is by no means obvious what rounding would be best, since the change made by rounding will be a large proportion of the original answer. In these 'awkward' cases, a rather different approach is necessary; this too is discussed later.

This chapter introduces three examples of problems that may be represented by linear models. It shows how the models may be formulated (although the task of finding solutions is postponed until the next chapter) and goes on to discuss some general features of such linear models.

1.2 A FIRST LINEAR MODEL: EXAMPLE A

For a first example (Example A), suppose that a manufacturer of a homogeneous commodity has stocks of 100 tons of the commodity at depot A and 250 tons at depot B; he has agreed to supply customers C, D and E (each at different locations) with 50, 120 and 180 tons respectively, as indicated in Table 1.1. The total stock (of 350 tons) exactly matches the total tonnage to be delivered to the customers, and thus all the stocks must be used up; all that the manufacturer has to decide is how many tons should be sent from each depot to each customer. Suppose that the manufacturer's aim is to arrange delivery at minimum transport cost, and that he finds that the cost of transport on any particular route is constant *per ton* irrespective of how many tons are sent along that route from a particular depot to a particular customer. Specific cost assumptions for each route are indicated in Table 1.1, and these complete the specification of the problem.

Model-building is begun by defining algebraic variables for the unknowns: let x_1 represent the number of tons to be sent from depot A to customer C, x_2 the number of tons from depot A to customer D, and so

Table 1.1 *The data for Example A – a transportation problem*

	Depot A	*Depot B*	
No. of tons available:	100	250	
	Customer C	*Customer D*	*Customer E*
No. of tons required:	50	120	180

Route	*Transport cost ($ per ton)*	*Variables (number of tons along each route)*
A to C	5	x_1
A to D	10	x_2
A to E	12	x_3
B to C	3	x_4
B to D	4	x_5
B to E	9	x_6