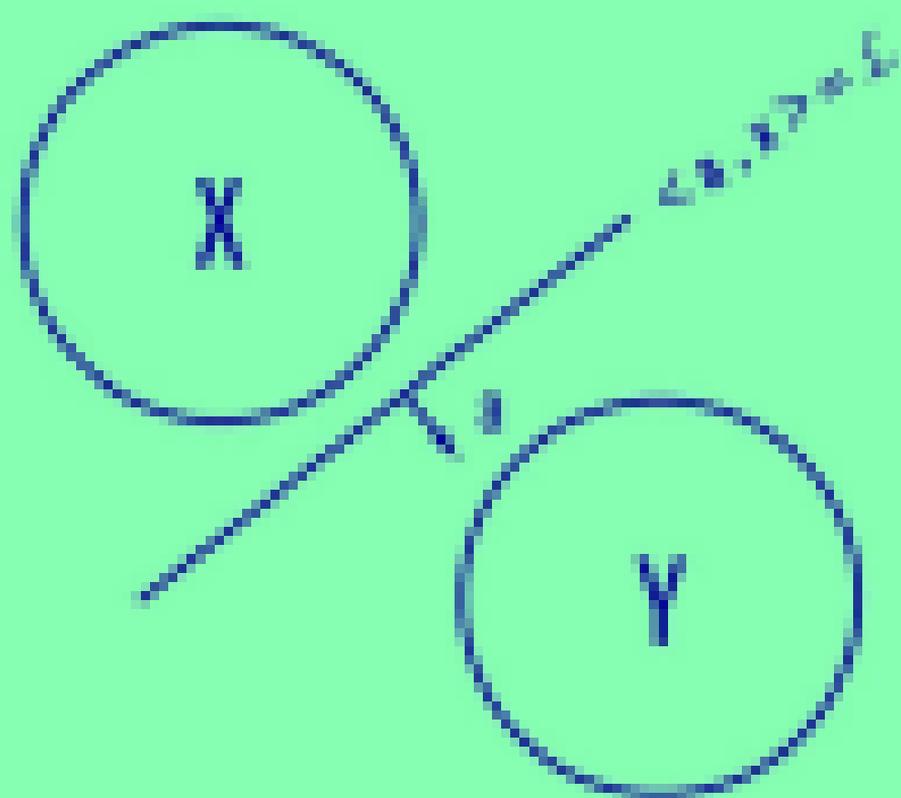


Pure and Applied Mathematics
A Wiley-Interscience Series of Texts, Monographs, and Surveys

Convexity and Optimization in \mathbb{R}^n

Leonard D. Berkovitz



CONVEXITY AND OPTIMIZATION IN \mathbb{R}^n

LEONARD B. ABSCOFFE
Yeshiva University



A Wiley-Interscience Publication
JOHN WILEY & SONS, INC.

CONVEXITY AND OPTIMIZATION IN \mathbb{R}^n

CONVEXITY AND OPTIMIZATION IN \mathbb{R}^n

LEONARD B. ABSCOFFE
Yonkers University



A Wiley-Interscience Publication
JOHN WILEY & SONS, INC.

This book is printed on acid-free paper. 

Copyright © 2002 by John Wiley & Sons, Inc., New York, NY. All rights reserved.

Published simultaneously in Canada.

No part of this publication may be reproduced, stored in a retrieval system or transmitted in any form or by any means, electronic, mechanical, photocopying, recording, scanning or otherwise, except as permitted under Sections 107 or 108 of the 1976 United States Copyright Act, without either the prior written permission of the Publisher, or authorization through payment of the appropriate per-copy fee to the Copyright Clearance Center, 222 Rosewood Drive, Danvers, MA 01923, (978) 7508400 for (978) 7508744. Requests to the Publisher for permission should be addressed to the Permissions Department, John Wiley & Sons, Inc., 485 Third Avenue, New York, NY 10158-0002, (212) 8506015, fax (212) 8506009, E-mail: PERMREQ @ WILEY.COM.

For ordering and customer service, call 1-800-CALL-WILEY.

Library of Congress Cataloging-in-Publication Data

Beikwa, Leonard David. EN-

Convexity and optimization in \mathbb{R}^n (superior ed.) / Leonard D. Beikwa.

p. cm.—Pure and applied mathematics

1. A Wiley-Interscience publication.

Includes bibliographical references and index.

ISBN 0-471-3320-0 (pbk.:alk. paper)

1. Convex sets. I. Mathematics—Optimization. II. Title. III. Pure and applied

mathematics (Wiley & Sons classification).

QA612.B45 2002

517.94—dc22

2001048391

Printed in the United States of America

10 9 8 7 6 5 4 3 2 1

To my wife, Anne

CONTENTS

Preface	ix
I Topics in Real Analysis	0
1. Introduction	0
2. Vectors in \mathbb{R}^n	0
3. Algebra of Sets	4
4. Metric Topology of \mathbb{R}^n	9
5. Limits and Continuity	9
6. Basic Property of Real Numbers	11
7. Completeness	14
8. Equivalent Norms and Cartesian Products	16
9. Fundamental Existence Theorem	18
10. Linear Transformations	21
11. Differentiation in \mathbb{R}^n	22
II Convex Sets in \mathbb{R}^n	26
1. Lines and Hyperplanes in \mathbb{R}^n	26
2. Properties of Convex Sets	39
3. Separation Theorems	42
4. Supporting Hyperplanes, Extreme Points	56
5. Systems of Linear Inequalities, Theorems of the Alternative	61
6. Affine Geometry	69
7. More on Separation and Support	80
III Convex Functions	87
1. Definition and Elementary Properties	87
2. Subgradients	102
3. Differentiable Convex Functions	109
4. Alternative Theorems for Convex Functions	113
5. Application to Game Theory	117

IV Optimization Problems	126
1. Introduction	126
2. Differentiable Unconstrained Problems	129
3. Optimization of Convex Functions	137
4. Linear Programming Problems	138
5. First-Order Conditions for Differentiable Nonlinear Programming Problems	143
6. Second-Order Conditions	153
V Convex Programming and Duality	179
1. Problem Statement	179
2. Necessary Conditions and Sufficient Conditions	181
3. Perturbation Theory	188
4. Lagrangian Duality	200
5. Geometric Interpretation	207
6. Quadratic Programming	210
7. Duality in Linear Programming	219
VI Simplex Method	222
1. Introduction	222
2. Extreme Points of Feasible Set	229
3. Preliminaries to Simplex Method	230
4. Phase II of Simplex Method	234
5. Termination and Cycling	249
6. Phase I of Simplex Method	251
7. Revised Simplex Method	255
Bibliography	263
Index	265

PREFACE

This book presents the mathematics of finite-dimensional optimization, focusing those aspects of convexity that are useful in this context. It provides a basis for the further study of convexity, of more general optimization problems, and of numerical algorithms for the solution of finite-dimensional optimization problems. The intended audience consists of beginning graduate students in engineering, economics, operations research, and mathematics and strong undergraduates in mathematics. This was the audience in a one-semester course at Purdue, MA 521, from which this book evolved.

Ideally, the prerequisites for reading this book are good introductory courses in real analysis and linear algebra. In teaching MA 521, I found that while the mathematics students had the real analysis prerequisites, many of the other students who took the course because of their interest in optimization did not have this prerequisite. Chapter I is for those students and readers who do not have the real analysis prerequisite; in it I present those concepts and results from real analysis that are needed. Except for the Weierstrass theorem on the existence of a minimum, the “heavy” or “deep” theorems are stated without proof. Students without the real variables prerequisite found the material difficult at first, but most managed to assimilate it at a satisfactory level. The advice to readers for whom this is the first encounter with the material in Chapter I is to make a serious effort to master it and to return to it as it is used in the sequel.

To address as wide an audience as possible, I have not always presented the most general result or argument. Thus, in Chapter II I chose the “clearest point” approach to separation theorems, rather than most generally valid arguments, because I believe it to be more intuitive and straightforward for the intended audience. Readers who wish to get the best possible separation theorems in finite dimensions should read Sections 6 and 7 of Chapter II. In proving the Farkas Lemma Theorem, I used a penalty function argument due to Motzkin rather than more technical arguments involving linearizations. I limited the discussion of duality to Lagrangian duality and did not consider Fenchel duality, since the latter would require the development of more mathematical machinery.

The numbering system and reference system for theorems, lemmas, remarks, and corollaries is the following. Within a given chapter, theorems, lemmas, and remarks are numbered consecutively in each section, preceded by the section number. Thus, the first theorem of Section 1 is Theorem 1.1, the second, Theorem 1.2, and so on. The same applies to lemmas and remarks. Corollaries are numbered consecutively within each section without a reference to the section number. Reference to a theorem in the same chapter is given by the theorem number. Reference to a theorem in a chapter different from the current one is given by the theorem number preceded by the chapter number in Roman numerals. Thus, a reference in Chapter IV to Theorem 4.1 in Chapter II would be Theorem II.4.1. References to lemmas and remarks are similar. References to corollaries within the same section are given by the number of the corollary. References to corollaries in a different section of the same chapter are given by prefixing the section number to the corollary number; references in a different chapter are given by prefixing the chapter number in Roman numerals to the preceding.

I thank Rita Stevens and John Gregory for reading the current notes version of this book and for their corrections and suggestions for improvement. I thank Terry Combe for preparing the figures. I also thank Betty Clark for typing numerous innumerable versions and revisions of the notes for MR 521. His skill and cooperation contributed greatly to the success of this project.

LEONARD D. BRADY

West Lafayette, Indiana

I

TOPICS IN REAL ANALYSIS

1. INTRODUCTION

The serious study of convexity and optimization problems in \mathbb{R}^n requires some background in real analysis and in linear algebra. In teaching a course based on notes from which this text evolved, the author and his colleagues assumed that the students had an undergraduate course in linear algebra but did not necessarily have a background in real analysis. The purpose of this chapter is to provide the reader with most of the necessary background in analysis. Not all statements will be proved. The reader, however, is expected to understand the definitions and the theorems and is expected to follow the proofs of statements whenever the proofs are given. The bad news is that many readers will find this chapter to be the most difficult one in the text. The good news is that careful study of this material will provide background for many other courses and that subsequent chapters should be easier. If necessary, the reader should return to this chapter when encountering this material later on.

2. VECTORS IN \mathbb{R}^n

By *column n -space*, or \mathbb{R}^n , we mean the set of all n -tuples $\mathbf{x} = (x_1, \dots, x_n)$, where the x_i , $i = 1, \dots, n$ are real numbers. Thus, \mathbb{R}^n is a generalization of the familiar two- and three-dimensional spaces \mathbb{R}^2 and \mathbb{R}^3 . The elements \mathbf{x} of \mathbb{R}^n are called *vectors* or *points*. We will often identify the vector $\mathbf{x} = (x_1, \dots, x_n)$ with the $n \times 1$ matrix

$$\begin{pmatrix} x_1 \\ \vdots \\ x_n \end{pmatrix}$$

and abuse the use of the equal sign to write

$$\mathbf{x} = \begin{pmatrix} x_1 \\ \vdots \\ x_n \end{pmatrix}.$$

In this case we shall call \mathbf{x} a column vector. We shall also identify \mathbf{x} with the $1 \times n$ matrix (x_1, \dots, x_n) and write $\mathbf{x} = (x_1, \dots, x_n)$. In this case we shall call \mathbf{x} a row vector. It will usually be clear from the context whether we consider \mathbf{x} to be a row vector or a column vector. When there is danger of confusion, we will identify \mathbf{x} with the column vector and use the transpose symbol, which is a superscript t , to denote the row vector. Thus $\mathbf{x}^t = (x_1, \dots, x_n)$.

We define two operations, vector addition and multiplication by a scalar. If $\mathbf{x} = (x_1, \dots, x_n)$ and $\mathbf{y} = (y_1, \dots, y_n)$ are two vectors, we define their sum $\mathbf{x} + \mathbf{y}$ to be

$$\mathbf{x} + \mathbf{y} = (x_1 + y_1, \dots, x_n + y_n).$$

For any scalar, or real number, α we define $\alpha\mathbf{x}$ to be

$$\alpha\mathbf{x} = (\alpha x_1, \dots, \alpha x_n).$$

We assume that the reader is familiar with the properties of these operations and knows that under these operations \mathbb{R}^n is a vector space over the reals.

Another important operation is the inner product, or dot product, of two vectors, denoted by $\langle \mathbf{x}, \mathbf{y} \rangle$ or $\mathbf{x} \cdot \mathbf{y}$ and defined by

$$\langle \mathbf{x}, \mathbf{y} \rangle = \sum_{j=1}^n x_j y_j.$$

Again, we assume that the reader is familiar with the properties of the inner product. We use the inner product to define the norm $\|\cdot\|$ or length of a vector \mathbf{x} as follows:

$$\|\mathbf{x}\| = \langle \mathbf{x}, \mathbf{x} \rangle^{1/2} = \left[\sum_{j=1}^n x_j^2 \right]^{1/2}.$$

This norm is called the euclidean norm. In \mathbb{R}^2 and \mathbb{R}^3 the euclidean norm reduces to the familiar length. It is straightforward to show that the norm has the following properties:

$$\|\mathbf{x}\| \geq 0 \quad \text{for all } \mathbf{x} \text{ in } \mathbb{R}^n, \quad (1)$$

$$\|\mathbf{x}\| = 0 \quad \text{if and only if } \mathbf{x} \text{ is the zero vector } \mathbf{0} \text{ in } \mathbb{R}^n, \quad (2)$$

$$\|c\mathbf{x}\| = |c| \|\mathbf{x}\| \quad \text{for all real numbers } c \text{ and vectors } \mathbf{x} \text{ in } \mathbb{R}^n. \quad (3)$$

The norm has two additional properties, which we will prove:

For all vectors \mathbf{x} and \mathbf{y} in \mathbb{R}^n

$$|\langle \mathbf{x}, \mathbf{y} \rangle| \leq \|\mathbf{x}\| \|\mathbf{y}\| \quad (4)$$

with equality holding if and only if one of the vectors is a scalar multiple of the other.

For all vectors \mathbf{x} and \mathbf{y} in \mathbb{R}^n

$$\|\mathbf{x} + \mathbf{y}\| \leq \|\mathbf{x}\| + \|\mathbf{y}\|, \quad (5)$$

with equality holding if and only if one of the vectors is a nonnegative scalar multiple of the other.

The inequality (4) is known as the Cauchy–Schwarz inequality. The reader should not be discouraged by the slickness of the proof that we will give or by the slickness of other proofs. Usually, the first time that a mathematical result is discovered, the proof is rather complicated. Often, many very smart people then look at it and constantly improve the proof until a relatively simple and clever argument is found. Therefore, do not be intimidated by clever arguments. New to the proof.

Let \mathbf{x} and \mathbf{y} be any pair of fixed vectors in \mathbb{R}^n . Let λ be any real scalar. Then

$$\begin{aligned} 0 &\leq \|\mathbf{x} + \lambda\mathbf{y}\|^2 = \langle \mathbf{x} + \lambda\mathbf{y}, \mathbf{x} + \lambda\mathbf{y} \rangle \\ &= \|\mathbf{x}\|^2 + 2\lambda\langle \mathbf{x}, \mathbf{y} \rangle + \lambda^2\|\mathbf{y}\|^2 \quad \text{for all } \lambda. \end{aligned} \quad (6)$$

This says that the quadratic in λ on the right either has a double real root or has no real roots. Therefore, by the quadratic formula, its discriminant must satisfy

$$4\langle \mathbf{x}, \mathbf{y} \rangle^2 - 4\|\mathbf{x}\|^2\|\mathbf{y}\|^2 \leq 0.$$

Transposing and taking square roots give

$$|\langle \mathbf{x}, \mathbf{y} \rangle| \leq \|\mathbf{x}\| \|\mathbf{y}\|.$$

Equality holds if and only if the quadratic has a double real root, say $\lambda = \alpha$. But then, from (6), $\|\mathbf{x} + \alpha\mathbf{y}\| = 0$. Hence by (2) $\mathbf{x} + \alpha\mathbf{y} = \mathbf{0}$, and $\mathbf{x} = -\alpha\mathbf{y}$.

To prove (5), we write

$$\begin{aligned} \|\mathbf{x} + \mathbf{y}\|^2 &= \langle \mathbf{x} + \mathbf{y}, \mathbf{x} + \mathbf{y} \rangle = \|\mathbf{x}\|^2 + 2\langle \mathbf{x}, \mathbf{y} \rangle + \|\mathbf{y}\|^2 \\ &\leq \|\mathbf{x}\|^2 + 2\|\mathbf{x}\| \|\mathbf{y}\| + \|\mathbf{y}\|^2 = (\|\mathbf{x}\| + \|\mathbf{y}\|)^2, \end{aligned} \quad (7)$$

where the inequality follows from the Cauchy–Schwarz inequality. If we now take square roots, we get the triangle inequality. We obtain equality in (7), and hence in the triangle inequality, if and only if $\langle \mathbf{x}, \mathbf{y} \rangle$ is nonnegative and either \mathbf{y} is a scalar multiple of \mathbf{x} or \mathbf{x} is a scalar multiple of \mathbf{y} . But, under those circumstances, for $\langle \mathbf{x}, \mathbf{y} \rangle$ to be nonnegative, the multiple must be nonnegative. The inequality (5) is called the triangle inequality because when it is applied to vectors in \mathbb{R}^2 or \mathbb{R}^3 it says that the length of a third side of a triangle is less than or equal to the sum of the lengths of the other two sides.

Exercise 1.1. For any vectors x and y in \mathbb{R}^2 , show that $\|x + y\|^2 + \|x - y\|^2 = 2\|x\|^2 + 2\|y\|^2$. Interpret this relation as a statement about parallelograms in \mathbb{R}^2 and \mathbb{R}^n .

In elementary courses it is shown that in \mathbb{R}^2 and \mathbb{R}^3 the cosine of the angle θ , $0 \leq \theta < \pi$, between two nonzero vectors x and y is given by the formula

$$\cos \theta = \frac{\langle x, y \rangle}{\|x\| \|y\|}.$$

Thus, x and y are orthogonal if and only if $\langle x, y \rangle = 0$. In \mathbb{R}^n , the Cauchy-Schwarz inequality allows us to give meaning to the notion of angle between nonzero vectors as follows. Since $|\langle x, y \rangle| \leq \|x\| \|y\|$, the absolute value of the quotient

$$\frac{\langle x, y \rangle}{\|x\| \|y\|} \tag{1}$$

is less than or equal to 1, and thus (1) is the cosine of an angle between zero and π . We define this angle to be the angle between x and y . We say that x and y are orthogonal if $\langle x, y \rangle = 0$.

8. ALGEBRA OF SETS

Given two sets A and B , the set A is said to be a subset of B , or to be contained in B , and written $A \subset B$, if every element of A is also an element of B . The set A is said to be a proper subset of B if A is a subset of B and there exists at least one element of B that is not in A . This is written as $A \subsetneq B$. Two sets A and B are said to be equal, written $A = B$, if $A \subset B$ and $B \subset A$. The empty set, or null set, is the set that has no members. The empty set will be denoted by \emptyset . The notation $x \in S$ will mean " x belongs to the set S " or " x is an element of the set S ."

Let $\{S_\alpha\}_{\alpha \in A}$ be a collection of sets indexed by an index set A . By the union of the sets $\{S_\alpha\}_{\alpha \in A}$, denoted by $\bigcup_{\alpha \in A} S_\alpha$, we mean the set consisting of all elements that belong to at least one of the S_α . Note that if not all of the S_α are empty, then $\bigcup_{\alpha \in A} S_\alpha$ is not empty. By the intersection of the sets $\{S_\alpha\}_{\alpha \in A}$, denoted by $\bigcap_{\alpha \in A} S_\alpha$, we mean the set of points x with the property that x belongs to every set S_α . Note that for a collection $\{S_\alpha\}_{\alpha \in A}$ of nonempty sets, the intersection can be empty.

Let X be a set, say \mathbb{R}^2 for example, and let S be a subset of X . By the complement of S (relative to X), written cS , we mean the set consisting of those points in X that do not belong to S . Note that $c\emptyset$ is the empty set. By convention, we take the complement of the empty set to be X . We then always have $c(cS) = S$. Also, if $A \subset B$, then $cB \subset cA$. The reader should convince himself or herself of the truth of the following lemma, known as de Morgan's

law, either by writing out a proof or drawing Venn diagrams for some simple cases.

Lemma 3.1. Let $\{S_\alpha\}_{\alpha \in A}$ be a collection of subsets of a set X . Then

$$\bigcup_{\alpha \in A} S_\alpha = c \left[\bigcap_{\alpha \in A} cS_\alpha \right],$$

$$\bigcap_{\alpha \in A} S_\alpha = c \left[\bigcup_{\alpha \in A} cS_\alpha \right].$$

4. METRIC TOPOLOGY OF \mathbb{R}^n

Let X be a set. A function d that assigns a real number to each pair of points (x, y) with $x \in X$ and $y \in X$ is said to be a *metric*, or *distance function*, on X if it satisfies the following:

$$d(x, y) \geq 0, \quad \text{with equality holding if and only if } x = y. \quad (1)$$

$$d(x, y) = d(y, x) \quad (2)$$

$$d(x, y) \leq d(x, z) + d(z, y) \quad \text{for all } x, y, z \text{ in } X. \quad (3)$$

The last inequality is called the *triangle inequality*.

In \mathbb{R}^n we define a function d on pairs of points x, y by the formula

$$d(x, y) = \|x - y\| = \left(\sum_{i=1}^n (x_i - y_i)^2 \right)^{1/2}. \quad (4)$$

Exercise 4b. Use the properties of the norm to show that the function d defined by (4) is a metric, or distance function, on \mathbb{R}^n .

A set X with metric d is called a *metric space*. Although metric spaces occur in many areas of mathematics and applications, we shall confine ourselves to \mathbb{R}^n . In \mathbb{R}^n and \mathbb{R}^2 the function defined in (4) is the ordinary euclidean distance.

We now present some important definitions to the reader, who should master them and try to improve his or her understanding by drawing pictures in \mathbb{R}^2 .

The (open) ball centered at x with radius $r > 0$, denoted by $B(x, r)$, is defined to be the set of points y whose distance from x is less than r . We write this in symbols as

$$B(x, r) = \{y : \|y - x\| < r\}.$$

The closed ball $\overline{B}(x, r)$ with center at x and radius $r > 0$ is defined by

$$\overline{B}(x, r) = \{y : \|y - x\| \leq r\}.$$

Let $S \subset \mathbb{R}^n$. A point \mathbf{x} is an *interior point* of S if there exists an $r > 0$ such that $B(\mathbf{x}, r) \subset S$. A set S need not have any interior points. For example, consider the set of points in the plane of the form $(x_1, 0)$ with $0 < x_1 < 1$. This is the interval $0 < x_1 < 1$ on the x_1 -axis. It has no interior points, considered as a set in \mathbb{R}^2 . On the other hand, if we consider this as a set in \mathbb{R}^1 , every point is an interior point. This leads to the important observation that for a given set, whether or not it has interior points may depend on in which Euclidean space \mathbb{R}^n we take the set to be lying, or embedded.

If the set of interior points of a set S is not empty, then we call the set of interior points the *interior* of S and denote it by $\text{int} S$.

A set S is said to be *open* if all points of S are interior points. Thus an equivalent definition of an open set is that it is equal to its interior. If we go back to the definition of interior point, we can restate the definition of an open set as follows. A set S is open if for every point \mathbf{x} in S there exists a positive number $r > 0$, which may depend on \mathbf{x} , such that the ball $B(\mathbf{x}, r)$ is contained in S . The last definition, while being wordy, is the one that the reader should picture constantly.

A point \mathbf{x} is said to be a *limit point* of a set S if for every $\epsilon > 0$ there exists a point $\mathbf{x}_\epsilon \neq \mathbf{x}$ such that \mathbf{x}_ϵ belongs to S and $\mathbf{x}_\epsilon \in B(\mathbf{x}, \epsilon)$. The point \mathbf{x}_ϵ will in general depend on ϵ . A set need not have any limit points, and a limit point of a set need not belong to the set. An example of a set without a limit point is the set \mathbb{N} of positive integers, considered as a set in \mathbb{R}^1 . For an example of a limit point that does not belong to a set, consider the set $S = \{1/n : n = 1, 2, 3, \dots\}$ in \mathbb{R}^1 . Zero is a limit point of the set, yet zero does not belong to S . We shall denote the set of limit points of a set S by S' .

Exercise 41. (a) Sketch the graph of $y = \sin(1/x)$, $x > 0$.

(b) Consider the graph as a set in \mathbb{R}^2 and find the limit points of this set.

The *closure* of a set S , denoted by \bar{S} , is defined to be the set $S \cup S'$; that is, $\bar{S} = S \cup S'$. A set S is said to be *closed* if S contains all of its limit points; that is, $S' \subset S$. A set S is closed if and only if $\bar{S} = S$. To see this, note that $\bar{S} = S$ and the definition $\bar{S} = S \cup S'$ imply that $S' \subset S \cup S'$. Hence $S' \subset S$. On the other hand, if $S' \subset S$, then $S \cup S' = S$ and so $\bar{S} = S \cup S' = S$.

A set can be neither open nor closed. In \mathbb{R}^2 consider $\text{Int}(1)$, the ball with center at the origin and radius 1. It should be intuitively clear that all points \mathbf{x} in $\text{Int}(1)$ with $\|\mathbf{x}\| = 1$ are limit points of $\text{Int}(1)$. Now consider the set

$$S = \text{Int}(1) \cup \{\mathbf{x} = (x_1, x_2) : \|\mathbf{x}\| = 1, x_1 \geq 0\}.$$

(The reader should sketch this set.) Points $\mathbf{x} = (x_1, x_2)$ with $\|\mathbf{x}\| = 1$ and $x_1 > 0$ are not interior points of S since for such an \mathbf{x} , no matter how small we choose $r > 0$, the ball $B(\mathbf{x}, r)$ will not belong to S . Hence S is not open. Also, S is not closed since points $\mathbf{x} = (x_1, x_2)$ with $\|\mathbf{x}\| = 1$ and $x_1 < 0$ are limit points of S , yet they do not belong to S .

It follows from our definitions that \mathbb{R}^n itself is both open and closed. By convention, we will also take the empty set to be open and closed.

THEOREM 4.1. *Complements of closed sets are open. Complements of open sets are closed.*

Proof. If $S = \emptyset$, where \emptyset denotes the empty set, or if $S = \mathbb{R}^n$, then there is nothing to prove. Let $S \neq \emptyset$, $S \neq \mathbb{R}^n$ and let S be closed. Let x be any point in cS . Then $x \notin S$ and x is not a limit point of S . If we recall the definition of “ x is a limit point of S ,” we see that the statement “ x is not a limit point of S ” means that there is an $\epsilon_0 > 0$ such that all points $y \neq x$ with $y \in B(x, \epsilon_0)$ are in cS . Since $x \in cS$, this means that $B(x, \epsilon_0) \subset cS$, and so cS is open. Now let $S = \emptyset$, $S \neq \mathbb{R}^n$ and let S be open. Let x be a limit point of cS . Then for every $\epsilon > 0$ there exists a point $x_\epsilon \neq x$ in cS such that $x_\epsilon \in B(x, \epsilon)$. Since $x_\epsilon \in cS$, the ball $B(x_\epsilon, \epsilon)$ does not belong to S . Thus x cannot belong to the open set S . For that we would be able to find an ϵ_0 such that $B(x, \epsilon_0)$ did belong to S . Hence $x \in cS$, so cS is closed. \square

Exercise 4.3. Show that for $x \in \mathbb{R}^n$ and $r > 0$ the set $B(x, r)$ is open; that is, show that an open ball is open.

Exercise 4.4. Show that for $x \in \mathbb{R}^n$ and $r > 0$ the closed ball $\overline{B(x, r)}$ is closed.

Exercise 4.5. Show that any finite set of points x_1, \dots, x_n in \mathbb{R}^n is closed.

Exercise 4.6. Show that in \mathbb{R}^n no point x with $\|x\| = 1$ is an interior point of $B(0, 1)$.

Exercise 4.7. Show that for any set S in \mathbb{R}^n the set \bar{S} is closed.

Exercise 4.8. Show that for any set S the closure \bar{S} is equal to the intersection of all closed sets containing S .

THEOREM 4.2. (i) Let $\{O_\alpha\}_{\alpha \in I}$ be a collection of open sets. Then $\bigcup_{\alpha \in I} O_\alpha$ is open.
 (ii) Let O_1, \dots, O_n be a finite collection of open sets. Then $\bigcap_{i=1}^n O_i$ is open.
 (iii) Let $\{F_\alpha\}_{\alpha \in I}$ be a collection of closed sets. Then $\bigcap_{\alpha \in I} F_\alpha$ is closed.
 (iv) Let F_1, \dots, F_n be a finite collection of closed sets. Then $\bigcup_{i=1}^n F_i$ is closed.

Note that an infinite collection of open sets need not have an intersection that is open. For example, in \mathbb{R}^1 , for each $n = 1, 2, 3, \dots$ let

$$O_n = (1 - 1/n, 1 + 1/n) = \{x \in \mathbb{R}^1 : 1 - 1/n < x < 1 + 1/n\}.$$

In \mathbb{R}^1 each O_n is open, and $\bigcap_{n=1}^{\infty} O_n = \{1\}$, a point. By Exercise 4.5, the set consisting of the point $x = 1$ is closed. Similarly, the union of an infinite number of closed sets need not be closed. To see this, for each $n = 1, 2, 3, \dots$ let $F_n = [0, 1 - 1/(n+1)] = \{x \in \mathbb{R}^1 : 0 \leq x \leq 1 - 1/(n+1)\}$. Each F_n is closed, and

$\bigcup_{i=1}^n \mathcal{F}_i = [x, y) = \{x \in \mathbb{R}^1 : 0 < x < 1\}$, which is not closed since it does not contain the limit point $x = 1$.

We now prove the theorem. To establish (i), we take an arbitrary element x in $\bigcup_{i=1}^n \mathcal{O}_i$. Since x is in the union, it must belong to at least one of the sets $\{\mathcal{O}_i\}_{i=1}^n$ say \mathcal{O}_{i_0} . Since \mathcal{O}_{i_0} is open, there exists an $r_0 > 0$ such that $\mathcal{B}(x, r_0) \subseteq \mathcal{O}_{i_0}$. Hence $\mathcal{B}(x, r_0) \subseteq \bigcup_{i=1}^n \mathcal{O}_i$, and thus the union is open. To establish (ii), let x be an arbitrary point in $\bigcap_{i=1}^n \mathcal{O}_i$. Then $x \in \mathcal{O}_i$ for each $i = 1, \dots, n$. Since each \mathcal{O}_i is open, for each $i = 1, \dots, n$ there exists an $r_i > 0$ such that $\mathcal{B}(x, r_i) \subseteq \mathcal{O}_i$. Let $r = \min\{r_i : i = 1, \dots, n\}$. Then $\mathcal{B}(x, r) \subseteq \mathcal{B}(x, r_i) \subseteq \mathcal{O}_i$ for each $i = 1, \dots, n$. Hence $\mathcal{B}(x, r) \subseteq \bigcap_{i=1}^n \mathcal{O}_i$, and thus the intersection is open.

The most efficient way to establish (ii) is to write

$$\bigcap_{i=1}^n \mathcal{F}_i = \bigcap_{i=1}^n \complement(\mathcal{F}_i) = \complement\left[\bigcup_{i=1}^n \mathcal{F}_i\right],$$

where the second equality follows from Lemma 3.1. Since \mathcal{F}_i is closed, $\complement \mathcal{F}_i$ is open and so is $\bigcup_{i=1}^n \complement(\mathcal{F}_i)$. Hence the complement of this set is closed, and therefore so is $\bigcap_{i=1}^n \mathcal{F}_i$. To establish (iv), write

$$\bigcup_{i=1}^n \mathcal{F}_i = \bigcup_{i=1}^n \complement(\mathcal{F}_i) = \complement\left[\bigcap_{i=1}^n \complement(\mathcal{F}_i)\right].$$

Again, $\complement \mathcal{F}_i$ is open for each $i = 1, \dots, n$. Hence the intersection $\bigcap_{i=1}^n \complement(\mathcal{F}_i)$ is open and the complement of the intersection is closed.

The reader who wishes to test his or her mastery of the relevant definitions should prove statements (iii) and (v) of the theorem directly.

8. LIMITS AND CONTINUITY

A sequence in \mathbb{R}^n is a function that assigns to each positive integer k a vector or point \mathbf{a}_k in \mathbb{R}^n . We usually write the sequence as $\{\mathbf{a}_k\}_{k=1}^{\infty}$ or $\{\mathbf{a}_k\}$, rather than in the usual function notation. Examples of sequences in \mathbb{R}^2 are

- (i) $\{\mathbf{a}_k\} = \{k\mathbf{i}, k\mathbf{j}\}$,
- (ii) $\{\mathbf{a}_k\} = \left\{ \left(\cos \frac{k\pi}{2}, \sin \frac{k\pi}{2} \right) \right\}$,
- (iii) $\{\mathbf{a}_k\} = \left\{ \left(\frac{1-1^k}{2}, \frac{1^k}{2} \right) \right\}$,
- (iv) $\{\mathbf{a}_k\} = \left\{ \left(1-1^k - \frac{1}{k}, 1-1^k - \frac{1}{k} \right) \right\}$.

The reader should plot the points in these sequences.

A sequence $\{x_n\}$ is said to have a limit y or to converge to y if for every $\epsilon > 0$ there is a positive integer $K(\epsilon)$ such that whenever $k > K(\epsilon)$, $x_k \in B_\epsilon(y)$. We write

$$\lim_{n \rightarrow \infty} x_n = y \quad \text{or} \quad x_n \rightarrow y.$$

The sequences (i), (ii), and (iv) in (1) do not converge, while the sequence (iii) converges to $0 = (0, 0)$. Sequences that converge are said to be convergent; sequences that do not converge are said to be divergent. A sequence $\{x_n\}$ is said to be bounded if there exists an $M > 0$ such that $|x_n| < M$ for all n . A sequence that is not bounded is said to be unbounded. The sequence (i) is unbounded; the sequences (ii), (iii), and (iv) are bounded.

The following theorem lists some basic properties of convergent sequences.

THEOREM 5.1. (i) The limit of a convergent sequence is unique.

(ii) Let $\lim_{n \rightarrow \infty} x_n = x_0$, let $\lim_{n \rightarrow \infty} y_n = y_0$, and let $\lim_{n \rightarrow \infty} r_n = r$, where $\{r_n\}$ is a sequence of scalars. Then $\lim_{n \rightarrow \infty} (x_n + y_n)$ exists and equals $x_0 + y_0$, and $\lim_{n \rightarrow \infty} r_n x_n$ exists and equals rx_0 .

(iii) A convergent sequence is bounded.

Exercise 5A. Prove Theorem 5.1.

Let $\{x_n\}$ be a sequence in \mathbb{R}^2 and let $r_1 < r_2 < r_3 < \dots < r_n < \dots$ be a strictly increasing sequence of positive integers. Then the sequence $\{x_{r_n}\}_{n=1}^{\infty}$ is called a subsequence of $\{x_n\}$. Thus, $\{x_{2k}\}_{k=1}^{\infty} = \{2k, 2k\}$ is a subsequence of (i). The sequence $\{x_{n^2}\}_{n=1}^{\infty} = \{1, 4, 9, 16, 25, \dots\}$ is a subsequence of (ii). Each of the sequences $\{x_{2n}\} = \{(1-1)^2 - 1/2k, (1-1)^2 - 1/2k\} = \{0 - 1/2k, 0 - 1/2k\}$ and $\{x_{2n+1}\} = \{(1-1)^2 - 1/2k + 1, (1-1)^2 - 1/2k + 1\}$ is a subsequence of (iv). The reader should plot the various subsequences.

We can now formulate a very useful criterion for a point to be a limit point of a set. (The reader should not confuse the notions of limit and limit point. They are different.)

LEMMA 5.1. A point x is a limit point of a set S if and only if there exists a sequence $\{x_n\}$ of points in S such that, for each k , $x_k \neq x$ and $x_k \rightarrow x$.

Proof. Let x be a limit point of S . Then for every positive integer k there is a point x_k in S such that $x_k \neq x$ and $x_k \in B_{1/k}(x)$. For every $\epsilon > 0$, there is a positive integer $K(\epsilon)$ satisfying $1/K(\epsilon) = \epsilon$. Hence for $k > K(\epsilon)$, we have $1/k < \epsilon$ and $B_{1/k}(x) \subset B_\epsilon(x)$. Hence the sequence $\{x_k\}$ converges to x . Conversely, let there exist a sequence $\{x_n\}$ of points in S with $x_n \neq x$ and $x_n \rightarrow x$. Let $\epsilon > 0$ be arbitrary. Since $x_n \rightarrow x$, there exists a positive integer $K(\epsilon)$ such that, for $k > K(\epsilon)$, $x_k \in B_\epsilon(x)$. Since by hypothesis $x_k \neq x$ for all k , it follows that we have satisfied the requirements of the definition that x is a limit point of S .

Remark 5.1. Let $\{x_k\}$ be a sequence of points belonging to a set S and converging to a point x . Then x must belong to \bar{S} . If x is in S , there is nothing to prove. If x were not in S , then since $x_k \in S$ for each k , it follows that $x_k \neq x$. Lemma 5.1 then implies that x is a limit point of S . Thus $x \in \bar{S}$. If S is a closed set, then since $\bar{S} = S$, the point x belongs to S .

Remark 5.2. Let S be a set in \mathbb{R}^n . If x belongs to \bar{S} , then there is a sequence of points $\{x_k\}$ in S such that $x_k \rightarrow x$. To see this, note that if x is in S , then x belongs either to S or to S' . If $x \in S'$, the assertion follows from Lemma 5.1. If $x \in S$ and $x \notin S'$, we may take $x_k = x$ for all positive integers k .

Let S be a subset (proper or improper) of \mathbb{R}^n . Let F be a function with domain S and range contained in \mathbb{R}^m . We denote this by $F: S \rightarrow \mathbb{R}^m$. Thus, $F = (f_1, \dots, f_m)$, where each f_i is a real-valued function.

Let a be a limit point of S . We say that "the limit of $f(x)$ as x approaches a exists and equals L " if for every $\epsilon > 0$ there exists a $\delta > 0$, which may depend on ϵ , such that for all $x \in B_\delta(a) \cap S$, $x \neq a$, we have $f(x) \in B_\epsilon(L)$. We write

$$\lim_{x \rightarrow a} f(x) = L.$$

It is not hard to show that, if $L = (L_1, \dots, L_m)$, $\lim_{x \rightarrow a} f(x) = L$ if and only if for each $i = 1, \dots, m$, $\lim_{x \rightarrow a} f_i(x) = L_i$.

It is also not hard to show that if a limit exists it is unique and that the algebra of limits, as stated in the next theorem, holds.

Theorem 5.2. Let $\lim_{x \rightarrow a} f(x) = L$ and let $\lim_{x \rightarrow a} g(x) = M$. Then

$$\lim_{x \rightarrow a} [f(x) + g(x)]$$

exists and equals $L + M$. Also $\lim_{x \rightarrow a} [c f(x)]$ exists and equals cL .

A useful sequential criterion for the existence of a limit is now given.

Theorem 5.3. $\lim_{x \rightarrow a} f(x) = L$ if and only if for every sequence $\{x_k\}$ of points in S such that $x_k \neq a$ for all k and $x_k \rightarrow a$ it is true that $\lim_{k \rightarrow \infty} f(x_k) = L$.

Proof. Let $\lim_{x \rightarrow a} f(x) = L$ and let $\{x_k\}$ be an arbitrary sequence of points in S with $x_k \neq a$ for all k and with $x_k \rightarrow a$. Let $\epsilon > 0$ be given. Since $\lim_{x \rightarrow a} f(x) = L$, there exists a $\delta(\epsilon) > 0$ such that, for all $x \neq a$ and in $B_\delta(a) \cap S$, $f(x) \in B_\epsilon(L)$. Since $x_k \rightarrow a$, there exists a positive integer $K(\delta)$ such that whenever $k > K(\delta)$, $x_k \in B_\delta(a)$. Thus, for $k > K(\delta)$, $f(x_k) \in B_\epsilon(L)$, and so $f(x_k) \rightarrow L$.

We now prove the implication in the other direction. Suppose that $\lim_{x \rightarrow a} f(x)$ does not equal L . Then there exists an $\epsilon_0 > 0$ such that for every $\delta > 0$ there exists a point x_δ in S satisfying the conditions $x_\delta \neq a$, $x_\delta \in B_\delta(a)$,

and $f(x_j) \in \mathbb{R} \setminus \{a\}$. Take j through the sequence of values $j = 1, 2, \dots$, where k runs through the positive integers. We then get a sequence $\{x_k\}$ with $x_k \in E$, $x_k \neq a$, and $x_k \rightarrow a$, with the further property that $f(x_k) \in \mathbb{R} \setminus \{a\}$. Thus $f(x_k)$ does not converge to L , and the theorem is proved.

Let x_0 be a point of S that is also a limit point of S ; that is, $x_0 \in S \cap F$. We say that f is *continuous at x_0* if

$$\lim_{x \rightarrow x_0} f(x) = f(x_0).$$

If $x_0 \in S$ and x_0 is not a limit point of S , we will take f to be continuous at x_0 by convention.

From the corresponding property of limits it follows that f is continuous at x_0 if and only if each of the real-valued functions f_1, \dots, f_n is continuous at x_0 . It follows from Theorem 5.2 that if f and g are continuous at x_0 , then so are $f + g$ and af .

The function f is said to be *continuous on S* if it is continuous at every point of S .

Exercise 5.2. Show that if $\lim_{x \rightarrow a} f(x) = L$, then L is unique.

Exercise 5.3. Prove Theorem 5.2.

Exercise 5.4. Let f be a real-valued function defined on a set S in \mathbb{R}^n . Show that if f is continuous at a point x_0 in S and if $f(x_0) < 0$, then there exists a $\delta > 0$ such that $f(x) < 0$ for all x in $\mathcal{B}(x_0, \delta) \cap S$.

Exercise 5.5. Show that the real-valued function $f: \mathbb{R}^n \rightarrow \mathbb{R}^n$ defined by $f(x) = \|x\|$ is continuous on \mathbb{R}^n (i.e., show that the norm is a continuous function). *Hint:* Use the triangle inequality to show that, for any pair of vectors x and y , $|\|x\| - \|y\|| \leq \|x - y\|$.

6. BASIC PROPERTY OF REAL NUMBERS

A set S in \mathbb{R}^n is said to be *bounded* if there exists a positive number M such that $\|x\| \leq M$ for all $x \in S$. Another way of stating this is $S \subset \mathcal{B}(0, M)$. The number M is said to be a *bound* for the set. Note that if a set is bounded, then there are infinitely many bounds.

We now shall deal with sets in \mathbb{R}^1 only. In \mathbb{R}^1 a set S being bounded means that, for every x in S , $|x| \leq M$, or $-M \leq x \leq M$. A set S in \mathbb{R}^1 is said to be *bounded above* if there exists a real number A such that $x \leq A$ for all x in S . A set S is *bounded below* if there exists a real number B such that $x \geq B$ for all x in S . The number A is said to be an *upper bound* of S , the number B a *lower bound*.

A number U is said to be a *least upper bound* (l.u.b.), or *supremum* (sup), of a set S

- (i) if U is an upper bound of S and
- (ii) if U' is another upper bound of S and then $U' \geq U$.

If a set has a least upper bound, then the least upper bound is unique. To see this, let U_1 and U_2 be two least upper bounds of a set S . Then since U_2 is an upper bound and U_1 is a least upper bound, $U_2 \geq U_1$. Similarly, $U_1 \geq U_2$ and so $U_1 = U_2$.

Condition (ii) states that no number $A < U$ can be an upper bound of S . Therefore condition (ii) can be replaced by the equivalent statement:

- (ii') For every $\epsilon > 0$ there is a number x_ϵ in S with $U - \epsilon < x_\epsilon < U$.

Let S be the set of numbers $1 - 1/n$, $n = 1, 2, 3, \dots$. Then one is the l.u.b. of this set. Note that one does not belong to S . Now consider the set $S_0 = S \cup \{1\}$. The number 1 is the l.u.b. of this set and belongs to S_0 . The reader should be clear about the distinction between an upper bound and a supremum, or l.u.b.

A number L is said to be a *greatest lower bound* (g.l.b.), or *infimum* (inf), of a set S

- (i) if L is a lower bound of S and
- (ii) if L' is another lower bound of S and then $L' \leq L$.

Observations analogous to those made following the definition of l.u.b. hold for the definitions of g.l.b. We leave their formulation to the reader.

Exercise 61. Let $L = \text{g.l.b.}$ of a set S . Show that there exists a sequence of points $\{x_k\}$ in S such that $x_k \rightarrow L$. Show that the sequence $\{x_k\}$ can be taken to be nonincreasing, that is, $x_{k+1} \leq x_k$ for every k . Does L have to be a limit point of S ?

Exercise 62. Let S be a set in \mathbb{R}^1 . We define $-S$ to be $\{-x \mid x \in S\}$. Thus $-S$ is the set that we obtain by replacing each element x in S by the element $-x$. Show that S is bounded below if and only if $-S$ is bounded above. Show that α is the g.l.b. of S if and only if $-\alpha$ is the l.u.b. of $-S$.

We can now state the basic property of the real numbers \mathbb{R}^1 , which is sometimes called the *completeness property*.

Basic Property of the Reals

Every set $S \subset \mathbb{R}^1$ that is bounded $\left(\begin{smallmatrix} \text{above} \\ \text{below} \end{smallmatrix}\right)$ has a $\left(\begin{smallmatrix} \text{l.u.b.} \\ \text{g.l.b.} \end{smallmatrix}\right)$.

It follows from Exercise 6.2 that every set in \mathbb{R}^1 that is bounded above has a l.u.b. if and only if every set that is bounded below has a g.l.b. Thus, one statement of the completeness property can be interpreted as two equivalent statements.

Not every number system has this property. Consider the rational numbers \mathbb{Q} . They possess all of the algebraic and order properties of the reals but do not possess the completeness property. We will outline the argument showing that the rationals are not complete. For details see Rudin [1976]. Recall that a rational number is a number that can be represented as a quotient of integers and recall that $\sqrt{2}$ is not rational. (The ancient Greeks knew that $\sqrt{2}$ is not rational.) Consider the set S of rational numbers x defined by $S = \{x \mid x \text{ rational, } x^2 < 2\}$. Clearly, 2 is an upper bound. It can be shown that for any x in S there is an x' in S such that $x' > x$. Thus no element of S can be an upper bound for S , let alone a l.u.b. It can also be shown that for any rational number y not in S that is an upper bound of S there is another rational number y' that is not in S , that is, also an upper bound of S and satisfies $y' < y$. Thus no rational number not in S can be a l.u.b. of S . Hence, there is no l.u.b. of S in the rationals. The number $\sqrt{2}$ in the reals is the candidate for the title of l.u.b., but $\sqrt{2}$ is not in the system of rational numbers.

Exercise 6.3. Let A and B be two bounded sets of real numbers with $A \subset B$. Show that

$$\begin{aligned} \sup\{a \mid a \in A\} &\leq \sup\{b \mid b \in B\}, \\ \inf\{a \mid a \in A\} &\geq \inf\{b \mid b \in B\}. \end{aligned}$$

Exercise 6.4. Let A and B be two sets of real numbers.

(i) Show that if A and B are bounded above, then

$$\sup\{a + b \mid a \in A, b \in B\} = \sup\{a \mid a \in A\} + \sup\{b \mid b \in B\}.$$

(ii) Show that if A and B are bounded below, then

$$\inf\{a + b \mid a \in A, b \in B\} = \inf\{a \mid a \in A\} + \inf\{b \mid b \in B\}.$$

Exercise 6.5. Let a and b be real numbers and let $J = \{x \mid a < x < b\}$. A real-valued function f is said to have a right-hand limit ρ at a point x_0 where $a < x_0 < b$, if for every $\epsilon > 0$ there exists a $\delta(\epsilon)$ such that whenever $0 < x - x_0$

< δ_0) the inequality $|f(x) - p| < \epsilon$ holds. We shall write

$$f(x_0+) = \lim_{x \rightarrow x_0^+} f(x) = p.$$

(ii) Formulate a definition for a left-hand limit at a point x_0 , where $a < x_0 < b$. If l is the left-hand limit, we shall write

$$f(x_0-) = \lim_{x \rightarrow x_0^-} f(x) = l.$$

(iii) Show that f has a limit at a point x_0 , where $a < x_0 < b$, if and only if f has right- and left-hand limits at x_0 and these limits are equal.

Exercise 6.8. Let f be as in Exercise 6.5. A multi-valued function f defined on I is said to be nondecreasing if for any pair of points x_1 and x_2 in I with $x_2 > x_1$ we have $f(x_1) \leq f(x_2)$.

- Formulate the definition of nonincreasing function.
- Show that if f is a nondecreasing function on I , then f has a right-hand limit at every point x_0 where $a < x_0 < b$, and left-hand limit at every point $a < x_0 < b$.
- Show that if I is an open interval $(a, b) = \{x \mid a < x < b\}$ and if f is nondecreasing and bounded below on I , then $\lim_{x \rightarrow a^+} f(x)$ exists. If f is bounded above, then $\lim_{x \rightarrow b^-} f(x)$ exists.

7. COMPACTNESS

A property that is of great importance in many contexts is that of compactness. The definition that we shall give is one that is applicable in very general contexts as well as in \mathbb{R}^n . We do this so that the reader who takes more advanced courses in analysis will not have to unlearn anything. After presenting the definition of compactness, we will state, without proof, two theorems that give necessary and sufficient conditions for a set X in \mathbb{R}^n to be compact. We frequently shall use the properties stated in these theorems when working with compact sets.

Let S be a set in \mathbb{R}^n . A collection of open sets $\{O_\alpha\}_{\alpha \in A}$ is said to be an open cover of S if every $x \in S$ is contained in some O_α . A set X in \mathbb{R}^n is said to be compact if for every open cover $\{O_\alpha\}_{\alpha \in A}$ of X there is a finite collection of sets O_1, \dots, O_m from the original collection $\{O_\alpha\}_{\alpha \in A}$ such that the finite collection O_1, \dots, O_m is also an open cover of X . In mathematical jargon this property is stated as "Every open cover has a finite subcover."

To illustrate, consider the set S in \mathbb{R}^1 defined by $S = \{x \mid 0 < x < 1\}$. For each positive integer k , let $O_k = (0, 1 - 1/k) = \{x \mid 0 < x < 1 - 1/k\}$. Then $\{O_k\}_{k=1}^\infty$ is an open cover of S , but no finite collection of sets in the cover will

cover \mathcal{S} . Thus, $(0, 1)$ is not compact. It is difficult to give a meaningful example of a compact set based on the definitions, since we must show that every open cover has a finite subcover.

THEOREM 7.1. In \mathbb{R}^n a set S is compact if and only if every sequence $\{x_k\}$ of points in S has a subsequence $\{x_{k_j}\}$ that converges to a point in S (Bolzano-Weierstrass property).

THEOREM 7.2. In \mathbb{R}^n a set S is compact if and only if S is closed and bounded.

We refer the reader to Rudin and Rudin ([1991] or Rudin [1976]) for proofs of these theorems.

We emphasize that the criteria for compactness given in Theorems 7.1 and 7.2 are not necessarily valid in contexts more general than \mathbb{R}^n .

Theorem 7.2 provides an easily applied criterion for compactness. Using Theorem 7.2, we see immediately that $S = (0, 1) = \{x: 0 < x < 1\}$ is not compact in \mathbb{R}^1 . (S is not closed.) We also see that $S_1 = [0, 1] = \{x: 0 \leq x \leq 1\}$ is compact, since it is closed and bounded.

We conclude with another characterization of compactness that is valid in more general contexts than \mathbb{R}^n . A set S is said to have the *finite-intersection property* if for every collection of closed sets $\{F_\alpha\}$ such that $\bigcap F_\alpha \cap S \neq \emptyset$ there is a finite subcollection $F_{\alpha_1}, \dots, F_{\alpha_n}$ such that $\bigcap_{i=1}^n F_{\alpha_i} \cap S \neq \emptyset$. We shall illustrate the definition by exhibiting a set that fails to have the finite-intersection property. The set $(0, \infty) = \{x: 0 < x < \infty\}$ in \mathbb{R}^1 does not have the finite-intersection property. Take $F_n = [1, 1/n] \cap (0, \infty)$, $n = 1, 2, 3, \dots$. Then $\bigcap_{n=1}^{\infty} F_n = [1, \infty) \cap (0, \infty)$, yet the intersection of any finite subcollection of the F_n has nonempty intersection with $(0, 1)$ and so is not contained in $(0, \infty)$.

THEOREM 7.3. A set S is compact if and only if it has the finite-intersection property.

The proof is an exercise in the use of Lemma 3.1. Let S have the finite-intersection property. Let (\mathcal{O}_α) be an open cover of S . Then $S \subset \bigcup \mathcal{O}_\alpha = \complement \left(\bigcap \complement \mathcal{O}_\alpha \right)$, and so $\complement \left(\bigcap \complement \mathcal{O}_\alpha \right) \cap S = \bigcap \complement \mathcal{O}_\alpha$. Each set $\complement \mathcal{O}_\alpha$ is closed, so there exist a finite number of sets $\mathcal{O}_{\alpha_1}, \dots, \mathcal{O}_{\alpha_n}$ such that $\complement \left(\bigcap_{i=1}^n \complement \mathcal{O}_{\alpha_i} \right) \supset S$. Hence $S \subset \complement \left(\bigcap_{i=1}^n \complement \mathcal{O}_{\alpha_i} \right) = \bigcup_{i=1}^n \mathcal{O}_{\alpha_i}$, and so $\{(\mathcal{O}_\alpha)\}$ has a finite subcover $\mathcal{O}_{\alpha_1}, \dots, \mathcal{O}_{\alpha_n}$. Thus we have shown that S is compact. Now let S be compact. Let $\{F_\alpha\}$ be a collection of closed sets such that $\bigcap F_\alpha \cap S \neq \emptyset$. Then $\complement \left(\bigcap F_\alpha \right) \cap S = \emptyset$, and so $\bigcap \complement (F_\alpha) \cap S = \emptyset$. Thus $(\complement F_\alpha)$ is an open cover of S . Since S is compact, there exist a finite subcollection of sets $F_{\alpha_1}, \dots, F_{\alpha_n}$ such that $\bigcap_{i=1}^n \complement (F_{\alpha_i}) \cap S = \emptyset$. Therefore $\bigcap_{i=1}^n F_{\alpha_i} \cap S = \complement \left(\bigcap_{i=1}^n \complement (F_{\alpha_i}) \right) \cap S \neq \emptyset$, and thus S has the finite-intersection property.

Remark 7.4. Note that the finite-intersection property can be stated in the following equivalent form. Every collection of closed sets $\{F_\alpha\}$ such that $\complement \left(\bigcap F_\alpha \right) \cap S = \emptyset$ has a finite subcollection $F_{\alpha_1}, \dots, F_{\alpha_n}$ such that $\bigcap_{i=1}^n F_{\alpha_i} \cap S \neq \emptyset$.

8. EQUIVALENT NORMS AND CARTESIAN PRODUCTS

A norm on a vector space \mathbb{R}^n is a function $\|\cdot\|$ that assigns a real number $\|\mathbf{x}\|$ to every vector \mathbf{x} in the space and has the properties (1)–(3) and (5) listed in Section 3, with (4) replaced by $\|\mathbf{x}\| \geq 0$. Once a norm has been defined, one can define a function ρ by the formula $\rho(\mathbf{x}, \mathbf{w}) = \|\mathbf{x} - \mathbf{w}\|$. Properties (1)–(3) and (5) of the norm imply that the function ρ satisfies the properties (1)–(3) of a metric, and therefore is a metric, or distance function. Thus a distance has been defined, open sets are defined in terms of the distance, as before. Convergence is defined in terms of open sets.

To illustrate these ideas, we give an example of another norm in \mathbb{R}^n , defined by

$$\|\mathbf{x}\| = \|\mathbf{x}\|_\infty = \max\{|x_1|, \dots, |x_n|\}.$$

The corresponding distance is defined by

$$\rho(\mathbf{x}, \mathbf{y}) = \|\mathbf{x} - \mathbf{y}\|_\infty = \max\{|x_1 - y_1|, \dots, |x_n - y_n|\}.$$

Since a norm determines the family of open sets, and since the family of open sets determines convergence, the question arises, "Given two norms, ν_1 and ν_2 , do they determine the same family of open sets?" The answer, which is not difficult to show, is, "If there exist positive constants m and M such that for all \mathbf{x} in \mathbb{R}^n ,

$$m\nu_2(\mathbf{x}) \leq \nu_1(\mathbf{x}) \leq M\nu_2(\mathbf{x}), \quad (1)$$

then a set is open under the ν_1 norm if and only if it is open under the norm ν_2 ." Two norms ν_1 and ν_2 that satisfy (1) are said to be *equivalent norms*.

An extremely important fact is that in \mathbb{R}^n all norms are equivalent. For a proof see Fleming [1977]. In optimization theory, and in other areas, it may happen that the analysis of a problem or the development of a computational algorithm is more convenient in one norm than in another. The equivalence of the norms guarantees that the convergence of an algorithm is intrinsic, that is, it is independent of the norm.

Let A_1, \dots, A_m be m sets with $A_i \subseteq \mathbb{R}^{n_i}$, $i = 1, \dots, m$. By the *cartesian product* of these sets, denoted by

$$\prod_{i=1}^m A_i \quad \text{or} \quad A_1 \times A_2 \times \cdots \times A_m,$$

we mean the set A in $\mathbb{R}^{n_1 + n_2 + \cdots + n_m}$ consisting of all possible points in $\mathbb{R}^{n_1 + n_2 + \cdots + n_m}$ obtained by taking the first n_1 components from a point in A_1 , the next n_2 components from a point in A_2 , the next n_3 components from a point in A_3 ,

and so forth, until the last \mathbf{x}_m components, which are obtained by taking a point in A_m . For example, \mathbb{R}^n can be considered to be the cartesian product of \mathbb{R}^1 taken n times with itself:

$$\mathbb{R}^n = \underbrace{\mathbb{R}^1 \times \mathbb{R}^1 \times \cdots \times \mathbb{R}^1}_n.$$

If A_1 is the set in \mathbb{R}^2 consisting of points on the circumference of the circle with center at the origin and radius 1 and if $A_2 = [0, 1]$ in \mathbb{R}^1 , then $A_1 \times A_2$ is the set in \mathbb{R}^3 that is the lateral surface of a cylinder of height 1 whose base is the closed disk in the (x_1, x_2) plane with center at the origin and radius equal to 1.

Let $X = \mathbb{R}^n = \cdots = \mathbb{R}^{2m}$. Let \mathbf{x} be a vector in X . Then $\mathbf{x} = (x_1, \dots, x_{2m})$ where $x_i = (x_{i1}, \dots, x_{i2})$, $i = 1, \dots, m$. If for vectors \mathbf{x} and \mathbf{y} in X and scalars α in \mathbb{R}^1 we define

$$(\mathbf{x} + \mathbf{y}) = (x_1 + y_1, \dots, x_m + y_m), \quad \alpha \mathbf{x} = (\alpha x_1, \dots, \alpha x_m)$$

and let $\mathbf{e} = \mathbf{e}_1 + \mathbf{e}_2 + \cdots + \mathbf{e}_{2m}$, then X can be identified with the usual vector space \mathbb{R}^n . The euclidean norm in X would then be given by

$$\|\mathbf{x}\| = \left[(x_1)^2 + \cdots + (x_m)^2 \right]^{1/2} = \left[\sum_{j=1}^m \sum_{i=1}^{2} (x_{ji})^2 \right]^{1/2}. \quad (2)$$

On the other hand, by keeping with the definitions used in more general contexts, the norm on X is defined by

$$\|\mathbf{x}\|_1 = \sum_{i=1}^m \|x_i\|, \quad (3)$$

where $\|x_i\|$ is the euclidean norm of x_i in \mathbb{R}^2 , defined by $\|x_i\| = (x_{i1}^2 + x_{i2}^2)^{1/2} = \left[\sum_{j=1}^2 (x_{ij})^2 \right]^{1/2}$. Since all norms on \mathbb{R}^2 are equivalent, the norm defined by (3) is equivalent to the euclidean norm defined by (2).

The following lemma is an immediate consequence of the definition of $\|\cdot\|_1$.

LEMMA 8.1. A sequence $\mathbf{x}_n = (x_{n1}, \dots, x_{nm})$ of points in X converges to a point $\mathbf{x}_0 = (x_{01}, \dots, x_{0m})$ of X in the metric determined by $\|\cdot\|_1$ if and only if for each $j = 1, \dots, m$, the sequence $\{x_{nj}\}$ converges to x_{0j} in the euclidean norm on \mathbb{R}^2 .

The next theorem can be proved by using Lemma 8.1 and Theorem 7.1 and by selecting appropriate subsequences.

THEOREM 8.1. Let A_1, \dots, A_m be a set with $A_j = \mathbb{R}^n$ and for each A_j be compact, $j = 1, \dots, m$. Then $\bigcap_{\epsilon > 0} A_\epsilon$ is compact in \mathbb{R}^n , where $\epsilon = \epsilon_1 + \epsilon_2 + \cdots + \epsilon_m$.

Exercise 8.1. (a) In \mathbb{R}^2 sketch the set $B_{\infty}(\mathbf{0}, 1) = \{\mathbf{x} : \|\mathbf{x}\|_{\infty} < 1\}$.

(b) Show that $\|\cdot\|_{\infty}$ is a norm (i.e., that it satisfies (1)–(3) and (5) of Section 2).

(c) Can you draw a sketch in \mathbb{R}^2 that would indicate that a set is open under the metric $d(\mathbf{x}, \mathbf{y}) = \|\mathbf{x} - \mathbf{y}\|$ if and only if it is open under the metric $\rho(\mathbf{x}, \mathbf{y}) = \|\mathbf{x} - \mathbf{y}\|_{\infty}$?

(d) Is it true that $\|\mathbf{x} + \mathbf{y}\|_{\infty} = \|\mathbf{x}\|_{\infty} + \|\mathbf{y}\|_{\infty}$ if and only if $\mathbf{y} = c\mathbf{x}$, $c > 0$?

Exercise 8.2. (a) Show that the function $\|\cdot\|_1$, defined by (3) is a norm on \mathbb{R}^n that is $\|\cdot\|_1$, satisfies (1)–(3) and (5) of Section 2.

(b) If $X = \mathbb{R}^n = \mathbb{R}^n \times \mathbb{R}^n$, sketch the set $B_0(\mathbf{0}, 1) = \{\mathbf{x} : \|\mathbf{x}\|_1 \leq 1\}$.

Exercise 8.3. Without using the fact that all norms in \mathbb{R}^n are equivalent, show that $\|\cdot\|$ and $\|\cdot\|_1$, defined by (2) and (3), respectively, are equivalent norms on \mathbb{R}^n .

Exercise 8.4. Show that the following functions are continuous:

(a) $f: \mathbb{R}^2 \times \mathbb{R}^2 \rightarrow \mathbb{R}^2$, where $f(\mathbf{x}, \mathbf{y}) = (\mathbf{x} + \mathbf{y})$ (addition is continuous);

(b) $f: \mathbb{R}^2 \times \mathbb{R}^2 \rightarrow \mathbb{R}^2$, where $f(\lambda, \mathbf{x}) = \lambda\mathbf{x}$ (scalar multiplication is continuous); and

(c) $f: \mathbb{R}^2 \times \mathbb{R}^2 \rightarrow \mathbb{R}^2$, where $f(\mathbf{x}, \mathbf{y}) = \langle \mathbf{x}, \mathbf{y} \rangle$ (scalar product is continuous).

Exercise 8.5. Prove Theorem 8.1.

Exercise 8.6. Show that if \mathcal{Q} is a positive-definite real symmetric matrix, then $\rho(\mathbf{x}) = \langle \mathbf{x}, \mathcal{Q}\mathbf{x} \rangle^{1/2}$ defines a norm on \mathbb{R}^n . *Hint:* Look at Section 1.

9. FUNDAMENTAL EXISTENCE THEOREM

The basic problem in optimization theory is as follows: Given a set S (not necessarily in \mathbb{R}^n) and a real-valued function f defined on S , does there exist an element x_0 in S such that $f(x_0) \leq f(x)$ for all x in S , and if so find it. The problem as stated is a minimization problem. The point x_0 is said to be a *minimizer* for f , or to *minimize* f . We also say that f attains a *minimum* on S at x_0 . The problems of “maximizing f over S ” or finding an x^0 in S such that $f(x^0) \geq f(x)$ for all x , can be reduced to the problem of minimizing the function $-f$ over S . To see this, observe that if $-f(x^0) \leq -f(x)$ for all x in S , then $f(x^0) \geq f(x)$ for all x in S , and conversely.

The first question that we address is, “Does a minimizer exist?” To point out some of the difficulties here, we look at some very simple examples in \mathbb{R}^1 .

Let $S = \{x \mid x \geq 1\}$ and let $f(x) = 1/x$. It is clear that $\inf\{f(x) \mid x \geq 1\} = 0$, yet there is no $x_0 \in S$ such that $f(x_0) = 0$. Thus a minimizer does not exist, or f does not attain its minimum on S . As a second example, let $S = [1, 2]$ and let $f(x) = 1/x$. Now, $\inf\{f(x) \mid 1 \leq x \leq 2\} = \frac{1}{2}$, yet for no x_0 in S do we have $f(x_0) = \frac{1}{2}$. Of course, if we modify the last example by taking $S = [1, 2] = \{x \mid 1 \leq x \leq 2\}$, then $x_0 = 2$ is the minimizer. For future reference, note that if $S = [1, 2] = \{x \mid 1 < x < 2\}$, then $x_0 = 2$ is still the minimizer.

We now state and prove the fundamental existence theorem for minimization of functions with domains in \mathbb{R}^n .

THEOREM 8.1. *Let S be a compact set in \mathbb{R}^n and let f be a real-valued continuous function defined on S . The f attains a maximum and a minimum on S .*

Before taking up the proof, we point out that the theorem gives a sufficient condition for the existence of a maximizer and minimizer. It is not a necessary condition, as the example in \mathbb{R}^1 , $S = [1, 2]$, $f(x) = 1/x$, shows. In more advanced courses the reader may encounter the notions of upper semicontinuity and lower semicontinuity and learn that continuity can be replaced by upper semicontinuity to ensure the existence of maximizers and that continuity can be replaced by lower semicontinuity to ensure the existence of minimizers. Finally, the reader should study the proof of this theorem very carefully and be sure to understand it. The proof of the theorem in \mathbb{R}^n is the model for the proofs of the existence of minimizers in more general contexts than \mathbb{R}^n .

Proof. If f is continuous on S , then so is $-f$. If $-f$ attains a minimum on S , then f attains a maximum on S . Thus we need only prove that f attains a minimum.

We first prove that the set of values that f assumes on S is bounded below, that is, that the set $\mathcal{R} = \{y \mid y = f(x), x \in S\}$ is bounded below. To show this, we suppose that \mathcal{R} were not bounded below. Then for each positive integer k there would exist a point \mathbf{x}_k in S such that $f(\mathbf{x}_k) < -k$. Since S is compact, the sequence $\{\mathbf{x}_k\}$ of points in S has a subsequence $\{\mathbf{x}_{k_j}\}$ that converges to a point \mathbf{x}_0 in S . The continuity of f implies that $\lim_{j \rightarrow \infty} f(\mathbf{x}_{k_j}) = f(\mathbf{x}_0)$. Thus, by the definition of limit (taking $\epsilon = 1$) we have that there exists a positive integer K such that, for $j > K$, $f(\mathbf{x}_{k_j}) > f(\mathbf{x}_0) - 1$. This contradicts the defining property of $\{\mathbf{x}_k\}$, namely that, for each k , $f(\mathbf{x}_k) = -k$.

Since the set \mathcal{R} is bounded below, it has an infimum, say L . Therefore, by definition of infimum, for each positive integer k there exists a point \mathbf{x}_k in S such that

$$L \leq f(\mathbf{x}_k) < L + \frac{1}{k}. \quad (1)$$

Since S is compact, the sequence $\{\mathbf{x}_k\}$ has a subsequence $\{\mathbf{x}_{k_j}\}$ that converges

to a point \mathbf{x}_0 in S . The continuity of f implies that $\lim_{\mathbf{x} \rightarrow \mathbf{x}_0} f(\mathbf{x}) = f(\mathbf{x}_0)$. Therefore, if we replace δ by δ_1 in (1) and let $\epsilon \rightarrow \infty$, we get that $L \leq f(\mathbf{x}_0) \leq L$. Hence

$$f(\mathbf{x}_0) = L \leq f(\mathbf{x}) \quad \text{for all } \mathbf{x} \in S.$$

This proves the theorem.

If the reader understands the proof of Theorem 9.1 and understands the use of compactness, he or she should have no trouble in proving the following theorem. Actually, Theorem 9.1 is a corollary of Theorem 9.2.

Theorem 9.2. Let S be a compact subset of \mathbb{R}^n and let $F = (f_1, \dots, f_m)$ be a continuous function defined on S with range in \mathbb{R}^m . Then the set

$$F(S) = \{ \mathbf{y} \in \mathbb{R}^m : \mathbf{y} = f(\mathbf{x}), \mathbf{x} \in S \}$$

is compact.

In mathematical jargon this result is stated as follows: The continuous image of a compact set is compact.

Exercise 9.1. Prove Theorem 9.2.

Exercise 9.2. Let A be a real $n \times n$ symmetric matrix with entries a_{ij} and let $Q(\mathbf{x})$ be the quadratic form

$$Q(\mathbf{x}) = \mathbf{x}^T A \mathbf{x} = Q(\mathbf{x}, A\mathbf{x}) = \sum_{i,j=1}^n a_{ij} x_i x_j.$$

Show that there exists a constant $M > 0$ such that, for all \mathbf{x} , $Q(\mathbf{x}) \leq M \|\mathbf{x}\|^2$. Show that if $Q(\mathbf{x}) > 0$ for all $\mathbf{x} \neq \mathbf{0}$, then there exists a constant $m > 0$ such that, for all \mathbf{x} , $Q(\mathbf{x}) \geq m \|\mathbf{x}\|^2$. *Hint:* First show that the conclusion is true for $n=2$, $1) = \mathbf{x}^T (\mathbf{x}) = x_1^2$.

Exercise 9.3. Let f be a real-valued function defined on \mathbb{R}^n such that

$$\lim_{\|\mathbf{x}\| \rightarrow \infty} f(\mathbf{x}) = +\infty.$$

By this is meant that for each $M > 0$ there exists an $R > 0$ such that if $\|\mathbf{x}\| \geq R$, then $f(\mathbf{x}) > M$. Such a function f is said to be coercive. Show that if f is continuous on \mathbb{R}^n and is coercive, then f attains a minimum on \mathbb{R}^n .

18. LINEAR TRANSFORMATIONS

A linear transformation from \mathbb{R}^n to \mathbb{R}^m is a function T with domain \mathbb{R}^n and range contained in \mathbb{R}^m such that for every x, y in \mathbb{R}^n and every pair of scalars α, β

$$T(\alpha x + \beta y) = \alpha T(x) + \beta T(y).$$

Example 18.1. $T: \mathbb{R}^2 \rightarrow \mathbb{R}^2$: Rotates each vector in \mathbb{R}^2 through an angle θ .

Example 18.2. $T: \mathbb{R}^3 \rightarrow \mathbb{R}^3$ defined by

$$T(x_1, x_2, x_3) = (x_1, x_2, 0).$$

This is the projection of \mathbb{R}^3 onto the x_1, x_2 plane.

Example 18.3. Let A be an $m \times n$ matrix. Define

$$Tx = Ax. \quad (1)$$

In the event that the range space is \mathbb{R}^1 , that is, we have a linear transformation from \mathbb{R}^n to \mathbb{R}^1 , we call the transformation a linear functional. The following is an example of a linear functional L . Let a be a fixed vector in \mathbb{R}^n . Then

$$L(x) = (a, x) \quad (2)$$

defines a linear functional.

We now show that (1) and (2) are essentially the only examples of a linear transformation and linear functional, respectively. Recall that the standard basis in \mathbb{R}^n is the basis consisting of the vectors e_1, \dots, e_n , where $e_j = (0, 0, \dots, 0, 1, 0, \dots, 0)$ and 1 is the j th entry.

THEOREM 18.1. Let T be a linear transformation from \mathbb{R}^n to \mathbb{R}^m . Then relative to the standard basis in \mathbb{R}^n and \mathbb{R}^m there exists an $m \times n$ matrix A such that $Tx = Ax$.

Let $x \in \mathbb{R}^n$, then $x = \sum_{j=1}^n x_j e_j$. Hence

$$Tx = T\left(\sum_{j=1}^n x_j e_j\right) = \sum_{j=1}^n x_j T(e_j). \quad (3)$$

For each $j = 1, \dots, n$, $T e_j$ is a vector in \mathbb{R}^m , so if e_1^m, \dots, e_m^m is the standard basis in \mathbb{R}^m ,

$$T e_j = \sum_{i=1}^m a_{ij} e_i^m. \quad (4)$$

Substituting (4) into (3) and changing the order of summation give

$$T\mathbf{x} = \sum_{j=1}^n a_j i \left(\sum_{i=1}^n a_{ij} x_i^j \right) = \sum_{i=1}^n \left(\sum_{j=1}^n a_{ij} x_j \right) x_i^i.$$

Therefore if A is the $n \times n$ matrix whose j th column is $(a_{1j}, \dots, a_{nj})^T$, then $T\mathbf{x}$ expressed as an n -vector $(y_1, \dots, y_n)^T$ is given by

$$\begin{pmatrix} y_1 \\ \vdots \\ y_n \end{pmatrix} = T\mathbf{x} = A\mathbf{x}.$$

From (4) we get that the j th column of A gives the coordinates of $T\mathbf{e}_j$ relative to $\mathbf{e}_1^i, \dots, \mathbf{e}_n^i$.

In the case of a linear functional L , the matrix A reduces to a row vector \mathbf{a} in \mathbb{R}^n , so that

$$L(\mathbf{x}) = (\mathbf{a}, \mathbf{x}).$$

Exercise 10.1. Show that a linear functional L is a continuous mapping from \mathbb{R}^n to \mathbb{R} .

11. DIFFERENTIATION IN \mathbb{R}^n

Let D be an open interval in \mathbb{R}^1 and let f be real-valued function defined on D . The function f is said to have a derivative or to be differentiable at a point x_0 in D if

$$\lim_{h \rightarrow 0} \frac{f(x_0 + h) - f(x_0)}{h} \quad (1)$$

exists. The limit is called the derivative of f at x_0 and is denoted by $f'(x_0)$.

The question now arises as to how one generalizes this concept to real-valued functions defined on open sets D in \mathbb{R}^n and to functions defined on open sets in \mathbb{R}^n with range in \mathbb{R}^m . In elementary calculus the notion of partial derivative for real-valued functions defined on open sets D in \mathbb{R}^n is introduced. For each $i = 1, \dots, n$ the partial derivative with respect to x_i at a point \mathbf{x}_0 is defined by

$$\frac{\partial f}{\partial x_i}(\mathbf{x}_0) = \lim_{h \rightarrow 0} \frac{f(x_0^1, \dots, x_0^{i-1}, x_0^i + h, x_0^{i+1}, \dots, x_0^n) - f(x_0^1, \dots, x_0^n)}{h} \quad (2)$$

provided the limit on the right exists. It turns out that the notion of partial derivative is not the correct generalization of the notion of derivative.

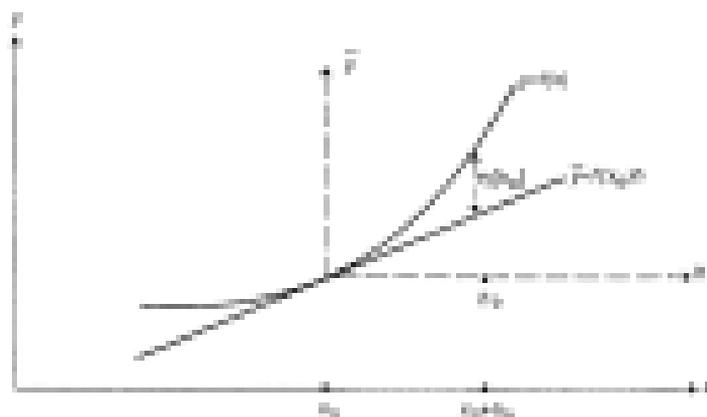


Figure 1.1.

To motivate the correct generalization of the notion of derivative to functions with domain and range in spaces of dimension greater than 1, we examine the notion of derivative for functions f defined on an open interval D in \mathbb{R}^1 with range in \mathbb{R}^1 .

We rewrite (1) as

$$f(x_0 + h) = \frac{f(x_0 + h) - f(x_0)}{h} + \phi(h),$$

where $\phi(h) \rightarrow 0$ as $h \rightarrow 0$. This in turn can be rewritten as

$$f(x_0 + h) - f(x_0) = f'(x_0)h + \phi(h), \quad (2)$$

where $\phi(h)/h \rightarrow 0$ as $h \rightarrow 0$. The term $f'(x_0)h$ is a linear approximation to f near x_0 and h is a "good approximation" for small h in the sense that $\phi(h)/h \rightarrow 0$ as $h \rightarrow 0$. See Figure 1.1.

If we consider h to be an arbitrary real number, then $f'(x_0)h$ defines a linear functional L on \mathbb{R}^1 by the formula $L(h) = f'(x_0)h$. Thus if f is differentiable at x_0 , there exists a linear functional (or linear transformation) on \mathbb{R}^1 such that

$$f(x_0 + h) - f(x_0) = L(h) + \phi(h), \quad (3)$$

where $\phi(h)/h \rightarrow 0$ as $h \rightarrow 0$. Conversely, let there exist a linear functional L on \mathbb{R}^1 such that (3) holds. Then $L(h) = ah$ for some real number a , and we may write

$$f(x_0 + h) - f(x_0) = ah + \phi(h).$$

If we divide by $h \neq 0$ and then let $h \rightarrow 0$, we get that $f'(x_0)$ exists and equals α . Thus we could have used (4) to define the notion of derivative and could have defined the derivative to be the linear functional L , which in this case is determined by the number α . The reader might feel that this is a very convoluted way of defining the notion of derivative. It has the merit, however, of being the definition that generalizes correctly to functions with domain in \mathbb{R}^n and range in \mathbb{R}^m for any $n \geq 1$ and any $m \geq 1$.

Let D be an open set in \mathbb{R}^n , $n \geq 1$, and let $F = (f_1, \dots, f_m)$ be a function defined on D with range in \mathbb{R}^m , $m \geq 1$. The function F is said to be differentiable at a point x_0 in D , or to be Fréchet differentiable at x_0 , if there exists a linear transformation $T(x_0)$ from \mathbb{R}^n to \mathbb{R}^m such that for all h in \mathbb{R}^n with h sufficiently small

$$F(x_0 + h) - F(x_0) = T(x_0)h + o(h), \quad (5)$$

where $\|o(h)\|/\|h\| \rightarrow 0$ as $\|h\| \rightarrow 0$.

We now represent $T(x_0)$ as a matrix relative to the standard bases e_1, \dots, e_n in \mathbb{R}^n and e_1^*, \dots, e_m^* in \mathbb{R}^m . Let ϵ be a real number and let $h = \epsilon e_j$ for a fixed j in the set $1, \dots, n$. Then $\epsilon e_j = (\epsilon, \dots, \epsilon, \dots, \epsilon)$, where the ϵ occurs in the j th component. From (5) we have

$$F(x_0 + \epsilon e_j) - F(x_0) = T(x_0)(\epsilon e_j) + o(\epsilon e_j).$$

Hence for $\epsilon \neq 0$

$$\frac{F(x_0 + \epsilon e_j) - F(x_0)}{\epsilon} = T(x_0)e_j + \frac{o(\epsilon e_j)}{\epsilon}.$$

Since $\|o(\epsilon e_j)\| = o$ and for $j = 1, \dots, n$,

$$\begin{aligned} f_j(x_0 + \epsilon e_j) - f_j(x_0) &= f_j(x_0, \dots, x_0, \epsilon, x_0, \dots, x_0) \\ &\quad - f_j(x_0, \dots, x_0, 0). \end{aligned}$$

It follows on letting $\epsilon \rightarrow 0$ that for $j = 1, \dots, n$ the partial derivatives $f_j'(x_0) = \partial f_j / \partial x_j$ exist and

$$T(x_0)e_j = \begin{pmatrix} \frac{\partial f_1}{\partial x_j}(x_0) \\ \frac{\partial f_2}{\partial x_j}(x_0) \\ \vdots \\ \frac{\partial f_m}{\partial x_j}(x_0) \end{pmatrix}.$$

Since the coordinates of $T(\mathbf{x}_0, \mathbf{e}_j)$, relative to the standard basis in \mathbb{R}^n are given by the j th column of the matrix representing $T(\mathbf{x}_0)$, the matrix representing $T(\mathbf{x}_0)$ is the matrix

$$\begin{pmatrix} \frac{\partial \xi}{\partial x_1}(\mathbf{x}_0) \\ \vdots \\ \frac{\partial \xi}{\partial x_n}(\mathbf{x}_0) \end{pmatrix}.$$

If we define \mathbf{v}_i to be the row vector

$$\mathbf{v}_i = \left(\frac{\partial \xi}{\partial x_1}, \dots, \frac{\partial \xi}{\partial x_n} \right), \quad i = 1, \dots, n,$$

and \mathbf{V} to be the matrix

$$\begin{pmatrix} \mathbf{v}_1 \\ \vdots \\ \mathbf{v}_n \end{pmatrix},$$

then we have shown that

$$T(\mathbf{x}_0) = \mathbf{V}\mathbf{U}(\mathbf{x}_0).$$

The matrix $\mathbf{V}\mathbf{U}(\mathbf{x}_0)$ is called the *Jacobian matrix* of f at \mathbf{x}_0 . It is sometimes also written as $J(f; \mathbf{x}_0)$. Note that we have shown that if f is differentiable at \mathbf{x}_0 , then each component function ξ has all first partial derivatives existing at \mathbf{x}_0 .

THEOREM 11.1. Let D be an open set in \mathbb{R}^n and let f be a function defined on D with range in \mathbb{R}^m . If f is differentiable at a point \mathbf{x}_0 in D , then f is continuous at \mathbf{x}_0 .

This follows immediately from (5) on letting $\mathbf{h} \rightarrow \mathbf{0}$.

We will now show that the existence of partial derivatives at a point does not imply differentiability. We shall consider real-valued functions f defined on \mathbb{R}^2 . For typographical convenience we shall represent points in the plane as (x, y) instead of (x_1, x_2) .

Example 11.1. Let

$$f(x, y) = \begin{cases} \frac{xy^2}{x^2 + y^2}, & (x, y) \neq (0, 0), \\ 0, & (x, y) = (0, 0). \end{cases}$$

Since $f(x, 0) = 0$ for all x , $f'_x(0, 0)$ exists and equals 0. Similarly $f'_y(0, 0) = 0$. Along the line $y = kx$, the function f becomes

$$f(x, kx) = \frac{kx^3}{x^2 + k^2x^2} = \frac{k}{1 + k^2}, \quad x \neq 0,$$

and so $\lim_{(x,y) \rightarrow (0,0)} f(x, y) = 0$ ($1 + 1^2$). Thus $\lim_{(x,y) \rightarrow (0,0)} f(x, y)$ does not exist. Therefore f is not continuous at the origin and so by Theorem 11.1 cannot be differentiable at the origin.

One might now ask, what if the function is continuous at a point and has partial derivatives at the point? Must the function be differentiable at the point? The answer is no, as the next example shows.

Example 11.2. Let

$$f(x, y) = \begin{cases} \frac{xy}{\sqrt{x^2 + y^2}} & (x, y) \neq (0, 0), \\ 0, & (x, y) = 0. \end{cases}$$

As in Example 11.1, $\partial f/\partial x$ and $\partial f/\partial y$ exist at the origin and are both equal to zero. Note that

$$0 \leq (x) = (x)^2 = x^2 = 2x|x| + (x)^2$$

implies that $x^2 + y^2 \geq 2|x||y|$. Hence for $(x, y) \neq 0$

$$|f(x, y)| = \frac{|xy|}{\sqrt{x^2 + y^2}} \leq \frac{1}{2} \frac{(x^2 + y^2)}{\sqrt{x^2 + y^2}} = \frac{1}{2} \sqrt{x^2 + y^2}.$$

Therefore $\lim_{(x,y) \rightarrow (0,0)} f(x, y) = 0$, and so f is continuous at $(0, 0)$.

If f were differentiable at $(0, 0)$, then for any $\mathbf{h} = (h_1, h_2)$

$$f(\mathbf{0} + \mathbf{h}) = f(\mathbf{0}) + (Df)(\mathbf{0}) \cdot \mathbf{h} + \eta(\mathbf{h}),$$

where $\eta(\mathbf{h})/|\mathbf{h}| \rightarrow 0$ as $\mathbf{h} \rightarrow \mathbf{0}$. Since $(Df)(\mathbf{0}) = (\partial f/\partial x)(\mathbf{0}), (\partial f/\partial y)(\mathbf{0}) = (0, 0)$ and since $f(\mathbf{0}) = 0$, we would have

$$f(\mathbf{h}) = \eta(\mathbf{h}), \quad \mathbf{h} \neq \mathbf{0}.$$

From the definition of f we get, writing (h_1, h_2) in place of (x, y) ,

$$\eta(\mathbf{h}) = \frac{h_1 h_2}{\sqrt{(h_1)^2 + (h_2)^2}}.$$

We require that $\eta(\mathbf{h})/|\mathbf{h}| \rightarrow 0$ as $\mathbf{h} \rightarrow \mathbf{0}$. But this is not true since

$$\frac{\eta(\mathbf{h})}{|\mathbf{h}|} = \frac{h_1 h_2}{(h_1)^2 + (h_2)^2},$$

and by Example 11.1, the right-hand side does not have a limit as $\mathbf{h} \rightarrow \mathbf{0}$.

A real-valued function f defined on an open set D in \mathbb{R}^n is said to be of class C^k on D if all of the partial derivatives up to and including those of order k exist and are continuous on D . A function \mathbf{F} defined on D with range in \mathbb{R}^m , $m > 1$, is said to be of class C^k on D if each of its component functions f_i , $i = 1, \dots, m$, is of class C^k on D .

We saw that if ϕ is a real-valued function defined on an open interval $D =]a - \epsilon, a + \epsilon[$ in \mathbb{R} , then $\phi(x_0 + h) - \phi(x_0)$ for small h at any point x_0 in D at which ϕ is differentiable. If ϕ is of class C^2 in D , then a better approximation can be obtained using Taylor's theorem from elementary calculus, one form of which states that for $x_0 \in D$ and h such that $x_0 + h \in D$

$$\phi(x_0 + h) - \phi(x_0) = \sum_{j=1}^{k-1} \frac{\phi^{(j)}(x_0)}{j!} h^j + \frac{\phi^{(k)}(\xi) h^k}{k!}, \quad (6)$$

where $\phi^{(j)}$ denotes the j th derivative of ϕ and ξ is a point lying between x_0 and $x_0 + h$. If $k = 1$, then the summation term in (6) is absent and we have a restatement of the mean-value theorem. Since $\phi^{(k)}$ is continuous,

$$\phi^{(k)}(\xi) = \phi^{(k)}(x_0) + o(h),$$

where $o(h) \rightarrow 0$ as $h \rightarrow 0$. Substituting this relation into (6) gives

$$\phi(x_0 + h) - \phi(x_0) = \sum_{j=1}^{k-1} \frac{\phi^{(j)}(x_0) h^j}{j!} + \eta(h), \quad (7)$$

where $\eta(h) = o(h)h^k/k!$. Thus $\eta(h)/h^k \rightarrow 0$ as $h \rightarrow 0$. For small h , the polynomial in (7) is thus a "good" approximation to $\phi(x_0 + h) - \phi(x_0)$ in the sense that the error committed in using the approximation tends to zero faster than h^k .

We now generalize (7) to the case in which f is a real-valued function of class C^{k+1} or C^{k+2} on an open set D in \mathbb{R}^n . We restrict our attention to functions of class C^{k+1} or C^{k+2} because for functions of class C^{k+1} with $k > 2$ the statement of the result is very cumbersome, and in this text we shall only need $k = 1, 2$.

THEOREM 11.2. Let D be an open set in \mathbb{R}^n , let x_0 be a point in D , and let f be a real-valued function defined on D . Let \mathbf{h} be a vector such that $x_0 + \mathbf{h}$ is in D .

(i) If f is of class C^{k+1} on D , then

$$f(x_0 + \mathbf{h}) - f(x_0) = \nabla f(x_0) \cdot \mathbf{h} + \eta_1(\mathbf{h}), \quad (8)$$

where $\eta_1(\mathbf{h})/|\mathbf{h}| \rightarrow 0$ as $|\mathbf{h}| \rightarrow 0$.

(ii) If f is of class C^{k+2} on D , then

$$f(x_0 + \mathbf{h}) - f(x_0) = \nabla f(x_0) \cdot \mathbf{h} + \frac{1}{2}(\mathbf{h} \cdot \mathbf{H} x_0 \mathbf{h}) + \eta_2(\mathbf{h}), \quad (9)$$

where $\alpha_k(\mathbf{h})|\mathbf{h}|^2 \rightarrow 0$ as $\mathbf{h} \rightarrow \mathbf{0}$ and

$$M(\mathbf{x}_0) = \left(\frac{\partial^2 f}{\partial x_j \partial x_k}(\mathbf{x}_0) \right). \quad (10)$$

Proof. Let $\phi(t) = f(\mathbf{x}_0 + t\mathbf{h})$. Then

$$\phi(0) = f(\mathbf{x}_0) \quad \text{and} \quad \phi(1) = f(\mathbf{x}_0 + \mathbf{h}). \quad (11)$$

It follows from the chain rule that ϕ is a function of class C^2 on an open interval containing the closed interval $[0, 1] = \{t \mid 0 \leq t \leq 1\}$ in its interior whenever f is of class C^2 . If f is of class C^2 , then using the chain rule, we get

$$\phi'(t) = \sum_{j=1}^n \frac{\partial f}{\partial x_j}(\mathbf{x}_0 + t\mathbf{h})_j = \nabla f(\mathbf{x}_0 + t\mathbf{h}) \cdot \mathbf{h}. \quad (12)$$

If f is of class C^3 , we get

$$\phi''(t) = \sum_{j=1}^n \sum_{k=1}^n \frac{\partial^2 f}{\partial x_j \partial x_k}(\mathbf{x}_0 + t\mathbf{h})_j \mathbf{h}_k = (\mathbf{h})^T M(\mathbf{x}_0 + t\mathbf{h}) \mathbf{h}, \quad (13)$$

where M is the matrix defined in (10).

Let f be of class C^3 . If we take $k = 1$, $\alpha_k = 0$, and $h = 1$ in (8), recall that the summation term in (8) is absent when $h = 1$, and use (11) and (12), we get

$$f(\mathbf{x}_0 + \mathbf{h}) - f(\mathbf{x}_0) = \phi(1) - \phi(0) = \nabla f(\mathbf{x}_0 + \theta\mathbf{h}) \cdot \mathbf{h}, \quad (14)$$

where $0 < \theta < 1$. Since f is of class C^3 , the function ∇f is continuous and thus

$$\nabla f(\mathbf{x}_0 + \theta\mathbf{h}) = \nabla f(\mathbf{x}_0) + \alpha_j(\mathbf{h}),$$

where $\alpha_j(\mathbf{h}) \rightarrow 0$ as $\mathbf{h} \rightarrow \mathbf{0}$. Substituting this relation into (14) and using the Cauchy-Schwarz inequality

$$|\langle \mathbf{x}, \mathbf{h} \rangle| \leq |\mathbf{x}_j(\mathbf{h})| |\mathbf{h}|$$

give (8), with $\alpha_j(\mathbf{h}) = \langle \alpha_j(\mathbf{h}), \mathbf{h} \rangle$.

If f is of class C^3 , we proceed in similar fashion. We take $k = 2$, $\alpha_k = 0$, and $h = 1$ in (8) and use (11), (12), and (13) to get

$$f(\mathbf{x}_0 + \mathbf{h}) - f(\mathbf{x}_0) = \nabla f(\mathbf{x}_0) \cdot \mathbf{h} + \frac{1}{2}(\mathbf{h})^T M(\mathbf{x}_0 + \theta\mathbf{h}) \mathbf{h}, \quad (15)$$

where $\theta < \bar{\epsilon} < 1$. Since f is of class $C^{2,1}$, each entry of D is continuous. Hence

$$f(\mathbf{x}_0 + \mathbf{h}) = f(\mathbf{x}_0) + Df(\mathbf{h}),$$

where $Df(\mathbf{h})$ is a matrix with entries $m_{ij}(\mathbf{h})$ such that $m_{ij}(\mathbf{h}) \rightarrow 0$ as $\mathbf{h} \rightarrow \mathbf{0}$. Substituting this relation into (15) and using the relation

$$\|(\mathbf{h}, Mf(\mathbf{h})\mathbf{h})\| \leq \|\mathbf{h}\| \|Mf(\mathbf{h})\| \leq m(\mathbf{h}) \|\mathbf{h}\|^2,$$

where $m(\mathbf{h}) \rightarrow 0$ as $\mathbf{h} \rightarrow \mathbf{0}$, gives (9) with $g_{ij}(\mathbf{h}) = (\mathbf{h}, Mf(\mathbf{h})\mathbf{h})$.

Corollary 1 (Taylor's Theorem). Let D be an open set in \mathbb{R}^n and let f be a real-valued function defined in D . If f is of class $C^{2,1}$ in D , then for each \mathbf{x}_0 in D and \mathbf{h} such that $\mathbf{x}_0 + \mathbf{h}$ is in D

$$f(\mathbf{x}_0 + \mathbf{h}) - f(\mathbf{x}_0) = \nabla f(\mathbf{x}_0) \cdot \mathbf{h} + g(\mathbf{h}), \quad (16)$$

where $\theta < \bar{\epsilon} < 1$, g is of class $C^{2,1}$, then

$$f(\mathbf{x}_0 + \mathbf{h}) - f(\mathbf{x}_0) = \nabla f(\mathbf{x}_0) \cdot \mathbf{h} + \frac{1}{2}(\mathbf{h}, Mf(\mathbf{x}_0) \cdot \mathbf{h}) + \eta(\mathbf{h}), \quad (17)$$

where $\theta < \bar{\epsilon} < 1$.

Proof. Equation (16) follows from (14) and equation (17) from (15).

We now give a condition under which the existence of partial derivatives does imply differentiability.

Theorem 11.3. Let D be an open set in \mathbb{R}^n , let f be a function defined on D with range in \mathbb{R}^m , and let \mathbf{x}_0 be a point in D . If f is of class $C^{1,1}$ on some open ball $B(\mathbf{x}_0, r)$ removed at \mathbf{x}_0 , then f is differentiable at \mathbf{x}_0 .

Proof. Since each component function f_i , $i = 1, \dots, m$, is of class $C^{1,1}$, from (8) of Theorem 11.2 we have that

$$f(\mathbf{x}_0 + \mathbf{h}) - f(\mathbf{x}_0) = \nabla f(\mathbf{x}_0) \cdot \mathbf{h} + \eta(\mathbf{h}),$$

where $\eta_i(\mathbf{h})/\|\mathbf{h}\| \rightarrow 0$ as $\mathbf{h} \rightarrow \mathbf{0}$.

Let $\eta(\mathbf{h}) = (\eta_1(\mathbf{h}), \dots, \eta_m(\mathbf{h}))$. Then

$$f(\mathbf{x}_0 + \mathbf{h}) - f(\mathbf{x}_0) = \nabla f(\mathbf{x}_0) \cdot \mathbf{h} + \eta(\mathbf{h}),$$

where $\eta(\mathbf{h})/\|\mathbf{h}\| \rightarrow \mathbf{0}$ as $\|\mathbf{h}\| \rightarrow 0$. Therefore f is differentiable at \mathbf{x}_0 .

II

CONVEX SETS IN \mathbb{R}^n

1. LINES AND HYPERPLANES IN \mathbb{R}^n

Convex sets and convex functions play an important role in many optimization problems. The study of convex sets in \mathbb{R}^n , in turn, involves the notions of line segments, lines, and hyperplanes in \mathbb{R}^n . We therefore begin with a discussion of these objects.

In \mathbb{R}^2 and \mathbb{R}^3 the vector equation of a line and in \mathbb{R}^n the vector equation of a plane are obtained from geometric considerations. In \mathbb{R}^n we shall reverse this process. We shall define a line to be the set of points satisfying an equation which in vector notation is the same as that of a line in \mathbb{R}^2 or \mathbb{R}^3 . We shall define a hyperplane in \mathbb{R}^n in similar fashion.

In \mathbb{R}^2 and \mathbb{R}^3 a vector equation of the line through two points \mathbf{x}_1 and \mathbf{x}_2 is given by (see Figure 2.1)

$$\mathbf{x} = \mathbf{x}_1 + t(\mathbf{x}_2 - \mathbf{x}_1), \quad -\infty < t < \infty. \quad (1)$$

The closed oriented line segment $[\mathbf{x}_1, \mathbf{x}_2]$ with orientation from initial point \mathbf{x}_1 to terminal point \mathbf{x}_2 corresponds to values of t in $[0, 1]$. The positive ray from \mathbf{x}_1 corresponds to values of $t \geq 0$, and the negative ray from \mathbf{x}_1 corresponds to values of $t \leq 0$.

In \mathbb{R}^n we define the line through two points \mathbf{x}_1 and \mathbf{x}_2 to be the set of points \mathbf{x} in \mathbb{R}^n that satisfy (1). Thus

Definition 1.1. The line in \mathbb{R}^n through two points \mathbf{x}_1 and \mathbf{x}_2 is defined to be the set of points \mathbf{x} such that $\mathbf{x} = \mathbf{x}_1 + c(\mathbf{x}_2 - \mathbf{x}_1)$, where c is any real number, or in set notation

$$\{\mathbf{x} \in \mathbb{R}^n = \mathbf{x}_1 + c(\mathbf{x}_2 - \mathbf{x}_1), \quad -\infty < c < \infty\}. \quad (2)$$

The set in (2) can also be written as

$$\{\mathbf{x} \in \mathbb{R}^n = (1 - t)\mathbf{x}_1 + t\mathbf{x}_2, \quad -\infty < t < \infty\}, \quad (3)$$

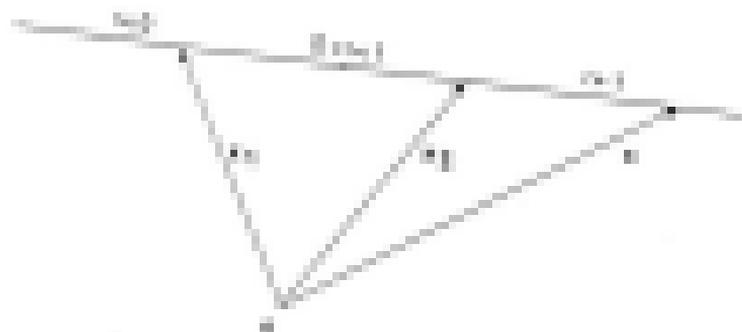


Figure 1.1.

and, recalling $\alpha + \beta + \gamma = 1$,

$$[x_1, x_2] = [x_1, \alpha x_1 + \beta x_2 + \gamma y] \quad (1)$$

Definition 1.1. The closed line segment joining the points x_1 and x_2 is denoted by $[x_1, x_2]$ and is defined by

$$[x_1, x_2] = \{x_1\alpha + \beta x_2 + \gamma y \mid \alpha + \beta + \gamma = 1\}$$

or equivalently,

$$[x_1, x_2] = \{x_1\alpha + \beta x_2 + \beta(x_2 - x_1) \mid \alpha + \beta = 1\}$$

Definition 1.2. The open line segment joining the points x_1 and x_2 is denoted by (x_1, x_2) and is defined by

$$(x_1, x_2) = \{x_1\alpha + \beta x_2 + \gamma y \mid \alpha + \beta + \gamma = 1\}$$

or equivalently,

$$(x_1, x_2) = \{x_1\alpha + \beta(x_2 - x_1) \mid \alpha + \beta = 1\}$$

The half-open segment which includes x_1 but not x_2 will be denoted by $[x_1, x_2)$ and is obtained by substituting α for the half-open interval $0 \leq \alpha < 1$, or equivalently by substituting α and β to satisfy $0 \leq \alpha + \beta < 1$, $\alpha + \beta + \gamma = 1$. Similarly, the half-open line segment which includes x_2 but not x_1 will be denoted by $(x_1, x_2]$ and is obtained by substituting α for the half-open interval $0 \leq \alpha < 1$, or equivalently by substituting α and β to satisfy $0 \leq \alpha + \beta < 1$, $\alpha + \beta + \gamma = 1$.

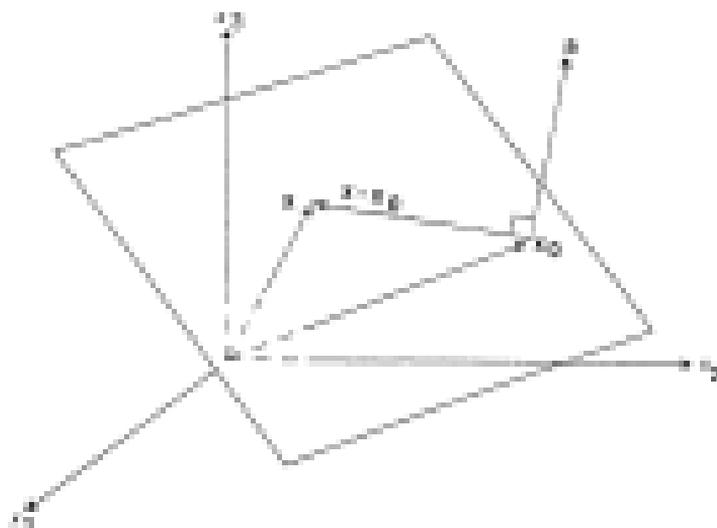


Figure 2.1

Let y be a point in the open line segment (x_1, x_2) . Then $y = \alpha x_1 + \beta x_2$, $\alpha > 0$, $\beta > 0$, $\alpha + \beta = 1$, and $y - x_1 = \beta(x_2 - x_1) = \beta(x_2 - x_1)$. Similarly $y - x_2 = \alpha(x_1 - x_2)$. Hence

$$\frac{\|y - x_1\|}{\|y - x_2\|} = \frac{\beta}{\alpha}. \quad (5)$$

This observation will be useful in the sequel.

In \mathbb{R}^3 a plane through the point x_0 with normal n consists of all points x such that $x - x_0$ is orthogonal to n . Thus, the plane is the set defined by (see Figure 2.2)

$$\langle n, x - x_0 \rangle = 0.$$

If $n = (n_1, n_2, n_3)$ and $x_0 = (x_{01}, x_{02}, x_{03})$, we can describe the plane as the set of $x = (x_1, x_2, x_3)$ satisfying

$$n_1(x_1 - x_{01}) + n_2(x_2 - x_{02}) + n_3(x_3 - x_{03}) = 0$$

or

$$n_1 x_1 + n_2 x_2 + n_3 x_3 = \beta,$$

where $\gamma = a_1x_{n+1} + a_2x_{n+2} + \dots + a_nx_{n+n}$. These are the familiar forms of the equation of a plane in \mathbb{R}^3 . In vector notation these equations can be written as

$$\langle \mathbf{a}, \mathbf{x} - \mathbf{x}_0 \rangle = 0 \quad \text{or} \quad \langle \mathbf{a}, \mathbf{x} \rangle = \gamma \quad (5)$$

where $\gamma = \langle \mathbf{a}, \mathbf{x}_0 \rangle$. Conversely, every equation of the form (5) is the equation of a plane with normal vector \mathbf{a} . To see this, let \mathbf{x}_0 be a point that satisfies (5). Then $\langle \mathbf{a}, \mathbf{x}_0 \rangle = \gamma$, so if \mathbf{x} is any other point satisfying (5), $\langle \mathbf{a}, \mathbf{x} \rangle = \langle \mathbf{a}, \mathbf{x}_0 \rangle$. Hence $\langle \mathbf{a}, \mathbf{x} - \mathbf{x}_0 \rangle = 0$ for all \mathbf{x} satisfying (5). But $\langle \mathbf{a}, \mathbf{x} - \mathbf{x}_0 \rangle = 0$ is an equation of the plane through \mathbf{x}_0 with normal \mathbf{a} .

In \mathbb{R}^n we take (6) to define a hyperplane.

Definition 1.4. A hyperplane H_α^n in \mathbb{R}^n is defined to be the set of points that satisfy the equation $\langle \mathbf{a}, \mathbf{x} \rangle = \alpha$. Thus

$$H_\alpha^n = \{ \mathbf{x} : \langle \mathbf{a}, \mathbf{x} \rangle = \alpha \}.$$

The vector \mathbf{a} is said to be a normal to the hyperplane.

As was the case in \mathbb{R}^3 , $H_\alpha^n = \{ \mathbf{x} : \langle \mathbf{a}, \mathbf{x} \rangle = \alpha \}$ can be written as

$$\langle \mathbf{a}, \langle \mathbf{a}, \mathbf{x} - \mathbf{x}_0 \rangle \rangle = 0$$

for any \mathbf{x}_0 satisfying $\langle \mathbf{a}, \mathbf{x}_0 \rangle = \alpha$, or equivalently for any \mathbf{x}_0 in H_α^n . Thus we may say that an equation of the hyperplane in \mathbb{R}^n through the point \mathbf{x}_0 with normal \mathbf{a} is

$$\langle \mathbf{a}, \mathbf{x} - \mathbf{x}_0 \rangle = 0. \quad (7)$$

Note that in \mathbb{R}^2 a hyperplane is a line.

Recall that an $(n-1)$ -dimensional subspace of \mathbb{R}^n can be represented as the set of vectors $\mathbf{x} = (x_1, \dots, x_n)$ whose coordinates satisfy $a_1x_1 + \dots + a_nx_n = 0$ for some nonzero vector $\mathbf{a} = (a_1, \dots, a_n)$. Thus, if $\alpha = 0$ and $\mathbf{a} \neq \mathbf{0}$, then the hyperplane H_0^n passes through the origin and is an $(n-1)$ -dimensional subspace of \mathbb{R}^n .

In Section 10 of Chapter I we saw that for a given linear functional L on \mathbb{R}^n there exists a unique vector \mathbf{a} such that $L(\mathbf{x}) = \langle \mathbf{a}, \mathbf{x} \rangle$ for all \mathbf{x} in \mathbb{R}^n . This representation of linear functionals and the definition of the hyperplane H_α^n show that hyperplanes are level surfaces of linear functionals. This interpretation will prove to be very useful.

Let A be any set in \mathbb{R}^n and let \mathbf{x}_0 be any vector in \mathbb{R}^n . The set $A + \mathbf{x}_0$ is defined to be

$$A + \mathbf{x}_0 = \{ \mathbf{y} : \mathbf{y} = \mathbf{x} + \mathbf{x}_0, \mathbf{x} \in A \}$$

and is called the *normal* of A by \mathbf{x}_0 .

LEMMA 1.1. For any \mathbf{a} in \mathbb{R}^n and scalar α , the hyperplane $H_{\mathbf{a}}^{\alpha}$ is the translate of $H_{\mathbf{a}}^0$, the hyperplane through the origin with normal \mathbf{a} , by any vector \mathbf{x}_0 in $H_{\mathbf{a}}^{\alpha}$.

Proof. Fix an \mathbf{x}_0 in $H_{\mathbf{a}}^{\alpha}$. Let \mathbf{x} be an arbitrary point in $H_{\mathbf{a}}^{\alpha}$ and let $\mathbf{u} = \mathbf{x} - \mathbf{x}_0$. Then

$$\alpha = \langle \mathbf{a}, \mathbf{x} \rangle = \langle \mathbf{a}, \mathbf{u} + \mathbf{x}_0 \rangle = \langle \mathbf{a}, \mathbf{u} \rangle + \alpha.$$

Thus $\langle \mathbf{a}, \mathbf{u} \rangle = 0$ and so $\mathbf{u} \in H_{\mathbf{a}}^0$. Since $\mathbf{x} = \mathbf{x}_0 + \mathbf{u}$, we have shown that

$$H_{\mathbf{a}}^{\alpha} = H_{\mathbf{a}}^0 + \mathbf{x}_0 \quad \mathbf{x}_0 \in H_{\mathbf{a}}^{\alpha}. \quad (8)$$

Now let \mathbf{u} be an arbitrary vector in $H_{\mathbf{a}}^0$. Then $\mathbf{u} + \mathbf{x}_0 \in H_{\mathbf{a}}^0 + \mathbf{x}_0$ and

$$\langle \mathbf{a}, \mathbf{u} + \mathbf{x}_0 \rangle = \langle \mathbf{a}, \mathbf{x}_0 \rangle = \alpha.$$

Thus we get the inclusion opposite to the one in (8), and the lemma is proved.

Two hyperplanes $H_{\mathbf{a}}^{\alpha}$ and $H_{\mathbf{b}}^{\beta}$ are said to be *parallel* if their normals are scalar multiples of each other. Thus two hyperplanes are parallel if and only if they are translates of the same hyperplane through the origin. The parallel hyperplane through the origin is called the *parallel subspace*.

Exercise 1.1. (a) Draw a sketch in \mathbb{R}^2 and \mathbb{R}^3 that illustrates Lemma 1.1.

(b) In \mathbb{R}^2 sketch the hyperplane $x + 2y = 1$ and the parallel subspace.

Exercise 1.2. (a) In \mathbb{R}^2 find an equation of the hyperplane through the points $(1, 1, 1)$, $(2, 0, 1)$, $(0, 1, 0)$, $(0, 2, 0, 1)$, and $(1, 1, -1, 0)$.

(b) Let $\mathbf{x}_1, \dots, \mathbf{x}_k$ be k points in \mathbb{R}^n . Find a sufficient condition for the existence of a unique hyperplane containing these points.

Exercise 1.3. Show that a hyperplane is a closed set.

Exercise 1.4. Given a hyperplane $H_{\mathbf{a}}^{\alpha}$ show that \mathbf{a} is indeed normal to $H_{\mathbf{a}}^{\alpha}$ in the sense that if \mathbf{x}_0 and \mathbf{x}_1 are any two points in $H_{\mathbf{a}}^{\alpha}$, then \mathbf{a} is orthogonal to $\mathbf{x}_1 - \mathbf{x}_0$.

Exercise 1.5 (Linear Algebra Review). Let V be an $(n - 1)$ -dimensional subspace of \mathbb{R}^n . Let \mathbf{p} be a vector in \mathbb{R}^n not in V . Show that every \mathbf{x} in \mathbb{R}^n has a unique representation $\mathbf{x} = \mathbf{v} + \alpha\mathbf{p}$ where $\mathbf{v} \in V$ and $\alpha \in \mathbb{R}$.

2. PROPERTIES OF CONVEX SETS

A subset C of \mathbb{R}^2 is said to be *convex* if given any two points $\mathbf{x}_1, \mathbf{x}_2$ in C the line segment $[\mathbf{x}_1, \mathbf{x}_2]$ joining these points is contained in C . We now formalize this definition analytically.

Definition 2.1. A subset C of \mathbb{R}^2 is *convex* if for every pair of points $\mathbf{x}_1, \mathbf{x}_2$ in C the line segment

$$[\mathbf{x}_1, \mathbf{x}_2] = \{\mathbf{x} : \mathbf{x} = \alpha\mathbf{x}_1 + \beta\mathbf{x}_2, \alpha \geq 0, \beta \geq 0, \alpha + \beta = 1\}$$

belongs to C .

At this point the reader may find it helpful to do the following exercise.

Exercise 2.1. Sketch the following sets in \mathbb{R}^2 and determine from your figure which sets are convex and which are not:

- (a) $\{(x, y) : x^2 + y^2 \leq 1\}$,
- (b) $\{(x, y) : 0 < x^2 + y^2 < 1\}$,
- (c) $\{(x, y) : x \geq x^2\}$,
- (d) $\{(x, y) : |x| + |y| \leq 1\}$, and
- (e) $\{(x, z) : z \geq 1/(1 + x^2)\}$.

In \mathbb{R}^2 a plane determines two sets, one on each side of the plane. We extend this notion to \mathbb{R}^n . The *positive half space* determined by \mathcal{H}_1^n and denoted by \mathcal{H}_1^{n+} is defined to be

$$\mathcal{H}_1^{n+} = \{\mathbf{x} : \langle \mathbf{n}, \mathbf{x} \rangle \geq c\}. \quad (1)$$

The *negative half space* \mathcal{H}_1^{n-} is defined to be

$$\mathcal{H}_1^{n-} = \{\mathbf{x} : \langle \mathbf{n}, \mathbf{x} \rangle \leq c\}. \quad (2)$$

The *closed positive half space*, denoted by $\bar{\mathcal{H}}_1^{n+}$, is defined to be the closure of \mathcal{H}_1^{n+} . The *closed negative half space* $\bar{\mathcal{H}}_1^{n-}$ is defined to be the closure of \mathcal{H}_1^{n-} .

Note that a half space is not a subspace of \mathbb{R}^n , since there exist points \mathbf{x} in a half space such that not all scalar multiples of \mathbf{x} are in the half space.

Exercise 2.2. Show that

$$\bar{\mathcal{H}}_1^{n+} = \{\mathbf{x} : \langle \mathbf{n}, \mathbf{x} \rangle \geq c\}, \quad \bar{\mathcal{H}}_1^{n-} = \{\mathbf{x} : \langle \mathbf{n}, \mathbf{x} \rangle \leq c\}.$$

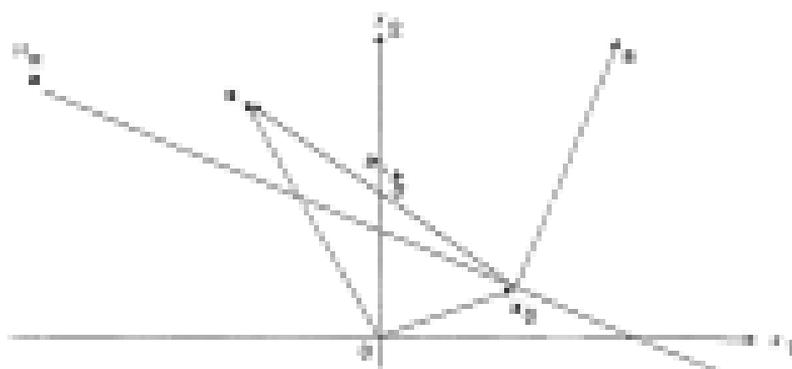


Figure 2.1

If $x_2 \in M_1^c$, then $\langle \mathbf{x}, \mathbf{x}_2 \rangle = \alpha$. Hence, for $\mathbf{x} \in H_1^+$, we have $\langle \mathbf{x}, \mathbf{x} \rangle = \alpha = \langle \mathbf{x}, \mathbf{x}_2 \rangle$. Therefore

$$\langle \mathbf{x}, \mathbf{x} - \mathbf{x}_2 \rangle = 0 \quad \text{for all } \mathbf{x} \in H_1^+.$$

Thus, vectors \mathbf{x} in H_1^+ are such that the angle between \mathbf{x} , the normal to M_1^c , and $\mathbf{x} - \mathbf{x}_2$ lies between 0 and $\pi/2$. In \mathbb{R}^2 and \mathbb{R}^3 this says that \mathbf{x} and $\mathbf{x} - \mathbf{x}_2$ point in the same direction or have the same orientation. See Figure 2.1.

Half-spaces are convex sets. Although this and other properties of convex sets are "geometrically obvious" in \mathbb{R}^2 and \mathbb{R}^3 , they require proof in \mathbb{R}^n . To show that H_1^+ is convex, we must show that if \mathbf{x}_1 and \mathbf{x}_2 are in H_1^+ , then so is $[\mathbf{x}_1, \mathbf{x}_2]$. Let $\mathbf{x} \in [\mathbf{x}_1, \mathbf{x}_2]$. Then there exist scalars $\lambda \geq 0$, $\mu \geq 0$ with $\lambda + \mu = 1$ such that $\mathbf{x} = \lambda \mathbf{x}_1 + \mu \mathbf{x}_2$. Hence

$$\langle \mathbf{x}, \mathbf{x} \rangle = \langle \lambda \mathbf{x}_1 + \mu \mathbf{x}_2, \lambda \mathbf{x}_1 + \mu \mathbf{x}_2 \rangle = \lambda \langle \mathbf{x}_1, \mathbf{x}_1 \rangle + \mu \langle \mathbf{x}_2, \mathbf{x}_2 \rangle \geq \lambda \alpha + \mu \alpha = \alpha,$$

and so $\mathbf{x} \in H_1^+$. Similar arguments show that M_1^+ , M_1^- , M_1^c and the hyperplane H_1^0 are convex sets.

In the series of lemmas that follow we list some elementary properties of convex sets. We shall provide proofs to show the reader how one "proves the obvious." When speaking of a convex set, we shall always assume that the convex set is not empty.

Lemma 2.1. Let $\{C_\alpha\}$ be a collection of convex sets such that $C = \bigcap_{\alpha} C_\alpha$ is not empty. Then C is convex.

Proof. Let \mathbf{x}_1 and \mathbf{x}_2 belong to C . Then for each α the points \mathbf{x}_1 and \mathbf{x}_2 belong to C_α . Since for each α the set C_α is convex, for each α the line segment $[\mathbf{x}_1, \mathbf{x}_2]$ belongs to C_α . Hence $[\mathbf{x}_1, \mathbf{x}_2] \subset C$, so C is convex. \square

Let A be an $m \times n$ matrix with i th row $a_i = (a_{i1}, \dots, a_{in})$ and let $\mathbf{b} = (b_1, \dots, b_m)$ be a vector in \mathbb{R}^m . Let S denote the solution set of $A\mathbf{x} \leq \mathbf{b}$. Then S is the intersection of half spaces $\{\mathbf{x} : \langle a_i, \mathbf{x} \rangle \leq b_i\}$, so if S is nonempty, then S is convex.

Note that if A and B are convex sets, then $A \cap B$ need not be convex.

If A and B are two sets in \mathbb{R}^n and if λ and μ are scalars, we define

$$\lambda A + \mu B = \{\lambda \mathbf{x} + \mu \mathbf{y} : \mathbf{x} \in A, \mathbf{y} \in B\}.$$

Lemma 2.2. If A and B are convex sets in \mathbb{R}^n and λ and μ are scalars, then $\lambda A + \mu B$ is convex.

Proof. Let \mathbf{x}_1 and \mathbf{x}_2 belong to $\lambda A + \mu B$. We must show that $\langle \mathbf{x}_1, \mathbf{x}_2 \rangle$ belongs to $\lambda A + \mu B$. We have $\mathbf{x}_1 = \lambda \mathbf{a}_1 + \mu \mathbf{b}_1$ for some $\mathbf{a}_1 \in A$ and $\mathbf{b}_1 \in B$ and $\mathbf{x}_2 = \lambda \mathbf{a}_2 + \mu \mathbf{b}_2$ for some $\mathbf{a}_2 \in A$ and $\mathbf{b}_2 \in B$. If $\alpha \in \langle \mathbf{x}_1, \mathbf{x}_2 \rangle$, then there exist scalars $\alpha > 0, \beta > 0$ with $\alpha + \beta = 1$ such that $\mathbf{x} = \alpha \mathbf{x}_1 + \beta \mathbf{x}_2$. Hence

$$\begin{aligned} \mathbf{x} &= \alpha \mathbf{x}_1 + \beta \mathbf{x}_2 = \alpha(\lambda \mathbf{a}_1 + \mu \mathbf{b}_1) + \beta(\lambda \mathbf{a}_2 + \mu \mathbf{b}_2) \\ &= \lambda(\alpha \mathbf{a}_1 + \beta \mathbf{a}_2) + \mu(\alpha \mathbf{b}_1 + \beta \mathbf{b}_2). \end{aligned}$$

Since A and B are convex, the first term in parentheses on the right is an element \mathbf{a}_3 in A and the second term in parentheses on the right is an element \mathbf{b}_3 in B . Hence for arbitrary $\alpha \in \langle \mathbf{x}_1, \mathbf{x}_2 \rangle$,

$$\mathbf{x} = \lambda \mathbf{a}_3 + \mu \mathbf{b}_3$$

so $\langle \mathbf{x}_1, \mathbf{x}_2 \rangle \subset \lambda A + \mu B$ as required.

Lemma 2.3. Let $A_1 \subset \mathbb{R}^n, A_2 \subset \mathbb{R}^n, \dots, A_k \subset \mathbb{R}^n$ and let A_1, \dots, A_k be convex. Then $A_1 \cap \dots \cap A_k$ is a convex set in $\mathbb{R}^{n \times \dots \times n}$.

We leave the proof as an exercise for the reader.

Lemma 2.4. If A is convex, then so is \bar{A} , the closure of A .

Proof. Let \mathbf{x} and \mathbf{y} be points of \bar{A} . Let $\alpha \in \langle \mathbf{x}, \mathbf{y} \rangle$, so $\alpha = \alpha \mathbf{x} + \beta \mathbf{y}$ for some $\alpha \geq 0, \beta \geq 0, \alpha + \beta = 1$. By Remark 1.3.2, there exist sequences $\{\mathbf{a}_k\}$ and $\{\mathbf{b}_k\}$ of points in A such that $\lim \mathbf{a}_k = \mathbf{x}$ and $\lim \mathbf{b}_k = \mathbf{y}$. Since A is convex, $\alpha_k = \alpha \mathbf{a}_k + \beta \mathbf{b}_k \in A$ for each k . Letting $k \rightarrow \infty$, we get that $\lim \alpha_k = \alpha$. From Remark 1.5.1 we have that $\alpha \in \bar{A}$.

The next lemma and its corollaries are fundamental results.

Lemma 2.5. Let C be a convex set with nonempty interior. Let \mathbf{x}_1 and \mathbf{x}_2 be two points with $\mathbf{x}_1 \in \text{int} C$ and $\mathbf{x}_2 \in \bar{C}$. Then the line segment $\langle \mathbf{x}_1, \mathbf{x}_2 \rangle$ is contained in the interior of C .



Figure 2.4.

We first assume that $x_1 \in C$. The proof of the lemma in this case is based on Figure 2.4, which illustrates how one uses geometry to motivate analytical proofs. In the figure we have normalized the dimensions so that $\|x_2 - x_1\| = 1$. Let x be an arbitrary point in $\|x_1, x_2\|$. Then $x = \alpha x_1 + \beta x_2$ with $\alpha > 0$, $\beta > 0$, and $\alpha + \beta = 1$. We must show that $x \in \text{int}(C)$. Since the point x_1 is in $\text{int}(C)$, there exists a circle of radius $r > 0$ centered at x_1 that is contained entirely within C . From x_2 draw tangents to the circle. Since $\|x_1 - x\| = \alpha \|x_2 - x_1\|$, by similar triangles we would expect the circle with center at x and radius αr to be in C . This would show that $x \in \text{int}(C)$.

We now verify analytically that the suggested argument is indeed valid. To prove the lemma, we must show that any point $x \in \|x_1, x_2\|$ is in $\text{int}(C)$. By hypothesis, $x_1 \in \text{int}(C)$, so we need only consider $x \in \|x_1, x_2\|$.

Since $x_1 \in \text{int}(C)$, there exists an $r > 0$ such that $B(x_1, r) \subset C$. If $x \in \|x_1, x_2\|$, then $x = \alpha x_1 + \beta x_2$ for some $\alpha > 0$, $\beta > 0$, $\alpha + \beta = 1$. We shall show that $B(x, \alpha r) \subset C$, and so $x \in \text{int}(C)$. Let $y \in B(x, \alpha r)$ and let

$$z = x_1 + \frac{y - x_1}{\alpha}.$$

Then,

$$\|z - x_1\| = \left\| \frac{y - x_1}{\alpha} \right\| \leq \frac{\alpha r}{\alpha} = r.$$

Hence $z \in B(x_1, r)$ and therefore in C . Now,

$$y = \alpha(z - x_1) + x = \alpha(z - x_1) + \alpha x_1 + \beta x_2 = \alpha z + \beta x_2.$$

Since C is convex, $y \in C$.

We now suppose that $x_2 \in C$ and $x_1 \notin C$. Figure 2.5 will illustrate the idea of the proof. In the figure we again take $\|x_2 - x_1\| = 1$.

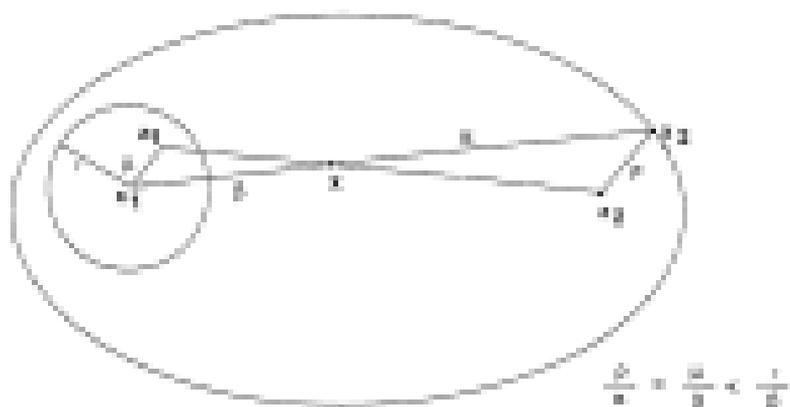


Figure 15

Since $x_1 \in \text{int}(C)$, there exists an $r > 0$ such that $B(x_1, r) \subset C$. To prove the lemma, we must show that if $x \in (x_1, x_2)$, then $x \in \text{int}(C)$. If $x \in (x_1, x_2)$, there exist $\alpha > 0$, $\beta > 0$, $\alpha + \beta = 1$ such that $x = \alpha x_1 + \beta x_2$. Also by relation (5) in Section 1, $\|x - x_1\|/\|x - x_2\| = \beta/\alpha$. Since $x_2 \in C$, there exists a r_2 in C such that $\|x_2 - x_1\| = r_2/\beta$. Define

$$r_1 = r_2 - \frac{\beta}{\alpha}(r_2 - r_1).$$

Hence $\|x_1 - x_2\| = r_1$, and so $x_1 \in \text{int}(C)$. Also,

$$\begin{aligned} x - x_1 + \beta r_1 &= \alpha \left(x_1 + \frac{\beta}{\alpha}(x_2 - x_1) \right) + \beta r_1 \\ &= \alpha x_1 + \beta x_2. \end{aligned}$$

Since $x_1 \in \text{int}(C)$ and $x_2 \in C$, by the case already proved, $x \in \text{int}(C)$.

COROLLARY 1. If C is convex, then $\text{int}(C)$ is either empty or convex.

COROLLARY 2. Let C be convex and let $\text{int}(C) \neq \emptyset$. Then (i) $\overline{\text{int}(C)} = C$ and (ii) $\text{int}(C) = \text{int}(\overline{C})$.

Before proving Corollary 2, we note that its conclusion need not be true if C is not convex. To see this, take $C = \{\text{rational points in } [0, 1] \cup [1, 2]\}$.

Proof. Since for any set C , $\text{int}(C) \subseteq C$, it follows that $\overline{\text{int}(C)} \subseteq C$. Let $x \in C$. Then for any y in $\text{int}(C)$, the line segment $[y, x]$ belongs to $\text{int}(C)$. Hence $x \in \text{int}(C)$, and so $C = \text{int}(C)$.

Now, for any set C , $C \subseteq \overline{C}$, it follows that $\text{int}(C) \subseteq \text{int}(\overline{C})$. Let $x \in \text{int}(\overline{C})$. Then there exists an $r > 0$ such that $B(x, r) \subseteq \overline{C}$. Now let $y \in \text{int}(C)$ and consider the line segment $[y, x]$, which is the extension of the segment $[y, x]$ a distance less than r beyond x . Then $x \in C$, and since $y \in \text{int}(C)$, we get that $[x, y] \subseteq \text{int}(C)$. In particular, $x \in \text{int}(C)$. Thus (ii) is proved.

For each positive integer n , define

$$P_n = \left\{ p = (p_1, \dots, p_n) : p_i \geq 0, \sum_{i=1}^n p_i = 1 \right\}.$$

For $n = 1$, P_1 is the point 1. For $n = 2$, P_2 is the closed line segment joining $(0, 1)$ and $(1, 0)$. For $n = 3$, P_3 is the closed triangle with vertices $(1, 0, 0)$, $(0, 1, 0)$, and $(0, 0, 1)$. It is easy to verify that, for each n , P_n is a closed convex set.

Definition 2.1. A point x in \mathbb{R}^n is a *convex combination* of points x_1, \dots, x_n if there exists a $p = (p_1, \dots, p_n) \in P_n$ such that

$$x = p_1 x_1 + p_2 x_2 + \dots + p_n x_n.$$

Lemma 2.6. A set C in \mathbb{R}^n is convex if and only if every convex combination of points in C is also in C .

Proof. If every convex combination of points in C is in C , then every convex combination of every two points in C is in C . But for any fixed pair of points x_1 and x_2 in C , the set of convex combinations of x_1 and x_2 is just $[x_1, x_2]$. In other words, for every pair of points x_1 and x_2 in C , the line segment $[x_1, x_2]$ belongs to C , and so C is convex.

We shall prove the "only if" statement by induction on k , the number of points in the convex combination. If C is convex, then the definition of convexity says that every convex combination of two points is in C . Suppose that we have shown that if C is convex, then every convex combination of k points is in C . We shall show that every convex combination of $k + 1$ points is in C . Let

$$x = p_1 x_1 + \dots + p_k x_k + p_{k+1} x_{k+1}, \quad p \in P_{k+1}.$$

If $p_{k+1} = 0$, then there is nothing to prove. If $p_{k+1} < 1$, then $\sum_{i=1}^k p_i > 0$. Hence,

$$x = \left(\sum_{i=1}^k p_i \right) \left(\frac{p_1}{\sum_{i=1}^k p_i} x_1 + \dots + \frac{p_k}{\sum_{i=1}^k p_i} x_k \right) + p_{k+1} x_{k+1}.$$

The term in square brackets is a convex combination of k points in C and so by the inductive hypothesis is in C . Therefore x is a convex combination of two points in C , and so since C is convex, x belongs to C . This proves the lemma.

For a given set A , let $K(A)$ denote the set of all convex combinations of points in A . It is easy to verify that $K(A)$ is convex. Clearly, $K(A) \supseteq A$.

Definition 2.4. The convex hull of a set A , denoted by $co(A)$, is the intersection of all convex sets containing A .

Since \mathbb{R}^n is convex, $co(A)$ is nonempty, and since the intersection of convex sets is convex, $co(A)$ is convex. Also, since we have $co(A) \subset C$ for any convex set C containing A , it follows that $co(A)$ is the smallest convex set containing A .

THEOREM 2.1. The convex hull of a set A is the set of all convex combinations of points in A ; that is, $co(A) = K(A)$.

Proof. Let $\{C_\alpha\}$ denote the family of convex sets containing A . Since $co(A) = \bigcap C_\alpha$ and $K(A)$ is a convex set containing A , $co(A) \subset K(A)$. To prove the opposite inclusion, note that $A \subseteq C_\alpha$ for each α and Lemma 2.3 imply that for each α all convex combinations of points in A belong to C_α . That is, for each α , $K(A) \subset C_\alpha$. Hence $K(A) \subset \bigcap C_\alpha = co(A)$.

Lemma 2.5 states that a set C is convex if and only if $K(C) \subset C$. Since $C \subseteq K(C)$ for any set C , Lemma 2.3 can be restated in the equivalent form that a set C is convex if and only if $K(C) = C$. Since, for any set A , $K(A) = co(A)$, we have the following corollary to Theorem 2.1.

COROLLARY 1. A set A is convex if and only if $A = co(A)$.

THEOREM 2.2 (CARATHÉODORY). Let A be a subset of \mathbb{R}^n and let $x \in co(A)$. Then there exist $n + 1$ points x_1, \dots, x_{n+1} in A and a point p in F_{n+1} such that

$$x = p_1 x_1 + \dots + p_{n+1} x_{n+1}.$$

Note that since some of the p_i 's may be zero, the result is often stated that a point $x \in co(A)$ can be written as a convex combination of at most $n + 1$ points of A .

Proof. If $x \in co(A)$, then by Theorem 2.1,

$$x = \sum_{i=1}^m q_i x_i, \quad x_i \in A, \quad i = 1, \dots, m, \quad \mathbf{q} = (q_1, \dots, q_m) \in F_m, \quad (2)$$

for some positive integer m and some \mathbf{q} in F_m . If $m < n + 1$, then there is

nothing to prove. If $m > n + 1$, we shall express x as a convex combination of at most $m - 1$ points. Repeating this argument, a finite number of times gives the desired result.

Let us then suppose that in (3) $m > n + 1$. Consider the $m - 1 > n$ vectors in \mathbb{R}^n , $(x_1 - x_m), (x_2 - x_m), \dots, (x_{m-1} - x_m)$. Since $m - 1 > n$, these vectors are linearly dependent. Hence there exist scalars $\lambda_1, \dots, \lambda_{m-1}$, not all zero such that

$$\lambda_1(x_1 - x_m) + \lambda_2(x_2 - x_m) + \dots + \lambda_{m-1}(x_{m-1} - x_m) = \theta.$$

Let $\lambda_m = -(\lambda_1 + \dots + \lambda_{m-1})$. Then

$$\lambda_1 x_1 + \dots + \lambda_{m-1} x_{m-1} + \lambda_m x_m = \theta \quad (4)$$

and

$$\sum_{i=1}^m \lambda_i = 0. \quad (5)$$

From (1) and (4) it follows that for any t

$$x = \sum_{i=1}^m t_i x_i - t_i x_m. \quad (6)$$

We shall show that there is a value of t such that the coefficients of the x_i are nonnegative, have sum equal to 1, and at least one of the coefficients is zero. This will prove the theorem.

Let $I = \{i \mid i = 1, \dots, m, \lambda_i > 0\}$. Since $\lambda_1, \dots, \lambda_m$ are not all zero, it follows from (5) that $I \neq \emptyset$. Let i_0 denote an index in I such that

$$\frac{\lambda_{i_0}}{\lambda_m} = \min\{i \in I, \lambda_i / \lambda_m\}.$$

Now take $t = \lambda_{i_0} / \lambda_m$. Then $t \geq 0$, and for $i \in I$

$$t_i(t - \lambda_i) = \lambda_i \left(\frac{\lambda_i}{\lambda_m} - \frac{\lambda_{i_0}}{\lambda_m} \right) \geq 0,$$

with equality holding for $i = i_0$. If $i \notin I$, then $\lambda_i \leq 0$, so $t_i - \lambda_i \geq 0$ whenever $t \geq 0$. Thus, if $t = \lambda_{i_0} / \lambda_m$, then

$$t_i - \lambda_i \geq 0, \quad i = 1, \dots, m, \quad \text{with } t_{i_0} - \lambda_{i_0} = 0.$$

Upon comparing this statement with (6), we see that we have written x as a nonnegative linear combination of at most $m - 1$ points. This combination is

also a convex combination since, using (3), we have

$$\sum_{j=1}^m (\lambda_j - \lambda_j') = \sum_{j=1}^m \lambda_j - 1 = \sum_{j=1}^m \lambda_j - \sum_{j=1}^m \lambda_j = 0.$$

Lemma 2.7. *If A is a compact subset of \mathbb{R}^n , then so is $\text{co}(A)$.*

If A is closed, then $\text{co}(A)$ need not be closed, as the following example shows. Let

$$A_1 = \{(x_1, x_2) : x_1 > 0, x_2 > 0, x_1 x_2 \geq 1\}$$

and let

$$A_2 = \{(x_1, x_2) : x_1 > 0, x_2 < 0, x_1 x_2 \leq -1\}.$$

Let $A = A_1 \cup A_2$. Then A is closed, but $\text{co}(A) = \{(x_1, x_2) : x_1 > 0\}$ is open.

We now prove the lemma. By Theorem 2.2

$$\text{co}(A) = \left\{ \mathbf{x} : \mathbf{x} = \sum_{j=1}^{m+1} \lambda_j \mathbf{x}_j, \lambda_j \geq 0, \mathbf{x}_j \in A, \mathbf{1} = (\lambda_1, \dots, \lambda_{m+1}) \in P_{m+1} \right\}.$$

Let

$$\pi = P_{m+1} = \left\{ \frac{\lambda_1, \dots, \lambda_{m+1}}{\lambda_1 + \dots + \lambda_{m+1}} \right\}.$$

Since A is compact and P_{m+1} is compact, by Theorem 18.1 the set π is compact in \mathbb{R}^m , where $m = (m+1) + (m+1) = (m+1)^2$. The mapping $\varphi: \mathbb{R}^m \rightarrow \mathbb{R}^n$ defined by

$$\varphi(\lambda_1, \dots, \lambda_{m+1}) = \sum_{j=1}^{m+1} \lambda_j \mathbf{x}_j$$

is continuous. Since π is compact, it follows from Theorem 18.2 that $\varphi(\pi)$ is compact. By Theorem 2.2, $\varphi(\pi) = \text{co}(A)$, so $\text{co}(A)$ is compact.

Lemma 2.8. *If \mathcal{C} is an open subset of \mathbb{R}^n , then $\text{co}(\mathcal{C})$ is also open.*

Proof. Since $O = \text{co}(\mathcal{C})$, the set $\text{co}(\mathcal{C})$ has nonempty interior. Therefore, by Corollary 1 of Lemma 2.5 $\text{int}(\text{co}(\mathcal{C}))$ is convex and $O \subset \text{int}(\text{co}(\mathcal{C}))$. (Why?) Since $\text{co}(\mathcal{C})$ is the intersection of all convex sets containing \mathcal{C} , we have $\text{co}(\mathcal{C}) \subset \text{int}(\text{co}(\mathcal{C}))$. Since we always have the reverse inclusion, we have that $\text{co}(\mathcal{C}) = \text{int}(\text{co}(\mathcal{C}))$.

Exercise 2.3. Using the definition of a convex set, show that (a) the non-negative orthant in $\mathbb{R}^n = \{x \in \mathbb{R}^n : x_1, \dots, x_n, x_i \geq 0 \forall i = 1, \dots, n\}$ is convex and (b) a hyperplane H_C^0 is convex.

Exercise 2.4. A mapping S from \mathbb{R}^n to \mathbb{R}^m is said to be affine if $Sx = Tx + k$, where T is a linear map from $\mathbb{R}^n \rightarrow \mathbb{R}^m$ and k is a fixed vector. (If $k = 0$, then S is linear.) Show that if C is a convex set in \mathbb{R}^n , then $S(C) = \{y \in \mathbb{R}^m : y = Tx + k, x \in C\}$ is a convex set in \mathbb{R}^m . In mathematical jargon, you are about to show that under an affine map convex sets are mapped into convex sets or under an affine map the image of a convex set is convex.

Exercise 2.5. Show that the sets P_i are compact and convex.

Exercise 2.6. Show that for any set A the set $co(A)$ of all convex combinations of points in A is convex.

Exercise 2.7. Show that the open ball $B(0, r)$ is convex.

Exercise 2.8. Consider the linear programming (LP) problem: Minimize (c, x) subject to $Ax = b, x \geq 0$. Let $S = \{x : x \text{ is a solution of the problem LP}\}$. Show that if S is not empty, then it is convex.

Exercise 2.9. Let $C \subset \mathbb{R}^n$. Show that C is convex if and only if $\lambda C + \mu C = C + \mu C$ for all $\lambda \geq 0, \mu \geq 0$.

Exercise 2.10. A set C is said to be a cone with apex at the origin, or simply a cone, if whenever $x \in C$, all vectors $\lambda x, \lambda \geq 0$, belong to C . If C is also convex, C is said to be a convex cone.

- Give an example of a cone that is not convex.
- Give an example of a cone that is convex.
- Let C be a nonempty set in \mathbb{R}^n . Show that C is a convex cone if and only if $\lambda_1 x_1$ and $\lambda_2 x_2 \in C$ implies that $\lambda_1 x_1 + \lambda_2 x_2 \in C$ for all $\lambda_1 \geq 0, \lambda_2 \geq 0$.

Exercise 2.11. Show that if C_1 and C_2 are convex cones, then so is $C_1 + C_2$ and that $C_1 + C_2 = co(C_1 \cup C_2)$.

Exercise 2.12. Show that if A is a bounded set in \mathbb{R}^n , then so is $co(A)$.

Exercise 2.13. Let X be a convex set in \mathbb{R}^n and let y be any vector not in X . Let $[y, X]$ denote the union of all line segments $[y, x]$ with x in X .

- Sketch a set $[y, X]$ in \mathbb{R}^2 .
- Show that $[y, X]$ is convex.

Exercise 2.14. Let C be a proper subset of \mathbb{R}^n . For each $\epsilon > 0$ define a set $\mathfrak{a}_\epsilon(C)$, the ϵ -neighborhood of C , by $\mathfrak{a}_\epsilon(C) = \bigcup_{x \in C} \mathcal{B}(x, \epsilon)$. Show that

$$\begin{aligned}\mathfrak{a}_\epsilon(C) &= \{x \mid \exists s - \delta \leq x \text{ for some } s \in C\} \\ &= \bigcup_{s \in C} (\mathcal{B}(0, \epsilon) + s).\end{aligned}$$

Show that $\mathfrak{a}_\epsilon(C)$ is open. Show that if C is convex, then so is $\mathfrak{a}_\epsilon(C)$.

Exercise 2.15. Show that if a vector \mathbf{x} in \mathbb{R}^n has distinct representations as a convex combination of a set of vectors $\mathbf{x}_0, \mathbf{x}_1, \dots, \mathbf{x}_n$, then the vectors $\mathbf{x}_0 - \mathbf{x}_1, \dots, \mathbf{x}_1 - \mathbf{x}_2, \dots, \mathbf{x}_{n-1} - \mathbf{x}_n$ are linearly dependent.

Exercise 2.16. Let N be a set contained in the closed negative half space determined by the hyperplane H_0^n . Show that if $\mathbf{x} \in \text{int}(N)$, then $\langle \mathbf{x}, \mathbf{x} \rangle < 0$.

Exercise 2.17. Let N be a convex set and let H_0^n be a hyperplane such that $X \cap H_0^n = \emptyset$. Show that X is contained in one of the open half spaces determined by H_0^n .

3. SEPARATION THEOREMS

This section is devoted to the establishment of separation theorems. In some sense, these theorems are the fundamental theorems of optimization theory. The validity of this assertion will be made evident in subsequent chapters of this book.

We have seen that a hyperplane H_0^n divides \mathbb{R}^n into two half spaces, one on each side of H_0^n . It therefore seems natural to say that two sets X and Y are separated by a hyperplane H_0^n if they are contained in different half spaces determined by H_0^n . There are various types of separation, which we will now define precisely and discuss. The reader should draw sketches in \mathbb{R}^2 . (Recall that in \mathbb{R}^2 a hyperplane is a line.)

Definition 3.1. Two sets X and Y are separated by the hyperplane H_0^n if for every $\mathbf{x} \in X$, $\langle \mathbf{x}, \mathbf{x} \rangle \geq \alpha$ and, for every $\mathbf{y} \in Y$, $\langle \mathbf{y}, \mathbf{y} \rangle \leq \alpha$.

This type of separation does not always correspond to what one intuitively considers a separation. For example, in \mathbb{R}^2 , let

$$X_1 = \{(x_1, x_2) \mid x_2 = 0, 0 \leq x_1 \leq 1\}$$

and let

$$Y_1 = \{(x_1, x_2) \mid x_2 = 0, 1 \leq x_1 \leq 2\}.$$

According to Definition 1.1, these sets are separated by the hyperplane $x_1 = 1$. On the other hand, one would not consider these sets as being separated in any reasonable way. This example illustrates that two sets which intuitively should not be considered as being separated can be separated by a hyperplane H_C^+ according to Definition 1.1 if the sets both lie in H_C^+ . To rule out the possibility just discussed, the notion of proper separation is introduced.

Definition 1.2. Two sets X and Y are properly separated by a hyperplane H_C^+ if for every $x \in X$, $\langle \mathbf{a}, \mathbf{x} \rangle = \alpha$, for every $y \in Y$, $\langle \mathbf{a}, \mathbf{y} \rangle < \alpha$, and at least one of the sets is not contained in H_C^+ .

Geometrically, the definition requires that X and Y lie in opposite closed half spaces and at least one of the sets not be contained in H_C^+ . Note that the sets X_1 and X_2 are not properly separated by the hyperplane $x_1 = 0$ and that they cannot be properly separated by any hyperplane.

Proper separation of two sets X and Y does not require that the sets be disjoint. The sets

$$X_2 = \{(x_1, x_2) : 0 \leq x_1 \leq 1, 0 \leq x_2 \leq 1\}$$

and

$$Y_1 = \{(x_1, x_2) : 0 \leq x_1 \leq 1, -1 \leq x_2 \leq 0\}$$

are not disjoint but are properly separated by the hyperplane $x_2 = 0$. A set may be a subset of another set, yet the two sets can be properly separated. To see this, let $X_2^* = X_2$ and let $X_1^* = \{(x_1, x_2) : 0 \leq x_1 \leq 1, x_2 = 0\}$. Then $X_1^* \subset X_2^*$ and the sets are properly separated by the hyperplane $x_2 = 0$. To rule out the possibility just described, the notion of strict separation is introduced.

Definition 1.3. Two sets X and Y are strictly separated by a hyperplane H_C^+ if for every $x \in X$, $\langle \mathbf{a}, \mathbf{x} \rangle > \alpha$ and, for every $y \in Y$, $\langle \mathbf{a}, \mathbf{y} \rangle < \alpha$.

Geometrically, the definition requires that X and Y lie in opposite open half spaces determined by H_C^+ . Note that the sets X_2 and Y_1 cannot be strictly separated by any hyperplane. The sets

$$X_2 = \{(x_1, x_2) : 0 \leq x_1 \leq 1, 0 < x_2 \leq 1\}$$

and

$$Y_1 = \{(x_1, x_2) : 0 \leq x_1 \leq 1, -1 < x_2 < 0\}$$

are strictly separated by the hyperplane $x_2 = 0$. Note that X_2 and Y_1 are properly separated but are not strictly separated by the hyperplane $x_2 = 0$.

Finally, note that the sets X_1 and T_1 are not strictly separated by the hyperplane $x_1 = 0$.

Definition 3.4. Two sets X and F are *strongly separated* by a hyperplane H_C^* if there exists an $\varepsilon > 0$ such that the sets $X + \varepsilon B(0, 1)$ and $F + \varepsilon B(0, 1)$ are strictly separated by H_C^* .

The sets X_0 and K_0 are not strongly separated by the hyperplane $x_0 = 0$ and cannot be strongly separated by any hyperplane. Let $0 < \eta < 1$ be fixed. Then for each such η the sets

$$X_{\eta 0} = \{(x_0, x_1) \in \mathbb{R}^2 \mid 0 \leq x_0 \leq 1, \eta \leq x_1 \leq 1\}$$

and

$$Y_{\eta 0} = \{(x_0, x_1) \in \mathbb{R}^2 \mid 0 \leq x_0 \leq 1, -1 \leq x_1 \leq -\eta\}$$

are strongly separated by the hyperplane $x_1 = 0$.

It is clear from the definitions that strong separation implies strict separation, which implies proper separation.

Exercise 3.4. Sketch the pairs of sets (X_0, Y_0) , $(X_{\eta 0}, Y_{\eta 0})$, (X_C^*, Y_C^*) , $(X_{\eta C}^*, Y_{\eta C}^*)$ and $(X_{\eta 0}, Y_{\eta 0})$.

The next lemma gives two conditions, each of which is equivalent to strong separation as given in Definition 3.4.

Lemma 3.1. The following statements are equivalent:

- The sets X and F are strongly separated by a hyperplane H_C^* .
- There exists an $\varepsilon > 0$ such that $\langle \alpha, x \rangle \geq \varepsilon + \varepsilon$ for all x in X and $\langle \alpha, y \rangle \leq \varepsilon - \varepsilon$ for all y in F .
- There exists an $\varepsilon' > 0$ such that

$$\inf\{\langle \alpha, x \rangle \mid x \in X\} \geq \varepsilon + \varepsilon'$$

and

$$\sup\{\langle \alpha, y \rangle \mid y \in F\} \leq \varepsilon - \varepsilon'.$$

Proof. The equivalence of (i) and (ii) is clear, so we need only prove the equivalence of (i) and (iii). We note that $x \in B(0, 1)$ if and only if $-x \in B(0, 1)$. Suppose X and F are strongly separated by H_C^* . Then there exists an $\varepsilon > 0$ such that, for all x in X and all v in $B(0, 1)$, $\langle \alpha, x - \varepsilon v \rangle \geq \varepsilon$. Thus for all x in X and

all $x \in \mathbb{R}(0, 1)$,

$$-c(\mathbf{a}, x) \geq \alpha - c(\mathbf{a}, \mathbf{x}).$$

If we take $x = (\beta\mathbf{a})/|\mathbf{a}|$, $0 < \beta < 1$, and then let $\beta \rightarrow 1$, we get $-c(\mathbf{a}) \geq \alpha - c(\mathbf{a}, \mathbf{x})$, or

$$c(\mathbf{a}, \mathbf{x}) \geq \alpha + c(\mathbf{a}).$$

for all \mathbf{x} in X . A similar argument shows that

$$c(\mathbf{a}, \mathbf{y}) \leq \alpha - c(\mathbf{a})$$

for all \mathbf{y} in Y . Hence (i) holds for any positive number α strictly less than $c(\mathbf{a})$.

We now suppose that there exists an $\alpha > 0$ such that $c(\mathbf{a}, \mathbf{x}) \geq \alpha + \eta$ for all \mathbf{x} in X . Let \mathbf{r} be an arbitrary element of $\mathbb{R}(0, 1)$. Since $|\mathbf{r}| = 1$, from the Cauchy-Schwarz inequality we get that $c(\mathbf{a}, \mathbf{r}) < |\mathbf{a}|$, and consequently

$$c(\mathbf{a}, \mathbf{x} + \mathbf{r}) = c(\mathbf{a}, \mathbf{x}) - c(\mathbf{a}, \mathbf{r}) \geq \alpha + \eta - c(\mathbf{a}).$$

Similarly, for all \mathbf{y} in Y and $\mathbf{r} \in \mathbb{R}(0, 1)$,

$$c(\mathbf{a}, \mathbf{y} + \mathbf{r}) < \alpha - \eta + c(\mathbf{a}).$$

Hence if we take $\alpha < \eta/|\mathbf{a}|$, we have that $c(\mathbf{a}, \mathbf{x}) \geq \alpha$ for all \mathbf{x} in $X + \alpha\mathbb{R}(0, 1) = X + \alpha\mathbb{R}(1)$ and that $c(\mathbf{a}, \mathbf{r}) < \alpha$ for all \mathbf{r} in $Y + \alpha\mathbb{R}(0, 1)$. This proves (i).

Remark 3.3. From the preceding lemma, it should be clear that a necessary and sufficient condition for two sets X and Y to be strongly separated is that there exists a vector \mathbf{a} such that

$$\inf\{c(\mathbf{a}, \mathbf{x}) : \mathbf{x} \in X\} > \sup\{c(\mathbf{a}, \mathbf{y}) : \mathbf{y} \in Y\}.$$

Our principal objective in this section is to prove Theorem 3.4. Although this theorem is not the best possible separation theorem in E^n , it is the one that is valid in infinite-dimensional spaces and it does cover the situations that occur in optimization problems. The best possible separation theorem in E^n will be discussed in Section 5. Our proof of Theorem 3.4 will not be the most economical one but will be one that is suggested by rather obvious geometric considerations.

The principal step in the proof is to show that if C is a convex set and \mathbf{y} is a point not in C , then \mathbf{y} and C can be properly separated. If C is not convex, then it may not be possible to separate \mathbf{y} and C , as can be seen by taking, in E^2 , $\mathbf{y} = \mathbf{0}$ and C to be the disjunction of any circle with center at the origin.



Figure 2A.

We first show that if C is closed and convex, then strong separation is possible.

THEOREM 1.1. Let C be a closed convex set and y a vector such that $y \notin C$. Then there exists a hyperplane L_C^y that strongly separates y and C .

To motivate the proof, we argue heuristically from Figure 2A, in which C is assumed to have a tangent at each boundary point.

Draw a line from y to x_0 , the point of C that is closest to y . The vector $p = x_0 - y$ will be perpendicular to C in the sense that $p = x_0$ is the normal to the tangent line at x_0 . The tangent line, whose equation is $(y - x_0, x - x_0) = 0$, properly separates y and C . The point x_0 is characterized by the fact that $(y - x_0, x - x_0) < 0$ for all $x \in C$. To obtain strong separation, we merely move the line parallel to itself so as to pass through a point $x_0 \in (x_0, y)$. We now justify these steps in a series of lemmas, some of which are of interest in their own right.

If X is a set and y is a vector, then we define the distance from y to X , denoted by $dy(X)$, to be

$$dy(X) = \inf\{|y - x| : x \in X\}.$$

A point x_0 in X is said to attain the distance from y to X , or to be closest to y , if $dy(X) = |y - x_0|$.

LEMMA 1.2. Let C be a convex subset of \mathbb{R}^n and let $y \notin C$. If there exists a point in C that is closest to y , then it is unique.

Proof. Suppose that there were two points x_1 and x_2 of C that were closest to y . Then since C is convex, $(x_1 + x_2)/2$ belongs to C , and so

$$\begin{aligned} d(y, C) &\leq \|(x_1 + x_2)/2 - y\| = \frac{1}{2}(\|x_1 - y\| + \|x_2 - y\|) \\ &\leq \frac{1}{2}(\|x_1 - y\| + \|x_2 - y\|) = d(y, C). \end{aligned}$$

Hence equality holds in the application of the triangle inequality, and so by (5) of Section 2 in Chapter I, we have that, for some $\alpha > 0$, $\|x_1 - y\| = \alpha\|x_2 - y\|$. Suppose $\alpha > 1$. Since $\|x_1 - y\| = \|x_2 - y\| = d(y, C)$, we get that $\alpha = 1$, and so $x_1 = x_2$. If $\alpha = 0$, we get $x_1 = y$, contradicting the assumption that $y \notin C$.

Lemma 3.3. *Let C be a closed subset of \mathbb{R}^n and let $y \notin C$. Then there exists a point x_0 in C that is closest to y .*

Proof. Let $x_0 \in C$ and let $r > \|x_0 - y\|$. Then $C_1 = \overline{B(y, r)} \cap C$ is nonempty, closed, and bounded and hence is compact. The function $x \rightarrow \|x - y\|$ is continuous on C_1 and so attains its minimum at some point x_0 in C_1 . Thus, for all $x \in C_1$, $\|x - y\| \geq \|x_0 - y\|$. For $x \in C$ and $x \notin C_1$, we have

$$\|x - y\| > r > \|x_0 - y\| \geq \|x_0 - y\|,$$

since $x_0 \in \overline{B(y, r)}$.

Lemmas 3.2 and 3.3 show that if C is a closed convex set and $y \notin C$, then there is a unique closest point in C to y .

The next lemma characterizes closest points.

Lemma 3.4. *Let C be a convex set and let $y \notin C$. Then $x_0 \in C$ is a closest point in C to y if and only if*

$$\langle y - x_0, x - x_0 \rangle \leq 0 \quad \text{for all } x \in C. \quad (1)$$

Note that if C is not convex, then the above characterization of closest point need not hold. To see this, let y be the origin in \mathbb{R}^2 and let C be the circumference of the unit circle. Let x_0 be a fixed point in C . Then it is not true that $\langle y - x_0, x - x_0 \rangle \leq 0$ for all x in C .

Proof. Let x_0 be a closest point to y and let x be any point in C . Since C is convex, the line segment $[x_0, x] = \{(t)x + (1-t)x_0 : 0 \leq t \leq 1\}$ belongs to C . Let

$$\varphi(t) = \|(t)x + (1-t)x_0 - y\|^2 = (x_0 + t(x - x_0) - y, x_0 + t(x - x_0) - y). \quad (2)$$

For $0 \leq t \leq 1$, $\varphi(t)$ is the square of the distance between the point $(t)x + (1-t)x_0$ and y . If $t = 0$, then $x = x_0$. Since φ is continuously differentiable on $(0, 1)$ and

\mathbf{x}_0 is a point in C closest to \mathbf{y} , we have $\rho(\delta) \geq 0$. Calculating $\rho(\delta)$ from (7) gives

$$\rho(\delta) = \delta - (\mathbf{y} - \mathbf{x}_0, \mathbf{y} - \mathbf{x}_0) + \delta (\mathbf{y} - \mathbf{x}_0)^2. \quad (8)$$

If we now let $\delta \rightarrow 0+$ and use $\rho(\delta) \geq 0$, we get (5).

We now suppose that (1) holds. Let \mathbf{x} be any other point in C . It follows from (8) that $\rho(\delta) \geq 0$ for $0 < \delta < 1$. This ρ is strictly increasing function on $[0, 1]$, and so for any point $\delta(\mathbf{x})$ in $(\mathbf{x}_0, \mathbf{x})$. We have $\|\delta(\mathbf{x}) - \mathbf{y}\| = \|\mathbf{x}_0 - \mathbf{y}\|$. In particular, this is true for $\mathbf{y} = \mathbf{x}$, and so \mathbf{x}_0 is a closest point to \mathbf{y} .

We can now complete the proof of Theorem 3.1. Let \mathbf{x}_0 be the closest point in C to \mathbf{y} and let $\mathbf{a} = \mathbf{y} - \mathbf{x}_0$. Then for all $\mathbf{x} \in C$, $\langle \mathbf{a}, \mathbf{x} - \mathbf{x}_0 \rangle \leq 0$, and so $\langle \mathbf{a}, \mathbf{x} \rangle \leq \langle \mathbf{a}, \mathbf{x}_0 \rangle$, with equality occurring when $\mathbf{x} = \mathbf{x}_0$. Therefore $\sup\{\langle \mathbf{a}, \mathbf{x} \rangle : \mathbf{x} \in C\} = \langle \mathbf{a}, \mathbf{x}_0 \rangle$. On the other hand, $\langle \mathbf{a}, \mathbf{y} - \mathbf{x}_0 \rangle = \|\mathbf{a}\|^2 > 0$, so

$$\langle \mathbf{a}, \mathbf{y} \rangle = \langle \mathbf{a}, \mathbf{x}_0 \rangle + \|\mathbf{a}\|^2 > \sup\{\langle \mathbf{a}, \mathbf{x} \rangle : \mathbf{x} \in C\}.$$

The conclusion of the theorem now follows from Remark 3.1 with $K = \{\mathbf{y}\}$ and $X = C$.

We now take up the separation of a point \mathbf{y} and an arbitrary convex set; that is, we shall no longer assume that C is closed. We shall obtain separation, but shall not be able to guarantee proper separation unless we assume that C has nonempty interior.

Theorem 3.2. Let C be a convex set and let $\mathbf{y} \notin C$. Then there exists a hyperplane $H_C^\mathbf{y}$ such that for all \mathbf{x} in C

$$\langle \mathbf{a}, \mathbf{x} \rangle < \alpha \quad \text{and} \quad \langle \mathbf{a}, \mathbf{y} \rangle = \alpha. \quad (9)$$

If $\text{int}(C) \neq \emptyset$, then for all \mathbf{x} in $\text{int}(C)$

$$\langle \mathbf{a}, \mathbf{x} \rangle < \alpha. \quad (10)$$

Proof. We first suppose that $\mathbf{y} = \mathbf{0}$. Then to establish (9), we must find an $\mathbf{a} \neq \mathbf{0}$ such that

$$\langle \mathbf{a}, \mathbf{x} \rangle < 0 \quad \text{for all } \mathbf{x} \text{ in } C. \quad (11)$$

For each \mathbf{x} in C define

$$N_{\mathbf{x}} = \{\mathbf{z} \mid \langle \mathbf{z}, \mathbf{x} \rangle = 0, \langle \mathbf{z}, \mathbf{x} \rangle \leq 0\}.$$

The set $N_{\mathbf{x}}$ is nonempty since it contains the element $-\mathbf{x}/\|\mathbf{x}\|$.

To establish (4), it suffices to show that

$$\bigcap_{\alpha \in I} N_\alpha \neq \emptyset.$$

Each N_α is closed and is contained in $S(\mathbb{R}, 1) = \{\mathbf{x} : \|\mathbf{x}\| = 1\}$. Thus we may write the last relation as

$$\bigcap_{\alpha \in I} N_\alpha = \left(\bigcap_{\alpha \in I} N_\alpha \right) \cap S(\mathbb{R}, 1) \neq \emptyset. \quad (5)$$

Since $S(\mathbb{R}, 1)$ is compact, it has the finite-intersection property. Thus, if the intersection in (5) were empty, there would exist a finite subcollection $N_{\alpha_1}, \dots, N_{\alpha_k}$ that is empty. Therefore, we may establish (5) by showing that every finite subcollection $N_{\alpha_1}, \dots, N_{\alpha_k}$ has nonempty intersection.

Now consider any finite collection of sets $N_{\alpha_1}, \dots, N_{\alpha_k}$ and the corresponding k points $\mathbf{x}_1, \dots, \mathbf{x}_k$ in C . Let $\text{co}\{\mathbf{x}_1, \dots, \mathbf{x}_k\}$ denote the convex hull of $\mathbf{x}_1, \dots, \mathbf{x}_k$. Then $\text{co}\{\mathbf{x}_1, \dots, \mathbf{x}_k\}$ is compact and convex. Since C is convex, $\text{co}\{\mathbf{x}_1, \dots, \mathbf{x}_k\} \subset C$. Thus, since $\mathbb{H} \cap C$, we have $\mathbb{H} \cap \text{co}\{\mathbf{x}_1, \dots, \mathbf{x}_k\}$. Hence, by Theorem 3.1, there exists a vector $\mathbf{w} \neq \mathbf{0}$ such that

$$\mathbf{0} = \langle \mathbf{w}, \mathbf{0} \rangle > \langle \mathbf{w}, \mathbf{x}_i \rangle, \quad \text{for } \text{co}\{\mathbf{x}_1, \dots, \mathbf{x}_k\}. \quad (6)$$

We may divide through by $\|\mathbf{w}\| \neq 0$ in (6) and so assume that $\|\mathbf{w}\| = 1$. Thus, (6) says that

$$\mathbf{w} \in \left(\bigcap_{i=1}^k N_{\alpha_i} \right) \cap S(\mathbb{R}, 1),$$

and we have shown that the intersection of any finite subcollection $N_{\alpha_1}, \dots, N_{\alpha_k}$ of the closed sets $\{N_\alpha\}_{\alpha \in I}$ has nonempty intersection with $S(\mathbb{R}, 1)$. Hence (5) holds.

We now remove the restriction that $p = 0$. We have

$$y \in C \quad \text{if and only if } \mathbb{H} \cap C - y = \{x' : x' = x - y, x \in C\}.$$

Therefore, there exists an $\mathbf{u} \neq \mathbf{0}$ such that $\langle \mathbf{u}, x' \rangle \leq 0$ for all x' in $C - y$. Hence for all x in C ,

$$\langle \mathbf{u}, x - y \rangle \leq 0,$$

and so

$$\langle \mathbf{u}, x \rangle \leq \langle \mathbf{u}, y \rangle \quad \text{for all } x \text{ in } C.$$

If we now let $x = (x, y)$, we get the first conclusion of the theorem. The second follows from Exercise 2.16.

Our first separation result for convex sets will follow from Theorem 3.2.

Theorem 3.3. Let X_0 and Y be two disjoint convex sets. Then there exists a hyperplane H_0^c that separates them.

Note that the theorem does not assert that proper separation can be achieved. Theorem 3.2 in the sequel will allow us to conclude that the separation is proper. Our proof of Theorem 3.3 does not yield this fact.

To prove the theorem, we first note that

$$X_0 \cap Y = \emptyset \quad \text{if and only if } \theta \notin X_0 - Y$$

Let

$$A = X_0 - Y = X_0 + (-1)Y$$

Since X_0 and Y are convex, by Lemma 2.2, so is A . Since X_0 and Y are disjoint, $\theta \notin A$. Hence by Theorem 2.2, there exists an $\alpha \neq \theta$ such that

$$\langle \alpha, x \rangle \leq 0 \quad \text{all } x \text{ in } A.$$

Let x be an arbitrary element of X_0 and let y be an arbitrary element of Y . Then $x - x - y$ is in A and

$$\langle \alpha, x \rangle \leq \langle \alpha, y \rangle. \quad (7)$$

Let $\beta = \sup\{\langle \alpha, x \rangle : x \in X_0\}$ and let $\gamma = \inf\{\langle \alpha, y \rangle : y \in Y\}$. Then from (7) we get that β and γ are finite, and for any number α such that $\beta \leq \alpha \leq \gamma$,

$$\langle \alpha, x \rangle \leq \alpha \leq \langle \alpha, y \rangle$$

for all x in X_0 and all y in Y . Thus the hyperplane H_0^c separates X_0 and Y .

Theorem 3.4. Let K and F be two convex sets such that $\text{int}(K)$ is not empty and $\text{int}(K)$ is disjoint from F . Then there exists a hyperplane H_0^c that properly separates K and F .

To illustrate the theorem, let $K = \{(x_1, x_2) : x_1 \leq 0, -1 \leq x_2 \leq 1\}$ and let $F = \{(x_1, x_2) : x_1 = 0, -1 \leq x_2 \leq 1\}$. The hypotheses of the theorem are fulfilled and $x_1 = 0$ properly separates K and F and hence K and F . Note that strict separation of K and F is not possible.

Proof. Since $\text{int}(X)$ is convex and disjoint from Y , we can apply Theorem 3.3 with $K_1 = \text{int}(X)$ and obtain the existence of an $\alpha \neq 0$ and an r such that

$$\langle \alpha, x \rangle \leq r \leq \langle \alpha, y \rangle \quad (8)$$

for all x in $\text{int}(X)$ and all y in Y .

Let $x \in \bar{X}$. By Corollary 2 to Lemma 2.5, $\bar{X} = \overline{\text{int}(X)}$. Hence $x \in \overline{\text{int}(X)}$. By Remark 1.2.2 there exists a sequence of points $\{x_n\}$ in $\text{int}(X)$ such that $x_n \rightarrow x$. It then follows from (8) and the continuity of the inner product that (8) holds for all x in \bar{X} and all $y \in Y$. Thus, the hyperplane H_r^α separates \bar{X} and Y . By Corollary 2 to Lemma 2.5, $\text{int}(\bar{X}) = \text{int}(X)$. By Exercise 2.16, for $x \in \text{int}(\bar{X})$, $\langle \alpha, x \rangle < r$, so the separation is proper.

Theorem 3.5. Let K be a compact convex set and let C be a closed convex set such that K and C are disjoint. Then K and C can be strongly separated.

Proof. Let

$$\begin{aligned} K_\epsilon &= K + B(0, \epsilon) = \{x + \alpha \mid x \in K, \|\alpha\| < \epsilon\}, \\ C_\epsilon &= C + B(0, \epsilon) = \{y + \beta \mid y \in C, \|\beta\| < \epsilon\}. \end{aligned}$$

Since $K_\epsilon = \bigcup_{\alpha \in B(0, \epsilon)} (K + \alpha)$, it is a union of open sets and hence is open. It is also readily verified that K_ϵ is convex. Similarly, the set C_ϵ is open and convex. (See Exercise 2.14.)

We now show that there exists an $\epsilon > 0$ such that $K_\epsilon \cap C_\epsilon = \emptyset$. If the assertion were false, there would exist a sequence $\{\epsilon_n\}$ with $\epsilon_n > 0$ and $\epsilon_n \rightarrow 0$ and a sequence $\{\mathbf{w}_n\}$ such that, for each n , $\mathbf{w}_n \in K_{\epsilon_n} \cap C_{\epsilon_n}$. Since $\mathbf{w}_n = \mathbf{u}_n + \mathbf{v}_n$ with $\mathbf{u}_n \in K$, $\|\mathbf{v}_n\| < \epsilon_n$ and $\mathbf{w}_n = \mathbf{y}_n + \mathbf{z}_n$ with $\mathbf{y}_n \in C$ and $\|\mathbf{z}_n\| < \epsilon_n$, we have a sequence $\{\mathbf{u}_n\}$ in K and a sequence $\{\mathbf{y}_n\}$ in C such that

$$\|\mathbf{u}_n - \mathbf{u}_m\| < \epsilon_n, \quad \|\mathbf{y}_n - \mathbf{y}_m\| < \epsilon_n.$$

Hence

$$\|\mathbf{u}_n - \mathbf{y}_n\| = \|(\mathbf{u}_n - \mathbf{u}_m) + (\mathbf{u}_m - \mathbf{y}_m) + (\mathbf{y}_m - \mathbf{y}_n)\| \leq \|\mathbf{u}_n - \mathbf{u}_m\| + \|\mathbf{y}_m - \mathbf{y}_n\| < 2\epsilon_n. \quad (9)$$

Since K is compact, there exists a subsequence $\{\mathbf{u}_{n_j}\}$ that converges to an element \mathbf{u}_0 in K . It follows from (9) that $\mathbf{y}_{n_j} \rightarrow \mathbf{u}_0$. Since C is closed, $\mathbf{u}_0 \in C$. This contradicts the assumption that C and K are disjoint, and so the assertion is true.

We have shown that there exists an $\epsilon > 0$ such that K_ϵ and C_ϵ are disjoint open convex sets. Hence by Theorem 3.4 there is a hyperplane that properly separates K_ϵ and C_ϵ . Since both K_ϵ and C_ϵ are open sets, it follows from Exercise 2.18 that K_ϵ and C_ϵ are strictly separated. According to Definition 3.4, this says that K and C are strongly separated.

We now apply the separation theorem to obtain an additional characterization of convex sets.

THEOREM 3.6. Let A be a set contained in some half space. Then the closure of the convex hull of A , $\text{co}(A)$, is the intersection of all closed half spaces containing A .

Proof. We select the hyperplanes that determine the closed half spaces containing A to be the negative closed half spaces. If H_0^+ is a closed half space containing A , then $\text{co}(A)$ is contained in H_0^+ and hence in the intersection of all such half spaces. Now let \mathbf{x} be a point in the intersection. If \mathbf{x} were not in $\text{co}(A)$, then by Theorem 3.1, there would exist a hyperplane H_0^+ such that

$$\langle \mathbf{n}, \mathbf{x} \rangle < \alpha < \langle \mathbf{n}, \mathbf{a} \rangle, \quad \mathbf{a} \in \overline{\text{co}(A)}.$$

Hence \mathbf{x} would not be in the closed half space by H_0^+ which contains A . Therefore \mathbf{x} is not in the intersection of all negative closed half spaces containing A , which contradicts the assumption that \mathbf{x} is in the intersection. Thus the intersection is contained in $\text{co}(A)$, and the theorem is proved.

The following corollary follows from the theorem and the fact that a closed convex set not equal to \mathbb{R}^n is contained in some closed half space.

COROLLARY 3.1. If A is a closed convex set not equal to \mathbb{R}^n , then A is the intersection of all closed half spaces containing A .

Note that the theorem is false if we replace $\overline{\text{co}(A)}$ by $\text{co}(A)$. To see this, consider the set A in \mathbb{R}^2 defined by

$$A = \{(x_1, x_2) : (x_1 \geq 0, x_2 \geq 1)\} \cup \{(x_1, x_2) : (x_1 \geq 0, x_2 \leq -1)\}.$$

Then $\text{co}(A) = \{(x_1, x_2) : x_1 \geq 0\}$, and the intersection of all closed half spaces is $\{(x_1, x_2) : (x_1 \geq 0)\}$.

Exercise 3.2. Let V be a linear subspace and let $\mathbf{y} \in V$. Show that \mathbf{x}_0 is V the closest point in V to \mathbf{y} if and only if $\mathbf{y} - \mathbf{x}_0$ is orthogonal to V ; that is, for every \mathbf{w} in V , $\mathbf{y} - \mathbf{x}_0$ is orthogonal to \mathbf{w} .

Exercise 3.3. Let C be a closed convex set and let $\mathbf{p} \in C$. Show that \mathbf{x}_0 is C is closest to \mathbf{p} if and only if $\langle \mathbf{x} - \mathbf{x}_0, \mathbf{x}_0 - \mathbf{p} \rangle \geq \|\mathbf{x}_0 - \mathbf{p}\|^2$ for all \mathbf{x} in C .

Exercise 3.4. In reference to Theorem 3.5, show that if K is assumed to be convex and closed, then it may not be possible to separate K and C strictly, let alone strongly.

Exercise 3A. Show that if F is closed and K is compact, then $K + F$ is closed. Give an example to which F and K are closed yet $F + K$ is not.

Exercise 3B. Let A and B be two compact sets. Show that A and B can be strongly separated if and only if $\text{co}(A) \cap \text{co}(B) = \emptyset$.

Exercise 3C. Prove Theorem 3.5 using Theorem 3.1 and Exercise 3A.

Exercise 3D. Let A be a bounded set. Show that $\text{co}(A)$ is the intersection of all closed half spaces containing A . Show that the statement is false if A is not bounded and a proper subset of \mathbb{R}^n .

Exercise 3E. Let A be a closed convex set such that $\text{co}(A)$ (the complement of A) is convex. Show that A is a closed half space.

Exercise 3F. Let C_1 and C_2 be two convex subsets of \mathbb{R}^n . Show that there exists a hyperplane that strongly separates C_1 and C_2 if and only if

$$\inf\{\|x - y\| : x \in C_1, y \in C_2\} > 0.$$

4. SUPPORTING HYPERPLANES; EXTREME POINTS

In studying convex sets we do not wish to restrict ourselves to sets with smooth boundaries, that is, sets such that every boundary point has a tangent plane. The concept that replaces that of tangent plane is that of supporting hyperplane. It is a concept that has many important consequences, as we shall see.

We first define what is meant by a boundary point of a set. A point x is said to be a boundary point of a set S if for every $\epsilon > 0$ the ball $B(x, \epsilon)$ contains a point $y \in S$ and a point $z \notin S$. Note that the definition allows us to have $y = z$ or $x = z$. For example, in \mathbb{R}^2 let $S = \{x = (x_1, x_2) : 0 \leq |x| \leq 1\}$. Then θ is a boundary point of S , and for every $\epsilon < 1 - |\theta|$ the only point in $B(\theta, \epsilon)$ that is not in S is θ . The other boundary points of S are the points x with $\|x\| = 1$. If there exists an $\epsilon > 0$ such that the only point in $B(x, \epsilon)$ that is in S is x itself, then we say that x is an isolated point of S .

Definition 4.1. A hyperplane H_C is said to be a supporting hyperplane to a set $S \subseteq \mathbb{R}^n$ for every $x \in S$, $\langle a, x \rangle \leq \alpha$ and there exists at least one point $x_0 \in S$ such that $\langle a, x_0 \rangle = \alpha$. The hyperplane H_C is said to support S at x_0 . The hyperplane H_C is a nontrivial supporting hyperplane if there exists a point $x_1 \in S$ such that $\langle a, x_1 \rangle < \alpha$.

To illustrate this definition, consider the set S in \mathbb{R}^2 defined by

$$S = \{x \in \mathbb{R}^2 : x_1^2 + x_2^2 < 1, x_2 = 0\}.$$

As a set in \mathbb{R}^2 , every point of S is a boundary point. The only supporting hyperplane at a point of S is the trivial one, $\alpha_1 = 0$. Every point in the set $E = \{(x, x^2) \mid x \in [-1, 1], \alpha_1 = 0\}$ is also a boundary point of S . At every point of E there is a nontrivial supporting hyperplane with equation $\alpha_1 + b\alpha_2 = 1$ for appropriate a and b .

Remark 4.1. Let K be a compact set in \mathbb{R}^n . Then for each \mathbf{a} in \mathbb{R}^n there exists an α such that the hyperplane $H_{\mathbf{a}}^{\alpha}$ with equation $\langle \mathbf{a}, \mathbf{x} \rangle = \alpha$ is a supporting hyperplane to K . To see this, note that the linear functional L defined by $L(\mathbf{x}) = \langle \mathbf{a}, \mathbf{x} \rangle$ is continuous on K and so attains its maximum at a point \mathbf{x}_0 in K . Then $\langle \mathbf{a}, \mathbf{x} \rangle \leq \langle \mathbf{a}, \mathbf{x}_0 \rangle$ for all $\mathbf{x} \in K$. The desired hyperplane is obtained by taking $\alpha = \langle \mathbf{a}, \mathbf{x}_0 \rangle$.

Theorem 4.1. Let C be a convex set and let \mathbf{p} be a boundary point of C . Then there exists a supporting hyperplane $H_{\mathbf{p}}^{\alpha}$ to C such that $\mathbf{p} \in H_{\mathbf{p}}^{\alpha}$, $H_{\mathbf{p}}^{\alpha} \cap \text{int}(C) = \emptyset$, the supporting hyperplane is nontrivial.

Proof. If $\mathbf{p} \in C$, then by Theorem 3.2 with $\mathbf{p} = \mathbf{p}$ there exists a hyperplane $H_{\mathbf{p}}^{\alpha}$ such that for all \mathbf{x} in C

$$\langle \mathbf{a}, \mathbf{x} \rangle \leq \alpha \quad \text{and} \quad \langle \mathbf{a}, \mathbf{p} \rangle = \alpha.$$

Since $\mathbf{p} \in C$ and \mathbf{p} is a boundary point of C , it follows that $\mathbf{p} \in \bar{C}$. Thus, $H_{\mathbf{p}}^{\alpha}$ is a supporting hyperplane to C with $\mathbf{p} \in H_{\mathbf{p}}^{\alpha}$.

If $\mathbf{p} \notin C$, then from the definition of boundary point it follows that there exists a sequence of points $\{\mathbf{x}_k\}$ with $\mathbf{x}_k \in C$ such that $\mathbf{x}_k \rightarrow \mathbf{p}$. It follows from Theorem 3.2 that for each positive integer k there exists a vector $\mathbf{a}_k \neq \mathbf{0}$ such that

$$\langle \mathbf{a}_k, \mathbf{x} \rangle \leq \langle \mathbf{a}_k, \mathbf{x}_k \rangle \quad \text{for all } \mathbf{x} \in C. \quad (1)$$

If we divide through by $\|\mathbf{a}_k\| \neq 0$ in (1), we see that we may assume that $\|\mathbf{a}_k\| = 1$. Since $\mathbb{S}(\mathbb{R}, 1)$ is compact, there exists a subsequence of $\{\mathbf{a}_k\}$ that we again label as $\{\mathbf{a}_k\}$ and a vector \mathbf{a} with $\|\mathbf{a}\| = 1$ such that $\mathbf{a}_k \rightarrow \mathbf{a}$. If we now let $k \rightarrow \infty$ in (1), we get that

$$\langle \mathbf{a}, \mathbf{x} \rangle \leq \langle \mathbf{a}, \mathbf{p} \rangle \quad \text{for all } \mathbf{x} \in C.$$

If we set $\alpha = \langle \mathbf{a}, \mathbf{p} \rangle$, we see that the hyperplane $H_{\mathbf{p}}^{\alpha}$ contains \mathbf{p} and is a supporting hyperplane to C .

It follows from Exercise 2.16 that the supporting hyperplane is nontrivial if $\text{int}(C) \neq \emptyset$.

Definition 4.2. A point \mathbf{x}_0 belonging to a convex set C is an *extreme point* of C if, for no pair of distinct points $\mathbf{x}_1, \mathbf{x}_2$ in C and for no $\alpha > 0, \beta > 0, \alpha + \beta = 1$, is it true that $\mathbf{x}_0 = \alpha\mathbf{x}_1 + \beta\mathbf{x}_2$.

In other words, a point x_0 in a convex set C is an extreme point of C if x_0 is interior to no closed interval $[x_1, x_2]$ in C . We leave it as an exercise for the reader to show that an extreme point is a boundary point.

Let

$$C_1 = \{(x_1, x_2) : x_1^2 + x_2^2 \leq 1\}.$$

The set of extreme points of C_1 is $\{(x_1, x_2) : x_1^2 + x_2^2 = 1\}$.

Let

$$C_2 = \{(x_1, x_2) : x_1^2 + x_2^2 \leq 1, x_1 \geq 0\},$$

and let $C_3 = C_1 \cup C_2$. Then the set of extreme points of C_3 is C_2 .

A convex set need not have any extreme points, as evidenced by C_1 . Another example of a convex set without any extreme points is $\{(x_1, x_2) : x_1 = 0\}$.

THEOREM 4.2. Let C be a compact convex set. Then the set of extreme points of C is not empty, and every supporting hyperplane of C contains at least one extreme point of C .

The following observation will be used in the proof of the theorem.

LEMMA 4.1. Let C be a convex set and let H be a supporting hyperplane of C . A point x_0 in $H \cap C$ is an extreme point of C if and only if x_0 is an extreme point of $H \cap C$.

Proof. Let x_0 in $H \cap C$ be an extreme point of C . If x_0 were not an extreme point of $H \cap C$, then x_0 would be an interior point of some closed interval $[x_1, x_2]$ contained in $H \cap C$. But then $[x_1, x_2]$ would be contained in C , and x_0 would not be an extreme point of C . Thus, if x_0 is an extreme point of C , then it is also an extreme point of $H \cap C$.

Let x_0 be an extreme point of $H \cap C$. If x_0 were not an extreme point of C , then there would exist points x_1 and x_2 in C with at least one of x_1 and x_2 not in H such that for some $\theta < \theta < 1$

$$x_0 = \theta x_1 + (1 - \theta)x_2.$$

Suppose that $H = H_{\alpha}^{\beta}$ that $\langle \alpha, x \rangle \leq \beta$ for x in C , and that $x_1 \notin H_{\alpha}^{\beta}$. Then $\langle \alpha, x_1 \rangle > \beta$ and

$$\theta = \langle \alpha, x_0 \rangle = \theta \langle \alpha, x_1 \rangle + (1 - \theta) \langle \alpha, x_2 \rangle \leq \theta \beta + (1 - \theta) \beta = \beta.$$

This contradiction shows that x_0 must be an extreme point of C .

We now show that the set of extreme points of a compact convex set is not empty.

LEMMA 4.2. *The set of extreme points of a compact convex set C is not empty.*

Proof. If C is a singleton, then there is nothing to prove. Therefore we may suppose that C is not a singleton. Since C is compact and the norm is a continuous function, there exists a point x^* in C of maximum norm. Thus $\|x\| \leq \|x^*\|$ for all x in C and $\|x^*\| > 0$.

We assert that x^* is an extreme point of C . If x^* were not an extreme point, there would exist distinct points x_1 and x_2 in C different from x^* and scalars $\alpha > 0, \beta > 0, \alpha + \beta = 1$ such that

$$x^* = \alpha x_1 + \beta x_2.$$

Since

$$\|x^*\| = \|\alpha x_1 + \beta x_2\| \leq \alpha\|x_1\| + \beta\|x_2\|,$$

we see that if either $\|x_1\|$ or $\|x_2\|$ was strictly less than $\|x^*\|$, we would have

$$\|x^*\| < (\alpha + \beta)\|x^*\| = \|x^*\|.$$

This contradiction implies that

$$\|x_1\| = \|x_2\| = \|x^*\|. \quad (2)$$

But then,

$$\begin{aligned} \|x^*\|^2 &= \|\alpha x_1 + \beta x_2\|^2 = \alpha^2\|x_1\|^2 + 2\alpha\beta\langle x_1, x_2 \rangle + \beta^2\|x_2\|^2 \\ &\leq \alpha^2\|x_1\|^2 + 2\alpha\beta\|x_1\|\|x_2\| + \beta^2\|x_2\|^2 \\ &= (\alpha + \beta)^2\|x^*\|^2 = \|x^*\|^2. \end{aligned}$$

Hence we must have

$$2\alpha\beta\langle x_1, x_2 \rangle = 2\alpha\beta\|x_1\|\|x_2\|.$$

Since $\alpha > 0, \beta > 0$, we have $\langle x_1, x_2 \rangle = \|x_1\|\|x_2\|$. Therefore $x_2 = \alpha x_1$ for some $\alpha > 0$. The relation (2) implies that $\alpha = 1$, and so $x_2 = x_1$. This contradicts the assumption that x_1 and x_2 are distinct. Thus x^* is an extreme point.

We now complete the proof of the theorem. Let M_1^0 be a supporting hyperplane to C . Since C is compact, the point x_0 in C such that $\langle x_0, x_0 \rangle = 0$ is also in C . Hence $M_1^0 \cap C$ is nonempty and is compact and convex. By Lemma 4.2, $M_1^0 \cap C$ contains an extreme point x_1 . By Lemma 4.1, x_1 is also an extreme point of C .

THEOREM 4.3. Let C be a compact convex set. Let C_e denote the set of extreme points of C . Then $C = \text{co}(C_e)$.

We first note that by Theorem 4.2 the set C_e is not empty.

Since $C_e \subset C$ and C is convex, we have that $\text{co}(C_e) \subset \text{co}(C) = C$. Since C is compact, we have that $\overline{\text{co}(C_e)} \cap C = C$.

We now assert that $\overline{\text{co}(C_e)} \subset \text{co}(C_e)$. If the last assertion were not true, then there would exist an $\mathbf{x}_0 \notin \text{co}(C_e)$ and in C . Since $\text{co}(C_e)$ is closed and convex, it follows from Theorem 3.1 that there exists a hyperplane H_1^* such that for all $\mathbf{x} \in \text{co}(C_e)$

$$\langle \mathbf{n}, \mathbf{x} \rangle = \alpha = \langle \mathbf{n}, \mathbf{x}_0 \rangle, \quad (3)$$

let

$$\beta = \max\{\langle \mathbf{n}, \mathbf{x} \rangle : \mathbf{x} \in C\}. \quad (4)$$

Since C is compact, the maximum is attained at some point \mathbf{x}_1 in C , and thus H_1^* is a supporting hyperplane to C at \mathbf{x}_1 . By Theorem 4.2, H_1^* contains an extreme point \mathbf{x}_2 of C . Thus, $\mathbf{x}_2 \in H_1^* \cap \overline{\text{co}(C_e)}$, and so by (3)

$$\beta = \langle \mathbf{n}, \mathbf{x}_2 \rangle = \alpha.$$

Since $\mathbf{x}_0 \in C$, it follows from (3) and (4) that

$$\alpha < \langle \mathbf{n}, \mathbf{x}_0 \rangle \leq \beta.$$

Combining the last two inequalities gives $\beta < \beta$. This contradiction shows that $\overline{\text{co}(C_e)} = \text{co}(C_e)$.

In more general contexts Theorem 4.3 is known as the *Krein–Milman theorem*. We shall see in Section 7 that in \mathbb{R}^n the conclusion of Theorem 4.3 can be strengthened to $C = \text{co}(C_e)$.

We conclude this section with an example of a compact convex set whose set of extreme points is not closed. In \mathbb{R}^2 let

$$C_1 = \{x : |x_1 - 1|^2 + x_2^2 \leq 1, x_1 \geq 0\},$$

$$C_2 = \{x : x_1 = 0, x_2 = 0, -1 \leq x_2 \leq 1\},$$

and let $C = \text{co}(C_1 \cup C_2)$. Then C_e is the union of the points $(0, 0, 1)$, $(0, 0, -1)$ and all points other than the origin that are on the circumference of the circle in the plane $x_1 = 0$ whose center is at $(1, 0, 0)$ and whose radius is 1. The origin is a limit point of this set yet is not in the set.

Exercise 4.1. Let S be a convex set. Show that a point x_0 in S is an extreme point of S if and only if the set $F = \{x \in S \mid x \neq x_0\}$ is convex.

Exercise 4.2. Show by examples that Theorem 4.3 fails if the convex set C has a nonempty set of extreme points and is merely assumed to be closed or is merely assumed to be bounded rather than closed and bounded.

Exercise 4.3. Let S be a set in \mathbb{R}^n . Show that every point of S is either an interior point or a boundary point. Show that if S is compact, then the set of boundary points of S is not empty. Show that if x_0 is an extreme point of S , then it is a boundary point of S .

Exercise 4.4. Let C be a convex subset of \mathbb{R}^n with nonempty interior. Show that $C_c = \text{int}(C)$ is open.

Exercise 4.5. Show directly, without using Lemma 4.2, that the set of extreme points of a closed ball is the set of its boundary points. (Hint: To simplify notation, take the center of the ball to be the origin.)

Exercise 4.6. Let C be a closed convex set in \mathbb{R}^n not equal to \mathbb{R}^n . Show that C is the intersection of all closed half spaces containing C that are determined by the supporting hyperplanes of C .

Exercise 4.7. Let x_0 be a point in \mathbb{R}^n . Let C_n denote the closed cube with center at x_0 and length of side equal to 1. Thus

$$C_n = \{x \mid |x_i - x_{0i}| \leq \frac{1}{2}, i = 1, \dots, n\}.$$

The vertices of the cube are the 2^n points of the form (x_1, x_2, \dots, x_n) where the x_i take on the values ± 1 .

- Show that the set of 2^n vertices is the set of extreme points of C_n .
- Show directly, without using Theorem 4.3, that C_n is the convex hull of the set of vertices. (Hint: Use induction on n .)

5. SYSTEMS OF LINEAR INEQUALITIES: THEOREMS OF THE ALTERNATIVE

In this section we shall use the separation theorem to obtain so-called theorems of the alternative. Theorems of the alternative have the following form: Given two systems of inequalities, I and II, either system I has a solution or system II has a solution but never both. Necessary conditions for many optimization problems follow from these theorems.

The following definitions and lemma are needed in our discussion. A vector $\mathbf{x} = (x_1, \dots, x_n)$ in \mathbb{R}^n is said to be nonnegative, or $\mathbf{x} \geq \mathbf{0}$, if for every $j = 1, 2, \dots, n$, $x_j \geq 0$. A vector \mathbf{x} in \mathbb{R}^n is said to be positive, or $\mathbf{x} > \mathbf{0}$, if for every $j = 1, 2, \dots, n$, $x_j > 0$.

Lemma 3.1. Let A be an $m \times n$ matrix and let

$$C = \{\mathbf{w} \mid \mathbf{w} = A\mathbf{x}, \mathbf{x} \in \mathbb{R}^n, \mathbf{x} \geq \mathbf{0}\}.$$

Then C is a closed convex cone.

Proof. A straightforward calculation verifies that C is a convex cone.

To show that C is closed, we first consider the case in which the columns of A are linearly independent. Let \mathbf{w}_0 be a limit point of C , and let $\{\mathbf{w}_k\}$ be a sequence of points in C converging to \mathbf{w}_0 . Then there exists a sequence of points $\{\mathbf{x}_k\}$ in \mathbb{R}^n with $\mathbf{x}_k \geq \mathbf{0}$ such that $A\mathbf{x}_k \rightarrow \mathbf{w}_0$.

We show that $\{\mathbf{x}_k\}$ is bounded. If $\{\mathbf{x}_k\}$ were unbounded, there would exist a subsequence, again denoted by $\{\mathbf{x}_k\}$, such that $\|\mathbf{x}_k\| \rightarrow \infty$. All points in the sequence $\{\mathbf{x}_k / \|\mathbf{x}_k\|\}$ have norm 1, so there exists a subsequence and a point \mathbf{u}_0 of norm 1 such that $\mathbf{u}_0 / \|\mathbf{u}_0\| = \mathbf{u}_0$. From

$$\|\mathbf{x}_k\| A(\mathbf{x}_k / \|\mathbf{x}_k\|) \rightarrow \mathbf{w}_0 \quad \|\mathbf{x}_k\| \rightarrow \infty$$

we see that we must have $A\mathbf{u}_0 / \|\mathbf{u}_0\| = \mathbf{0}$. On the other hand, $\mathbf{u}_0 / \|\mathbf{u}_0\| = \mathbf{u}_0$ implies that $A\mathbf{u}_0 / \|\mathbf{u}_0\| = A\mathbf{u}_0$. Hence $A\mathbf{u}_0 = \mathbf{0}$. Recall that $\|\mathbf{u}_0\| = 1$, so that $\mathbf{u}_0 \neq \mathbf{0}$. Therefore the relation $A\mathbf{u}_0 = \mathbf{0}$ contradicts the linear independence of the columns of A .

Since $\{\mathbf{x}_k\}$ is bounded, $\{\mathbf{x}_k\}$ has a convergent subsequence, which we denote again by $\{\mathbf{x}_k\}$. Let $\mathbf{x}_0 = \lim \mathbf{x}_k$. Then $\mathbf{x}_0 \geq \mathbf{0}$. Also, $A\mathbf{x}_k \rightarrow A\mathbf{x}_0$, so that $A\mathbf{x}_0 = \mathbf{w}_0$. Thus $\mathbf{w}_0 \in C$, and C is closed.

If the columns of A are not linearly independent, we first show that any point \mathbf{x} in C can be expressed as a nonnegative linear combination of $p < n$ linearly independent columns of A . Let \mathbf{A}_j , $j = 1, \dots, n$, denote the j th column of A . Then there exists an $\mathbf{x} \in \mathbb{R}^n$ such that

$$\mathbf{x} = \sum_{j=1}^n x_j \mathbf{A}_j, \quad x_j \geq 0.$$

Since the columns of A are linearly dependent, there exists a $\mu = (\mu_1, \dots, \mu_p) \neq \mathbf{0}$ such that

$$\mu_1 \mathbf{A}_1 + \mu_2 \mathbf{A}_2 + \dots + \mu_p \mathbf{A}_p = \mathbf{0}.$$

Therefore, for any $\rho \in \mathbb{R}$,

$$\mathbf{x} = \sum_{j=1}^n (x_j - \rho \mu_j) \mathbf{A}_j = \sum_{j=1}^p (x_j - \rho \mu_j) \mathbf{A}_j + \sum_{j=p+1}^n x_j \mathbf{A}_j.$$

Since $\mu \neq 0$, the first case on the right exists. If $\mu_j < 0$, then $x_j - \mu x_j \geq 0$ whenever $\mu \geq \alpha_j/\beta_j$. In particular, since $\alpha_j \geq 0$, we have $x_j - \mu x_j \geq 0$ whenever $\mu \leq 0$. If $\mu_j > 0$, then $\alpha_j/\beta_j \leq 0$ and $x_j - \mu x_j \geq 0$ if and only if $\mu \leq \alpha_j/\beta_j$. If there exist indices j such that $\mu_j > 0$, set

$$\mu = \min \left\{ \frac{\alpha_j}{\beta_j} : \mu_j > 0 \right\}.$$

If $\mu_j < 0$ for all indices j for which $\mu_j \neq 0$, set

$$\mu = \max \left\{ \frac{\alpha_j}{\beta_j} : \mu_j < 0 \right\}.$$

Then $x_j - \mu x_j \geq 0$ for all $j = 1, \dots, n$ and $x_j - \mu x_j = 0$ for at least one value of j . We have now expressed x as a nonnegative linear combination of $q < n$ columns of A . We continue the process until we express x as a nonnegative linear combination of $p < n$ linearly independent columns of A .

Let v denote a choice of $p < n$ linearly independent columns of A and let $A(v)$ denote the corresponding $m \times p$ matrix. There are a finite number of such choices. Let

$$C(v) = \{w : w = A(v)y, w \in \mathbb{R}^m, y \in \mathbb{R}^p, y \geq 0\}.$$

We have shown that each set $C(v)$ is closed and that

$$C \subseteq \bigcup_v C(v).$$

We now show the opposite inclusion. Let $w \in C(v)$ for some v . By relabeling columns of A , we can assume without loss of generality that v selects the first p columns of A . Then there exists a $y \in \mathbb{R}^p, y \geq 0$ such that

$$w = \sum_{j=1}^p y_j A_j = \sum_{j=1}^p y_j A_j + \sum_{j=p+1}^n 0 A_j, \quad y_j \geq 0, \quad j = 1, \dots, p.$$

Let $x = (y_1, \dots, y_p, 0, \dots, 0)$. Then $w = Ax$, and so $w \in C$.

We have expressed C as the union of a finite number of closed sets. Hence C is closed and Lemma 5.1 is proved.

THEOREM 5.1 (FARKAS'S LEMMA). Let A be an $m \times n$ matrix and let b be a vector in \mathbb{R}^m . Then one and only one of the systems

$$\begin{array}{ll} \text{I.} & Ax \leq b, \quad \langle b, b \rangle > 0, \quad x \in \mathbb{R}^n, \\ \text{II.} & Ay = b, \quad y \geq 0, \quad y \in \mathbb{R}^n \end{array}$$

has a solution.

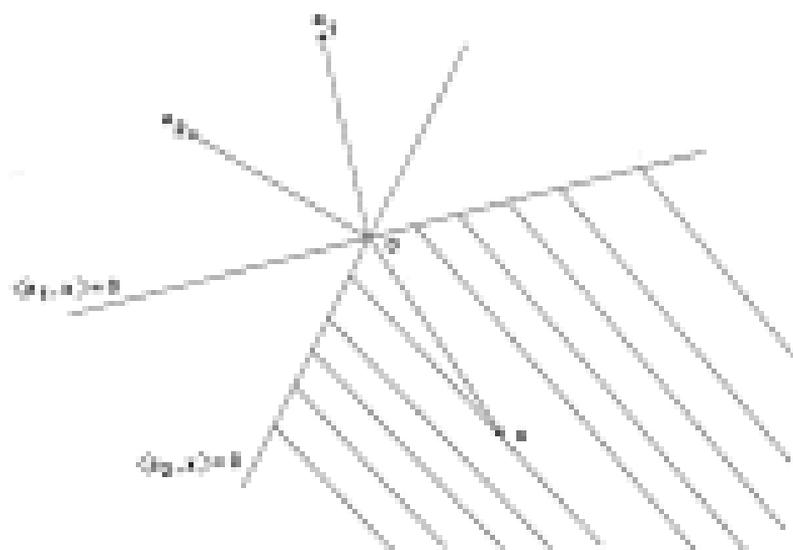


Figure 2.1.

The conclusion of the theorem is sometimes stated in the logically equivalent form: Either I has a solution or II has a solution but never both. It can also be stated in the following form: System I has a solution if and only if system II has no solution.

Figure 2.1 illustrates Farkas's lemma when A is a 2×2 matrix with rows a_1 and a_2 . Any vector x in the unshaded region satisfies $Ax \leq b$. Thus, if b does not lie in the acute angle determined by a_1 and a_2 , then I has a solution. (Draw the hyperplane $(b, x) = 0$ for such a b .) If we write II in the equivalent form $y'A = b$, $y' \geq 0$, we see that for II to have a solution b must lie in the angle determined by a_1, a_2 . The vector b is given, so its position is fixed. It either lies in the angle or does not. Thus, precisely one of the systems I and II has a solution.

We now prove the theorem. The reader will get an appreciation of the efficiency and aesthetic appeal of convexity arguments by looking up the original proof [Farkas, 1953].

We first suppose that II has a solution y_0 . We shall show that I has no solution. If $Ax > b$ for all x in \mathbb{R}^n , then I has no solution. Suppose there exists an x_0 in \mathbb{R}^n such that $Ax_0 \leq b$. Then

$$(b, x_0) = b'x_0 = (A'y_0)'x_0 = y_0'Ax_0 = y_0'(Ax_0).$$

Since $y_0 > 0$ and $Ax_0 \in \mathbf{R}$, it follows that $(b, x_0) \in \mathbf{R}$. Hence I has no solution.

We now suppose that II has no solution. Let

$$C = \{x \mid x = A'y, y \in \mathbb{R}^n, y \geq 0, ax \in \mathbb{R}\}.$$

The statement that II has no solution is equivalent to the statement that $b \notin C$. By Lemma 5.1, C is closed and convex. Therefore, by Theorem 3.1 there exists a hyperplane H_C^b that strongly separates b and C . Thus $\alpha_0 \notin \mathbf{R}$ and

$$(\alpha_0, b) > \alpha > (\alpha_0, x)$$

for all x in C . Since $b \in C$, we get that $\alpha = 0$. Thus

$$(\alpha_0, b) > 0. \quad (1)$$

For all x in C we have

$$(\alpha_0, x) = (\alpha_0, A'y) = y'A\alpha_0 = \alpha,$$

the last inequality now holding for all y in \mathbb{R}^n satisfying $y \geq 0$.

We claim that the last inequality implies that

$$A\alpha_0 \in \mathbf{R}. \quad (2)$$

If (2) were false, there would exist at least one component of $A\alpha_0$, say the j th component, that is positive. Let i_j denote this component. Let $y_j = \alpha y_1 = \alpha(0, \dots, 0, 1, 0, \dots, 0)$, where $1 \geq 0$. Then $(y_j, A\alpha_0) = i_j > \alpha$. Since $i_j > 0$, if we take i to be sufficiently large, we get a contradiction. Thus (2) is true. But (1) and (2) say that α_0 is a solution of I.

THEOREM 5.2 [GONTER 1973]. *Let A be an $m \times n$ matrix. Then one and only one of the following systems has a solution:*

$$\begin{array}{ll} \text{I} & Ax = b, \quad x \in \mathbb{R}^n \\ \text{II} & A'y = 0, \quad y \in \mathbb{R}^m, \quad y \neq 0, \quad y \geq 0. \end{array}$$

The reader should draw a figure which illustrates this theorem when A is a 3×2 matrix with rows $\alpha_1, \alpha_2, \alpha_3$.

Proof. We shall show that the system I has a solution if and only if the system

$$\text{III} \quad Ax \leq 0, \quad \zeta = 0,$$

has a solution, where ζ is the vector in \mathbb{R}^m all of whose entries equal 1. Let

(x_0, ζ_0) be a solution of F . Then

$$\Delta x_0 = \zeta_0 \delta < \mathbf{0},$$

so I has a solution. Now let x_0 be a solution of I and let $\zeta = \Delta x_0$. If $\zeta = (\zeta_1, \dots, \zeta_m)$, then $\zeta_i < 0$ for $i = 1, \dots, m$. Let $\zeta_0 = \max\{\zeta_1, \dots, \zeta_m\}$. Then (x_0, ζ_0) is a solution of F .

The system F has a solution if and only if the system

$$F: (A_1 - \sigma) \begin{pmatrix} x \\ \zeta \end{pmatrix} < \mathbf{0}, \quad \langle \mathbf{0}_{m-1}, (x, \zeta) \rangle > 0,$$

where $\mathbf{0}_k$ is the zero vector in \mathbb{R}^k , has a solution. By Farkas's lemma, either F (and hence I) has a solution or

$$F': \begin{pmatrix} A' \\ -\sigma' \end{pmatrix} y = \begin{pmatrix} \mathbf{0} \\ -1 \end{pmatrix}, \quad y \geq \mathbf{0}, \quad y \in \mathbb{R}^n,$$

has a solution but never both. System F' can be written as

$$F': A'y = \mathbf{0}, \quad \langle a, y \rangle = 1, \quad y \geq 0.$$

From $\langle a, y \rangle = 1$ it follows that $y \neq \mathbf{0}$. Hence if F' has a solution, so does

$$F'': A'y = \mathbf{0}, \quad y \geq \mathbf{0}, \quad y \neq \mathbf{0}.$$

If F'' has a solution, then $\langle a, y \rangle < 0$. If so, set $y' = y/\langle a, y \rangle$, then F' will have a solution. Hence either I has a solution or F'' has a solution but never both.

Theorem 5.3 [MIRMAN 1963]. Let A, B, C be given matrices with n columns and let $\mathbf{a} \neq \mathbf{0}$. Then one and only one of the following systems has a solution:

$$I: Ax > \mathbf{0}, \quad Bx \geq \mathbf{0}, \quad Cx = \mathbf{0}, \quad x \in \mathbb{R}^n,$$

$$II: \lambda'y_1 + \mu'y_2 + \nu'y_3 = \mathbf{0}, \quad y_1 \geq \mathbf{0}, \quad y_2 \neq \mathbf{0}, \quad y_3 \geq \mathbf{0}.$$

Proof. Let $\mathbf{e} = (1, 1, 1, \dots, 1)$ be a vector all of whose entries equal 1 and whose dimension is the row-dimension of A . We assert that the system I has a solution if and only if the system F has a solution, where

$$F: Ax \geq \zeta \mathbf{e}, \quad Bx \geq \mathbf{0}, \quad Cx = \mathbf{0}, \quad \zeta > 0.$$

Here ζ is a real scalar. It is clear that if F has a solution, then I has a solution. If I has a solution x_0 , let $\xi = (\xi_1, \dots, \xi_m) = Ax_0$. Let $\zeta = \min\{\xi_1, \dots, \xi_m\}$. Then $\zeta > 0$ and $\Delta x_0 \geq \zeta \mathbf{e}$, and so F has a solution.

Since I has a solution if and only if II has a solution, I has no solution if and only if II has no solution. Now, I has no solution if and only if the system

$$I): \begin{pmatrix} -A & B \\ -B & B \\ -C & B \\ C & B \end{pmatrix} \begin{pmatrix} x \\ z \end{pmatrix} \leq \begin{pmatrix} 0 \\ 0 \\ 0 \\ 0 \end{pmatrix}, \quad (0, 1) \leq (x, z) \leq 0,$$

where 0 , is the zero vector in E , has no solution.

By Farkas's lemma I has no solution if and only if the system II' has a solution, where

$$II'): \begin{pmatrix} -A' & -B' & -C' & C' \\ A' & B' & B' & 0 \end{pmatrix} \begin{pmatrix} y_1 \\ y_2 \\ y_3 \\ y_4 \end{pmatrix} = \begin{pmatrix} 0 \\ 0 \\ 1 \end{pmatrix}, \quad (0, y_2, y_3, y_4) \geq 0.$$

But II' having a solution is equivalent to the system

$$II): -A'y_1 + B'y_2 + C'y_3 - y_4 = 0, \quad (x, y_2) = 1,$$

having a solution $y_1 \geq 0$, $y_2 \geq 0$, $y_3 \geq 0$, $y_4 \geq 0$. If we let $y_2 = (y_3 = y_4) = 0$, we see that if II has a solution as indicated, then

$$III): -A'y_1 + B'y_2 + C'y_3 = 0, \quad y_1 \geq 0, \quad y_2 \geq 0, \quad y_3 \geq 0,$$

has a solution. In summary, we have shown that if I has no solution, then II has a solution.

Now suppose that II has a solution (y_1, y_2, y_3) . Then $(x, y_2) > 0$. Hence if we set

$$y'_1 = \frac{y_1}{(x, y_2)}, \quad y'_2 = \frac{y_2}{(x, y_2)}, \quad y'_3 = \frac{y_3}{(x, y_2)},$$

then (y'_1, y'_2, y'_3) is a solution of II with $(x, y'_2) = 1$. Thus we may assume that II has a solution (y_1, y_2, y_3) with $(x, y_2) = 1$.

Let η_1 be a vector whose dimension is that of y_1 and whose i th component is equal to y_{1i} , the i th component of y_1 , if $y_{1i} > 0$ and zero otherwise. Let η_2 be a vector whose dimension is that of y_2 and whose i th component is equal to y_{2i} if $y_{2i} > 0$ and zero otherwise. Then $\eta_1 \geq 0$, $\eta_2 \geq 0$ and $y_1 = (y_2 = y_3)$. Hence if II has a solution, then so does III. Therefore, I has no solution.

Remark 3.1. Theorem 5.3 is valid if either B or C or both are zero. In fact, if both B and C are zero, the lemma becomes Gordan's theorem.

THEOREM 5.4 [GAL 1965]. Let A be an $m \times n$ matrix and let $a, b \in \mathbb{R}$ be a vector in \mathbb{R}^m . Then one and only one of the following systems has a solution:

$$\begin{aligned} \text{I: } & Ax \leq a, \\ \text{II: } & A^T y = b, \quad (a, y) = -1, \quad y \geq 0 \end{aligned}$$

Proof. The system $Ax \leq a$ has a solution if and only if the system $\zeta \geq 0$, $Ax \leq \zeta$ has a solution. Thus the system I is equivalent to the system

$$\text{I': } (A, -a) \begin{pmatrix} x \\ \zeta \end{pmatrix} \leq 0, \quad \langle \theta_0, (x, \theta_0, \zeta) \rangle \geq 0,$$

where θ_0 is the zero vector in \mathbb{R}^n . By Farkas's lemma I' has a solution if and only if the system II'

$$\begin{pmatrix} A^T \\ -a \end{pmatrix} y = \begin{pmatrix} \theta_0 \\ 0 \end{pmatrix}, \quad y \geq 0, \quad y \in \mathbb{R}^m,$$

has no solution. But system II' is the same as

$$(a, y) = -1, \quad A^T y = \theta_0, \quad y \geq 0,$$

which is system II).

COROLLARY 1. Let A be an $m \times n$ matrix and let $a \neq 0$ in \mathbb{R}^m be such that, for all $x \in \mathbb{R}^n$, $Ax - a \geq 0$. Then there exists a vector $y \in \mathbb{R}^m$ such that $y \geq 0$, $A^T y = 0$, and $(a, y) = -1$.

To prove the corollary, we need only note that if $Ax - a \geq 0$ for all x in \mathbb{R}^n , then $Ax \leq a$ has no solution in \mathbb{R}^n .

Exercise 5.1. Show that the set C of Lemma 5.1 is a convex cone.

Exercise 5.2. Let A be an $m \times n$ matrix and b a vector in \mathbb{R}^m . Show that one and only one of the following systems has a solution:

$$\begin{aligned} \text{I: } & Ax \geq b, \quad x \geq 0, \quad (b, x) \geq 0 \\ \text{II: } & A^T y \geq b, \quad y \leq 0 \end{aligned}$$

Exercise 5.3. Let A be an $m \times n$ matrix and let $a \neq 0$ be a vector in \mathbb{R}^m . Show that one and only one of the following systems has a solution:

$$\begin{aligned} \text{I: } & Ax = a, \quad x \in \mathbb{R}^n, \\ \text{II: } & A^T y = 0, \quad (a, y) = 1. \end{aligned}$$

6. AFFINE SUBSETS

Consider a plane Π in \mathbb{R}^3 and a line A in \mathbb{R}^3 . Let C be a subset of Π and let S be a subset of A . Although C and S are subsets of \mathbb{R}^3 , it is natural to consider C as a two-dimensional object and S as a one-dimensional object. In this section we shall make these ideas precise for subsets of \mathbb{R}^n .

Let V_r denote an r -dimensional vector subspace of \mathbb{R}^n . When we do not wish to emphasize the dimension of the subspace, we shall write V . The reader will recall from linear algebra that an $(n - 1)$ -dimensional subspace V_{n-1} of \mathbb{R}^n can be expressed as

$$V_{n-1} = \left\{ \mathbf{x} : \sum_{j=1}^n a_j x_j = 0 \right\} = \{ \mathbf{x} : (\mathbf{a}, \mathbf{x}) = 0 \},$$

where \mathbf{a} is a nonzero vector. Thus V_{n-1} is a hyperplane through the origin.

Also, given a subspace V_{n-k} , there exist k linearly independent vectors $\mathbf{a}_1, \dots, \mathbf{a}_k$ such that

$$V_{n-k} = \{ \mathbf{x} : (\mathbf{a}_i, \mathbf{x}) = 0, i = 1, \dots, k \}.$$

Thus V_{n-k} , a subspace of dimension $n - k$, is the intersection of k hyperplanes through the origin.

Definition 6.1. A linear manifold M is a subset of \mathbb{R}^n of the form

$$M = V + \mathbf{h} = \{ \mathbf{x} : \mathbf{x} = \mathbf{v} + \mathbf{h}, \mathbf{v} \in V \},$$

where V is a subspace and \mathbf{h} is a vector in \mathbb{R}^n .

Note that since $\mathbf{h} = \mathbf{0} + \mathbf{h}$ and $\mathbf{0} \in V$, it follows that $\mathbf{h} \in M$.

Thus a linear manifold is a translate of a subspace. Since subspaces are closed sets in \mathbb{R}^n , linear manifolds are closed sets. Note that \mathbb{R}^n is a linear manifold, since $\mathbb{R}^n = \mathbb{R}^n + \mathbf{0}$, and that a subspace V is a linear manifold, since $V = V + \mathbf{0}$. In Section 1 we saw that a hyperplane is the translate of a hyperplane through the origin and thus the translate of an $(n - 1)$ -dimensional vector space. Therefore, hyperplanes are linear manifolds. In \mathbb{R}^3 the linear manifolds are planes, lines, and points. (A point is the translate of the zero-dimensional subspace consisting of the origin.)

Linear manifolds are also called *flat* or *affine varieties* or *affine sets*.

Lemma 6.1. Let M be a linear manifold, $M = V + \mathbf{h}$. Then the subspace V is unique.

To prove the lemma, we suppose that M' has another representation $M' = V' + \mathbf{k}'$ and show that $V' = V$. Since $\mathbf{k}_0 \in M$, from $M' = V' + \mathbf{k}'$, we have that $\mathbf{k}_0 = \mathbf{v}' + \mathbf{k}'$ for some $\mathbf{v}' \in V'$. Hence $\mathbf{k}' - \mathbf{k}_0 = -\mathbf{v}' \in V'$. Now let $\mathbf{x}_0 \in V$. Then $\mathbf{x}_0 = \mathbf{m} - \mathbf{k}_0$ for some $\mathbf{m} \in M$. But $\mathbf{m} = \mathbf{y} + \mathbf{k}'$ for some $\mathbf{y} \in V'$, and so

$$\mathbf{x}_0 = (\mathbf{y} + \mathbf{k}') - \mathbf{k}_0 = \mathbf{y} + (\mathbf{k}' - \mathbf{k}_0) \in V'.$$

Thus $V \subseteq V'$. A similar argument shows that $V' \subseteq V$.

If $M = V + \mathbf{k}$, then the vector \mathbf{k} is not unique. In fact, $M = V + \mathbf{k}'$ for any $\mathbf{k}' \in M$. To see this, let \mathbf{k}' be any vector in M and let $M' = V + \mathbf{k}'$. Let $\mathbf{m} \in M$, so $\mathbf{m} = \mathbf{v} + \mathbf{k}$ for some $\mathbf{v} \in V$. Since $\mathbf{k}' \in M$, we have $\mathbf{k}' = \mathbf{v}' + \mathbf{k}$ for some $\mathbf{v}' \in V$. Hence $\mathbf{m} = (\mathbf{v} - \mathbf{v}') + \mathbf{k}'$, which is an element of M' . Thus, $M' = M$. A similar argument shows that $M' \subseteq M$.

On the other hand, if \mathbf{k} is fixed and $M = V + \mathbf{k}$, then the representation of an element \mathbf{x} in M as $\mathbf{x} = \mathbf{v} + \mathbf{k}$ is unique. For if $\mathbf{x} = \mathbf{v}' + \mathbf{k}$, then $\mathbf{v} = \mathbf{x} - \mathbf{k}$ and $\mathbf{v}' = \mathbf{x} - \mathbf{k}$.

If M is a linear manifold, then the unique subspace V such that $M = V + \mathbf{k}$ for some \mathbf{k} is called the *parallel subspace* of M .

Definition 4.2. The *dimension* of a linear manifold M is defined to be the dimension of the parallel subspace V .

If M_1 and M_2 are two linear manifolds with the same parallel subspace, then M_1 and M_2 are said to be *parallel*.

The next lemma characterizes linear manifolds as sets having the property that if \mathbf{x}_1 and \mathbf{x}_2 belong to the set, then the entire line determined by these points lies in the set.

Lemma 4.3. A set M is a linear manifold if and only if, for every $\mathbf{x}_1, \mathbf{x}_2$ in M and λ_1, λ_2 real with $\lambda_1 + \lambda_2 = 1$, the point $\lambda_1 \mathbf{x}_1 + \lambda_2 \mathbf{x}_2$ belongs to M .

Proof. Let M be a linear manifold, $M = V + \mathbf{k}$. Let \mathbf{x}_1 and \mathbf{x}_2 be in M . Then $\mathbf{x}_1 = \mathbf{v}_1 + \mathbf{k}$ and $\mathbf{x}_2 = \mathbf{v}_2 + \mathbf{k}$, where \mathbf{v}_1 and \mathbf{v}_2 are in V . Hence for λ_1, λ_2 with $\lambda_1 + \lambda_2 = 1$,

$$\lambda_1 \mathbf{x}_1 + \lambda_2 \mathbf{x}_2 = \lambda_1(\mathbf{v}_1 + \mathbf{k}) + \lambda_2(\mathbf{v}_2 + \mathbf{k}) = (\lambda_1 \mathbf{v}_1 + \lambda_2 \mathbf{v}_2) + \mathbf{k}.$$

Since V is a subspace, $\lambda_1 \mathbf{v}_1 + \lambda_2 \mathbf{v}_2 \in V$, and so $\lambda_1 \mathbf{x}_1 + \lambda_2 \mathbf{x}_2 \in M$.

We now prove the reverse implication. Let M be a set such that if \mathbf{x}_1 and \mathbf{x}_2 are in M and λ_1, λ_2 are scalars such that $\lambda_1 + \lambda_2 = 1$, then $\lambda_1 \mathbf{x}_1 + \lambda_2 \mathbf{x}_2$ is in M . Let $\mathbf{x}_0 \in M$ and let $W = M - \mathbf{x}_0$. We shall show that W is a subspace, and then show that M is a linear manifold, since $M = W + \mathbf{x}_0$. Let \mathbf{y}_1 and \mathbf{y}_2

belong to W and let α and β be arbitrary scalars. Then

$$\alpha(\gamma_1 + \beta\gamma_2 + x_0) + \beta(\gamma_1 + x_0) + (1 - \alpha)(x_0) + [\beta(\gamma_1 + x_0) + (1 - \beta)x_0] = x_0.$$

Since $\gamma_1 + x_0$ and $\gamma_2 + x_0$ belong to M , each of the expressions in square brackets belongs to M . If we denote the expression in the first bracket by m_1 and the expression in the second bracket by m_2 , we may write

$$\alpha(\gamma_1 + \beta\gamma_2 + x_0) + \beta(\gamma_1 + x_0) + (1 - \alpha)(x_0) = \alpha m_1 + \beta m_2 + x_0,$$

where m_1 denotes the expression in curly braces. Since m_1 and m_2 are in M and the sum of the coefficients on the extreme right is 1, we get that $\alpha(\gamma_1 + \beta\gamma_2 + x_0) \in M$. Thus, $\alpha\gamma_1 + \beta\gamma_2 \in W$ and so W is a subspace.

COROLLARY 1. If M_1 and M_2 are linear manifolds, then so is $M_1 + M_2 = \{\alpha x = m_1 + m_2, m_1 \in M_1, m_2 \in M_2\}$.

COROLLARY 2. If $\{M_i\}$ is a family of linear manifolds, then so is $\bigcup_i M_i$.

We leave the proofs of these corollaries as exercises for the reader.

Definition 4.4. A point x is said to be an affine combination of points x_1, \dots, x_n if there exists a vector $q = (q_1, \dots, q_n)$ in \mathbb{R}^n such that

$$x = q_1 x_1 + \dots + q_n x_n \quad \text{and} \quad \sum_{i=1}^n q_i = 1.$$

COROLLARY 3. A set X is a linear manifold if and only if every affine combination of points in X is in X .

The proof is similar to that of Lemma 2.6.

Definition 4.5. A finite set of points x_1, \dots, x_n is affinely dependent if there exist real numbers $\lambda_1, \dots, \lambda_n$, not all zero, such that $\lambda_1 + \dots + \lambda_n = 0$ and $\lambda_1 x_1 + \dots + \lambda_n x_n = 0$. If the points x_1, \dots, x_n are not affinely dependent, then they are affinely independent.

Lemma 4.6. A set of points x_1, x_2, \dots, x_n is affinely dependent if and only if the vectors $(x_2 - x_1), (x_3 - x_1), \dots, (x_n - x_1)$ are linearly dependent.

Proof. Let x_1, \dots, x_n be affinely dependent. Then there exist real numbers $\lambda_1, \dots, \lambda_n$, not all zero, such that

$$\lambda_1 x_1 + \lambda_2 x_2 + \dots + \lambda_n x_n = 0, \quad \sum_{i=1}^n \lambda_i = 0. \quad (1)$$

Hence

$$\lambda_1 = -(\lambda_2 + \cdots + \lambda_r). \quad (7)$$

Not all of $\lambda_2, \dots, \lambda_r$ are zero, for if they were, λ_1 would also be zero, contradicting the assumption that not all of $\lambda_1, \dots, \lambda_r$ are zero. Substituting (7) into the first relation in (1) gives

$$\lambda_2(\mathbf{x}_1 - \mathbf{x}_2) + \lambda_3(\mathbf{x}_1 - \mathbf{x}_3) + \cdots + \lambda_r(\mathbf{x}_1 - \mathbf{x}_r) = \mathbf{0},$$

where $\lambda_2, \dots, \lambda_r$ are not all zero. Thus the vectors $(\mathbf{x}_1 - \mathbf{x}_2), \dots, (\mathbf{x}_1 - \mathbf{x}_r)$ are linearly dependent.

To prove the reverse implication, we reverse the argument.

Remark 6.3. Clearly, an equivalent statement of Lemma 6.2 is the following. A set of points $\mathbf{x}_1, \dots, \mathbf{x}_r$ is affinely independent if and only if the vectors $\mathbf{x}_2 - \mathbf{x}_1, \dots, \mathbf{x}_r - \mathbf{x}_1$ are linearly independent.

Lemma 6.4. Let M be a linear manifold. Then the following statements are equivalent:

- The dimension of M equals r .
- There exist $r + 1$ points $\mathbf{x}_0, \mathbf{x}_1, \dots, \mathbf{x}_r$ in M that are affinely independent and any set of $r + 2$ points in M is affinely dependent.
- There exist points $\mathbf{x}_0, \mathbf{x}_1, \dots, \mathbf{x}_r$ that are affinely independent and such that each point \mathbf{x} in M has a unique representation

$$\mathbf{x} = \sum_{i=0}^r \lambda_i \mathbf{x}_i, \quad \sum_{i=0}^r \lambda_i = 1 \quad (8)$$

in terms of the \mathbf{x}_i 's.

Proof. We first show that (1) implies (2). If (1) holds, then we may write $M = V_1 + \mathbf{x}_0$, where V_1 is a subspace of dimension r and \mathbf{x}_0 is an arbitrary point in M . Let $\mathbf{x}_1, \dots, \mathbf{x}_r$ be a basis for V_1 and let

$$\mathbf{x}_i = \mathbf{x}_0 + \mathbf{v}_i, \quad i = 1, \dots, r.$$

Then the vectors $\mathbf{x}_i - \mathbf{x}_0$, $i = 1, \dots, r$, are linearly independent and the $r + 1$ points $\mathbf{x}_0, \mathbf{x}_1, \dots, \mathbf{x}_r$ are affinely independent.

Let $\mathbf{x}_0, \mathbf{x}_1, \dots, \mathbf{x}_r, \mathbf{x}_{r+1}$ be any set of $r + 2$ points in M . If we write $M = V_1 + \mathbf{x}_0$, then each of the vectors $\mathbf{x}_i - \mathbf{x}_0$, $i = 1, \dots, r + 1$, is in V_1 . Therefore the $r + 1$ vectors in this set are linearly dependent, and hence $\mathbf{x}_0, \mathbf{x}_1, \dots, \mathbf{x}_r, \mathbf{x}_{r+1}$ are affinely dependent.

Next we show that (C) implies (B). Since M is a linear manifold and $\mathbf{x}_0 \in M$, we may write $M = F + \mathbf{x}_0$, where F is a subspace. We assert that F has dimension r . To see this, let

$$\mathbf{v}_i = \mathbf{x}_i - \mathbf{x}_0 \quad (i = 1, \dots, r).$$

Since the points $\mathbf{x}_0, \mathbf{x}_1, \dots, \mathbf{x}_r$ are affinely independent, the vectors $\mathbf{v}_1, \dots, \mathbf{v}_r$ are linearly independent. Now let $\mathbf{w}_1, \dots, \mathbf{w}_{r+1}$ be a set of $r+1$ distinct nonzero vectors in F . Then there exist points $\mathbf{p}_1, \dots, \mathbf{p}_{r+1}$ in M such that

$$\mathbf{w}_j = \mathbf{p}_j - \mathbf{x}_0 \quad (j = 1, \dots, r+1).$$

Since the \mathbf{w}_j are distinct and none of the vectors \mathbf{w}_j is zero, the $r+2$ points $\mathbf{x}_0, \mathbf{p}_1, \dots, \mathbf{p}_{r+1}$ are distinct. By hypothesis, they are affinely dependent. Hence the vectors $\mathbf{w}_1, \dots, \mathbf{w}_{r+1}$ are linearly dependent. Thus, F has dimension r , we may write $F = U_r$, and the vectors $\mathbf{v}_1, \dots, \mathbf{v}_r$ are a basis for U_r .

For arbitrary \mathbf{x} in M we have $\mathbf{x} = \mathbf{x}_0 + \mathbf{u}$, uniquely for some \mathbf{u} in U_r . Hence there exist unique real numbers $\lambda_1, \dots, \lambda_r$ such that

$$\mathbf{x} = \sum_{i=1}^r \lambda_i \mathbf{v}_i + \mathbf{x}_0 = \sum_{i=1}^r \lambda_i (\mathbf{x}_i - \mathbf{x}_0) + \mathbf{x}_0 = \left(1 - \sum_{i=1}^r \lambda_i\right) \mathbf{x}_0 + \sum_{i=1}^r \lambda_i \mathbf{x}_i.$$

We obtain the representation (1) by setting $\lambda_0 = (1 - \lambda_1 - \dots - \lambda_r)$.

To show that the representation is unique in terms of the \mathbf{x}_i 's, we suppose that there were another representation

$$\mathbf{x} = \sum_{i=0}^r \lambda'_i \mathbf{x}_i, \quad \sum_{i=0}^r \lambda'_i = 1.$$

Then

$$\mathbf{0} = \sum_{i=0}^r (\lambda'_i - \lambda_i) \mathbf{x}_i, \quad \sum_{i=0}^r (\lambda'_i - \lambda_i) = 0$$

with not all of the coefficients $(\lambda'_i - \lambda_i)$ equal to zero. But then the points \mathbf{x}_i , $i = 0, \dots, r$, would be affinely dependent, contrary to assumption.

We conclude the proof by showing that (B) implies (C). Let \mathbf{x} be an arbitrary point in M . Then by virtue of (1)

$$\mathbf{x} = \sum_{i=0}^r \lambda_i \mathbf{x}_i = \left(1 - \sum_{i=1}^r \lambda_i\right) \mathbf{x}_0 + \sum_{i=1}^r \lambda_i \mathbf{x}_i = \sum_{i=1}^r \lambda_i (\mathbf{x}_i - \mathbf{x}_0) + \mathbf{x}_0. \quad (2)$$

Since $\mathbf{x}_0, \mathbf{x}_1, \dots, \mathbf{x}_r$ are affinely independent, the vectors

$$\mathbf{v}_1 = \mathbf{x}_1 - \mathbf{x}_0, \mathbf{v}_2 = \mathbf{x}_2 - \mathbf{x}_0, \dots, \mathbf{v}_r = \mathbf{x}_r - \mathbf{x}_0$$

are linearly independent.

Let V_1 denote the $(r-1)$ -dimensional vector space spanned by $\mathbf{x}_1 - \mathbf{x}_r, \dots, \mathbf{x}_{r-1} - \mathbf{x}_r$. It follows from (4) that $\mathbf{x}_r \in V_1 + \mathbf{x}_r$. Thus $M \subseteq V_1 + \mathbf{x}_r$.

To prove the opposite inclusion, we let $W = M - \mathbf{x}_r$. We saw in the proof of Lemma 6.2 that W is a subspace. Moreover, since $\mathbf{x}_1 - \mathbf{x}_r, \dots, \mathbf{x}_{r-1} - \mathbf{x}_r$ belong to W , we have that $V_1 \subseteq W$. Hence

$$V_1 + \mathbf{x}_r \subseteq W + \mathbf{x}_r = M$$

and so $M = V_1 + \mathbf{x}_r$. This concludes the proof of Lemma 6.4.

Points $\mathbf{x}_0, \mathbf{x}_1, \dots, \mathbf{x}_r$ as in Lemma 6.4 are said to be an *affine basis* for M . The numbers λ_i , $i = 0, 1, \dots, r$, are called the *barycentric coordinates* of \mathbf{x} relative to the affine basis $\mathbf{x}_0, \mathbf{x}_1, \dots, \mathbf{x}_r$.

Definition 6.5. Let A be a subset of \mathbb{R}^n . The *affine hull* of A , denoted by $A[A]$, is defined to be the intersection of all linear manifolds containing A . The *dimension* of A is the dimension of $A[A]$.

Since $A \subseteq \mathbb{R}^n$, since \mathbb{R}^n is a linear manifold, and since the intersection of linear manifolds is a linear manifold, it follows that, for any nonempty set A , $A[A]$ is not empty. Thus the dimension of A is well defined.

For any linear manifold M , the relations $M \subseteq A[M]$ and $A[M] \subseteq M$ hold, so $M = A[M]$. Thus, for a linear manifold M , the dimension of M as defined in Definition 6.2 coincides with that of Definition 6.5.

We shall often write $\dim A$ to refer to the dimension of A .

Lemma 6.6. Let A be a subset of \mathbb{R}^n and let $M(A)$ denote the set of all affine combinations of points in A . Then

- $A[A] = M(A)$ and
- if there exist $p+1$ points $\mathbf{x}_0, \mathbf{x}_1, \dots, \mathbf{x}_p$ in A that are affinely independent and if any set of $p+2$ points in A is affinely dependent, then $\dim A = p$.

Remark 6.2. Statement (i) is the analog for affine hulls of Theorem 2.1 for convex hulls.

Proof. It is readily verified that if \mathbf{x} and \mathbf{y} are two points in $M(A)$, then each affine combination of these points is in $M(A)$. Hence, by Lemma 6.2, $M(A)$ is a linear manifold. Since $M(A) \supseteq A$, it follows from the definition of $A[A]$ that $A[A] \subseteq M(A)$. On the other hand, it follows from Corollary 5 to Lemma 6.2 that $M(A)$ is contained in every linear manifold containing A . Thus $A[A] = M(A)$. Hence $A[A] = M(A)$.

We now establish (ii). Let $\mathbf{x}_0, \mathbf{x}_1, \dots, \mathbf{x}_p$ be $p+1$ affinely independent points in A and suppose that any set of $p+2$ points is affinely dependent. Since the points $\mathbf{x}_0, \mathbf{x}_1, \dots, \mathbf{x}_p$ are in $A[A]$, it follows from Lemma 6.4 that $\dim A[A] \geq p$.

On the other hand, for arbitrary \mathbf{x} in A , the vectors

$$\mathbf{x}_1 - \mathbf{x}_0, \mathbf{x}_2 - \mathbf{x}_0, \dots, \mathbf{x}_\rho - \mathbf{x}_0, \mathbf{x} - \mathbf{x}_0$$

are linearly dependent, and the first ρ vectors are linearly independent. Hence for arbitrary \mathbf{x} in A there exist unique scalars μ_0, \dots, μ_ρ such that

$$\mathbf{x} - \mathbf{x}_0 = \sum_{j=1}^{\rho} \mu_j (\mathbf{x}_j - \mathbf{x}_0)$$

On setting $\mu_0 = 1 - (\mu_1 + \mu_2 + \dots + \mu_\rho)$, we get, uniquely,

$$\mathbf{x} = \sum_{j=0}^{\rho} \mu_j \mathbf{x}_j, \quad \sum_{j=0}^{\rho} \mu_j = 1.$$

Hence A is contained in the linear manifold $N[\mathbf{x}_0, \mathbf{x}_1, \dots, \mathbf{x}_\rho]$ and so $\dim A \leq \dim N[\mathbf{x}_0, \mathbf{x}_1, \dots, \mathbf{x}_\rho]$.

To complete the proof of (ii), it suffices to show that $N[\mathbf{x}_0, \mathbf{x}_1, \dots, \mathbf{x}_\rho]$ has dimension ρ . For then $\dim A \leq \rho$, and since we have already shown that $\dim A \geq \rho$, statement (ii) will follow.

We now show that $\dim N[\mathbf{x}_0, \mathbf{x}_1, \dots, \mathbf{x}_\rho] = \rho$. Since $\mathbf{x}_0, \mathbf{x}_1, \dots, \mathbf{x}_\rho$ are affinely independent to show that $\dim N[\mathbf{x}_0, \mathbf{x}_1, \dots, \mathbf{x}_\rho] = \rho$, we must show that any set $\mathbf{r}_0, \mathbf{r}_1, \dots, \mathbf{r}_\rho, \mathbf{r}_{\rho+1}$ of $\rho + 2$ vectors in $N[\mathbf{x}_0, \mathbf{x}_1, \dots, \mathbf{x}_\rho]$ is affinely dependent. From (i), we have that $N[\mathbf{x}_0, \dots, \mathbf{x}_\rho] = M[\mathbf{x}_1, \dots, \mathbf{x}_\rho]$. Thus for $i = 0, 1, \dots, \rho + 1$ we have

$$\mathbf{r}_i = \sum_{j=1}^{\rho} \alpha_{ij} \mathbf{x}_j, \quad \sum_{j=1}^{\rho} \alpha_{ij} = 1.$$

Since $\alpha_{00} = 1 - \alpha_{01} - \alpha_{02} - \dots - \alpha_{0\rho}$, we have

$$\mathbf{r}_0 = \sum_{j=1}^{\rho} \alpha_{0j} (\mathbf{x}_j - \mathbf{x}_0) + \mathbf{x}_0, \quad i = 0, 1, 2, \dots, \rho + 1.$$

Hence,

$$\mathbf{r}_i - \mathbf{r}_0 = \sum_{j=1}^{\rho} (\alpha_{ij} - \alpha_{0j}) (\mathbf{x}_j - \mathbf{x}_0), \quad i = 1, \dots, \rho + 1. \quad (5)$$

Thus, the $\rho + 1$ vectors $\mathbf{r}_i - \mathbf{r}_0$ in (5) are in the vector space spanned by the ρ linearly independent vectors $\mathbf{x}_1 - \mathbf{x}_0, \dots, \mathbf{x}_\rho - \mathbf{x}_0$. Hence the vectors $\mathbf{r}_i - \mathbf{r}_0$ in (5) are linearly dependent, and the corresponding points $\mathbf{r}_0, \mathbf{r}_1, \dots, \mathbf{r}_{\rho+1}$ are affinely dependent.

COROLLARY 4. The points $\mathbf{x}_0, \mathbf{x}_1, \dots, \mathbf{x}_n$ are an affine basis for $A[C]$.

Proof. In the proof of Lemma 6.5 we showed that $A \subset A[\{\mathbf{x}_0, \mathbf{x}_1, \dots, \mathbf{x}_n\}]$. Hence $A[C] \subset A[\{\mathbf{x}_0, \mathbf{x}_1, \dots, \mathbf{x}_n\}]$. Since the opposite inclusion clearly holds, we have $A[C] = A[\{\mathbf{x}_0, \mathbf{x}_1, \dots, \mathbf{x}_n\}]$. The corollary then follows from the uniqueness of the representation (ii) of a point \mathbf{x} in $A[\{\mathbf{x}_0, \mathbf{x}_1, \dots, \mathbf{x}_n\}]$.

Let $\mathbf{x}_0, \mathbf{x}_1, \dots, \mathbf{x}_n$ be affinely independent. We call $\text{co}\{\mathbf{x}_0, \mathbf{x}_1, \dots, \mathbf{x}_n\}$ the simplex spanned by $\mathbf{x}_0, \mathbf{x}_1, \dots, \mathbf{x}_n$. The points $\mathbf{x}_0, \mathbf{x}_1, \dots, \mathbf{x}_n$ are called the vertices of the simplex.

We leave it as an exercise for the reader to show that any $n + 2$ points in the simplex are affinely dependent. Thus the dimension of the simplex $\text{co}\{\mathbf{x}_0, \mathbf{x}_1, \dots, \mathbf{x}_n\}$ is n .

COROLLARY 5. The dimension of a convex set C is the maximum dimension of simplices in C .

This corollary is a restatement of Lemma 6.5 when C is a convex set.

Intuitively, in \mathbb{R}^3 one would expect that a convex set that is not contained in a plane or in a line to have nonempty interior. In terms of our definition of dimension of a set, we would expect a three-dimensional convex set to have nonempty interior. This is indeed the case, not only in \mathbb{R}^3 , but also in \mathbb{R}^n , as the following theorem shows.

THEOREM 6.1. A convex set C in \mathbb{R}^n has dimension n if and only if C has nonempty interior.

Suppose C has nonempty interior. Let \mathbf{x}_0 be an interior point of C and let $\mathbf{e}_1, \dots, \mathbf{e}_n$ denote the standard basis vectors in \mathbb{R}^n . Then there exists an $\epsilon > 0$ such that each of the $n + 1$ points $\mathbf{x}_0, \mathbf{x}_0 + \epsilon\mathbf{e}_1, \dots, \mathbf{x}_0 + \epsilon\mathbf{e}_n$ lies in C . These points are affinely independent. Since in \mathbb{R}^n any $n + 2$ points are affinely dependent, it follows from (ii) of Lemma 6.5 that C has dimension n .

Now suppose that C has dimension n . Then $A[C] = \mathbb{R}^n$, and by Corollary 4 of Lemma 6.5 there exist points $\mathbf{x}_0, \mathbf{x}_1, \dots, \mathbf{x}_n$ in C that form an affine basis for $A[C] = \mathbb{R}^n$. Thus each \mathbf{x} in \mathbb{R}^n can be written uniquely as

$$\mathbf{x} = \sum_{i=0}^n \lambda_i \mathbf{x}_i, \quad \sum_{i=0}^n \lambda_i = 1. \quad (6)$$

Since C is convex, all points \mathbf{x} in \mathbb{R}^n with $\lambda_i \geq 0, i = 0, 1, \dots, n$, lie in C . In particular,

$$\mathbf{x}_0 = \frac{1}{n+1}(\mathbf{x}_0 + \mathbf{x}_1 + \dots + \mathbf{x}_n)$$

the centroid of the simplex $\text{co}\{\mathbf{x}_0, \mathbf{x}_1, \dots, \mathbf{x}_n\}$, is in C .

We can also write \mathbf{x} uniquely as

$$\mathbf{x} = \sum_{j=1}^n \lambda_j (\mathbf{u}_j - \mathbf{u}_0),$$

with the λ_j as in (14). Since the vectors $\mathbf{u}_1 - \mathbf{u}_0, \dots, \mathbf{u}_n - \mathbf{u}_0$ are linearly independent, they form a basis for \mathbb{R}^n , and the linear transformation T from \mathbb{R}^n to \mathbb{R}^n defined by

$$T(\mathbf{x}) = (\lambda_1, \dots, \lambda_n)$$

is nonsingular and continuous. Let f be the mapping from \mathbb{R}^n to \mathbb{R}^{n+1} defined by

$$f(\mathbf{x}) = f\left(\sum_{j=1}^n \lambda_j \mathbf{u}_j\right) = (\lambda_0, \lambda_1, \dots, \lambda_n),$$

where the λ_j are as in (15). The mapping f can be viewed as a composite map $\mathbf{h} \circ T$, where

$$\mathbf{h}(\lambda_1, \dots, \lambda_n) = (\lambda_0, \lambda_1, \dots, \lambda_n), \quad \lambda_0 = 1 - \sum_{j=1}^n \lambda_j.$$

Thus, f is continuous.

Since f is continuous and since

$$f(\mathbf{x}_0) = \left(\frac{1}{n+1}, \frac{1}{n+1}, \dots, \frac{1}{n+1}\right),$$

there exists a $\delta > 0$ such that for $\mathbf{x} \in B(\mathbf{x}_0, \delta)$ all components λ_j of $f(\mathbf{x})$ will be strictly positive. Thus, all $\mathbf{x} \in B(\mathbf{x}_0, \delta)$ lie in C , and so \mathbf{x}_0 is an interior point of C .

By experimenting with a sheet of cardboard, if necessary, the reader should find it plausible that an arbitrary plane in \mathbb{R}^3 can be put in coincidence with the plane $x_3 = 0$ by means of an appropriate sequence of rotations followed by a possible translation. The plane $x_3 = 0$ consists of those points in \mathbb{R}^3 of the form $(x_1, x_2, 0)$ and thus is essentially \mathbb{R}^2 . Thus, in some sense a plane in \mathbb{R}^3 is equivalent to \mathbb{R}^2 . Similarly, any line in \mathbb{R}^3 can be made to coincide with any one of the coordinate axes by means of an appropriate sequence of rotations followed by a possible translation.

We shall now show that any n -dimensional manifold M in \mathbb{R}^n is equivalent, in a sense to be made precise, to \mathbb{R}^n .

Definition 6.6. An affine transformation T from \mathbb{R}^n to \mathbb{R}^n is a mapping of the form $T\mathbf{x} = S\mathbf{x} + \mathbf{a}$, where \mathbf{a} is a fixed vector in \mathbb{R}^n and S is a linear transformation from \mathbb{R}^n to \mathbb{R}^n .

Since relative to appropriate bases in \mathbb{R}^n and \mathbb{R}^n a linear transformation S from \mathbb{R}^n to \mathbb{R}^n is given by $S\mathbf{x} = A\mathbf{x}$, where A is an $n \times n$ matrix, we have $T\mathbf{x} = A\mathbf{x} + \mathbf{a}$.

Lemma 6.6. A mapping T from \mathbb{R}^n to \mathbb{R}^n is affine if and only if for each $\lambda \in \mathbb{R}$ and each pair of points \mathbf{x}_1 and \mathbf{x}_2 in \mathbb{R}^n

$$T[(1 - \lambda)\mathbf{x}_1 + \lambda\mathbf{x}_2] = (1 - \lambda)T\mathbf{x}_1 + \lambda T\mathbf{x}_2. \quad (7)$$

It is a straightforward calculation to show that if T is affine, then (7) holds. Now suppose that (7) holds. Let $\mathbf{a} = T\mathbf{0}$, and let $S\mathbf{x} = T\mathbf{x} - \mathbf{a}$. To show that T is affine, we must show that S is linear.

For any α in \mathbb{R} and \mathbf{x} in \mathbb{R}^n ,

$$\begin{aligned} S(\alpha\mathbf{x}) &= T[\alpha\mathbf{x} + (1 - \alpha)\mathbf{0}] - \mathbf{a} = \alpha T\mathbf{x} + (1 - \alpha)T\mathbf{0} - \mathbf{a} \\ &= \alpha(T\mathbf{x} - \mathbf{a}) = \alpha S(\mathbf{x}). \end{aligned}$$

For any \mathbf{x}_1 and \mathbf{x}_2 in \mathbb{R}^n we have

$$\begin{aligned} S(\mathbf{x}_1 + \mathbf{x}_2) &= S[\frac{1}{2}(\mathbf{x}_1 + \mathbf{x}_2) + \frac{1}{2}(\mathbf{x}_1 + \mathbf{x}_2)] = 2S[\frac{1}{2}(\mathbf{x}_1 + \mathbf{x}_2)] = 2T[\frac{1}{2}(\mathbf{x}_1 + \mathbf{x}_2)] - 2\mathbf{a} \\ &= [T(\mathbf{x}_1) - \mathbf{a}] + [T(\mathbf{x}_2) - \mathbf{a}] = S(\mathbf{x}_1) + S(\mathbf{x}_2). \end{aligned}$$

Thus S is linear.

From (7) we see that under an affine map lines are mapped onto lines. Thus under an affine map linear manifolds are mapped onto linear manifolds. By taking $0 \leq \lambda \leq 1$ in (7), we see that under an affine map convex sets are mapped onto convex sets.

Theorem 6.7. Let $\mathbf{x}_0, \mathbf{x}_1, \dots, \mathbf{x}_n$ and $\mathbf{y}_0, \mathbf{y}_1, \dots, \mathbf{y}_n$ be two affinely independent sets in \mathbb{R}^n . Then there exists a one-to-one affine transformation T from \mathbb{R}^n onto itself such that $T\mathbf{x}_i = \mathbf{y}_i$, $i = 0, 1, \dots, n$, if $\mathbf{y}_0 = \mathbf{x}_0$, then T is unique.

If $\mathbf{y}_0 \neq \mathbf{x}_0$, obtain two affinely independent sets $\mathbf{x}_0, \dots, \mathbf{x}_n$, $\mathbf{x}_{0+1}, \dots, \mathbf{x}_n$ and $\mathbf{y}_0, \dots, \mathbf{y}_n$, $\mathbf{y}_{0+1}, \dots, \mathbf{y}_n$ of $n + 1$ points by adjoining appropriate points to the original sets. The vectors $\mathbf{x}_1 - \mathbf{x}_0, \dots, \mathbf{x}_n - \mathbf{x}_0$ constitute a basis for \mathbb{R}^n , as do the vectors $\mathbf{y}_1 - \mathbf{y}_0, \dots, \mathbf{y}_n - \mathbf{y}_0$. Thus, there exists a unique one-to-one linear transformation S of \mathbb{R}^n onto itself such that $S(\mathbf{x}_i - \mathbf{x}_0) = \mathbf{y}_i - \mathbf{y}_0$, $i = 1, \dots, n$. Let $T\mathbf{x} = S\mathbf{x} + \mathbf{a}$ for a fixed \mathbf{a} to be determined. Thus for $i = 0, 1, \dots, n$,

$$T(\mathbf{x}_i) = S(\mathbf{x}_i) + \mathbf{a} = S(\mathbf{x}_i - \mathbf{x}_0) + S(\mathbf{x}_0) + \mathbf{a} = \mathbf{y}_i - \mathbf{y}_0 + S(\mathbf{x}_0) + \mathbf{a}.$$

Therefore, if we take $\mathbf{x} = \mathbf{y}_i - \beta\mathbf{y}_n$, we will have $T\mathbf{x} = \mathbf{y}_i$, $i = 0, 1, \dots, n$.

COROLLARY 6. Let M_1 and M_2 be two linear manifolds of the same dimension. Then there exists a one-to-one affine transformation T of \mathbb{R}^n onto itself such that $T M_1 = M_2$.

By Lemma 6A, M_1 has an affine basis $\mathbf{x}_0, \mathbf{x}_1, \dots, \mathbf{x}_n$ and M_2 has an affine basis $\mathbf{y}_0, \mathbf{y}_1, \dots, \mathbf{y}_n$. By Theorem 6.1 there exists a one-to-one transformation T from \mathbb{R}^n onto itself such that $T\mathbf{x}_i = \mathbf{y}_i$, $i = 1, \dots, n$. The corollary now follows from the facts that each \mathbf{x} in M_1 is a unique affine combination of $\mathbf{x}_0, \mathbf{x}_1, \dots, \mathbf{x}_n$, each \mathbf{y} in M_2 is a unique affine combination of $\mathbf{y}_0, \mathbf{y}_1, \dots, \mathbf{y}_n$, and T maps affine combinations into affine combinations.

Remark 6.1. If M is a linear manifold of dimension r , then there exists a one-to-one affine transformation T of \mathbb{R}^n onto itself that maps M onto the manifold

$$U^r = \{\mathbf{y} \mid \mathbf{y} = (y_1, \dots, y_r, 0, \dots, 0), y_1 + y_2 + \dots + y_r = 0\}.$$

The manifold U^r can be identified with \mathbb{R}^r by means of the one-to-one mapping \mathbb{R}^r from U^r onto \mathbb{R}^r given by

$$\mathbb{R}^r \mathbf{y} = (y_1, \dots, y_r, 0, \dots, 0) \mapsto (y_1, \dots, y_r).$$

Denote the restriction of T to M by T_M , and define a mapping τ from M to \mathbb{R}^r by the formula

$$\mathbf{x} \in M \mapsto \tau(\mathbf{x}) = \mathbb{R}(T_M \mathbf{x}).$$

Then τ is one-to-one onto, is continuous, and has a continuous inverse, τ^{-1} . Note that τ and τ^{-1} map convex sets onto convex sets and linear manifolds onto linear manifolds of the same dimension.

Exercise 6A. Let k be a positive integer, $1 \leq k \leq n - 1$. Show that for a given subspace V_{n-k} of \mathbb{R}^n , there exist k linearly independent vectors $\mathbf{a}_1, \dots, \mathbf{a}_k$ such that

$$V_{n-k} = \{\mathbf{x} \in \mathbb{R}^n \mid (\mathbf{a}_i, \mathbf{x}) = 0, i = 1, \dots, k\}.$$

Exercise 6B. Show that a set of points $\mathbf{a}_0, \dots, \mathbf{a}_n$ is affinely dependent if and only if one of the points is an affine combination of the other points.

Exercise 6C. In \mathbb{R}^3 consider the plane $2x + y + z = 5$.

- (a) Find an affine basis for this plane that does not include the point $\mathbf{x} = (-1, -2, 9)$.
- (b) Express $\mathbf{x} = (-1, -2, 9)$ in barycentric coordinates relative to your basis.

Exercise 4.4. Let $\mathbf{x}_0, \mathbf{x}_1, \dots, \mathbf{x}_r$ be affinely independent. Show that any $r + 2$ points $\mathbf{p}_0, \mathbf{p}_1, \dots, \mathbf{p}_r, \mathbf{x}_{r+1}$ in $\text{aff}\{\mathbf{x}_0, \mathbf{x}_1, \dots, \mathbf{x}_r\}$ are affinely dependent.

Exercise 4.5. Show that an affine map T maps affine combinations of points into affine combinations of points.

Exercise 4.6. In Theorem 4.2, why is T unique if $r = n$?

5. SEPARATION, SEPARATION AND SUPPORT

Although the separation and support theorems of Sections 3 and 4 are adequate for most applications to optimization problems, they are not applicable in some situations where proper separation and nontrivial support of convex sets do occur. To illustrate this, let

$$C_1 = \{\mathbf{x} \in \mathbb{R}^3 : x_1 \geq 0, x_2 \geq 0, x_3 = 0\}$$

$$C_2 = \{\mathbf{x} \in \mathbb{R}^3 : x_1 \geq 0, x_2 \geq 0, x_3 = 0\}.$$

These two sets are properly separated by the plane (hyperplane in \mathbb{R}^3) $x_3 = x_4 = 0$. These sets are convex, are not disjoint, and have empty interiors so that neither Theorem 3.3 nor Theorem 3.4 can be applied to deduce separation.

Let $C = \{\mathbf{x} \in \mathbb{R}^2 : x_1^2 + x_2^2 \leq 1, x_3 = 0\}$. At each boundary point \mathbf{x}_0 with $x_1^2 + x_2^2 = 1$ there exist a nontrivial supporting hyperplane, namely the plane that is orthogonal to the plane $x_3 = 0$ and whose intersection with this plane is the tangent line to the circle $x_1^2 + x_2^2 = 1, x_3 = 0$ at \mathbf{x}_0 . Again, since the convex set C has empty interior, Theorem 3.1 does not give us the existence of this nontrivial supporting hyperplane.

The sets $C_1, C_2,$ and C , considered as sets in \mathbb{R}^3 , do have interiors. It turns out that this observation is the key to theorems that are applicable to examples such as those described above. We introduce the notion of relative interior of a convex set C of dimension r . The relative interior is the interior of a convex set C considered as a set in $\mathcal{N}(C)$. In other words we consider the set C as a set in \mathbb{R}^r rather than a set in \mathbb{R}^n .

Definition 5.1. Let A be a set of dimension $r < n$ and let $M = \mathcal{N}_r(A)$. A point \mathbf{x}_0 in A is said to be a relative interior point of A if there exists an $\epsilon > 0$ such that $B(\mathbf{x}_0, \epsilon) \cap M \subset A$. The set of relative interior points of A is called the

relative interior of A and will be denoted by $\text{ri}(A)$. A set A is said to be relatively open if $\text{ri}(A) = A$, that is, if every point of A is a relative interior point of A .

Note that if $\dim M = n$, then a relative interior point is an interior point and a relatively open set is an open set.

Since a linear manifold M is a closed set, any limit point of A will also be in M . Thus the closure of A as a subset of M is the same as the closure of A as a subset of \mathbb{R}^n .

If A is contained in M and $\dim M < n$, then every point of A is a boundary point of A . We define the relative boundary of A to be the set $\bar{A} - \text{ri}(A)$, that is, those points of \bar{A} that are not relative interior points of A . If $\dim M = n$, then the relative boundary of A is the boundary of A .

Let M be a linear manifold of dimension $r < n$ and let τ be the mapping defined in Remark 3.1. We noted there that τ and τ^{-1} are one-to-one and continuous.

LEMMA 3.1. If O is an open set in \mathbb{R}^n , then $\tau^{-1}(O)$ is a relatively open set in M . If O is a relatively open set in M , then $\tau(O)$ is an open set in \mathbb{R}^n .

Proof. Let O be an open set in \mathbb{R}^n and let $x_0 \in \tau^{-1}(O)$. Then there is a unique $y_0 \in O$ such that $\tau^{-1}(y_0) = x_0$, or equivalently, $x_0 = \tau(y_0)$. Since $x_0 \in O$, there exists an $\epsilon > 0$ such that the open ball $B_\epsilon(x_0, \epsilon)$ in \mathbb{R}^n is contained in O . Since τ is continuous, there exists a $\delta > 0$ such that for all $B_{\delta}(y_0, \delta) \cap M$ we have $\tau(B_{\delta}(y_0, \delta)) \subset O$. But this says that $B_{\delta}(x_0, \delta) \cap M \subset \tau^{-1}(O)$. Hence x_0 is a relative interior point of $\tau^{-1}(O)$, and so $\tau^{-1}(O)$ is relatively open. Since $\tau = (\tau^{-1})^{-1}$, a similar argument with τ replaced by τ^{-1} shows that if O is relatively open, then $\tau(O)$ is open in \mathbb{R}^n .

Let C be a convex set in \mathbb{R}^n of dimension $r < n$. As a consequence of the one-to-one relationship between $M = \text{aff}(C)$ and \mathbb{R}^n that preserves open sets, convexity, dimension of convex sets, and linear manifolds, we can obtain many properties of C by establishing these properties for convex sets of full dimension. These properties then hold for the set $\tau(C)$ considered as a set of full dimension in \mathbb{R}^n and hence for $C = \tau^{-1}(\tau(C))$. In particular, we obtain the following results in this fashion.

THEOREM 3.1. A convex set of dimension r has nonempty relative interior.

This theorem follows from Theorem 6.1 and Lemma 3.1.

LEMMA 3.2. Let C be a convex set and let x_1 and x_2 be two points with $x_1 \in \text{ri}(C)$ and $x_2 \in C$. Then the line segment $[x_1, x_2]$ is contained in $\text{ri}(C)$.

COROLLARY 1. If C is convex, then $\text{ri}(C)$ is convex.

COROLLARY 2. Let C be convex. Then

- (i) $\overline{\text{ri}(C)} = \bar{C}$ and
 (ii) $\text{ri}(C) = \text{int}(C)$.

LEMMA 7.2. Corollary 1 and (i) of Corollary 2 follow from Lemma 7.1 and the corresponding corollaries. In proving (ii) of Corollary 2 to Lemma 7.1, we argued that since $C \subseteq \bar{C}$, it follows that $\text{int}(C) \subseteq \text{int}(\bar{C})$. In general, it is not true that if A and B are convex sets with $A \subseteq B$ that $\text{ri}(A) = \text{ri}(B)$, as the following example shows. Take B to be a closed unit cube in \mathbb{R}^n . Let A be any closed face of B . Then $A \subseteq B$, but it is not true that $\text{ri}(A) \subseteq \text{ri}(B)$. In fact, $\text{ri}(A)$ and $\text{ri}(B)$ are disjoint. Since $\bar{A}(C) = \bar{A}(\bar{C})$, it does follow that $\text{ri}(C) = \text{ri}(\bar{C})$. The proof of the opposite inclusion given in Section 2 is valid if we replace interiors by relative interiors.

We now state and prove the best possible separation theorem in \mathbb{R}^n .

THEOREM 7.1. Let X and Y be two convex subsets of \mathbb{R}^n . Then X and Y can be properly separated if and only if $\text{ri}(X)$ and $\text{ri}(Y)$ are disjoint.

To prove the "only if" part of the theorem, we shall need to define what is meant by hyperplane cutting a convex set and to develop conditions under which cutting occurs.

DEFINITION 7.1. A hyperplane M_C^* is said to cut a convex set C if there exist points x_1 and x_2 in C such that $\langle x, n_C \rangle < \alpha$ and $\langle x, n_C \rangle > \alpha$. A hyperplane that does not cut C is said to bound C .

LEMMA 7.1. A hyperplane M_C^* cuts a convex set C if and only if the following two conditions hold:

- (i) M_C^* does not contain C and
 (ii) $M_C^* \cap \text{ri}(C) \neq \emptyset$.

PROOF. Let M_C^* cut C . Then there exist points x_1 and x_2 in C such that $\langle x, n_C \rangle < \alpha$ and $\langle x, n_C \rangle > \alpha$. Since for $y \in M_C^*$ we have $\langle y, n_C \rangle = \alpha$, conclusion (i) follows. Now let $x \in \text{ri}(C)$. Then by Lemma 7.2 the line segments $[x, x_1]$ and $[x, x_2]$ belong to $\text{ri}(C)$. Let

$$f(t) = \begin{cases} x + t(x_1 - x), & -1 \leq t \leq 1, \\ x + t(x_2 - x), & 0 \leq t \leq 1. \end{cases}$$

Then $f(t) \in C$ for $-1 \leq t \leq 1$ and $f(t) \in \text{ri}(C)$ for $-1 < t < 1$. Let $g(t) = \langle x, g(t) \rangle$. Since g is continuous on $[-1, 1]$ with $g(-1) < \alpha$ and $g(1) > \alpha$, there exists a t_0 with $-1 < t_0 < 1$ such that $\langle x, g(t_0) \rangle = \alpha$. But then $g(t_0) \in \text{ri}(C)$ and $\langle x, g(t_0) \rangle = \alpha$, which establishes (ii).

Now suppose that (i) and (ii) hold. Then there exist a point $x_0 \in \text{ri}(C) \cap A_1^c$ and a point $x_1 \in C$ and not in A_1^c . Suppose that $\langle a, x_1 \rangle = \alpha$. Since $x_0 \in \text{ri}(C)$, there exists a $\delta > 0$ such that $x_2 = x_0 - \delta x_1 = x_0' \in C$. Then

$$\langle a, x_2 \rangle = \alpha - \delta \langle a, x_1 \rangle = \alpha' < \alpha.$$

If $\langle a, x_1 \rangle > \alpha$, then $\langle a, x_2 \rangle < \alpha$.

We now prove Theorem 3.2.

If a hyperplane properly separates the convex sets X and Y then it certainly properly separates $\text{ri}(X)$ and $\text{ri}(Y)$. Conversely, if a hyperplane properly separates $\text{ri}(X)$ and $\text{ri}(Y)$, then since $\text{ri}(X) = \overset{\circ}{X}$ and $\text{ri}(Y) = \overset{\circ}{Y}$ it follows that the hyperplane properly separates X and Y . Thus a hyperplane properly separates X and Y if and only if it properly separates $\text{ri}(X)$ and $\text{ri}(Y)$. Thus we may assume without loss of generality that X and Y are disjoint relatively open convex sets. The proof of separation will be carried out by a descending induction on $d = \max\{\dim X, \dim Y\}$.

If $\dim X = n$, then the separation follows from Theorem 3.1. We now assume that proper separation occurs if $d \geq k$, where $1 = k \leq n$. We now suppose that $\dim X = k - 1$ and $\dim Y \leq k - 1$. To simplify the notation, we assume that $0 \in X$. There is no loss of generality in assuming this, since this can always be achieved by a translation of coordinates. Since $0 \in X$, the manifold $A(X)$ is a subspace of dimension $k - n - 1$. Let w be a point in \mathbb{R}^n that is not in $A(X)$. Let

$$A = \text{co}\{X \cup \{w\}\}, \quad B = \text{co}\{X \cup \{-w\}\}.$$

We assert that Y cannot intersect both A and B . See Figure 23.

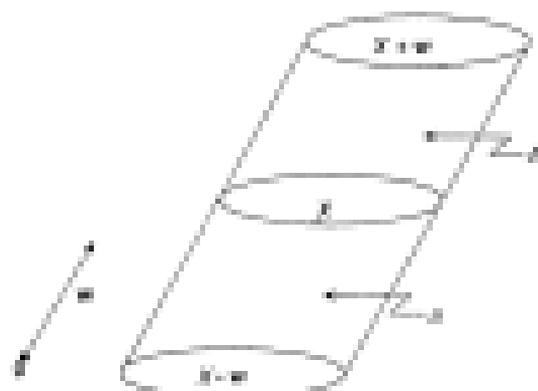


Figure 23.

By Theorem 2.2 point ξ in A can be written

$$\xi = \sum_{i=1}^{n+1} \lambda_i x_i + \lambda_{n+2} w,$$

where $x_i \in X$ for $i = 1, \dots, n+1$, the vector $w = (x_1, \dots, x_{n+1}) \in F_{n+1,1}$, and λ_i is either zero or 1. Thus, if $\xi \in A$, then there exists an $\alpha \in X$ such that

$$\xi = \alpha + \theta w, \quad 0 \leq \theta \leq 1. \quad (1)$$

Since

$$\xi = \alpha + \theta w = (1 - \theta)\alpha + \theta(\alpha + w),$$

every point ξ given by (1) is in $A \cap B$; similar argument shows that

$$\xi = (\eta; \eta = \alpha + \theta w, \alpha \in X, 0 \leq \theta \leq 1).$$

Thus if $A \cap F \neq \emptyset$ and $B \cap F \neq \emptyset$, then there exist points y_1 and y_2 in F such that

$$y_1 = \alpha_1 + \theta_1 w, \quad y_2 = \alpha_2 + \theta_2 w,$$

where $\alpha_1, \alpha_2 \in X$, $0 \leq \theta_1 \leq 1$, $0 \leq \theta_2 \leq 1$. The inequalities $\theta_1 = 0$ and $\theta_2 = 0$ follow from $X \cap F = \emptyset$.

Let $\alpha = \theta_1 \theta_2 \alpha_1 + \theta_2 \alpha_2$ and let $\beta = \theta_2 (\theta_1 + \theta_2) w$. Then $\alpha \in X$, $\beta \in B$, $\alpha + \beta = y_1$, and

$$\theta y_1 + \theta y_2 = \alpha + \beta. \quad (2)$$

Since X and B are convex, the left-hand side of (2) belongs to X and the right-hand side belongs to B . This contradicts $X \cap B = \emptyset$, and so F cannot intersect both A and B .

For the sake of definiteness, suppose that $A \cap F = \emptyset$. Since also $A \neq \emptyset$ (why?) and $\dim F \leq k-1$, by the inductive hypothesis there exists a hyperplane H that properly separates A and F . Since $X \subset A$, the hyperplane H properly separates X and F .

To prove the "only if" assertion, suppose that there exists a hyperplane H that properly separates X and F and that there exists a point α in $\text{ri}(X) \cap \text{ri}(F)$. Then $\alpha \in H$, and thus $H \cap \text{ri}(X) \neq \emptyset$ and $H \cap \text{ri}(F) \neq \emptyset$. Since the separation is proper, at least one of the sets, say X , is not contained in H . Hence by Lemma 7.3, H cuts X , which contradicts the assumption that H separates X and F .

We conclude this section with results that are stronger than Theorems 4.3 and 4.4.

THEOREM 7.3. Let C be a convex set and let D be a convex subset of the relative boundary of C . Then there exists a nontrivial supporting hyperplane to C that contains D .

Proof. Since $\text{ri}(C) \cap D = \emptyset$, we have $\text{ri}(C) \cap \text{ri}(D) = \emptyset$. Hence by Theorem 7.2 there exists a hyperplane M_0^C that properly separates C and D . Thus

$$(\mathbf{x}, \mathbf{y}) < \alpha \quad \text{for } \mathbf{x} \in C \quad \text{and} \quad (\mathbf{x}, \mathbf{y}) \geq \alpha \quad \text{for } \mathbf{y} \in D. \quad (3)$$

Since each \mathbf{y} in D is a relative boundary point of C , for each \mathbf{y} in D there exists a sequence (\mathbf{x}_n) of points in C such that $\mathbf{x}_n \rightarrow \mathbf{y}$. For this sequence we have $(\mathbf{x}_n, \mathbf{y}) < \alpha$ and $(\mathbf{x}_n, \mathbf{y}) \rightarrow (\mathbf{y}, \mathbf{y})$. Thus $(\mathbf{y}, \mathbf{y}) < \alpha$. From this and from (3) we get that $(\mathbf{y}, \mathbf{y}) = \alpha$. Thus, each \mathbf{y} in D is in $\text{ri}(C)$. Since $D \subseteq C$, it follows that M_0^C is a supporting hyperplane to C that contains D . Since the separation of C and D is proper and since $D \subseteq M_0^C$, there is an \mathbf{x}_0 in C such that $(\mathbf{x}_0, \mathbf{y}_0) < \alpha$. Thus the supporting hyperplane is nontrivial.

The next theorem is stronger than Theorem 4.3 but is only valid in finite dimensions.

THEOREM 7.4. Let C be a compact convex subset of \mathbb{R}^k . Then C is the convex hull of its extreme points.

Proof. Let C_e denote the set of extreme points of C . By Theorem 4.2 the set $C_e \neq \emptyset$. Since $C_e \subseteq C$ and $\text{co}(C_e) \subseteq \text{co}(C) = C$, to prove the theorem, we must show that

$$C = \text{co}(C_e). \quad (4)$$

The proof of (4) proceeds by induction on the dimension k of C .

If $k = 0$, then C is a point, and if $k = 1$, then C is a closed line segment. Thus (4) holds in these cases. Now suppose that (4) holds if $\dim C = k - 1$, where $1 \leq k \leq n$.

Let $\dim C = k$ and let $\mathbf{x} \in C$. By Theorem 7.1, either $\mathbf{x} \in \text{ri}(C)$ or \mathbf{x} belongs to the relative boundary of C . If \mathbf{x} belongs to the relative boundary of C , then by Theorem 7.3, there exists a supporting hyperplane to C that contains \mathbf{x} . This hyperplane will intersect $\text{ri}(C)$ in a manifold M_{k-1} that is a hyperplane of dimension $k - 1$ in \mathbb{R}^k and that supports C at \mathbf{x} . The set $M_{k-1} \cap C$ is compact and convex and has dimension $\leq k - 1$. By the inductive hypothesis \mathbf{x} is a convex combination of the extreme points of $M_{k-1} \cap C$. By Lemma 4.1 the extreme points of $M_{k-1} \cap C$ are extreme points of C . Thus $\mathbf{x} \in \text{co}(C_e)$ whenever \mathbf{x} is a relative boundary point of C .

If $x \in \text{int} C$, let L be a line in $A[C]$ through the point x . Then $L \cap \text{co}(C)$ is a line segment (p, q) whose end points p and q are in the relative boundary of C . Since C is compact, $y \in C$ and $z \in C$. Since x is a convex combination of y and z and since p and q are in $\text{co}(C)_\Delta$, it follows that $x \in \text{co}(C)_\Delta$, and (8) is established.

III

CONVEX FUNCTIONS

1. DEFINITION AND ELEMENTARY PROPERTIES

The simplest class of functions are real-valued functions that are translates of linear functions, or affine functions, f defined by the formula $f(x) = (a, x) + b$. One could consider the next class of functions in order of complexity to be the translates of the functions $y = |x|^p$, $p \geq 1$, in $\mathbb{R}^2 = (x, y)$, where $x \in \mathbb{R}^1$, the graph of $y = |x|$ is the upper nappe of a right circular cone with vertex at the origin and the graph of $y = |x|^2$ is a paraboloid of revolution. The graphs of such functions are axially symmetric about a line parallel to the y -axis. Intuitively, it appears that all of the above functions have the property that the set of points (x, y) in \mathbb{R}^{n+1} that lie on or above the graph is a convex set. It is therefore reasonable to consider the next level of functions to be the functions with this property. Such functions are called convex functions. It turns out that convex functions are important in optimization problems as well as in other areas of mathematics. We shall define convex functions using a different property and then show in Lemma 1.1 below that our definition is equivalent to the one suggested.

Definition 1.1. A real-valued function f defined in a convex set C in \mathbb{R}^n is said to be convex if for every x_1 and x_2 in C and every $\alpha \geq 0$, $\beta \geq 0$, $\alpha + \beta = 1$,

$$f(\alpha x_1 + \beta x_2) \leq \alpha f(x_1) + \beta f(x_2). \quad (1)$$

The function f is said to be strictly convex if strict inequality holds in (1).

The geometric interpretation of this definition is that a convex function is a function such that if x_1 and x_2 are any two points in \mathbb{R}^{n+1} on the graph of f then points of the line segment $[x_1, x_2]$ joining x_1 and x_2 lie on or above the graph of f .

Figure 1.1 illustrates the definition. Let $C = (x_\alpha, x_\beta)$ where $x_\beta = \alpha x_1 + \beta x_2$, $\alpha \geq 0$, $\beta \geq 0$, $\alpha + \beta = 1$, be a point on the line segment joining the points $A = (x_1, f(x_1))$ and $B = (x_2, f(x_2))$ on the graph of f . Let $D = (x_\beta, f(x_\beta))$ and let

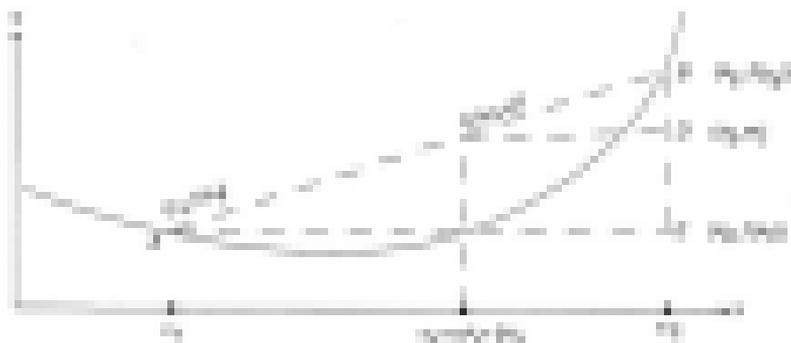


Figure 1.1.

$f^{-1}(y_2) = x_2$. Then

$$\frac{f(x_2) - y_2}{f(x_1) - y_1} = \frac{f(x_2) - f(x_1)}{f(x_1) - f(x_2)} = \frac{f(x_2) - f(x_1)}{f(x_1) - y_1}.$$

From the extreme terms of these equalities we get that

$$y_1 < f(x_2) < y_2 \implies f(x_2) < f(x_1) < f(x_2).$$

Thus $f(x_2) < f(x_1) < f(x_2)$ holds only if f is *not* injective.

Note that if the domain of f is not convex, then the left-hand side of (1.1) need not be defined. Therefore, when defining convex functions, it will make sense to consider functions defined on convex sets. In addition, we shall not repeat the term *convexity* when considering convex functions; it shall be tacitly assumed that f is not convex.

Definition 1.1. A function f defined on a convex set J is said to be *convex* if $-f$ is concave.

Definition 1.2. Let f be a real-valued function defined on a set J . The *epigraph* of f , denoted by $\text{epi}(f)$, is the set

$$\text{epi}(f) = \{(x, y) \in \mathbb{R}^2 \mid x \in J, y \geq f(x)\}.$$

In Figure 1.2 the bounded curve represents the epigraph of $f(x) = x^2 - 4$ ($x \in \mathbb{R}$), and $y = 0$, $-4 \leq x \leq 4$.

The epigraph of a function is the set of points $(x, y) \in \mathbb{R}^2$ that lie on or above the graph of f .



Figure 1.1

Lemma 1.1. A function f defined on a convex set C is convex if and only if $\text{epi } f$ is convex.

Proof. Suppose that $\text{epi } f$ is convex. Let x_1 and x_2 belong to C . Then $(x_1, f(x_1))$ and $(x_2, f(x_2))$ belong to $\text{epi } f$. Since $\text{epi } f$ is convex, for any $(\alpha, \beta) \in P_1$

$$\alpha(x_1, f(x_1)) + \beta(x_2, f(x_2)) = (\alpha x_1 + \beta x_2, \alpha f(x_1) + \beta f(x_2))$$

belongs to $\text{epi } f$. By the definition of $\text{epi } f$,

$$\alpha f(x_1) + \beta f(x_2) \geq f(\alpha x_1 + \beta x_2),$$

and so f is convex.

Now suppose f is convex. Let (x_1, y_1) and (x_2, y_2) be any two points in $\text{epi } f$. Then $y_1 \geq f(x_1)$ and $y_2 \geq f(x_2)$. For any $(\alpha, \beta) \in P_1$ let

$$(x, y) = \alpha(x_1, y_1) + \beta(x_2, y_2) = (\alpha x_1 + \beta x_2, \alpha y_1 + \beta y_2).$$

Since (x_1, y_1) and (x_2, y_2) are in $\text{epi } f$, and since f is convex, we have

$$y = \alpha y_1 + \beta y_2 \geq \alpha f(x_1) + \beta f(x_2) \geq f(\alpha x_1 + \beta x_2) = f(x).$$

Thus, $(x, y) \in \text{epi } f$, and so $\text{epi } f$ is convex.

Lemma 1.2. If f is a convex function defined on a convex set C in \mathbb{R}^n , then for any real γ the set $\{x \in C \mid f(x) \leq \gamma\}$ is either empty or convex.

We leave the proof of this lemma as an exercise for the reader.

The converse to this lemma is false, as can be seen from the function $f(x) = x^2$.

Lemma 1.3. Let $\{f_\alpha\}$ be a family of convex functions defined on a convex set C in \mathbb{R}^n and let there exist a real-valued function M with domain C such that, for each α , $f_\alpha(x) \leq M(x)$ for all $x \in C$. Then the function f defined by $f(x) = \sup\{f_\alpha(x) \mid \alpha\}$ is convex.

Proof. Since for each α in C , $f_\alpha(x) \leq M(x)$ for all x , it follows that $f(x) \leq M(x)$, and so f is real valued. Also $\text{epi}(f) = \bigcap_\alpha \text{epi}(f_\alpha)$, so $\text{epi}(f)$, being the intersection of convex sets, is convex. Therefore f is convex.

Theorem 1.4 (Jensen's Inequality). Let f be a real-valued function defined on a convex set C in \mathbb{R}^n . Then f is convex if and only if for every finite set of points x_1, \dots, x_k in C and every $\lambda = (\lambda_1, \dots, \lambda_k)$ in P_k

$$f(\lambda_1 x_1 + \dots + \lambda_k x_k) \leq \lambda_1 f(x_1) + \dots + \lambda_k f(x_k) \quad (1)$$

Proof. If (1) holds for all positive integers k , then it holds for $k = 2$. But for $k = 2$, the relation (1) is the definition of convex function.

Now suppose that f is convex. We shall prove (1) by induction on k . For $k = 1$ the relation is trivial. For $k = 2$ the relation follows from the definition of a convex function. Suppose that $k = 3$ and that (1) has been established for $k = 1$. We shall show that (1) holds for k . If $\lambda_k = 0$, then there is nothing to prove. If $\lambda_k \neq 0$, set $A = \sum_{i=1}^{k-1} \lambda_i$. Then

$$A > 0, \quad A + \lambda_k = 1, \quad \sum_{i=1}^{k-1} \frac{\lambda_i}{A} = 1,$$

and we may write (compare the proof of Lemma B.2.6)

$$\begin{aligned} f\left(\sum_{i=1}^k \lambda_i x_i\right) &= f\left(A \left[\sum_{i=1}^{k-1} \left(\frac{\lambda_i}{A}\right) x_i\right] + \lambda_k x_k\right) \\ &\leq A f\left(\sum_{i=1}^{k-1} \left(\frac{\lambda_i}{A}\right) x_i\right) + \lambda_k f(x_k) \\ &\leq \lambda_1 f(x_1) + \dots + \lambda_{k-1} f(x_{k-1}) + \lambda_k f(x_k) \end{aligned}$$

where the first inequality follows from the convexity of f and the second inequality follows from the inductive hypothesis.

The next lemma in conjunction with other criteria is sometimes useful in determining whether a given function is convex.

Lemma 1.5. Let f be a convex function defined on a convex set C in \mathbb{R}^n . Let g be a nondecreasing convex function defined on an interval I in \mathbb{R} . Let $f(x) \in I$. Then the composite function $g \circ f$ defined by $(g \circ f)(x) = g(f(x))$ is convex on C .

We leave the proof as an exercise.

Let C be a convex set. Choose an \mathbf{x} in C and a \mathbf{v} in \mathbb{R}^n . Let

$$A = \{\lambda \mid \lambda \in \mathbb{R}, \mathbf{x} + \lambda\mathbf{v} \in C\}.$$

Let $l_1 = \sup\{\lambda \mid \lambda \in A\}$ and let $l_2 = \inf\{\lambda \mid \lambda \in A\}$. Then $l_2 \leq 0 \leq l_1$. If $l_1 \neq l_2$, then since C is convex, $\mathbf{x} + \lambda\mathbf{v} \in C$ for all λ in the open interval (l_2, l_1) . Note that, depending on the structure of the set C , $\mathbf{x} + \lambda\mathbf{v}$ can belong to C for all λ in $[l_2, l_1]$ or $[l_2, l_1)$ or $(l_2, l_1]$.

Let f be a convex function defined on a convex set C . From the discussion in the preceding paragraph it follows that for each \mathbf{x} in C and each \mathbf{v} in \mathbb{R}^n there is an interval $J \subset \mathbb{R}$ that depends on \mathbf{x} and \mathbf{v} such that the function $\varphi(\lambda; \mathbf{x}, \mathbf{v})$ defined by

$$\varphi(\lambda; \mathbf{x}, \mathbf{v}) = f(\mathbf{x} + \lambda\mathbf{v}) \quad (3)$$

has J as its domain of definition. The interval J always includes the origin and can degenerate to the single point $\{0\}$; it can be open, half open, or closed. The function $\varphi(\lambda; \mathbf{x}, \mathbf{v})$ gives the one-dimensional section of the function f through the point \mathbf{x} in the direction of \mathbf{v} . The next lemma states that f is convex if and only if all the nondegenerate one-dimensional functions $\varphi(\lambda; \mathbf{x}, \mathbf{v})$ are convex.

Lemma 1.5. *A function f defined on a convex set C is convex if and only if for each \mathbf{x} in C and \mathbf{v} in \mathbb{R}^n such that $\varphi(\lambda; \mathbf{x}, \mathbf{v})$ is defined on a nondegenerate interval J the function $\varphi(\lambda; \mathbf{x}, \mathbf{v})$ is convex on J .*

Proof. Suppose that f is convex. Then for any λ_1, λ_2 in J and (α, β) in \mathcal{P}_2

$$\begin{aligned} \varphi(\alpha\lambda_1 + \beta\lambda_2; \mathbf{x}, \mathbf{v}) &= f(\mathbf{x} + (\alpha\lambda_1 + \beta\lambda_2)\mathbf{v}) \\ &= f(\alpha(\mathbf{x} + \lambda_1\mathbf{v}) + \beta(\mathbf{x} + \lambda_2\mathbf{v})) \\ &\stackrel{1}{\leq} \alpha f(\mathbf{x} + \lambda_1\mathbf{v}) + \beta f(\mathbf{x} + \lambda_2\mathbf{v}) \\ &\stackrel{1}{\leq} \alpha\varphi(\lambda_1; \mathbf{x}, \mathbf{v}) + \beta\varphi(\lambda_2; \mathbf{x}, \mathbf{v}) \\ &\stackrel{1}{\leq} \varphi(\alpha\lambda_1 + \beta\lambda_2; \mathbf{x}, \mathbf{v}). \end{aligned}$$

This proves the convexity of $\varphi(\lambda; \mathbf{x}, \mathbf{v})$.

Now suppose that for each \mathbf{x} in C and \mathbf{v} in \mathbb{R}^n such that $J \neq \{0\}$, $\varphi(\lambda; \mathbf{x}, \mathbf{v})$ is convex. Let \mathbf{x}_1 and \mathbf{x}_2 be in C and let (α, β) be in \mathcal{P}_2 . Then

$$\begin{aligned} f(\alpha\mathbf{x}_1 + \beta\mathbf{x}_2) &= f(\alpha(\mathbf{x}_1 - \mathbf{x}_2) + \beta(\mathbf{x}_1 + \mathbf{x}_2)) \\ &= f(\alpha(\mathbf{x}_1 - \mathbf{x}_2) + \beta(\mathbf{x}_2 - \mathbf{x}_2)) = \varphi(\beta; \mathbf{x}_1, \mathbf{x}_2 - \mathbf{x}_2) \\ &= \varphi(\alpha + \beta; \mathbf{x}_1, \mathbf{x}_2 - \mathbf{x}_2). \end{aligned} \quad (4)$$

Since

$$\varphi(t)(x_2 - x_1) = f(x_1 + t(x_2 - x_1))$$

and since C is convex, $\varphi(t)(x_2 - x_1)$ is defined for $0 \leq t \leq 1$ and

$$\varphi(0)(x_2 - x_1) = f(x_1), \quad \varphi(1)(x_2 - x_1) = f(x_2). \quad (5)$$

From the convexity of $\varphi(t)(x_2 - x_1)$ and (5) we get that the rightmost side of (4) is less than or equal to

$$\varphi(0)(x_2 - x_1) + \beta(\varphi(1)(x_2 - x_1)) = \varphi'(x_1) + \beta f'(x_2).$$

Hence f is convex.

Remark 1.1. An examination of the proof above shows that f is strictly convex (if and only if, for each x and x_1 in (x_0, x) is strictly convex.

We now present some results for convex functions defined on real intervals. The first result is the so-called three-point property.

Theorem 1.2. Let g be a convex function defined on an interval $I \subset \mathbb{R}$. Then if $x < y < z$ are three points in I ,

$$\frac{g(y) - g(x)}{y - x} \leq \frac{g(z) - g(x)}{z - x} \leq \frac{g(z) - g(y)}{z - y}. \quad (6)$$

If g is strictly convex, then strict inequality holds.

Figure 1.1 illustrates the theorem and shows why the conclusion of the theorem is called the three-point property.

Proof. Since $x < y < z$, there exists a θ ($0 < \theta < 1$) such that $y = x + \theta(z - x)$. Since g is convex,

$$\begin{aligned} g(y) - g(x) &= g(\theta(z - x) + (1 - \theta)x) - g(x) \\ &\leq \theta(g(z) + (1 - \theta)g(x)) - g(x) \\ &= \theta(g(z) - g(x)). \end{aligned} \quad (7)$$

If in (7) we now divide through by $\theta(z - x) > 0$ and note that $y - x = \theta(z - x)$, we get that

$$\frac{g(y) - g(x)}{y - x} \leq \frac{g(z) - g(x)}{z - x}.$$



Figure 1.1.

To obtain the second inequality, we write

$$\begin{aligned} g(2) - g(1) &= g(2) - g(1) + (2) - (1) \\ &= g(2) - (2) + f(2) - f(1) \\ &= 0 + f(2) - f(1). \end{aligned}$$

If we now divide through by $x - y = 2 - 1$ and note that $x - y = (2) - (1) = 1$, we obtain the second inequality.

DEFINITION 1. For each side of the digeneration of f we define by

$$d_{\text{convex}}(f) = \frac{f(2) - f(1)}{2 - 1}, \quad \text{or } d_{\text{conv}}, \quad \text{and}$$

if an increasing function f is **CONVEX**, then d_{conv} is strictly increasing.

Proof. We must show that if $x_1 < x_2$, then

$$d_{\text{conv}}(f) < d_{\text{conv}}(f). \quad (1)$$

For two points x_1 and x_2 such that $x - y = x_2$, inequality (1) follows from (1) by taking $x = x_1$, $y = x_2$, and $x = x_1$. For two points x_1 and x_2 such that $x_1 < x_2$, i.e., inequality (1) follows from (1) by taking $x = x_1$, $y = x_2$, and $x = x_2$. Lastly, if $x_1 < x_2 < x_3$, the inequality (1) follows from (1) by taking $x = x_1$, $y = x_2$, and $x = x_3$.

Let f be a real-valued function defined on an interval I . If at a point c in I

$$\lim_{h \rightarrow 0^+} \frac{f(c+h) - f(c)}{h}$$

exists, then f is said to have a *right-hand derivative* at c whose value $f'_+(c)$ is the value of the above limit. Similarly, the *left-hand derivative* of f at c , denoted by $f'_-(c)$, is

$$\lim_{h \rightarrow 0^-} \frac{f(c+h) - f(c)}{h}$$

provided the limit exists.

COROLLARY 1. At each point c in $\text{int}(I)$ the right- and left-hand derivatives $f'_+(c)$ and $f'_-(c)$ exist and $f'_-(c) \leq f'_+(c)$. Moreover, f'_+ and f'_- are increasing functions.

Proof. In (1) take $y = c$. Then for $a < c < b$ we have

$$\frac{g(b) - g(c)}{b - c} \leq \frac{g(c) - g(a)}{c - a}.$$

It follows from this relation, from Corollary 1, and from Exercise 155, that $g'_+(c)$ and $g'_-(c)$ exist and $g'_-(c) \leq g'_+(c)$.

Now let $a = d$ be any point in $\text{int}(I)$ and let $h = 0$ be such that $c + h = d$ and $d - h > a$. Then by Theorem 1.2

$$\frac{g(c+h) - g(c)}{h} \leq \frac{g(d) - g(c)}{d - c} \leq \frac{g(d - h) - g(d)}{-h}.$$

If we now let $h \rightarrow 0^+$, we get $g'_+(c) \leq g'_-(d)$. Combining this with $g'_-(c) \leq g'_+(c)$ for any a in $\text{int}(I)$, we get

$$g'_-(c) \leq g'_+(c) \leq g'_-(d) \leq g'_+(d),$$

which gives the desired monotonicity property.

Remark 1.1. If c is a left-hand end point of I and g is defined at c , then from the monotonicity of $g'_-(c)$ it follows that $g'_-(c)$ always exists if we permit the value $-\infty$.

The next corollary is an immediate consequence of Corollary 2 and is a result that the reader encountered in elementary calculus.

COROLLARY 1. *If f is convex and twice differentiable on I , then f' is an increasing function on I and $f''(x) \geq 0$ for all x in I .*

We can now return to functions defined on \mathbb{R}^n .

Definition 1.4. Let f be a real-valued function defined on a set S in \mathbb{R}^n . Let \mathbf{x} be a point in S and let \mathbf{v} be a vector in \mathbb{R}^n such that $\mathbf{x} + \lambda\mathbf{v}$ belongs to S for λ in some interval $(0, \delta)$. Then f is said to have a directional derivative at \mathbf{x} in the direction \mathbf{v} if

$$\lim_{\lambda \rightarrow 0^+} \frac{f(\mathbf{x} + \lambda\mathbf{v}) - f(\mathbf{x})}{\lambda}$$

exists. We denote the directional derivative by $f'(\mathbf{x}; \mathbf{v})$.

Lemma 1.6. *Let f be a convex function defined on a convex set C and let \mathbf{x} be a point of C . If \mathbf{x} is an interior point, then $f'(\mathbf{x}; \mathbf{v})$ exists for every \mathbf{v} in \mathbb{R}^n . If \mathbf{x} is a boundary point and \mathbf{v} is a vector in \mathbb{R}^n such that $\mathbf{x} + \lambda\mathbf{v}$ belongs to C for sufficiently small $\lambda > 0$, then $f'(\mathbf{x}; \mathbf{v})$ exists if we allow the value $-\infty$.*

Proof. If \mathbf{x} is an interior point of C , then for each \mathbf{v} in \mathbb{R}^n there exists a positive number $\lambda_0 = \lambda_0(\mathbf{v})$ such that for $|\lambda| \leq \lambda_0$ the point $\mathbf{x} + \lambda\mathbf{v}$ is in C . Hence the function $\varphi(\lambda; \mathbf{x}, \mathbf{v})$ of (3) is defined for $|\lambda| \leq \lambda_0$. By Lemma 1.5, $\varphi(\cdot; \mathbf{x}, \mathbf{v})$ is convex, and so by Corollary 2 of Theorem 1.1, $\varphi'(\lambda; \mathbf{x}, \mathbf{v})$ exists. But

$$\frac{\varphi'(\lambda; \mathbf{x}, \mathbf{v}) - \varphi'(0; \mathbf{x}, \mathbf{v})}{\lambda} = \frac{f(\mathbf{x} + \lambda\mathbf{v}) - f(\mathbf{x})}{\lambda}$$

and so $f'(\mathbf{x}; \mathbf{v})$ exists and equals $\varphi'(0; \mathbf{x}, \mathbf{v})$. If \mathbf{x} is a boundary point of C , the assertion follows from similar arguments and Remark 1.1.

In the next lemma we record a useful inequality.

Lemma 1.7. *Let f be a convex function defined on a convex set C . Let \mathbf{x} and \mathbf{y} be points in C and let $\mathbf{z} = \alpha\mathbf{x} + (1 - \alpha)\mathbf{y}$, $0 < \alpha < 1$. Then*

$$f(\mathbf{z}) - f(\mathbf{x}) \geq \frac{f(\mathbf{z}) - f(\mathbf{y})}{1 - \alpha}. \quad (4)$$

Proof. The function $\varphi(\lambda; \mathbf{x}, \mathbf{y} - \mathbf{x})$ in (3) is defined for $0 \leq \lambda < 1$, with $\varphi(0; \mathbf{x}, \mathbf{y} - \mathbf{x}) = f(\mathbf{x})$, $\varphi(1; \mathbf{x}, \mathbf{y} - \mathbf{x}) = f(\mathbf{y})$ and $\varphi(\alpha; \mathbf{x}, \mathbf{y} - \mathbf{x}) = f(\mathbf{z})$ for $0 <$

$t < 1$. By Corollary 1 to Theorem 1.2,

$$\varphi(t)(y - x) - \varphi(0)(y - x) \geq \frac{\varphi(t)(y - x) - \varphi(0)(y - x)}{t},$$

which is the same as (9).

We note that (8) can also be established directly from the definition of a convex function.

We next show that a convex function is continuous on the interior of its domain of definition.

THEOREM 1.3. Let f be a convex function defined on a convex set C in \mathbb{R}^n . If x is an interior point of C , then f is continuous at x .

If x is not an interior point of C , then f need not be continuous at x , as the following example shows. Let $C = [-1, 1] \subset \mathbb{R}$ and let $f(x) = x^2$ for $-1 < x < 1$ and $f(-1) = f(1) = 2$. The function f is convex on $[-1, 1]$ since $\text{epi}(f)$ is convex, but f is discontinuous at ± 1 .

We now prove the theorem.

If x is an interior point of C , then there exists a $\rho > 0$ such that $\overline{B(x, \rho)} \subset C$. Let $\delta = \rho/\sqrt{n}$. Let C_δ denote the cube with center at x and length of side equal to 2δ ; that is,

$$C_\delta = \{y \mid |y_i - x_i| \leq \delta, i = 1, \dots, n\}.$$

Since

$$\|y - x\| = \left[\sum_{i=1}^n (y_i - x_i)^2 \right]^{1/2},$$

it follows that $\|y - x\| \leq \rho/\sqrt{n}$ for each y in C_δ . Thus $C_\delta \subset \overline{B(x, \rho)}$. On the other hand, if $\|y - x\| \leq \delta$, then we cannot have $|y_i - x_i| > \delta$ for any index i in set $\{1, \dots, n\}$. Hence $B(x, \delta) \subset C_\delta$.

The first step in the proof is to show that f is bounded above on $B(x, \delta)$. It is easier to show that f is bounded above on the larger set C_δ than to work directly with $B(x, \delta)$. Let $C_{2\delta}$ denote the set of 2^n vertices of C_δ . By Exercise II.4.7(d), the set $C_{2\delta}$ is the set of extreme points of C_δ . By Theorem II.4.1, $C_\delta = \text{co}(C_{2\delta})$. But $C_{2\delta}$ is compact, so by Lemma II.1.7, $\text{co}(C_{2\delta})$ is compact. Hence

$$C_\delta = \overline{\text{co}(C_{2\delta})} = \text{co}(C_{2\delta}).$$

[We could also conclude this directly from Exercise II.4.7 (c).] Thus, for each

p is in C_1 , then there is a point P_1 such that

$$P_1 = \sum_{j=1}^n \alpha_j P_j,$$

where the α_j are the vertices of C_1 .

Let

$$M = \max\{\|P_j\| : j = 1, \dots, n\}.$$

Then by the convexity of f and Theorem 1.1 (Jensen's inequality), because p is in C_1

$$f(P_1) \leq \sum_{j=1}^n \alpha_j f(P_j) \leq \sum_{j=1}^n \alpha_j M = M.$$

Thus, since $\|P_1\| \in C_1$,

$$f(\|P_1\|) \leq M \quad \text{for } p \text{ in } \text{int}(C_1). \quad (10)$$

Let x be an arbitrary point in $\text{int}(C)$, and consider the line determined by x and 0 . This line will intersect the surface of $\text{int}(C)$ at two points, $x = \alpha$ and $x = -\alpha$, where $\alpha = \alpha_1 x_1 + \dots + \alpha_n x_n$ and $\alpha_j \geq 0$ and $\sum \alpha_j = 1$ (see Figure 1a).

We shall consider first α , and the convexity of f becomes

$$f(\alpha) \leq f(\| \alpha \|)$$

in terms of $\{\alpha_1, \dots, \alpha_n\}$ and this gives the coefficients of f at α . On the other, we

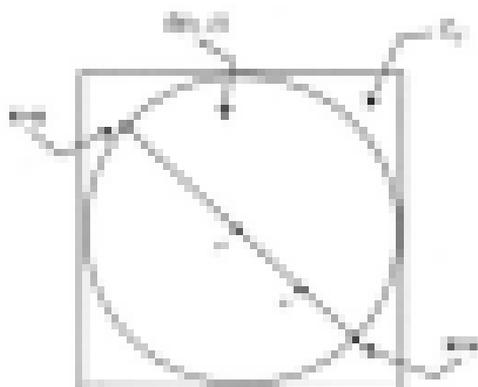


Figure 1a.

first express x as a convex combination of x and $x + u$. Thus

$$x = (1 - \lambda)(x + \lambda(x + u)) = x + \lambda u, \quad 0 < \lambda < 1. \quad (11)$$

Therefore, $\lambda = (x - u)^{-1}$. We next express x as a convex combination of u and $x - u$. Thus

$$x = (1 - \alpha)(x - u) + \alpha u, \quad 0 < \alpha < 1.$$

Invoking relation (1) in Section 1 of Chapter II and using the relation $\lambda = (x - u)^{-1}$, we see that

$$\lambda(1 - \lambda)^{-1} = \lambda(x - u)^{-1} = \lambda^{-2},$$

and so $\alpha = (1 + \lambda)^{-1}$. Therefore,

$$x = \frac{1}{1 + \lambda}u + \frac{\lambda}{1 + \lambda}(x - u). \quad (12)$$

Since f is convex, we have, using (10), (11), and (12),

$$f(x) \leq (1 - \lambda)f(x) + \lambda f(x + u) \leq (1 - \lambda)f(x) + \lambda M,$$

$$f(x) \leq \frac{1}{1 + \lambda}f(u) + \frac{\lambda}{1 + \lambda}f(x - u) \leq \frac{1}{1 + \lambda}f(u) + \frac{\lambda}{1 + \lambda}M.$$

From these inequalities we get

$$-K[M - f(x)] \leq f(u) - f(x) \leq \lambda[M - f(x)].$$

Hence, for $u \in B(x, \delta)$

$$|f(u) - f(x)| \leq \delta^{-1}K[M - f(x)]\|u - x\|, \quad (13)$$

from which the continuity of f at x follows.

Remark 1.1. A function f defined on a set D in \mathbb{R}^n with range in \mathbb{R}^m is said to be Lipschitz continuous on D if there exists a constant $K > 0$ such that

$$\|f(x) - f(y)\| \leq K\|x - y\|$$

for all x, y in D . Since a real-valued continuous function defined on a compact set is bounded above and below, relation (13) implies that a convex function defined on an open convex set C is Lipschitz continuous on every compact subset contained in C .

Remark 1.1. The proof of Theorem 1.3 can easily be modified to show that if f is convex on a convex set C , then f is continuous on $\text{int} C$.

In Section 11 of Chapter I we saw that the existence of partial derivatives at a point does not imply differentiability. For convex functions, however, the existence of partial derivatives at a point does imply differentiability.

THEOREM 1.4. Let f be a convex function defined on a convex set C . If the partial derivatives of f with respect to each variable exist at an interior point \mathbf{x} of C , then f is differentiable at \mathbf{x} .

Proof. Define the function η on \mathbb{R}^n by

$$\eta(\mathbf{h}) = f(\mathbf{x} + \mathbf{h}) - f(\mathbf{x}) - \sum_{j=1}^n \frac{\partial f}{\partial x_j}(\mathbf{x}) h_j.$$

To prove that f is differentiable at \mathbf{x} , we must show that

$$\frac{\eta(\mathbf{h})}{\|\mathbf{h}\|} \rightarrow 0 \quad \text{as } \|\mathbf{h}\| \rightarrow 0. \quad (14)$$

Let

$$\eta_j = \frac{f(\mathbf{x} + h_j \mathbf{e}_j) - f(\mathbf{x})}{h_j} - \frac{\partial f}{\partial x_j}(\mathbf{x}).$$

Since f is convex and a linear function is concave and convex, the function η is convex. Thus

$$\begin{aligned} \eta(\mathbf{h}) &= \eta\left(\sum_{j=1}^n \frac{h_j \mathbf{e}_j}{\|\mathbf{h}\|} \|\mathbf{h}\|\right) \leq \frac{1}{\|\mathbf{h}\|} \sum_{j=1}^n \eta(h_j \mathbf{e}_j) \\ &= \sum_{j=1}^n \eta_j h_j \leq \sum_{j=1}^n |h_j| |\eta_j| \leq \left(\sum_{j=1}^n |\eta_j|\right) \|\mathbf{h}\| \end{aligned} \quad (15)$$

for $j = 1, \dots, n$. The existence of the partial derivatives implies that each of the $\eta_j \rightarrow 0$ as $\|\mathbf{h}\| \rightarrow 0$. Also,

$$0 = \eta(\mathbf{0}) = \eta\left(\frac{\mathbf{h} + (-\mathbf{h})}{2}\right) \leq \frac{\eta(\mathbf{h}) + \eta(-\mathbf{h})}{2}.$$

Hence

$$\eta(\mathbf{h}) \geq -\eta(-\mathbf{h}) \geq -\left(\sum_{j=1}^n |\eta_j|\right) \|\mathbf{h}\|. \quad (16)$$

where, for each $i = 1, \dots, n$, $c_i^2 \rightarrow 0$ as $\|h\| \rightarrow 0$. We now obtain (14) from (15) and (16).

Exercise 1.1. Show that if f and g are convex functions defined on a convex set C , then for any $\lambda > 0$, $\mu > 0$, the function $\lambda f + \mu g$ is convex.

Exercise 1.2. Let $\{f_n\}$ be a sequence of convex functions defined on a convex set C . Show that if $f(x) = \lim_{n \rightarrow \infty} f_n(x)$ exists for all $x \in C$, then f is convex.

Exercise 1.3. Prove Lemma 1.2.

Exercise 1.4. Prove Lemma 1.4.

Exercise 1.5. Let $J = [a, b]$ be a real interval and let C be a convex set in \mathbb{R}^n . Let $f(x, \alpha) = f(x, \alpha)$ be a real-valued function defined on $J \times C$ that is continuous on J for each fixed x in C and is convex on C for each fixed x in J . Show that

$$F(x) = \int_a^b f(x, \alpha) d\alpha$$

is defined for each x in C and that F is convex on C .

Exercise 1.6. (a) Show that any norm $\|\cdot\|$ on \mathbb{R}^n is convex on \mathbb{R}^n .

(b) Let S be a nonempty convex set, and let $\|\cdot\|$ denote the euclidean norm in \mathbb{R}^n . Let f be the distance function to S defined by

$$f(x) = \inf\{\|y - x\| : y \in S\}.$$

Show that f is convex on \mathbb{R}^n .

(c) Show that f is Lipschitz continuous on compact subsets of \mathbb{R}^n .

Exercise 1.7. Let f be a convex function defined on \mathbb{R}^n . Let x_0 be a fixed point in \mathbb{R}^n . Show that the function $h(\lambda) : \mathbb{R}^+ \rightarrow \mathbb{R}$ defined by $h(\lambda) = f(x_0 + \lambda v)$ is (a) positively homogeneous [i.e., for $\lambda > 0$, $h(\lambda v) = \lambda h(v)$] and (b) convex.

Exercise 1.8. Let g be a convex function defined on \mathbb{R}^n . Let $S = \{x : g(x) > 0\}$ be nonempty.

(a) Show that S is convex.

(b) Show that if $f(x) = 1/g(x)$ for x in S , then f is convex on S .

Exercise 1.8. Let C be a convex cone (see Exercise II.2.10 for the definition). A real-valued function g defined on C is said to be a *gauge function* on C if

- (a) $g(\alpha x) = \alpha g(x)$ for all x in C and $\alpha \geq 0$ and
 (b) $g(x + y) \leq g(x) + g(y)$ for all x, y in C .

Show that g is convex.

Exercise 1.9. Let C be a convex set. The *support function* s of C is the real-valued function defined by

$$s(\mathbf{n}) = \sup\{(\mathbf{n}, \mathbf{x}) : \mathbf{x} \in C\}$$

for all \mathbf{n} such that the supremum is finite.

- (a) Show that if the domain D of s is nonempty, then D is in a convex cone.
 (b) Show that s is a gauge function on D .

Exercise 1.11. Let C be a compact convex set with $0 \in \text{int}(C)$. The *Minkowski distance function* p determined by C is defined by

$$p(\mathbf{x}) = \inf\{\lambda : \lambda \geq 0, \mathbf{x} \in \lambda C\}.$$

Show that p is a gauge function.

Exercise 1.12. Let C be a compact convex set, and let s be its support function. Show that the following statements hold:

- (a) The domain of s is \mathbb{R}^n .
 (b) For each $\mathbf{n} \neq \mathbf{0}$ there exists a point \mathbf{x}_n in C such that $s(\mathbf{n}) = (\mathbf{n}, \mathbf{x}_n)$.
 (c) The hyperplane $H_n = \{\mathbf{x} : (\mathbf{n}, \mathbf{x}) = s(\mathbf{n})\}$ supports C at \mathbf{x}_n .
 (d) The distance from H_n to $\mathbf{0}$ is equal to $s(\mathbf{n})/\|\mathbf{n}\|$.

Exercise 1.13. Let C be a nonempty closed convex set $C \neq \mathbb{R}^n$. Let s be the support function of C with nonempty domain D . Show that

$$C = \{\mathbf{x} : (\mathbf{n}, \mathbf{x}) \leq s(\mathbf{n}) \text{ for all } \mathbf{n} \text{ in } D\}.$$

Hint: An alternate definition of C is

$$C = \bigcap_{\mathbf{n} \in D} \{\mathbf{x} : (\mathbf{n}, \mathbf{x}) \leq s(\mathbf{n}), \mathbf{n} \text{ fixed in } D\},$$

and see Exercise II.6.6.

2. SUBGRADIENTS

If f is a real-valued function defined on a set C in \mathbb{R}^n , then the graph of f is the set of points (\mathbf{x}, z) in \mathbb{R}^{n+1} of the form $(\mathbf{x}, f(\mathbf{x}))$, where $\mathbf{x} \in C$. If f is a differentiable function, then given a point \mathbf{x}_0 in the interior of C , the graph of f has a tangent plane at $(\mathbf{x}_0, f(\mathbf{x}_0))$ whose normal vector is $(\nabla f(\mathbf{x}_0), -1)$.

If C is convex and f is convex, f need not be differentiable at an interior point of C , as the function f defined by $f(x) = |x|$ shows. In this section we shall introduce a generalization of the gradient vector, called the subgradient, that exists at points $(\mathbf{x}_0, f(\mathbf{x}_0))$ of the graph of a convex function f , where \mathbf{x}_0 is an interior point of C . We shall show that if f is a convex function defined on C and \mathbf{x} is an interior point of C , then $\text{epi}(f)$ has a supporting hyperplane at the point $(\mathbf{x}, f(\mathbf{x}))$. The supporting hyperplanes need not be unique. In the next section we shall show that f is differentiable at \mathbf{x} if and only if f has a unique subgradient at \mathbf{x} . This subgradient turns out to be $\nabla f(\mathbf{x})$. Thus, if f is differentiable at \mathbf{x} , there is a unique supporting hyperplane to $\text{epi}(f)$ at $(\mathbf{x}, f(\mathbf{x}))$. This hyperplane is the tangent plane to the graph at $(\mathbf{x}, f(\mathbf{x}))$. If $(\mathbf{q}, -1)$ is a normal vector to a supporting hyperplane, then \mathbf{q} has some of the properties that the gradient vector would have.

We now define the notion of subgradient.

Definition 2.1. A function f defined on a set C is said to have a subgradient at a point \mathbf{x}_0 in C if there exists a vector \mathbf{q} in \mathbb{R}^n such that

$$f(\mathbf{x}) \geq f(\mathbf{x}_0) + (\mathbf{q}, \mathbf{x} - \mathbf{x}_0)$$

for all $\mathbf{x} \in C$. The vector \mathbf{q} is called a subgradient.

Geometrically, \mathbf{q} is a subgradient of f at \mathbf{x}_0 if the graph of f in \mathbb{R}^{n+1} lies on or above the graph of the hyperplane $z = f(\mathbf{x}_0) + (\mathbf{q}, \mathbf{x} - \mathbf{x}_0)$. Since $(\mathbf{x}_0, f(\mathbf{x}_0))$ is in this hyperplane, this hyperplane will be a supporting hyperplane to $\text{epi}(f)$ at $(\mathbf{x}_0, f(\mathbf{x}_0))$. Thus, the existence of a subgradient is a statement about the existence of a nonvertical supporting hyperplane to $\text{epi}(f)$. In Figure 2.1 we illustrate the definition with the function $f(x) = |x|$ at $x = 0$.

Any line through the origin with slope ζ satisfying $-1 \leq \zeta \leq 1$ has the property that

$$|x| \geq \zeta x = \zeta(0) + \zeta(x - 0)$$

for all x . Thus any ζ in the interval $[-1, 1]$ is a subgradient of $|x|$ at $x = 0$. If $\mathbf{x}_0 > 0$, then $\zeta = 1$ is the only subgradient of $|x|$ at \mathbf{x}_0 ; if $\mathbf{x}_0 < 0$, then $\zeta = -1$ is the only subgradient of $|x|$ at \mathbf{x}_0 .

Definition 2.2. The set of subgradients of a function f at a point \mathbf{x}_0 is called the subdifferential of f at \mathbf{x}_0 and is denoted by $\partial f(\mathbf{x}_0)$.

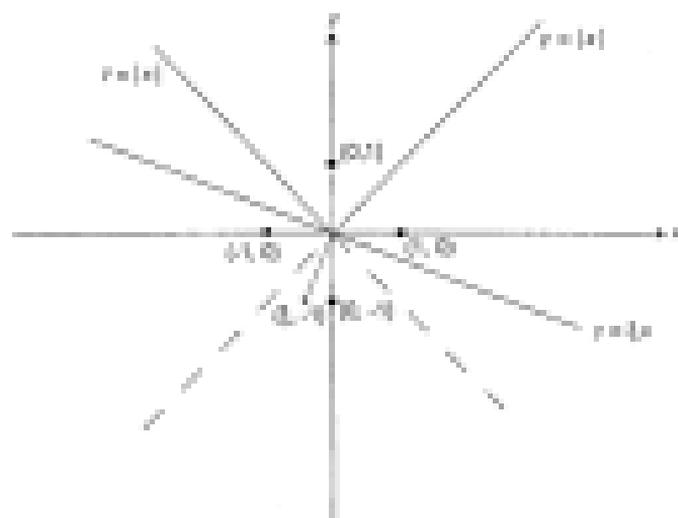


Figure 15

The subdifferential of $|x|$ at $x = 0$ is the interval $[-1, 1]$. The subdifferential of $|x|$ at $x_0 > 0$ is the singleton $\{1\}$, and the subdifferential at $x_0 < 0$ is $\{-1\}$.

Example 1.7. The subdifferential will be the empty set if no subgradient exists. This may happen, as the following example shows. Let $C = [-1, 1]$ and let $f(x) = -\sqrt{1-x^2}$. (Sketch the graph.) The graph of f has a nonvertical tangent line at any point x_0 in $(-1, 1)$. From the graph it appears that for such x_0 the graph lies above the tangent line, and so $\partial f(x_0) = \{f'(x_0)\}$. Although this can be verified in this example by relatively straightforward calculations, we shall not do so. The validity of the statement will follow from general theorems that we will prove. At the points $x = \pm 1$, the graph has a vertical tangent line, but the tangent line does not lie below the graph. Thus it appears that f does not have a subgradient at these points. We now show this analytically.

A number ζ is a subgradient of f at $x = -1$ if and only if

$$-\sqrt{1-x^2} \geq \zeta - \zeta(x+1) \quad \text{for all } x \text{ in } [-1, 1].$$

For x in $(-1, 1]$ this inequality is equivalent to

$$-\left(\frac{1-x^2}{1+x}\right)^{1/2} \geq \zeta - \zeta(x+1), \quad x \text{ in } (-1, 1].$$

If we let x in $(-1, 1]$ tend to -1 , the left-hand side of this inequality tends to $-\infty$. Hence there is no $\zeta \in \mathbb{R}$ such that the subgradient inequality holds for all x in $(-1, 1]$.

A similar argument shows that $\partial f(1) = \emptyset$.

THEOREM 2.1. Let f be a convex function defined on a convex set C . Then at each point $x_0 \in \text{int}(C)$ there exists a vector $\xi \in \mathbb{R}^n$ such that

$$f(x) \geq f(x_0) + \langle \xi, x - x_0 \rangle \quad (1)$$

for all $x \in C$. In other words, $\partial f(x_0) \neq \emptyset$ at every interior point x_0 of C .

Example 2.1 shows that the requirement $x_0 \in \text{int}(C)$ is essential.

Proof. We shall show that the conclusion of the theorem is a consequence of the existence of a supporting hyperplane to $\text{epi} f$ at $(x_0, f(x_0))$.

The point $(x_0, f(x_0))$ is a boundary point of $\text{epi} f$. Since $\text{epi} f$ is convex, we may apply Theorem II.4.1 with $C = \text{epi} f$ and get that there exists a vector $(\alpha, \beta) \neq (0, 0)$ in \mathbb{R}^{n+1} such that

$$\langle \alpha, x \rangle + \beta y \leq \langle \alpha, x_0 \rangle + \beta f(x_0)$$

for all $(x, y) \in \text{epi} f$. This relation is equivalent to

$$\langle \alpha, x \rangle + \beta y \leq \langle \alpha, x_0 \rangle + \beta f(x_0) \quad (2)$$

for all $(x, y) \in \text{epi} f$, that is, for all $x \in C$ and $y \geq f(x)$.

Since the right-hand side of (2) is a constant, if β were strictly positive, we could obtain a contradiction by letting y become arbitrarily large. Hence $\beta \leq 0$.

If $\beta = 0$, then $\alpha \neq 0$ and (2) becomes

$$\langle \alpha, x \rangle \leq \langle \alpha, x_0 \rangle \quad \text{all } x \in C. \quad (3)$$

Since $x_0 \in \text{int}(C)$, there exists an $\epsilon > 0$ such that $x_\epsilon = x_0 + \epsilon \alpha$ belongs to C . Setting $x = x_\epsilon$ in (3) gives $\epsilon \langle \alpha, \alpha \rangle \leq 0$, which is a contradiction. Thus, $\beta < 0$.

In (2) we now divide through by $\beta < 0$ and set $y = f(x)$ to get

$$f(x) \geq f(x_0) + \left\langle \frac{\alpha}{\beta}, x_0 - x \right\rangle$$

for all $x \in C$. Now set $\xi = -\alpha/\beta$ to get (1).

COROLLARY 1. If f is strictly convex, then strict inequality holds in (1).

We leave the proof as an exercise.

The converse of Theorem 2.1 is false. To see this, let C be the square in \mathbb{R}^2 , $[x_1 \geq 0, x_2 \leq 1, 0 \leq x_1 \leq 1, 0 \leq x_2 \leq 1]$. Let $f(x) = 0$ for x such that $0 \leq x_1 \leq 1$ and $0 < x_2 \leq 1$. Let $f(x) = \frac{1}{2} - |x_1 - \frac{1}{2}|$ for x such that $0 \leq x_1 \leq 1$ and $x_2 = 0$. The function f has a unique subgradient $\xi = 0$ at every point in $\text{int}(C)$, yet f is not convex on C . The function ξ however, is convex on $\text{int}(C)$, and this phenomenon illustrates the general state of affairs.

Theorem 2.2. Let C be a convex set in \mathbb{R}^n . If at each point x_0 in C f has a subgradient, then f is convex on C .

Remark 2.1. If we merely assume that at each point in $\text{int}(C)$ the subdifferential is not empty, then we can only conclude that f is convex on $\text{int}(C)$.

Proof. Let x and y be two points in C and let $\alpha = \alpha x + \beta y$, $\alpha > 0$, $\beta > 0$, $\alpha + \beta = 1$. Since f has a subgradient at x , we get

$$f(x) \geq f(x) + \langle \xi, x - x \rangle = f(x) + \beta \langle \xi, x - y \rangle,$$

$$f(y) \geq f(y) + \langle \eta, y - y \rangle = f(y) - \alpha \langle \eta, x - y \rangle.$$

If we now multiply the first inequality by $\alpha > 0$, the second inequality by $\beta > 0$, and add the results, we get

$$\alpha f(x) + \beta f(y) \geq f(\alpha x + \beta y).$$

Hence f is convex.

Remark 2.2. The proof of the theorem shows that if at each x_0 in C there exists a subgradient ξ such that $f(x) \geq f(x_0) + \langle \xi, x - x_0 \rangle$, for $x \in C$, then f is strictly convex on C .

Exercise 2.1. Show that if $\partial f(x_0)$ is not empty, then $\partial f(x_0)$ is closed and convex.

Exercise 2.2. Prove Corollary 1 to Theorem 2.1.

Exercise 2.3. Let f be a convex function defined on a convex set C and let $x_0 \in \text{int}(C)$. Show that $\xi \in \partial f(x_0)$ if and only if $f(x_0 + v) \geq \langle \xi, v \rangle$ for all v in \mathbb{R}^n .

Exercise 2.4. Let f be convex and concave on \mathbb{R}^n . Show that f is affine, that is, $f(x) = \langle a, x \rangle + b$.

Exercise 2.5. Let $f(x) = |x|$. Show that $\partial f(0) = \{\xi \in \mathbb{R}^n \mid \|\xi\| \leq 1\}$. (From Theorem 2.1, we will get that, for $x \neq 0$, $\partial f(x) = \{\xi \mid x \in \xi\}$.)

Exercise 14. Let $f(x_1, x_2) = x_1 + x_2$. Show that f is convex, and calculate $\partial f(x, 0)$.

5. DIFFERENTIABLE CONVEX FUNCTIONS

We first show that if f is a convex function and differentiable [in the sense of relation (2) in Section 11 of Chapter I] at an interior point of its domain, then there is a unique subgradient at the point and the subgradient is the gradient vector, and conversely, if f has a unique subgradient at x , then f is differentiable at x .

Theorem 2.1. Let f be a convex function defined on a convex set C in E^n , and let x be an interior point of C . Then f is differentiable at x if and only if f has a unique subgradient at x . Moreover, the unique element in the subgradient is $\nabla f(x)$.

Proof. Let f be differentiable at x . Since f is convex and $x \in \text{int}(C)$, by Theorem 2.1 the subgradient $\partial f(x)$ is not empty. Let $\xi \in \partial f(x)$. Then for any h in E^n and sufficiently small $\epsilon > 0$

$$f(x + \epsilon h) \geq f(x) + \langle \xi, \epsilon h \rangle.$$

Since f is differentiable at x , we have that for $h \neq 0$

$$f(x + \epsilon h) - f(x) + \langle \nabla f(x), \epsilon h \rangle = o(\epsilon),$$

where $o(\epsilon)/\epsilon \rightarrow 0$ as $\epsilon \rightarrow 0$. If we subtract the second relation from the first, we get

$$0 \geq \epsilon \langle \xi, h \rangle - \langle \nabla f(x), \epsilon h \rangle - o(\epsilon)(h)/\epsilon.$$

If we now divide by $\epsilon > 0$ and then let $\epsilon \rightarrow 0$, we get that $\langle \xi - \nabla f(x), h \rangle \leq 0$ for all $h \neq 0$ in E^n . In particular, if $\xi - \nabla f(x) \neq 0$, then this inequality holds for $h = \xi - \nabla f(x)$, whence $\langle \xi - \nabla f(x), \xi - \nabla f(x) \rangle < 0$. Thus, $\xi = \nabla f(x)$.

We now suppose that the subgradient $\partial f(x)$ consists of a unique element ξ . We will show that all the first-order partial derivatives of f exist at x . The differentiability of f at x will then follow from Theorem 1.4.

To simplify the notation, we assume that $x = 0$ and that $f(x) = 0$. There is no loss of generality in making this assumption, since this amounts to a translation of the origin in $E^n \times E$ to $(x, f(x))$. Let us now suppose that at least one partial derivative of f at 0 does not exist. For definiteness let us assume that it is the partial with respect to x_1 . Let

$$g(t) = f(0 + t e_1),$$

where e_1 is the first standard basis vector $(1, 0, \dots, 0)$. Since f is convex and $\mathcal{D}_1 \in \mathcal{C}$, the function g is defined and convex on some maximal interval J containing $\mathbf{0}$ in its interior. Moreover, $f'(x)$ does not exist at $\mathbf{0}$ if and only if $g'(0)$ does not exist. Since g is convex, it follows from Corollary 1 to Theorem 1.2 that $g'(0)$ and $g'_+(0)$ exist and that $g'_+(0) \leq g'_-(0)$. Therefore, if $g'(0)$ does not exist, then $g'_+(0) < g'_-(0)$.

Let d_1 be any number satisfying

$$g'_-(0) = d_1 < g'_+(0). \quad (11)$$

We shall show that corresponding to each such d_1 we can find a ξ in \mathcal{D}_1 with distinct d_1 determining distinct ξ . This will contradict the hypothesis that \mathcal{D}_1 consists of a unique element, and the proof will be accomplished.

It follows from (1) and Corollary 1.1 that for all $h > 0$

$$\frac{g(-h) - g(0)}{-h} \leq g'_-(0) = d_1 < g'_+(0) \leq \frac{g(h) - g(0)}{h}.$$

Since $g(0) = f(0) = 0$, we get from the preceding inequalities that for all $h > 0$

$$f(h e_1) = g(h) > d_1 h, \quad (12)$$

with equality holding only if $h = 0$. Let V_1 denote the one-dimensional subspace spanned by e_1 . Thus $V_1 = \{x; x = \alpha e_1, \alpha \in \mathbb{R}\}$. Define a linear functional L_1 on V_1 by the formula

$$L_1(x) = L_1(\alpha e_1) = d_1 \alpha. \quad (13)$$

Since the representation $x = \alpha e_1$ is unique for x in V_1 , the linear functional L_1 is well defined. Note that L_1 depends on d_1 . Combining (2) and (13), we get that for x in $V_1 \cap \mathcal{C}$

$$f(x) \geq L_1(x) = d_1 x. \quad (14)$$

If $n = 1$, we are through. Otherwise, let V_2 denote the two-dimensional subspace spanned by e_1 and e_2 . Thus

$$V_2 = \{y; y = \alpha e_1 + \beta e_2, \alpha, \beta \in \mathbb{R}\} = \{y; y = x + \alpha e_2, x \in V_1, \alpha \in \mathbb{R}\}.$$

We shall extend L_1 to a linear functional L_2 defined on V_2 and shall find a number d_2 such that for all y in $V_2 \cap \mathcal{C}$

$$f(y) \geq L_2(y) = d_2 y. \quad (15)$$

Let \mathbf{x} and \mathbf{y} belong to $V_1 \cap C$ and let $\alpha > 0$, $\beta > 0$. It then follows from the linearity of $L_{\mathbf{y}}$, from the convexity of $V_1 \cap C$, from (4), and from the convexity of f that

$$\begin{aligned} \frac{\alpha}{\alpha+\beta} L_{\mathbf{y}}(\mathbf{x}) + \frac{\beta}{\alpha+\beta} L_{\mathbf{y}}(\mathbf{y}) &= L_{\mathbf{y}}\left(\frac{\alpha}{\alpha+\beta}\mathbf{x} + \frac{\beta}{\alpha+\beta}\mathbf{y}\right) \\ &\leq f\left(\frac{\alpha}{\alpha+\beta}\mathbf{x} + \frac{\beta}{\alpha+\beta}\mathbf{y}\right) \\ &= f\left(\frac{\alpha}{\alpha+\beta}(\mathbf{x} - \beta\mathbf{w}_1) + \frac{\beta}{\alpha+\beta}(\mathbf{y} + \alpha\mathbf{w}_2)\right) \\ &\leq \frac{\alpha}{\alpha+\beta} f(\mathbf{x} - \beta\mathbf{w}_1) + \frac{\beta}{\alpha+\beta} f(\mathbf{y} + \alpha\mathbf{w}_2). \end{aligned}$$

Multiplying through by $\alpha + \beta$ gives

$$\alpha\mathbf{w}_1(\mathbf{x}) + \beta L_{\mathbf{y}}(\mathbf{x}) \leq \alpha(\mathbf{x} - \beta\mathbf{w}_1) + \beta(\mathbf{y} + \alpha\mathbf{w}_2)$$

and so

$$\frac{L_{\mathbf{y}}(\mathbf{x}) - f(\mathbf{x} - \beta\mathbf{w}_1)}{\beta} \leq \frac{f(\mathbf{y} + \alpha\mathbf{w}_2) - L_{\mathbf{y}}(\mathbf{y})}{\alpha}. \quad (5)$$

Denote the left-hand side of (5) by $\lambda(\mathbf{x}, \beta)$ and denote the right-hand side of (5) by $\mu(\alpha, \mathbf{y})$. Hence

$$\sup\{\lambda(\mathbf{x}, \beta) : \mathbf{x} \in V_1 \cap C, \beta > 0\} \leq \inf\{\mu(\alpha, \mathbf{y}) : \mathbf{y} \in V_1 \cap C, \alpha > 0\}$$

and both the supremum and infimum are finite. Hence there exists a number $L_{\mathbf{y}}$ such that

$$\frac{L_{\mathbf{y}}(\mathbf{x}) - f(\mathbf{x} - \beta\mathbf{w}_1)}{\beta} \leq L_{\mathbf{y}} \leq \frac{f(\mathbf{y} + \alpha\mathbf{w}_2) - L_{\mathbf{y}}(\mathbf{y})}{\alpha} \quad (7)$$

for all $\mathbf{x} \in V_1 \cap C$, $\alpha > 0$, $\beta > 0$.

Now let \mathbf{y} be an element in $V_1 \cap C$ that is not in V_2 . Then $\mathbf{y} = \mathbf{x} + \alpha\mathbf{w}_1$ with $\mathbf{x} \in \mathcal{O}$ and $\alpha > 0$. If $\mathbf{x} \in \mathcal{O}$, take $\alpha = \alpha$ in (7). If $\mathbf{x} \in \mathcal{O}$, take $\beta = -\alpha$ in (7). Then from (7) and (5) we get that for all \mathbf{y} in $V_2 \cap C$

$$f(\mathbf{y}) - f(\mathbf{x} + \alpha\mathbf{w}_1) \geq L_{\mathbf{y}}(\mathbf{y}) + L_{\mathbf{y}}\alpha - L_{\mathbf{y}}\alpha + L_{\mathbf{y}}\alpha,$$

which is (3).

Proceeding inductively, we get that for each i , satisfying (1) there exists a vector $\xi = (\xi_1, \xi_2, \dots, \xi_n)$ such that if $\mathbf{x} = (x_1, \dots, x_n)$ is in C , then

$$f(\mathbf{x}) \geq \xi_0 + \xi \cdot \mathbf{x}$$

Since $f(\mathbf{0}) = 0$, this says that $\xi_0 \leq f(\mathbf{0})$.

The next theorem is an immediate consequence of Theorems 3.1, 3.2, and 3.3.

THEOREM 3.2. Let C be an open convex set in \mathbb{R}^n and let f be real valued and differentiable on C . Then f is convex if and only if for each \mathbf{x}_0 in C

$$f(\mathbf{x}) \geq f(\mathbf{x}_0) + (Df)(\mathbf{x}_0) \cdot (\mathbf{x} - \mathbf{x}_0)$$

for all \mathbf{x} in C . Also, f is strictly convex if and only if for each \mathbf{x}_0 in C and all $\mathbf{x} \neq \mathbf{x}_0$ in C

$$f(\mathbf{x}) > f(\mathbf{x}_0) + (Df)(\mathbf{x}_0) \cdot (\mathbf{x} - \mathbf{x}_0).$$

REMARK 3.2. If $n = 2$ and f is differentiable at \mathbf{x}_0 , then the surface in \mathbb{R}^3 defined by $y = f(\mathbf{x})$ has a tangent plane at \mathbf{x}_0 whose equation can be written as

$$y - f(\mathbf{x}_0) = (Df)(\mathbf{x}_0) \cdot (\mathbf{x} - \mathbf{x}_0).$$

Thus, the geometric interpretation of Theorem 3.2 is that a function f that is differentiable on C is convex on C if and only if for each point \mathbf{x}_0 in C the surface defined by $y = f(\mathbf{x})$ lies above the tangent plane to the surface at \mathbf{x}_0 . This is not surprising since we obtained the result in Theorem 3.2 by considering the supporting hyperplanes to $\text{epi}(f)$, and the differentiable f thus are the tangent planes to the surface.

Let f be a real-valued function defined and of class C^2 on an open set D in \mathbb{R}^n (see Section 11 in Chapter I for the definition of class C^2). For such functions we have a computationally feasible criterion for convexity. In the ensuing discussion recall that, by Theorem 3.1.1.3, f is differentiable at each point \mathbf{x}_0 in D .

We first develop our criterion if D is an interval $I \subseteq \mathbb{R}$.

LEMMA 3.1. Let f be of class C^2 on an open interval $I \subseteq \mathbb{R}$. Then f is convex on I if and only if $f''(x) \geq 0$ for all x in I . If $f''(x) > 0$ or all x in I , then f is strictly convex.

Note that if f is strictly convex, then we cannot conclude that $f''(x) > 0$ for all x , as the example $f(x) = x^4$ shows. The function f is strictly convex, yet $f''(0) = 0$.

Proof. If f is convex on J , then by Corollary 3 to Theorem 3.2, $f''(x) \geq 0$ on J . Conversely, suppose $f''(x) \geq 0$ on J . Let x_0 be a fixed point in J and let x be any other point in J . Then by Taylor's theorem

$$f(x) - f(x_0) - f'(x_0)(x - x_0) = \frac{1}{2}f''(x_1)(x - x_0)^2 \geq 0,$$

where x_1 is a point between x_0 and x . It now follows from Theorem 3.2 that f is convex on J . If $f''(x) > 0$ on J it follows from Theorem 3.2 that f is strictly convex.

Recall that a symmetric $n \times n$ matrix A with entries (a_{ij}) is said to be positive semidefinite if the quadratic form

$$\mathbf{x}'A\mathbf{x} = (\mathbf{x}, A\mathbf{x}) = \sum_{i=1}^n a_{ii}x_i^2$$

is nonnegative for all \mathbf{x} in \mathbb{R}^n . The matrix A is said to be positive definite if $(\mathbf{x}, A\mathbf{x})$ is positive for all $\mathbf{x} \neq \mathbf{0}$ in \mathbb{R}^n .

THEOREM 3.3. Let f be of class $C^{2,1}$ on an open convex set D . Then f is convex on D if and only if the Hessian matrix

$$H(\mathbf{x}) = \left(\frac{\partial^2 f}{\partial x_i \partial x_j}(\mathbf{x}) \right)$$

is positive semidefinite at each point \mathbf{x} in D . If $H(\mathbf{x})$ is positive definite at each \mathbf{x} , then f is strictly convex.

Proof. Let \mathbf{x} be a point in D . Since D is open, for each vector \mathbf{v} in \mathbb{R}^n , $\mathbf{v} \neq \mathbf{0}$, there exists an $\epsilon > 0$ such that the line segment $(\mathbf{x} - \epsilon\mathbf{v}, \mathbf{x} + \epsilon\mathbf{v})$ is in D . Hence the function $\varphi(t, \mathbf{v})$ in relation (8) in Section 1 is defined on $(-\epsilon, \epsilon)$. Also, since f is of class $C^{2,1}$ on D , $\partial^2 f(\mathbf{x})/\partial x_i \partial x_j = \partial^2 f(\mathbf{x})/\partial x_j \partial x_i$ at each \mathbf{x} in D , and so $H(\mathbf{x})$ is symmetric.

If f is convex, then by Lemma 1.5, the function $\varphi(t, \mathbf{v})$ is convex on the interval $(-\epsilon, \epsilon)$. By Lemma 5.1, $\varphi''(t, \mathbf{v}) \geq 0$. If $\mathbf{v} = (v_1, \dots, v_n)$, then straightforward calculations as in relations (12) and (13) in Section 11 of Chapter I give

$$\begin{aligned} \varphi''(t, \mathbf{v}) &= \frac{d}{dt} f(\mathbf{x} + t\mathbf{v}) = \sum_{i=1}^n \frac{\partial f}{\partial x_i}(\mathbf{x} + t\mathbf{v})v_i \\ \varphi''(t, \mathbf{v}) &= \sum_{i=1}^n \sum_{j=1}^n \frac{\partial^2 f}{\partial x_i \partial x_j} f(\mathbf{x} + t\mathbf{v})v_j = (\mathbf{v}, H(\mathbf{x} + t\mathbf{v})\mathbf{v}). \end{aligned} \quad (9)$$

Setting $v = 0$ gives

$$\langle x, H(x)x \rangle = \varphi''(x)\langle x, x \rangle \geq 0. \quad (9)$$

Since (9) holds for all x , $H(x)$ is positive semidefinite.

Conversely, suppose now that $H(x)$ is positive semidefinite for each x in D . Then, from (8), for each x in D and for each v in \mathbb{R}^n , $\varphi''(x)\langle x, v \rangle \geq 0$ for all t in $(-\epsilon, \epsilon)$. Hence by Lemma 3.1, φ' is convex on D . Since x and v are arbitrary, it follows from Lemma 1.5 that f is convex on D . The argument also shows that if $H(x)$ is positive definite for all x in D , then f is strictly convex.

We assume that the reader is familiar with the following facts. Given a symmetric matrix A , there exists an orthogonal matrix P such that $PAF = D$, where D is a diagonal matrix whose diagonal elements are the eigenvalues of A . Thus A is positive definite if and only if all the diagonal entries of D (eigenvalues of A) are positive and A is positive semidefinite if and only if they are nonnegative. For a survey of numerical methods for obtaining the eigenvalues of a symmetric matrix see Golub and Van Loan [1996].

Let A be an $n \times n$ matrix with entries a_{ij} . Let

$$A_k = a_{kk}, \quad A_k = \det \begin{pmatrix} a_{11} & \cdots & a_{1k} \\ \vdots & \ddots & \vdots \\ a_{k1} & \cdots & a_{kk} \end{pmatrix}, \quad k = 1, 2, 3, \dots, n.$$

The determinants A_k are called the principal minors of A . Another criterion for positive definiteness of a symmetric matrix A with entries a_{ij} is the following.

Lemma 3.2. The symmetric matrix A is positive definite if and only if $A_k > 0$ for $k = 1, \dots, n$ and is negative definite if and only if $(-1)^k A_k > 0$ for $k = 1, \dots, n$.

We shall prove this lemma in Exercise IV.2.16.

If we assume that $A_k > 0$ for all $k = 1, 2, \dots, n$ and not all the $a_{ij} = 0$, then there is difference between $n \geq 3$ and $n = 2$.

For $n \geq 3$, if $A_k > 0$, $k = 1, \dots, n$, and not all $a_{ij} = 0$, then A need not be positive semidefinite. To see this, consider the matrix

$$A = \begin{pmatrix} 1 & 0 & \cdots & 0 \\ 0 & 0 & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & -1 \end{pmatrix}.$$

Then $\langle x, Ax \rangle = x_1^2 - x_n^2$, which is not nonnegative for all x .

For $n = 2$, however, if $A_{11} > 0$, $A_{22} > 0$, and not both A_{11} and A_{22} are zero, then the symmetric matrix A is positive semidefinite. To see this, let

$$A = \begin{pmatrix} A_{11} & a \\ a & A_{22} \end{pmatrix}.$$

If $A_{11} = 0$, then $A_{22} = -a^2$. Then to have $A_{22} > 0$, we must have $a = 0$. But then $A_{22} = 0$. Hence we must have $A_{11} > 0$. Thus $A_{11} > 0$. Since $A_{11} > 0$,

$$A_{11}A_{22} \geq a^2. \quad (18)$$

Since $A_{11} > 0$, it follows that $A_{22} \geq 0$. For $\mathbf{x} = (x_1, x_2)$

$$\langle \mathbf{x}, A\mathbf{x} \rangle = A_{11}x_1^2 + 2ax_1x_2 + A_{22}x_2^2.$$

If $\mathbf{x} = (x_1, 0)$ with $x_1 \neq 0$, we get

$$\langle \mathbf{x}, A\mathbf{x} \rangle = A_{11}x_1^2 > 0.$$

If $\mathbf{x} = (x_1, x_2)$ with $x_2 \neq 0$, we have $x_1 = cx_2$ for some $-\infty < c < \infty$ and

$$\langle \mathbf{x}, A\mathbf{x} \rangle = [A_{11}c^2 + 2ac + A_{22}]x_2^2.$$

Denote the term in square brackets by $q(c)$. From (18) we get that the discriminant of q , which equals $4a^2 - 4A_{11}A_{22}$, is nonpositive. Thus the quadratic q either has a double root or has no real roots. Since $A_{11} > 0$, it follows that $q(c) \geq 0$ for all c . Hence, $\langle \mathbf{x}, A\mathbf{x} \rangle \geq 0$ for all \mathbf{x} .

Exercise 82. Determine whether or not the following functions are convex on \mathbb{R}^2 :

(i) $f(x, y, z) = x^2 + 2xy + 4xz + 3y^2 + yz + 7z^2$,

(ii) $f(x, y, z) = x^2 + 4xy + 4y^2 + 2xz + 4yz$.

Exercise 83. Determine for convex set in \mathbb{R}^2 on which the function

$$f(x, y) = x^2 - 2xy + \frac{1}{2}y^2 - 2y$$

is (i) convex and (ii) strictly convex.

Exercise 84. For what values of r is the function $f(x) = x^r - r \ln x$ convex on $(0, \infty)$?

Exercise 85. Show that $f(\mathbf{x}) = x_1^2 + \cdots + x_n^2$, $n \geq 1$, is convex on $\{\mathbf{x} : \mathbf{x} \geq \mathbf{0}\}$.

Exercise 3.4. Show that the n -dimensional ellipsoid

$$\left\{ \mathbf{x} : \sum_{i=1}^n \left(\frac{x_i}{a_i} \right)^2 \leq 1 \right\}$$

is convex. *Hint:* There is a very easy solution.

Exercise 3.6. For what values of p is $f(\mathbf{x}) = \|\mathbf{x}\|^p$ convex on \mathbb{R}^n . *Hint:* Use Lemma 3.4.

Exercise 3.7. Show that if $c_i > 0$, $i = 1, \dots, k$, the function

$$f(x) = \exp(c_1 x_1^2) + c_2 x_1^2 + c_3 x_2^2 + c_4 x_3^2$$

is convex on $S = \{x_1, x_2\} : x_1 > 0, x_2 > 0\}$.

Exercise 3.8. Show that the function $f(\mathbf{x}) = 1 + \|\mathbf{x}\|^{2/p}$, $p \geq 1$, is convex on \mathbb{R}^n .

Exercise 3.9. Show that the function $f(x) = -(x_1 + x_2 + \dots + x_n)^{1/p}$, $n \geq 1$, is convex on $S = \{x : x \geq \mathbf{0}\}$.

Exercise 3.10. A function f that is positive on a convex set C in \mathbb{R}^n is said to be logarithmically convex if the function $g(\mathbf{x}) = \ln f(\mathbf{x})$ is convex on C .

(a) Show that if f is logarithmically convex, then f is convex.

(b) Is it true that every positive convex function is logarithmically convex?

Exercise 3.11. Let f be differentiable on an open convex set X . Show that if f is strictly convex, then $\nabla f(\mathbf{x}) \neq \nabla f(\mathbf{y})$ for all distinct pairs of points \mathbf{x}, \mathbf{y} in X .

Exercise 3.12. Let C be an open convex set and let f be real valued and differentiable on C . Prove that f is convex on C if and only if for each pair of points \mathbf{x}_1 and \mathbf{x}_2 in C

$$\langle \nabla f(\mathbf{x}_1) - \nabla f(\mathbf{x}_2), \mathbf{x}_1 - \mathbf{x}_2 \rangle \geq 0.$$

Show that f is strictly convex if and only if strict inequality holds.

4. ALTERNATIVE THEOREMS FOR CONVEX FUNCTIONS

The theorems of this section are alternative theorems for systems of inequalities involving convex and affine functions. They find application in optimization problems involving such functions.

Definition 4.1. A vector-valued function $\mathbf{f} = (f_1, \dots, f_m)$ with domain a convex set C in \mathbb{R}^n and range in \mathbb{R}^m is said to be convex if each of the real-valued component functions f_i , $i = 1, \dots, m$, is convex.

Affine functions or affine transformations were defined in Definition B.3.3. It was shown there that an affine function \mathbf{h} with domain \mathbb{R}^n and range \mathbb{R}^m has the form

$$\mathbf{h}(x) = Ax + \mathbf{b},$$

where A is an $m \times n$ matrix and \mathbf{b} is a vector in \mathbb{R}^m .

Note that $\mathbf{b} = (b_1, \dots, b_m)$, where $b_i = (\mathbf{a}_i, \mathbf{0}) + b_i$, and \mathbf{a}_i is the i th row of A .

Theorem 4.1. Let C be a convex set in \mathbb{R}^n and let $\mathbf{f} = (f_1, \dots, f_m)$ be defined and convex on C . Let $\mathbf{h} = (h_1, \dots, h_m)$ be an affine mapping from \mathbb{R}^n to \mathbb{R}^m . If

$$\mathbf{h}(x) < \mathbf{0}, \quad \mathbf{h}(x) = \mathbf{0} \quad (1)$$

for no x in C , then there exist a vector \mathbf{p} in \mathbb{R}^m and a vector \mathbf{q} in \mathbb{R}^n such that $\mathbf{p} \geq \mathbf{0}$, $(\mathbf{p}, \mathbf{q}) \neq \mathbf{0}$, and

$$\langle \mathbf{p}, \mathbf{f}(x) \rangle + \langle \mathbf{q}, \mathbf{h}(x) \rangle \geq 0 \quad (2)$$

for all x in C .

Proof. For each x in C define the set

$$S(x) = \{(\mathbf{p}, \mathbf{q}) \mid \mathbf{p} \in \mathbb{R}^m, \mathbf{q} \in \mathbb{R}^n, \mathbf{p} = \mathbf{f}(x), \mathbf{q} = \mathbf{h}(x)\}.$$

For each x in C , the set $S(x)$ is the set of solutions (\mathbf{p}, \mathbf{q}) in \mathbb{R}^{m+n} of the system

$$\mathbf{p} = \mathbf{f}(x), \quad \mathbf{q} = \mathbf{h}(x). \quad (3)$$

Thus for each x , the origin in \mathbb{R}^{m+n} is not in $S(x)$. Let

$$S = \bigcup_{x \in C} S(x).$$

Since the origin in \mathbb{R}^{m+n} belongs to none of the sets $S(x)$, it does not belong to S . It is readily verified that S is convex. Hence by Theorem B.3.2 there is a hyperplane that separates S and $\mathbf{0}$ in \mathbb{R}^{m+n} . That is, there exist vectors \mathbf{p} in \mathbb{R}^m and \mathbf{q} in \mathbb{R}^n such that $(\mathbf{p}, \mathbf{q}) \neq \mathbf{0}$ and

$$\langle \mathbf{p}, \mathbf{f}(x) \rangle + \langle \mathbf{q}, \mathbf{h}(x) \rangle \geq 0 \quad (4)$$

for all (\mathbf{p}, \mathbf{q}) in S .

If (p, q) belongs to \mathcal{K} , then $p > f(x)$ for some x in C . Hence if some component of q , say q_i , were negative, we could take x_i to be arbitrarily large and obtain a contradiction to (4). Hence $q \geq 0$.

Now let $\varepsilon > 0$ and let ε denote the vector in \mathbb{R}^m all of whose components are equal to ε . Then for $x \in C$ the vector $f(x) + \varepsilon \langle q, f(x) \rangle$ belongs to \mathcal{K} , and so

$$\langle q, f(x) \rangle + \langle q, \varepsilon f(x) \rangle \geq -\varepsilon \langle q, \varepsilon \rangle.$$

We obtain (2) by letting $\varepsilon \rightarrow 0$.

Remark 4.1. From the proof it is clear that if the affine function h is absent, then the conclusion is that there exists a $p \neq 0$, $p \geq 0$ such that $\langle q, h(x) \rangle \geq 0$ for all x in C .

Remark 4.2. This theorem is not a true alternative theorem, as it does not assert that either the inequality (1) has a solution or the inequality (2) has a solution, but *neither* holds. If we strengthen the hypotheses of the theorem by assuming $C = \mathbb{R}^n$ and the rows of A to be linearly independent, then we do get a true alternative theorem. This is taken up in Exercise 4.2.

The next two theorems are corollaries of Theorem 4.1.

Theorem 4.2. Let f be a vector-valued convex function defined on a convex set C in \mathbb{R}^n and with range in \mathbb{R}^m . Then either

$$I: f(x) = 0$$

has a solution x in C or there exists a $p \neq 0$, $p \geq 0$, such that

$$II: \langle q, f(x) \rangle \geq 0$$

for all x in C , but never both.

Proof. Suppose I has a solution x_0 . Then for all $p \neq 0$, $p \geq 0$, we have $\langle p, f(x_0) \rangle \leq 0$, so II cannot hold.

Suppose now that I has no solution. Then by Theorem 4.1 and Remark 4.1 following the theorem, there exists a $p \neq 0$, $p \geq 0$ such that II holds for all x .

Theorem 4.2 is a true generalization of Gordan's theorem, Theorem B.5.2. To see that this is so, let $C = \mathbb{R}^n$ and let $f(x) = Ax$. Then $Ax = 0$ has no solution x_0 if and only if there exists a vector p in \mathbb{R}^m , $p \neq 0$, $p \geq 0$, such that $\langle p, Ax \rangle \geq 0$ for all x . If we take $x = A^*p$, we get $\langle A^*p, A^*p \rangle \geq 0$. If we take $x = -A^*p$, we get $\langle A^*p, A^*p \rangle \leq 0$. Hence $A^*p = 0$. In other words we have shown that $Ax = 0$ has no solution if and only if there exists a $p \neq 0$, $p \geq 0$ such that $A^*p = 0$, which is Gordan's theorem.

THEOREM 4.1. Let C be a compact convex set in \mathbb{R}^n and let $\{f_{\alpha}\}_{\alpha \in A}$ be a possibly infinite family of real-valued convex functions that are continuous on C . If the system

$$f_{\alpha}(x) \leq 0, \quad \alpha \in A, \quad (5)$$

has no solution in C , then there exists a finite subfamily $f_{\alpha_1}, \dots, f_{\alpha_m}$ and a vector $\mathbf{p} = (p_1, \dots, p_m)$ with $\mathbf{p} \neq \mathbf{0}$, $\mathbf{p} \geq \mathbf{0}$ and such that the real-valued function F defined by

$$F(x) = \sum_{i=1}^m p_i f_{\alpha_i}(x) \quad (6)$$

satisfies

$$\min_{x \in C} F(x) > 0. \quad (7)$$

The proof will be carried out by translating the statement of the nonexistence of solutions to the system (5) to a statement about the empty intersection of a collection of closed subsets of a compact set. The finite-intersection property will then be used to obtain a statement about the empty intersection of a finite collection of sets. The last statement will then be translated to a statement about the nonexistence of a solution to the system (1). An application of Theorem 4.1 will then yield (7).

Proof. Since for each $\alpha \in A$ the function f_{α} is continuous, it follows that for each $\alpha \in A$ and $\epsilon > 0$ the set

$$C_{\alpha, \epsilon} = \{x \in C, f_{\alpha}(x) \leq \epsilon\}$$

is closed and is contained in C . If the intersection of the closed sets $\{C_{\alpha, \epsilon}\}$, where $\alpha \in A$ and $\epsilon > 0$, were not empty, there would exist an x_0 in C such that $f_{\alpha}(x_0) \leq \epsilon$ for all $\alpha \in A$ and all $\epsilon > 0$. Since $f_{\alpha}(x_0) \leq 0$ for all $\alpha \in A$, contradicting the hypothesis that $\{f_{\alpha}(x) \leq 0\}$ has no solution in C . Thus the intersection of the sets $C_{\alpha, \epsilon}$ is empty. Since C is compact, by Remark 1.1 and Theorem 1.1 there exists a finite subcollection $C_{\alpha_1, \epsilon_1}, C_{\alpha_2, \epsilon_2}, \dots, C_{\alpha_m, \epsilon_m}$ of $\{C_{\alpha, \epsilon}\}$ whose intersection is empty. That is, the system $f_{\alpha_j}(x) = \epsilon_j < 0$, $j = 1, \dots, m$, has no solution x in C . Therefore the system

$$f_{\alpha_j}(x) - \epsilon_j < 0, \quad j = 1, \dots, m,$$

has no solution x in C . Therefore, by Theorem 4.1 and Remark 1, there exists a vector $\mathbf{p} = (p_1, \dots, p_m)$ with $\mathbf{p} \neq \mathbf{0}$ and $p_i \geq 0$, $i = 1, \dots, m$, such that

$$\sum_{i=1}^m p_i f_{\alpha_i}(x) \geq \sum_{i=1}^m p_i \epsilon_i > 0 \quad (8)$$

for all x in C . The function J defined in (9) is continuous on the compact set C and hence attains a minimum on C . From this and from (8) the conclusion follows.

Remark 4.1. The reader who is familiar with the notion of semicontinuity will see that the requirement that the functions f_i be continuous can be replaced by the less restrictive requirement that they be lower semicontinuous.

Exercise 4.1. Verify that the set S in Theorem 4.1 is convex.

Exercise 4.2. Prove the following corollary to Theorem 4.1. Let $f = (f_1, \dots, f_n)$ be defined and convex on \mathbb{R}^n . Let $h = Ax - b$ be an affine mapping from \mathbb{R}^n to \mathbb{R}^m and let the rows of A be linearly independent. Then either

$$\text{I: } f(x) < 0, \quad Ax - b = 0$$

has a solution x in \mathbb{R}^n or

$$\text{II: } \exists p, (p, q) \neq 0, p \in \mathbb{R}^m, p \geq 0, q \in \mathbb{R}^n,$$

but a solution $(p, q) \neq 0, p \in \mathbb{R}^m, p \geq 0, q \in \mathbb{R}^n$, but never both.

5. APPLICATION TO GAME THEORY

In this section we use Theorem 4.1 to prove the von Neumann minimum theorem and then apply the minimum theorem to prove the existence of value and saddle points in mixed strategies for matrix games. Our proof will require the concept of uniform continuity and a sufficient condition for uniform continuity.

Definition 5.1. A function f defined on a set S in \mathbb{R}^n with range in \mathbb{R}^m is said to be uniformly continuous on S if for each $\epsilon > 0$ there exists a $\delta > 0$ such that if x and y are points in S satisfying $\|x - y\| < \delta$, then $\|f(x) - f(y)\| < \epsilon$.

Theorem 5.1. Let S be a compact set in \mathbb{R}^n and let f be a mapping defined and continuous on S with range in \mathbb{R}^m . Then f is uniformly continuous on S .

For a proof we refer the reader to Bartle and Sherbert [1980] or Rudin [1976].

Theorem 5.2 (Von Neumann's Minimum Theorem). Let $X = \mathbb{R}^n$ and $Y = \mathbb{R}^m$ be compact convex sets. Let $f: X \times Y \rightarrow \mathbb{R}$ be continuous on $X \times Y$. For each x in X let $f(x, \cdot): Y \rightarrow \mathbb{R}$ be convex and for each y in Y let $f(\cdot, y): X \rightarrow \mathbb{R}$ be concave.

Then

$$\max_{x \in X} \min_{y \in Y} f(x, y) = \min_{y \in Y} \max_{x \in X} f(x, y), \quad (1)$$

and there exists a pair (x_0, y_0) with $x_0 \in X$ and $y_0 \in Y$ such that for all x in X and all y in Y

$$f(x, y_0) \leq f(x_0, y_0) \leq f(x_0, y). \quad (2)$$

Moreover, if v denotes the common value in (1), then $f(x_0, y_0) = v$.

A pair (x_0, y_0) as in (2) is called a *saddle point* of the function f .

The relation (1) is not true in general for continuous functions f defined in $X \times Y$ where X and Y are compact and convex. To see this, let $X = [0, 1]$, let $Y = [0, 1]$, and let $f(x, y) = (x - y)^2$. Then for fixed x , $\min\{f(x, y) : y \in Y\}$ is achieved at $y = x$ and equals zero. Hence in this example the left-hand side of (1) equals zero. On the other hand, for fixed y in the interval $[0, \frac{1}{2}]$

$$\max_{x \in X} \{f(x, y) : x \in X\} = (1 - y)^2,$$

and for fixed y in the interval $(\frac{1}{2}, 1]$

$$\max_{x \in X} \{f(x, y) : x \in X\} = y^2.$$

Hence in this example the right-hand side of (1) equals $\frac{1}{4}$.

The convexity in X and concavity in Y is a sufficient condition but not a necessary condition, as the example $X = [0, 1]$, $Y = [0, 1]$, and $f(x, y) = e^x e^y$ shows.

Before we take up the proof of Theorem 5.2, which will be carried out in several steps, we observe the following: If A and B are arbitrary sets and g is a real-valued bounded function defined on $A \times B$, then the quantities

$$\sup_{x \in A} \inf_{y \in B} g(x, y) \quad \text{and} \quad \inf_{y \in B} \sup_{x \in A} g(x, y)$$

are well defined, while the quantities

$$\max_{x \in A} \min_{y \in B} g(x, y) \quad \text{and} \quad \min_{y \in B} \max_{x \in A} g(x, y)$$

need not be. For a simple example, take $A = [0, 1]$, $B = [0, 1]$, and $g(x, y) = xy$. Thus, implicit in (1) is the statement, which must be proved, that the left- and right-hand sides of (1) exist. The proof of this statement will constitute the second step of our proof. The first step is the general observation given in Lemma 5.1.

LEMMA 5.1. Let A and B be arbitrary sets and let $g: A \times B \rightarrow \mathbb{R}$ be bounded. Then

$$\sup_{a \in A} \inf_{b \in B} g(a, b) \leq \inf_{b \in B} \sup_{a \in A} g(a, b).$$

Proof. Choose an element $b_0 \in B$. Then for each $a \in A$,

$$g(a, b_0) \leq \sup_{b \in B} g(a, b)$$

Since the choice of b_0 was arbitrary, this relationship holds for each a in A and each b in B . Hence, for each $a \in A$

$$\inf_{b \in B} g(a, b) \leq \inf_{b \in B} \sup_{a \in A} g(a, b).$$

The right-hand side is a fixed real number, so

$$\sup_{a \in A} \inf_{b \in B} g(a, b) \leq \inf_{b \in B} \sup_{a \in A} g(a, b).$$

We now proceed to the second step of the proof and show that the left- and right-hand sides of (1) exist.

Fix fixed x in X , the function $f(x, \cdot): Y \rightarrow \mathbb{R}$ defined by $f(x, y)$ is continuous on the compact set Y . Hence

$$\phi(x) = \min_{y \in Y} f(x, y)$$

is defined for each x in X . We assert that the function ϕ is continuous on X . To prove the assertion, let $\epsilon > 0$ be given. Since f is continuous on the compact set $X \times Y$, it is uniformly continuous on $X \times Y$. Hence there exists a $\delta > 0$ such that if $\|x_1 - x_2\| = \delta$, then

$$|f(x_1, y) - f(x_2, y)| < \frac{1}{2}\epsilon$$

for all y in Y . Upon taking the minimum over Y we get that

$$\phi(x_1) - \frac{1}{2}\epsilon \leq \phi(x_2) \leq \phi(x_1) + \frac{1}{2}\epsilon,$$

or $|\phi(x_1) - \phi(x_2)| < \epsilon$ whenever $\|x_1 - x_2\| = \delta$. This proves the continuity of ϕ .

Since ϕ is continuous on the compact set X , there exists and x_0 in X such that

$$\phi(x_0) = \max_{x \in X} \phi(x) = \max_{x \in X} \left[\min_{y \in Y} f(x, y) \right]. \quad (3)$$

It follows from (3) that the left-hand side of (1) is well defined.

From the definition of ϕ we have

$$\phi(x_0) = \min_{y \in Y} f(x_0, y). \quad (4)$$

Similarly, for each y in Y we can set

$$\psi(y) = \max_{x \in X} f(x, y)$$

and then obtain the existence of a y_0 in Y such that

$$\psi(y_0) = \min_{y \in Y} \psi(y) = \min_{y \in Y} \left[\max_{x \in X} f(x, y) \right] \quad (5)$$

and

$$\psi(y_0) = \max_{x \in X} f(x, y_0). \quad (6)$$

Thus the right-hand side of (1) is well defined.

The third step in the proof is to establish the equality in (1). Let

$$\alpha = \max_{x \in X} \min_{y \in Y} f(x, y) \quad \beta = \min_{y \in Y} \max_{x \in X} f(x, y).$$

By Lemma 5.1 with $X = A$, $Y = B$, $f = g$ and by the second step, we have that $\alpha \leq \beta$. Hence to prove (1), we must show that $\alpha \geq \beta$.

Let $\epsilon > 0$ be arbitrary. We assert that there does not exist a y_0 in Y such that for all x in X

$$f(x, y_0) \leq \beta - \epsilon.$$

For if such a y_0 did exist, we would have

$$\beta = \min_{y \in Y} \max_{x \in X} f(x, y) \leq \max_{x \in X} f(x, y_0) \leq \beta - \epsilon$$

which is not possible.

Let $\{g_\alpha\}_{\alpha \in K}$ be the family of functions defined on Y by the formula

$$g_\alpha(y) = f(x_\alpha, y), \quad x \in K.$$

Then for each $x \in K$, the function g_x is continuous and convex on Y . Moreover, for arbitrary $\epsilon > 0$ the system

$$g_x(y) - \beta + \epsilon < 0, \quad x \in K, \quad (7)$$

has no solution y in Y .

Since (7) has no solution, it follows from Theorem 4.3 that there exist points x_1, \dots, x_n in X , a vector $p \in E^n$, $p \geq 0$, $p \neq 0$, and a $\delta > 0$ such that

$$G(y) = \sum_{i=1}^n p_i [f(x_i, y) - \beta + \delta] \geq \delta > 0 \quad (8)$$

for each y in Y . Since $\sum_{i=1}^n p_i > 0$, we may divide through by this quantity in (8) and assume that $p \in E^n$. Therefore, since X is convex, $\sum_{i=1}^n p_i x_i$ is in X . Using (8) and the concavity of $f(\cdot, y)$ for each fixed y , we may write

$$f\left(\sum_{i=1}^n p_i x_i, y\right) \geq \sum_{i=1}^n p_i f(x_i, y) \geq \beta + \delta$$

for each y in Y . Therefore,

$$\begin{aligned} \alpha &= \max_{x \in X} \min_{y \in Y} f(x, y) \geq \min_{y \in Y} f\left(\sum_{i=1}^n p_i x_i, y\right) \\ &\geq \min_{y \in Y} \sum_{i=1}^n p_i f(x_i, y) = \beta + \delta. \end{aligned}$$

Since $\delta > 0$ is arbitrary, we get that $\alpha \geq \beta$, and (1) is established.

The final step in the proof is to show that $\alpha = f(x_\alpha, y_\alpha)$ and that (2) holds. It follows from (1), (3), and (5) that

$$f(x_\alpha) - \beta(x_\alpha) = \alpha.$$

It now follows from (4) and (5) that

$$f(x_\alpha, y_\alpha) \leq \alpha \leq f(x_\alpha, y) \quad (9)$$

for all x in X and y in Y . In particular, (9) holds for $x = x_\alpha$ and $y = y_\alpha$, so that $\alpha = f(x_\alpha, y_\alpha)$. The relation (2) now follows from (9).

The simplest two-person zero-sum games are matrix games. These games are played by two players that we shall call Blue and Red. Blue selects a positive integer i from a finite set I of positive integers $\{1, \dots, m\}$. Red, simultaneously and in ignorance of Blue's choice, selects a positive integer j from a finite set J of positive integers $\{1, \dots, n\}$. The "payoff" or "reward" to Blue in the event that he chooses i and Red chooses j is a predetermined real number a_{ij} . The payoff to Red is $-a_{ij}$. Hence the name zero-sum game. A strategy for Blue is a rule which tells Blue which positive integer i in I to select. A strategy for Red is a rule which tells Red which positive integer j in J to select. Blue's objective is to choose i in I so as to maximize a_{ij} . Red's objective is to choose j in J so as to minimize a_{ij} . That is, to maximize $-a_{ij}$. If the strategy for Blue is to choose the number i_0 in I , then we call the strategy a pure strategy for Blue. Similarly, if the strategy for Red is to choose the number j_0 in J , then we call the strategy a pure strategy for Red. Later we shall consider strategies that are not pure.

Another way of describing a matrix game is as follows. An $m \times n$ matrix A with entries a_{ij} , $i = 1, \dots, m$, $j = 1, \dots, n$, is given. Blue selects a row i and Red selects a column j . The choices are made simultaneously and without knowledge of the opponent's choice. If Blue selects row i and Red selects column j , then the payoff to Blue is a_{ij} and the payoff to Red is $-a_{ij}$. Blue's objective is to maximize the payoff, while Red's objective is to minimize the payoff.

At this point, the formulation of a matrix game may strike the reader as rather artificial and uninteresting. It turns out, however, that many decision problems in conflict situations can be modeled as matrix games. As an example, we shall formulate a tactical naval problem as a matrix game. The reader will find many additional interesting examples elsewhere [Dresher, 1951]. The naval game appears there as a land conflict game called Colonel Blotto.

Two Red cargo ships are protected by three escort frigates under the command of Captain Plover. The entire convoy is subject to attack by a pack of four Blue submarines under the command of Captain Nemo. The ships carry a hazardous cargo and thus maintain a separation such that one submarine cannot attack both ships. Captain Nemo's problem is to allocate his submarines to attack the two ships. Thus, he may allocate all four of his submarines to attack one of the ships and none to attack the second ship. Or, he may allocate three submarines to attack one ship and one submarine to attack the second ship. Or, he may allocate two submarines to attack each of the ships. We may summarize the choices, or strategies, available to Captain Nemo by five ordered pairs, (4,0), (0,4), (3,1), (1,3), (2,2), where (x, y) indicates that he has allocated x submarines to attack ship 1 and y submarines to attack ship 2. Captain Plover's problem is to allocate his frigates to defend the ships. He has four choices, or strategies, each of which can be denoted by an ordered pair (p, q), $0 \leq p \leq 3$, $0 \leq q \leq 3$, $p + q = 3$, where p frigates are allocated to defend ship 1 and q frigates are allocated to defend ship 2. We assume that conditions

one such that when Captain Nemo allocates his submarines, he cannot determine the disposition of the escort frigates, and that when Captain Plato allocates his escort frigates, he cannot determine the disposition of the attacking submarines.

If at a given ship the number of attacking submarines exceeds the number of defending frigates, then the payoff to Blue at that ship equals the number of Red frigates at that ship plus 1. The "plus 1" is awarded for the cargo ship that is assumed to be sunk. If at a given ship the number of attacking submarines is less than the number of defending frigates, then the payoff to Blue at the ship is

$$-(\text{Number of attacking submarines} + 1).$$

The "plus 1" is deducted for allowing the cargo ship to proceed to its destination. If at a given ship the number of attacking submarines equals the number of defending ships, then the encounter is a draw, and the payoff to Blue at that ship is zero. The total payoff to Blue resulting from given choices by the two antagonists is the sum of the payoffs at each ship. The payoff to Red is the negative of the payoff to Blue.

The game just described can be cast as a matrix game with payoff matrix

		Captain Plato's strategies				
		(3, 0)	(3, 2)	(2, 1)	(1, 2)	
Captain Nemo's strategies	(4, 0)	4	0	2	1	
	(5, 4)	0	4	1	2	
	(3, 1)	1	-1	3	0	
	(1, 2)	-1	1	4	3	
		(2, 2)	-2	-2	2	2

a strategy for Blue (Captain Nemo) is a choice of row; a strategy for Red (Captain Plato) is a choice of column.

If Captain Nemo chooses row i , he is guaranteed a payoff of at least

$$\min\{a_{ij}\}; j = 1, \dots, 4.$$

If he picks the row that will give the payoff

$$v'' = \max_i \min_j a_{ij},$$

then he will be guaranteed to obtain v'' . Captain Nemo can thus pick either row 1 or row 2 and be guaranteed a payoff of at least zero. Thus $v'' = 0$. On the other hand, if Captain Plato chooses column j , he will lose at most

$\max\{a_i, j\} (i = 1, \dots, 5)$. If Captain Photo picks the columns that give Blue the payoff

$$v^* = \min_i \max_j a(i, j),$$

then he is guaranteed to lose no more than v^* . Captain Photo can thus pick either column 3 or 4 and be guaranteed to lose no more than three. Thus $v^* = 3$.

Note that if Captain Nemo knows beforehand that Captain Photo will pick column 3, then he can pick row 3 and obtain a payoff of 3, which is greater than $v^* = 3$. On the other hand, if Captain Photo knows beforehand that Captain Nemo will pick row 2, then he can pick column 1 and lose zero units, which is less of a loss than three units ($v^* = 3$). An examination of the payoff matrix shows that if one of the commander's choices is known in advance by the second commander, then the second commander can usually improve his payoff if he bases his choice on this information. Thus, in the submarine-convoys game victory of choice is essential. The open question, however, in the submarine-convoys game is, "What is the best strategy for each player?"

In contrast to the submarine-convoys game, consider the game with payoff matrix

$$\begin{pmatrix} 3 & 7 & 4 & 3 \\ 12 & 10 & 7 & 8 \\ 10 & 4 & 1 & 8 \\ 3 & 5 & 6 & 3 \end{pmatrix}.$$

Here

$$\max_i \min_j a(i, j) = \min_j \max_i a(i, j) = 7.$$

The strategies $i = i_0$, where $i_0 = 2$, and $j = j_0$, where $j_0 = 3$, have the properties that $a(i_0, j_0) = 7$ and

$$a(i, j_0) \leq a(i_0, j_0) \leq a(i_0, j)$$

for all $i = 1, \dots, 4$ and $j = 1, \dots, 4$. We say that the value of the game is 7, that $i_0 = 2$ is an optimal (pure) strategy for Blue, and that $j_0 = 3$ is an optimal (pure) strategy for Red. If Blue chooses his optimal strategy $i_0 = 2$, then the best that Red can do is to also choose his optimal strategy $j_0 = 3$. Also, if Red chooses his optimal strategy $j_0 = 3$, then the best that Blue can do is to choose his optimal strategy $i_0 = 2$. Thus, if either player announces beforehand that he will choose an optimal strategy, the opponent cannot take advantage of this information, other than to choose an optimal strategy.

We now consider the general matrix game with $m \times n$ payoff matrix A with entries a_{ij} , $i = 1, \dots, m$, $j = 1, \dots, n$. We always have

$$v^* = \max_i \min_j a_{ij} \leq \min_j \max_i a_{ij} = v^*.$$

If the reverse inequality $v^* \geq v^*$ holds, then we set

$$v = \min_j \max_i a_{ij} = \max_i \min_j a_{ij}$$

and say that the matrix game has a value v in pure strategies. It is easy to show that if the matrix game has a value v in pure strategies, then there exists a pair (i_0, j_0) such that

$$v = a_{i_0, j_0}, \quad (10)$$

and for all $i = 1, \dots, m$, $j = 1, \dots, n$

$$a_{i, j_0} \leq a_{i_0, j_0} \leq a_{i_0, j}. \quad (11)$$

The strategies i_0 and j_0 are called *optimal (pure) strategies*. We point out that optimal strategies need not be unique. A pair of strategies (i_0, j_0) such that (11) holds is called a *saddle point in pure strategies*. It is also easy to show that if a matrix game has a saddle point in pure strategies, then the game has value v and (10) holds.

If the matrix game has a value, then each player should use an optimal strategy. For then, Blue, the maximizer, is guaranteed a payoff of v and Red, the minimizer, is guaranteed to lose at most $-v$. If one player, say Blue, uses an optimal strategy and Red does not, then the payoff to Blue, who plays optimally, will, in general, be more favorable than it would be had Red played optimally. This is the content of (12).

The question that still thrives itself upon us is, "How does one play if the game does not have a saddle point, as in the submarine-coercing game?" To answer this question, we introduce the notion of mixed strategy.

Definition 3.2. A mixed strategy for Blue in the game with $m \times n$ payoff matrix A is a probability distribution on the rows of A . Thus, a mixed strategy is a vector $x = (x_1, \dots, x_m)$ with $x_i \geq 0$, $i = 1, \dots, m$, and $\sum_{i=1}^m x_i = 1$.

A mixed strategy for Red is a probability distribution on the columns of A . Thus, a mixed strategy for Red is a vector $y = (y_1, \dots, y_n)$ with $y_j \geq 0$, $j = 1, \dots, n$, and $\sum_{j=1}^n y_j = 1$.

The set of pure strategies for a player is a subset of the set of mixed strategies. For example, the pure strategy for Blue in which row i is always

chosen is also the mixed strategy in which row i is chosen with probability 1, that is, the vector $a_i = (0, \dots, 0, 1, 0, \dots, 0)$, where the i th entry is 1 and all other entries are zero.

In actual play, Blue will select his row using a random device which produces the value i with probability a_i , $i = 1, \dots, m$. Blue chooses row i if the random device produces the value i . Red chooses column j in similar fashion. If mixed strategies are used, neither player tries to outguess his opponent and both tend to luck. Moreover, if Blue uses a mixed strategy, Red cannot determine Blue's choice beforehand and take advantage of this information. A similar statement holds if Red uses a mixed strategy.

If mixed strategies are allowed, then it is appropriate to define the payoff of the matrix game with mixed strategies as the expected value of the payoff. Thus, if we denote the payoff resulting from the choice of strategies (x, y) by $F(x, y)$, we have

$$F(x, y) = \sum_i x_i a_{ij} = (x, Ay). \quad (12)$$

We define the matrix game with matrix A in which mixed strategies are used as follows. Blue selects a vector x in F_m , and Red selects a vector y in F_n . The payoff to Blue is given by (12), and the payoff to Red is $-F(x, y)$. We say that the matrix game with mixed strategies has a value if

$$\max_{x \in F_m} \min_{y \in F_n} F(x, y) = \min_{y \in F_n} \max_{x \in F_m} F(x, y). \quad (13)$$

We denote the number in (13) by v and call it the value of the game. We say that the game has a saddle point (x_0, y_0) , that x_0 is an optimal mixed strategy for Blue, and that y_0 is an optimal mixed strategy for Red if

$$F(x, y_0) \leq F(x_0, y_0) \leq F(x_0, y) \quad (14)$$

for all x in F_m and y in F_n .

THEOREM 3.3. *A matrix game has a value and a saddle point in mixed strategies.*

This theorem is a corollary of Theorem 3.2. Take $S = F_m$, $T = F_n$, and $f = F$. The function F is continuous on $F_m \times F_n$. The sets F_m and F_n are compact and convex. Since F is bilinear, it is concave in x and convex in y . Theorem 3.2 follows by applying Theorem 3.2 to F .

Note that the saddle point condition (14) implies that

$$v = F(x_0, y_0).$$

Exercise 8.1. Verify that in the subgame-perfect game the value in mixed strategies is $\frac{1}{2}$ and that

$$X_1 = \left\{ \frac{1}{2} \begin{pmatrix} 1 \\ 1 \end{pmatrix}, 0, \frac{1}{2} \begin{pmatrix} 1 \\ 1 \end{pmatrix} \right\}, \quad X_2 = \left\{ \frac{1}{2} \begin{pmatrix} 1 \\ 1 \end{pmatrix}, \frac{1}{2} \begin{pmatrix} 1 \\ 1 \end{pmatrix} \right\}$$

is a saddle point.

Exercise 8.2. Show that if a matrix game has a value in pure strategies, then (i) and (iii) hold. Conversely, show that if a matrix game has a saddle point in pure strategies, then the game has a value v and (ii) holds.

Exercise 8.3. Let X_1 denote the set of optimal mixed strategies for Blue and let X_2 denote the set of optimal mixed strategies for Red. Show that X_1 and X_2 are compact convex sets.

IV

OPTIMIZATION PROBLEMS

1. INTRODUCTION

In Section 9 of Chapter I we stated the basic problem in optimization theory as follows: Given a set S (not necessarily in \mathbb{R}^n) and a real-valued function f defined on S , does there exist an element x_0 in S such that $f(x_0) \leq f(x)$ for all x in S , and if so find it. We also showed that the problem of maximizing f over S is equivalent to the problem of minimizing $-f$ over S . Thus, there is no loss of generality in only considering minimization problems, as we shall do henceforth.

If S is a set in \mathbb{R}^n and f is a real-valued function on S , we say that a point x_0 in S is a *local minimizer* or that f has a *local minimum* at x_0 if there exists a $\delta > 0$ such that $f(x_0) \leq f(x)$ for all x in $B(x_0, \delta) \cap S$. If the strict inequality $f(x_0) < f(x)$ holds, then x_0 is said to be a *strict local minimizer*. If S is open, then $B(x_0, \delta) \cap S$ is an open set contained in S . Hence there exists a $\delta' > 0$ such that for all x in $B(x_0, \delta')$, $f(x_0) \leq f(x)$. Therefore, if S is open, for x_0 to be a local minimizer, we need to find a $\delta > 0$ such that $f(x_0) \leq f(x)$ for all x in $B(x_0, \delta)$ and can omit the intersection with S . A point x_0 that is a minimizer is also a local minimizer. A point x_0 that is a local minimizer need not be a minimizer.

In this chapter we shall develop necessary conditions and sufficient conditions for a point x_0 to be a minimizer or local minimizer for various classes of problems. The determination of the set E of points that satisfy the necessary conditions does not, of course, solve the optimization problem. All we learn is that the set of solutions is to be found in this set. The determination of those points in E that are solutions to the problem requires further analysis. Sufficient conditions, if applicable to the problem at hand, are useful in this connection.

If E is an open set in \mathbb{R}^n , the optimization problem is often called an *unconstrained problem*.

A *programming problem* is a problem of the following type. A set X_0 in \mathbb{R}^n is given as are functions f and g with domains X_0 and ranges in \mathbb{R}^1 and \mathbb{R}^m

respectively. Let

$$X = \{x : x \in X_0, g(x) \leq 0\}.$$

Problem P

Minimize f over the set X . That is, find an x_0 in X such that $f(x_0) \leq f(x)$ for all x in X .

The set X is called the *feasible set*, and elements of X are called *feasible points* or *feasible vectors*.

The problem is sometimes stated as follows: Minimize $f(x)$ subject to $x \in X_0$ and $g(x) \leq 0$.

If $X_0 = \mathbb{R}^n$, A is an $m \times n$ matrix, and

$$f(x) = c - k \cdot x, \quad g(x) = \begin{pmatrix} Ax - r \\ -b \end{pmatrix}, \quad (1)$$

the problem is said to be a *linear programming problem*. We have written $-k$ rather than k because in the linear programming literature the problem is usually stated as follows: Minimize (k, x) subject to $Ax \leq r$ and $x \geq 0$.

If X_0 is convex and the functions f and g are convex, the problem is said to be a *convex programming problem*.

2. DIFFERENTIABLE UNCONSTRAINED PROBLEMS

We will now denote the set on which the function f is defined by X rather than S . We first assume that X is an open interval in \mathbb{R}^1 and that f is differentiable on X . We will derive necessary conditions and sufficient conditions for a point x_0 in X to be a local minimizer. Although the reader surely knows these conditions from elementary calculus, it is instructive to review them.

Let x_0 in X be a local minimizer of f and let f be differentiable at x_0 . Then there exists a $\delta > 0$ such that $f(x) \geq f(x_0)$ for x in $(x_0 - \delta, x_0 + \delta)$. Therefore, for $x_0 - \delta \leq x \leq x_0 + \delta$

$$\frac{f(x) - f(x_0)}{x - x_0} \geq 0, \quad (2)$$

and for $x_0 - \delta \leq x \leq x_0$

$$\frac{f(x) - f(x_0)}{x - x_0} \leq 0. \quad (3)$$

Since f is differentiable at x_0 ,

$$\lim_{x \rightarrow x_0} \frac{f(x) - f(x_0)}{x - x_0} = \lim_{x \rightarrow x_0} \frac{f(x) - f(x_0)}{x - x_0} = f'(x_0).$$

From this and from (7) we get $f'(x_0) \geq 0$, while from (2) we get $f'(x_0) \leq 0$. Hence $f'(x_0) = 0$.

Let us now assume further that f is of class $C^{(2)}$ on $(x_0 - \delta, x_0 + \delta)$. By Taylor's theorem, for each $x \neq x_0$ in $(x_0 - \delta, x_0 + \delta)$ there exists a point ξ that lies between x_0 and x such that

$$f(x) = f(x_0) + f'(x_0)(x - x_0) + \frac{1}{2}f''(\xi)(x - x_0)^2. \quad (8)$$

Since x_0 is a local minimum, $f(x_0) = 0$ and therefore

$$\frac{1}{2}f''(\xi)(x - x_0)^2 = f(x) - f(x_0) \geq 0.$$

Thus

$$f''(\xi) = 2 \frac{f(x) - f(x_0)}{(x - x_0)^2} \geq 0.$$

If we now let $x \rightarrow x_0$, then $\xi \rightarrow x_0$ and $f''(\xi) = f''(x_0)$. Hence $f''(x_0) \geq 0$.

We have proved the following theorem.

THEOREM 2.1. Let X be an open interval in \mathbb{R} , let f be a real-valued function defined on X , and let x_0 in X be a local minimum for f . If f is differentiable at x_0 , then $f'(x_0) = 0$. If f is of class $C^{(2)}$ on some interval about x_0 , then $f''(x_0) \geq 0$.

A point c in X such that $f'(c) = 0$ is called a *critical point* of f . We have shown that for a differentiable function a local minimum is a critical point.

If we replace the necessary condition $f'(x_0) = 0$ by the stronger condition $f''(x_0) > 0$, we obtain a sufficient condition for a critical point to be a local minimum.

THEOREM 2.2. Let X be an interval in \mathbb{R} , and let f be real-valued and of class $C^{(2)}$ on some interval about a point x_0 in X . If $f'(x_0) = 0$ and $f''(x_0) > 0$, then x_0 is a strict local minimum for f .

Proof. Since f' is continuous on some interval about x_0 and since $f''(x_0) > 0$, there exists a $\delta > 0$ such that $f'(x) > 0$ for all x in $(x_0 - \delta, x_0 + \delta)$. (See Exercise 1.5.4.) By Taylor's Theorem, for each $x \neq x_0$ in $(x_0 - \delta, x_0 + \delta)$ there exists an θ lying between x_0 and x such that (3) holds. Since $f'(x_0) = 0$, we get that for each $x \neq x_0$ in $(x_0 - \delta, x_0 + \delta)$

$$f(x) - f(x_0) = \frac{1}{2}f''(\theta)(x - x_0)^2 > 0. \quad (9)$$

Thus x_0 is a strict local minimum.

In elementary calculus texts Theorem 2.2 is often called the *second derivative test*.

We must have $f''(x_0) > 0$ rather than $f''(x_0) \neq 0$ for the conclusion of Theorem 2.2 to hold. To see this, take $X = \mathbb{R}$ and take $f(x) = x^4$. Then $f'(0) = 0$, $f''(0) = 0$, but zero is not a local minimum. If, however, we assume that $f''(x) \neq 0$ for all x in X , then all the statements in the proof of Theorem 2.2 hold for each x in X , except that (H) now reads

$$f(x) - f(x_0) = \frac{1}{2}f''(x_0)(x - x_0)^2 + o.$$

Thus x_0 is a minimum. If we assume that $f''(x) \geq 0$ on some interval about x_0 , then x_0 is a local minimum. Also, if we assume that $f(x) > 0$ on all of X , then x_0 is a strict minimum.

We summarize this discussion in the following corollary to Theorem 2.2.

COROLLARY 1. Let $f'(x_0) = 0$. If $f''(x) \geq 0$ on X , then x_0 is a minimum. If $f''(x) > 0$ on X , then x_0 is a strict minimum. If $f''(x) \geq 0$ on some interval about x_0 , then x_0 is a local minimum.

We now take X to be an open set in \mathbb{R}^n . The next two theorems are the n -dimensional generalizations of Theorems 2.1 and 2.2.

THEOREM 2.3. Let X be an open set in \mathbb{R}^n , let f be a real-valued function defined on X , and let x_0 be a local minimum. If f has first-order partial derivatives at x_0 , then

$$\frac{\partial f}{\partial x_i}(x_0) = 0, \quad i = 1, \dots, n. \quad (5)$$

If f is of class C^{2k} on some open ball $B(x_0, \delta)$ centered at x_0 , then

$$B(x_0) = \left(\frac{\partial^{2k} f}{\partial x_1^{2k} \partial x_2^{2k}}(x_0) \right)$$

is positive semidefinite.

Proof. Since X is open, there exists a $\delta > 0$ such that for each $i = 1, \dots, n$ the points

$$x_0 + \theta e_i = (x_{01}, \dots, x_{0i-1}, x_{0i} + \theta, x_{0i+1}, \dots, x_{0n})$$

are in X whenever $|\theta| < \delta$. Therefore for each $i = 1, \dots, n$ the function φ_i defined by

$$\varphi_i(\theta) = f(x_0 + \theta e_i), \quad |\theta| < \delta,$$

is well defined and has a local minimum at $t = 0$. Since the partial derivatives of f exist at \mathbf{x}_0 , the chain rule implies that for each $i = 1, \dots, n$ the derivative $\varphi_i'(0)$ exists and is given by

$$\varphi_i'(0) = \frac{\partial f}{\partial x_i}(\mathbf{x}_0).$$

Since each φ_i has a local minimum at $t = 0$, $\varphi_i'(0) = 0$. Thus (5) holds.

Now suppose that f is of class $C^{2,1}$ on some open ball $B(\mathbf{x}_0, \delta)$. Then for each unit vector \mathbf{u} in \mathbb{R}^n the function $\varphi(\cdot)$ defined by

$$\varphi(t) = f(\mathbf{x}_0 + t\mathbf{u}), \quad |t| < \delta,$$

is well defined, is of class $C^{2,1}$, and has a local minimum at $t = 0$. By Theorem 2.1, $\varphi'(0) = 0$ and $\varphi''(0)\mathbf{u} \geq 0$. By a straightforward calculation as in relation (11) in Section 1.1 of Chapter I, we get

$$\varphi''(0)\mathbf{u} = \langle \mathbf{u}, H(\mathbf{x}_0 + t\mathbf{u})\mathbf{u} \rangle.$$

Setting $t = 0$, we get that $\langle \mathbf{u}, H(\mathbf{x}_0)\mathbf{u} \rangle \geq 0$. Since \mathbf{u} is an arbitrary unit vector, we get that $H(\mathbf{x}_0)$ is positive semidefinite.

If \mathbb{R}^n points \mathbf{x} such that $\nabla f(\mathbf{x}) = \mathbf{0}$ are called *critical points* of f .

THEOREM 2.6. Let X be an open set in \mathbb{R}^n and let f be a real-valued function of class $C^{2,1}$ on some ball $B(\mathbf{x}_0, \delta)$. If $H(\mathbf{x}_0)$ is positive definite and $\nabla f(\mathbf{x}_0) = \mathbf{0}$, then \mathbf{x}_0 is a strict local minimizer. If X is convex, f is of class $C^{2,1}$ on X , and $H(\mathbf{x})$ is positive semidefinite for all \mathbf{x} in X , then \mathbf{x}_0 is a minimizer for f . If $H(\mathbf{x})$ is positive definite for all \mathbf{x} in X , then \mathbf{x}_0 is a strict minimizer.

Proof. We shall only prove the local statement and leave the proof of the last two to the reader.

If $H(\mathbf{x}_0)$ is positive definite at \mathbf{x}_0 , then it follows from the continuity of the second partial derivatives of f and Lemma III.1.7 that there exists a $\delta > 0$ such that $H(\mathbf{x})$ is positive definite for all \mathbf{x} in $B(\mathbf{x}_0, \delta)$. Let \mathbf{x} be any point in $B(\mathbf{x}_0, \delta)$ and let $\mathbf{v} = \mathbf{x} - \mathbf{x}_0$. Then the function $\varphi(\cdot)$ defined by

$$\varphi(t) = f(\mathbf{x}_0 + t\mathbf{v})$$

is well defined and is $C^{2,1}$ on some real interval $(-1 + \eta, 1 + \eta)$, where $\eta > 0$. Moreover

$$\varphi'(0) = f'(\mathbf{x}_0), \quad \varphi''(0) = H(\mathbf{x}). \quad (6)$$

By Taylor's theorem, there exists a \bar{t} in the interval $(0, 1)$ such that

$$\varphi(x) - \varphi(x_0) = \varphi'(x_0) \cdot \tau + \frac{1}{2} \tau^2 \varphi''(\bar{x}). \quad (7)$$

Again, straightforward calculations as in relations (17) and (18) in Section 13 of Chapter 3 give

$$\begin{aligned} \varphi'(x) &= (\nabla f)(x_0) + \tau \nabla g(x), \\ \varphi''(x) &= (x, \mathcal{H}(x_0) + \tau \mathcal{H}(x)), \quad -1 - \eta \leq \tau \leq 1 + \eta. \end{aligned}$$

If we set $\tau = 0$ in the first of these relations, we get that $\varphi'(x_0) = (\nabla f)(x_0)$. If we then set $\tau = \bar{t}$ in the second of these relations, use the hypothesis $\nabla f(x_0) = 0$, and substitute into equation (7), we get

$$\varphi(x) - \varphi(x_0) = \frac{1}{2} (x, \mathcal{H}(x_0) + \bar{t} \mathcal{H}(x)).$$

From this relation, from (5), and from the fact that $\mathcal{H}(x)$ is positive-definite for all x in $B(x_0, \delta)$, we get that

$$f(x) - f(x_0) = \frac{1}{2} (x, \mathcal{H}(x_0) + \bar{t} \mathcal{H}(x)) > 0$$

for all x in $B(x_0, \delta)$. Hence x_0 is a strict local minimizer.

Remark 2.1. Theorems 2.1 and 2.3 suggest the well-known strategy of solving unconstrained optimization problems by first finding the critical points. Some of these critical points can then be eliminated by considering the second derivative or Hessian at these points. By then considering strengthened conditions on the second derivative or the Hessian as in Theorems 2.2 and 2.4, one can show in some cases that a critical point is indeed a local minimizer or maximizer.

Exercise 2.1. Let K be a closed interval $[a, b]$ in \mathbb{R} and let f be a real-valued function defined on K . Show that if $x = a$ is a local minimizer for f and if the right-hand derivative $f'_+(a)$ exists, then $f'_+(a) \geq 0$. Show that if $x = b$ is a local minimizer and the left-hand derivative $f'_-(b)$ exists, then $f'_-(b) \leq 0$.

Exercise 2.2. Find local minimizers, local maximizers, minimizers, and maximizers, if they exist, for the following functions. Exercise 15.3 may be useful in some cases:

- $f(x_1, x_2) = x_1^2 - 4x_1 + 2x_2^2 + 7$
- $f(x_1, x_2, x_3) = 3x_1^2 + 2x_2^2 + 2x_3^2 + 2x_1x_2 + 2x_2x_3 + 2x_1x_3$
- $f(x_1, x_2) = x_1x_2 + 1/x_1 + 1/x_2$, $(x_1, x_2) \neq (0, 0)$, and
- $f(x_1, x_2) = (x_1 + x_2)(x_1, x_2 + 1)$

Exercise 2.3. Let

$$f(x_1, x_2) = x_1^2 + x_2^2 \quad g(x_1, x_2) = x_1^2 + x_2^2.$$

Show that $\nabla f(\mathbf{0}) = \nabla g(\mathbf{0}) = \mathbf{0}$ and that both f and g have positive semidefinite Hessians at $\mathbf{0}$, but $\mathbf{0}$ is a local minimizer for g but not for f . This shows that the condition of positive definiteness of $\text{Hess}_g(\mathbf{0})$ in Theorem 2.4 cannot be replaced by positive semidefiniteness.

Exercise 2.4. Complete the proof of Theorem 2.4 by proving the last two statements.

Exercise 2.5. Let

$$f(x, y) = x^2 + y^2 - 3xy^2.$$

Minimize f on \mathbb{R}^2 . (Hint: see Exercise 18.1.)

Exercise 2.6. (a) Show that for no value of a does the function

$$f(x, y) = x^2 - 3axy + y^2$$

have a minimizer.

(b) For each value of a find local minimizers and local maximizers, if any exist, of this function.

Exercise 2.7. Let f be defined on \mathbb{R}^2 by

$$f(x, y) = x^2 + e^{2y} - 3xy^2.$$

Show that there is a unique point (x_0, y_0) at which $\nabla f(x, y) = \mathbf{0}$ and that this point is a local minimizer but not a maximizer.

Exercise 2.8 (Linear Regression). In the linear regression problem n points $(x_1, y_1), (x_2, y_2), \dots, (x_n, y_n)$ are given in the xy -plane and it is required to "fit" a straight line $y = ax + b$ to these points in such a way that the sum of the squares of the vertical distances of the given points from the line is minimized. That is, a and b are to be chosen so that

$$f(a, b) = \sum_{i=1}^n (ax_i + b - y_i)^2$$

is minimized. The resulting line is called the linear regression line for the given points.

Show that the coefficients a and b of the linear regression line are given by

$$b = \bar{y} - a\bar{x}, \quad a = \frac{n\bar{xy} - \sum_{i=1}^n x_i y_i}{n(\bar{x})^2 - \sum_{i=1}^n x_i^2},$$

where

$$\bar{x} = \frac{1}{n} \sum_{i=1}^n x_i, \quad \bar{y} = \frac{1}{n} \sum_{i=1}^n y_i.$$

Exercise 2.8 (Best Mean-Square Approximation by Orthogonal Functions). Let $\varphi_1, \dots, \varphi_r$ be continuous and not identically zero on an interval $[a, b]$ and such that

$$\int_a^b \varphi_j(x)\varphi_k(x) dx = 0 \quad \text{if } j \neq k.$$

Given a piecewise continuous function f on $[a, b]$ show that the values of c_1, \dots, c_r that minimize

$$F(c) = \int_a^b \left| f - \sum_{j=1}^r c_j \varphi_j \right|^2 dx$$

are given by the formulas

$$c_i = \frac{\langle f, \varphi_i \rangle}{\langle \varphi_i, \varphi_i \rangle}, \quad i = 1, \dots, r,$$

where

$$\langle \varphi_i, \varphi_j \rangle = \int_a^b \varphi_i(x)\varphi_j(x) dx.$$

Exercise 2.10. In this exercise we shall establish a series of results, some of which are of interest in their own right, that lead to a proof of Lemma III.1.2.

(a) Let $Q(x) = \langle x, Ax \rangle$, where A is a symmetric matrix. Show that if Q is positive definite, then A is nonsingular. Also: If A were singular, there would exist an $x_0 \neq 0$ such that $Ax_0 = 0$.

(b) Show that if Q is positive semidefinite and A is nonsingular, then Q is positive definite. Also: If A is positive semidefinite, then there exists an $x_0 \neq 0$ that minimizes Q .

(c) If Q is positive definite, then $\det A > 0$. Also: Let I denote the identity matrix and let $0 < \lambda < 1$. For each $0 < \lambda < 1$ consider the quadratic form

$$F(\mathbf{x}, \lambda) = (1 - \lambda)\|\mathbf{x}\|^2 + \lambda g(\mathbf{x}) = \langle \mathbf{x}, \mathbf{1} \rangle - \lambda \|\mathbf{x}\|^2 + \lambda \langle \mathbf{x}, \mathbf{c} \rangle.$$

Let $d(\lambda) = \det[(1 - \lambda)I + \lambda A]$. Show that d is a continuous function, that $d(\lambda) \neq 0$, and that $d(0) = 1$.

(ii) Let the function f with domain \mathbb{R}^n be defined by

$$f(\mathbf{x}) = \langle \mathbf{x}, A\mathbf{x} \rangle + 2\langle \mathbf{b}, \mathbf{x} \rangle + c,$$

where A is an $n \times n$ positive-definite matrix and \mathbf{b} is an n -vector. Show that $\mathbf{x}_0 = -A^{-1}\mathbf{b}$ is the unique minimizer of f and that

$$f(\mathbf{x}_0) = \langle \mathbf{b}, \mathbf{x}_0 \rangle + c = -\frac{\det \begin{pmatrix} A & \mathbf{b} \\ \mathbf{b}^T & c \end{pmatrix}}{\det A}.$$

Hint

$$\det \begin{pmatrix} A & \mathbf{b} \\ \mathbf{b}^T & c \end{pmatrix} = \det \begin{pmatrix} A & A\mathbf{x}_0 + \mathbf{b} \\ \mathbf{b}^T & \langle \mathbf{b}, \mathbf{x}_0 \rangle + c \end{pmatrix}.$$

Why?

(i) Let

$$g(\mathbf{x}, y) = \langle \mathbf{x}, A\mathbf{x} \rangle + 2y\langle \mathbf{b}, \mathbf{x} \rangle + cy^2 = \langle \mathbf{x}, y^2 B \mathbf{x}, y \rangle,$$

where

$$B = \begin{pmatrix} A & \mathbf{b} \\ \mathbf{b}^T & c \end{pmatrix}.$$

Show that the quadratic form g is positive definite if and only if A is positive definite and $\det B > 0$. Hint:

(i) Since $g(\mathbf{x}, 0) = \langle \mathbf{x}, A\mathbf{x} \rangle$, if g is positive definite, then A is positive definite.

(ii) For $y \neq 0$,

$$g(\mathbf{x}, y) = y^2 g \left(\frac{\mathbf{x}}{y}, 1 \right).$$

(iii) If A is positive definite, then

$$g(\mathbf{x}, 1) = f(\mathbf{x}) \geq f(\mathbf{x}_0).$$

(c) Prove Lemma III.5.2. (Hint: If A is positive definite, use (i) and induction. If A is negative definite, then $-A$ is positive definite.)

3. OPTIMIZATION OF CONVEX FUNCTIONS

The points at which a convex function attains a maximum or minimum have important properties not shared by all functions.

THEOREM 3.1. *Let f be a convex function defined on a convex set C .*

- (i) *If f attains a local minimum at x_0 , then f attains a minimum at x_0 .*
- (ii) *The set of points at which f attains a minimum is either empty or convex.*
- (iii) *If f is strictly convex and f attains a minimum at x_0 , then x_0 is unique.*
- (iv) *If f is not a constant function and f attains a maximum at some point x in C , then x must be a boundary point of C .*

Proof: (i) Since f attains a local minimum at x_0 , there exists a $\delta > 0$ such that for all x in $B(x_0, \delta) \cap C$

$$f(x) - f(x_0) \geq 0. \quad (1)$$

Let y be an arbitrary point in C . Then $[x_0, y]$ belongs to C . Hence for $0 < t < \delta/(y - x_0)$ the point $x = x_0 + ty = x_0 + t(y - x_0)$ belongs to $B(x_0, \delta) \cap C$. Hence, from Lemma III.1.7 and (1) we get

$$f(y) - f(x_0) \geq \frac{f(x) - f(x_0)}{t} \geq 0.$$

(ii) If f attains a minimum at a unique point, then there is nothing to prove. Suppose that f attains a minimum at two distinct points x_1 and x_2 . Let $\mu = f(x_1) = f(x_2) = \min\{f(x) \mid x \in C\}$. Then for any (α, β) in \mathcal{P}_2 with $\alpha + \beta = 1$, $\beta > 0$,

$$\mu \leq f(\alpha x_1 + \beta x_2) \leq \alpha f(x_1) + \beta f(x_2) = \alpha\mu + \beta\mu = \mu.$$

Hence f attains a minimum at each point of $[x_1, x_2]$.

(iii) Suppose that f attains a minimum at two distinct points x_1 and x_2 . Let μ be as in (ii). Then for any (α, β) in \mathcal{P}_2

$$\mu \leq f(\alpha x_1 + \beta x_2) = \alpha f(x_1) + \beta f(x_2) = \mu.$$

Thus, we get a contradiction, $\mu < \mu$, and so x_1 and x_2 cannot be distinct.

(iv) Suppose f attains a maximum at x . If x were not a boundary point, then x would be an interior point. Since f is not a constant function, there exists a point y in C such that $f(y) < f(x)$. Since C is convex, the line segment $[x, y]$ belongs to C . Since x is an interior point, there is a line segment $[x, y]$ such that x is an interior point of $[x, y]$. Then $x = \alpha z + \beta y$ for some (α, β) in

F_1 with $\alpha > 0$, $\beta > 0$, and

$$f(x) = \alpha f(x) + \beta f(x) = \alpha f(x) + \beta f(x) = f(x).$$

Thus we have the contradiction $f(x) = f(x)$.

The next theorem and its corollary characterize points x_0 at which a function attains a minimum.

THEOREM 3.2. Let f be a real-valued function defined on a set C . Then f attains a minimum at $x_0 \in C$ if and only if $0 \in \partial f(x_0)$.

Proof. If f attains a minimum at x_0 , then for all x in C ,

$$f(x) \geq f(x_0) = f(x_0) + \langle 0, x - x_0 \rangle. \quad (2)$$

Hence $0 \in \partial f(x_0)$. Conversely, if $0 \in \partial f(x_0)$, then (2) holds and f attains a minimum at x_0 .

Note that the theorem does not require f or C to be convex. If, however, f is not convex, it can fail to have a subgradient at all points of its domain. Such a function therefore cannot attain a minimum. The logarithm function in x , or any other logarithm function, is an example. The next corollary does require f and C to be convex.

COROLLARY 1. Let f be convex and differentiable on an open convex set C . Then f attains a minimum at $x_0 \in C$ if and only if $\nabla f(x_0) = 0$.

Proof. Let f attain a minimum at x_0 . That $\nabla f(x_0) = 0$ follows from Theorem 3.2, whether or not f and C are convex. If f is convex and differentiable on an open convex set C and $\nabla f(x_0) = 0$, then by Theorem III.3.1

$$\{0\} = [\nabla f(x_0)] = \partial f(x_0).$$

Thus, by the theorem, f attains a minimum at x_0 .

Remark 3.1. In geometric terms, Theorem 3.2 states that for a convex function f defined on a convex set C , a point x_0 in C furnishes a minimum if and only if the epigraph of f has a horizontal supporting hyperplane at $(x_0, f(x_0))$. The corollary states that for a differentiable convex function defined on an open convex set C , a point x_0 furnishes a minimum if and only if the tangent plane to the graph of f at x_0 is horizontal.

The corollary also points out the true content of Theorem 2.4. By Theorem III.3.3, the positive definiteness or semidefiniteness of $\partial^2 f(x)$ on E implies that

f is convex. It is the convexity of f that is the important element in Theorem 2.4. The alert reader may have observed that much of the proof of Theorem 2.4 was a repetition of the proof of Theorem III.3.

Exercise 31. Show that $f(x) = x_1^2 + x_2^2 - 2x_1 - 7x_2$ attains a minimum on \mathbb{R}^2 at a unique point and find the minimum.

Exercise 32. Minimize $f(x) = x_1^2 + u_1x_1 + x_2^2$ subject to

$$-2x_1 - 2x_2 + 6 \leq 0, \quad -u_1 + x_1 - 3 \leq 0, \quad x_1 - 2 \leq 0.$$

- Sketch the feasible set.
- Show that a solution exists.
- Find the solution.

4. LINEAR PROGRAMMING PROBLEMS

Linear programming is used extensively in resource allocation problems. To illustrate, we formulate a simple job-scheduling problem.

An office furniture manufacturer makes a different types of desks. His profit per desk of type j is b_j . Thus, if each week he produces x_j desks of type j , then his profit for the week is

$$b_1x_1 + \cdots + b_nx_n.$$

The manufacturer's objective is to maximize his profit. The number of desks that he can produce per week is not unlimited but is constrained by the capacity of his plant. Suppose that the manufacture of a type 1 desk requires a_{11} hours per week on machine 1, that the manufacture of a type 2 desk requires a_{21} hours of time on machine 1, and the manufacture of a type j desk requires a_{j1} hours per week of time on machine 1. If $a_{j1}, j = 1, \dots, n$, denotes the number of desks of type j manufactured per week, then the number of hours that machine 1 is used per week is

$$a_{11}x_1 + a_{21}x_2 + \cdots + a_{n1}x_n.$$

Because of maintenance requirements, readjustments, and the number of shifts employed, the total time that machine 1 is in operation cannot exceed some fixed preassigned number c_1 . Thus we have a constraint

$$a_{11}x_1 + \cdots + a_{n1}x_n \leq c_1.$$

Similarly, if the manufacture of one desk of type j requires a_{j2} hours per week

on machine i and the total number of hours per week that machine i can operate is c_i , then

$$a_{1j}x_j + \cdots + a_{ij}x_j \leq c_i, \quad i = \bar{1}, \bar{2}, \dots, \bar{n}.$$

Also the number of disks of type j made is nonnegative, so $x_j \geq 0$.

The problem for the disk manufacturer is to manufacture x_j disks of type j for as little as possible:

$$\sum_{j=1}^n b_j x_j$$

subject to

$$\sum_{j=1}^n a_{ij} x_j \leq c_i, \quad i = 1, \dots, m, \quad x_j \geq 0, \quad j = 1, \dots, n.$$

This is a linear programming problem in standard form.

For a brief account of the origins and early days of linear programming by a pioneer in the field, see Dantzig [1963]. Some early applications are also discussed there.

In this section we shall derive a necessary and sufficient condition that a point x_0 be a solution to a linear programming problem. In practice the use of this condition to solve the problem is precluded by the dimension of the space E^n . Instead, a direct method, the simplex method, is used. The simplex method will be discussed in Chapter VI. Nevertheless, we derive the necessary conditions for linear programming problems to illustrate the simplest case of the use of Farkas's theorem (and hence separation theorems) to derive necessary conditions in constrained optimization problems. The basic idea in this proof is a theme that runs throughout much of optimization theory.

The linear programming problem in standard form (SLP) is given as

$$\begin{aligned} & \text{Minimize } \zeta = \mathbf{k} \cdot \mathbf{x} \\ & \text{subject to } A\mathbf{x} \leq \mathbf{c}, \quad \mathbf{x} \geq \mathbf{0} \end{aligned}$$

where $\mathbf{x} \in E^n$, A is an $m \times n$ matrix, $\mathbf{k}, \mathbf{c} \in E^n$ is a vector in E^n , and \mathbf{c} is a vector in E^m .

The linear programming problem SLP can also be written as follows: Minimize $\zeta = \mathbf{k} \cdot \mathbf{x}$ subject to

$$\begin{pmatrix} A \\ -I \end{pmatrix} \mathbf{x} \leq \begin{pmatrix} \mathbf{c} \\ \mathbf{0} \end{pmatrix} \quad (1)$$

where I is the $n \times n$ identity matrix and $\mathbf{0}$ is the zero vector in \mathbb{R}^n . If we denote the matrix on the left in (2) by A and the vector on the right by \mathbf{b} and then drop the tildes, we can state the linear programming problem (LPI) as

$$\begin{aligned} & \text{minimize } (-\mathbf{b}, \mathbf{x}) \\ & \text{subject to } A\mathbf{x} \leq \mathbf{a}, \end{aligned}$$

where A is an $m \times n$ matrix, $\mathbf{b} \in \mathbb{R}^n$, and $\mathbf{a} \in \mathbb{R}^m$. Here, we have incorporated the constraint $\mathbf{x} \geq \mathbf{0}$ in the matrix A and the vector \mathbf{a} .

Theorem 4.1. Let \mathbf{x}_0 be a solution of LPI. Then there exists a $\mathbf{k} \neq \mathbf{0}$ in \mathbb{R}^m such that $(\mathbf{x}_0, \mathbf{k})$ satisfies the following:

$$\begin{aligned} \text{(i)} \quad & A\mathbf{x}_0 \leq \mathbf{a}, & \text{(ii)} \quad A\mathbf{k} = \mathbf{b}, \\ \text{(iii)} \quad & \mathbf{k} \geq \mathbf{0}, & \text{(iv)} \quad (-A\mathbf{x}_0 - \mathbf{a}, \mathbf{k}) = 0, \\ \text{(v)} \quad & (\mathbf{b}, \mathbf{x}_0) = (\mathbf{a}, \mathbf{k}). \end{aligned} \quad (3)$$

If there exists an $\mathbf{x}_0 \in \mathbb{R}^n$ and a $\mathbf{k} \neq \mathbf{0}$ in \mathbb{R}^m that satisfy (i)–(v), then \mathbf{x}_0 is a solution of LPI.

Proof. The feasible set for LPI is $X = \{\mathbf{x} \mid A\mathbf{x} \leq \mathbf{a}\}$. Since \mathbf{x}_0 is a solution to LPI, $\mathbf{x}_0 \in X$, and so (i) holds.

We now prove (ii) and (iii). If \mathbf{x}_0 in X is a solution to LPI, then $(-\mathbf{b}, \mathbf{x}_0) \leq (-\mathbf{b}, \mathbf{x})$ for all \mathbf{x} in X . Hence $(\mathbf{b}, \mathbf{x}_0) \leq (\mathbf{b}, \mathbf{x})$ for all \mathbf{x} in X . If we set $\mathbf{a} = (\mathbf{b}, \mathbf{x}_0)$, then we have that X is contained in the closed negative half space determined by the hyperplane $H_{\mathbf{a}}$. By Exercise B.2.15, for \mathbf{x} in $\text{int}(X)$, we have $(\mathbf{b}, \mathbf{x}) < \mathbf{a}$, and so \mathbf{x}_0 cannot be an interior point of X .

Let \mathbf{a}_i denote the i th row of the matrix A . Let

$$K = \{i \mid (\mathbf{a}_i, \mathbf{x}_0) = \mathbf{a}_i\}, \quad I = \{i \mid (\mathbf{a}_i, \mathbf{x}_0) < \mathbf{a}_i\}.$$

Since \mathbf{x}_0 is not an interior point of X , the set $K \neq \emptyset$. Let A_2 denote the submatrix of A consisting of those rows \mathbf{a}_i with i in K , and let A_1 denote the submatrix of A consisting of those rows \mathbf{a}_i with i in I .

We claim that the system

$$A_2 \mathbf{x} \leq \mathbf{0}, \quad (-\mathbf{b}, \mathbf{x}) = 0, \quad (4)$$

has no solution \mathbf{x} in \mathbb{R}^n .

Intuitively this is clear, for if \mathbf{x}_0 would be a solution, then we would move, for $\alpha > 0$, and we would have $A_2(\mathbf{x}_0 + \alpha \mathbf{x}) \leq \mathbf{a}_2$. Also for α sufficiently small, by continuity, we would have $A_1(\mathbf{x}_0 + \alpha \mathbf{x}) < \mathbf{a}_1$. Thus, if (4) had a solution, we would be able to move into the feasible set and at the same time increase $(\mathbf{b}, \mathbf{x}_0)$ or decrease $(-\mathbf{b}, \mathbf{x}_0)$. This would contradict the optimality of \mathbf{x}_0 .

We now make this argument precise. Again, if x_0 were a solution of (1), then $x_0 + \alpha x_1$ would be a solution of (3) for all $\alpha > 0$, and

$$\langle -k, x_0 + \alpha x_1 \rangle = \langle -k, x_0 \rangle + \alpha \langle k, x_1 \rangle < \langle -k, x_0 \rangle \quad (8)$$

for all $\alpha > 0$. Also,

$$A_1(x_0 + \alpha x_1) = x_0 + \alpha A_1 x_1 < x_0 \quad \alpha > 0, \quad (9)$$

and

$$A_2(x_0 + \alpha x_1) = A_2 x_0 + \alpha A_2 x_1 = \beta_1 - \beta_2 + \alpha A_2 x_1,$$

where $\beta > 0$. (Recall that $\beta > 0$ means that all components of β are strictly positive.) By taking $\alpha > 0$ to be sufficiently small, we can make $\alpha A_2 x_1 - \beta < 0$. Hence for $\alpha > 0$ and sufficiently small

$$A_2(x_0 + \alpha x_1) = x_1 + (\alpha A_2 x_1 - \beta) < x_1. \quad (10)$$

In summary, from (5) and (6), we get that for sufficiently small $\alpha > 0$ the vector $x_0 + \alpha x_1$ are in X . From this and from (4), we get that x_0 is not a solution to the linear programming problem. This contradiction shows that x_0 cannot be a solution of (3).

Since the system (3) has no solution, it follows, from Theorem B5.1 (Farkas's lemma) that there exists a vector y in \mathbb{R}^m , where m is the row dimension of A_1 , such that

$$A_1^T y = k, \quad y \geq 0.$$

Therefore, if m_1 is the row dimension of A_1 , then

$$A_1^T y + A_2^T \theta = k,$$

where θ is the zero vector in \mathbb{R}^m . If we now take $\lambda = (y, \theta)$, then λ is in \mathbb{R}^n and satisfies

$$A\lambda = (A_1^T, A_2^T) \begin{pmatrix} y \\ \theta \end{pmatrix} = k, \quad \lambda \geq 0. \quad (7)$$

Thus, (ii) and (iii) are established.

From (7) we see that those components of λ corresponding to indices in J are zero. By the definition of J , $A_1 \lambda = x_0 - \theta$. Hence

$$\langle Ax_0 - b, \lambda \rangle = \langle A_1 \lambda - x_0, A_2 \lambda - x_1 \rangle = \langle x_0 - \theta, x_1 \rangle = 0, \quad (8)$$

which is (iv).

From (ii) and (i) we get

$$\langle \mathbf{a}, \mathbf{k} \rangle = \langle A\mathbf{x}_0, \mathbf{k} \rangle = \langle \mathbf{x}_0, A^T\mathbf{k} \rangle = \langle \mathbf{x}_0, \mathbf{b} \rangle,$$

which is (v).

The proof of sufficiency is elementary and does not involve separation theorems. Let \mathbf{x} be any feasible point. Then

$$\begin{aligned} \langle -\mathbf{b}, \mathbf{x} \rangle - \langle -\mathbf{b}, \mathbf{x}_0 \rangle &= \langle -A^T\mathbf{k}, \mathbf{x} \rangle + \langle \mathbf{x}, \mathbf{k} \rangle \\ &= \langle \mathbf{k}, -A\mathbf{x} + \mathbf{x} \rangle = -\langle \mathbf{k}, A\mathbf{x} - \mathbf{b} \rangle \geq 0. \end{aligned}$$

Hence \mathbf{x}_0 is a solution to L.P.

Remark 4.1. An alternate proof of (ii) and (v) is to establish (v) first and then use (v) to establish (ii). Thus

$$\begin{aligned} \langle \mathbf{b}, \mathbf{x}_0 \rangle &= \langle A^T\mathbf{k}, \mathbf{x}_0 \rangle = \langle \mathbf{k}, A\mathbf{x}_0 \rangle = \left\langle \mathbf{y}, \mathbf{b}_0, \begin{pmatrix} A_0 \\ A_1 \end{pmatrix} \mathbf{x}_0 \right\rangle \\ &= \langle \mathbf{y}, A_0\mathbf{x}_0 \rangle + \langle \mathbf{y}, \mathbf{b}_0 \rangle = \langle \mathbf{y}, \mathbf{b}_0 \rangle = \langle \mathbf{y}, \mathbf{b}_0 \rangle = \langle \mathbf{y}, \mathbf{b} \rangle \\ &= \langle \mathbf{b}, \mathbf{c} \rangle, \end{aligned}$$

which is (v). To establish (ii), we write

$$\begin{aligned} 0 &= \langle \mathbf{b}, \mathbf{x}_0 \rangle - \langle \mathbf{a}, \mathbf{k} \rangle = \langle A^T\mathbf{k}, \mathbf{x}_0 \rangle - \langle \mathbf{a}, \mathbf{k} \rangle \\ &= \langle \mathbf{k}, A\mathbf{x}_0 \rangle - \langle \mathbf{a}, \mathbf{k} \rangle = \langle A\mathbf{x}_0 - \mathbf{a}, \mathbf{k} \rangle. \end{aligned}$$

Remark 4.2. It follows from (i) and (ii) that (iv) holds if and only if for each $i = 1, \dots, m$

$$\langle \mathbf{a}_i, \mathbf{x}_0 \rangle - c_i x_i = 0. \quad (iv')$$

That is, (iv') is equivalent to (iv). The condition (iv') is sometimes called the complementary slackness condition. It states that if there is "slack" in the i th constraint at \mathbf{x}_0 , that is, if $\langle \mathbf{a}_i, \mathbf{x}_0 \rangle < c_i$, then $\lambda_i = 0$.

Remark 4.3. Let the linear programming problem be reformulated with X_0 taken to be an arbitrary open set instead of \mathbb{R}^n . It is clear from the proof that if \mathbf{x}_0 is a solution to this problem, then the necessary conditions hold. Conversely, if there is a point \mathbf{x}_0 in X_0 at which the necessary conditions hold, then \mathbf{x}_0 is feasible and the minimum is achieved at \mathbf{x}_0 .

The inequality constraints $Ax \leq r$ can be converted to equality constraints by introducing a nonnegative vector ξ whose components are called slack variables and writing

$$Ax + \xi = r, \quad \xi \geq \theta.$$

Thus the linear programming problem formulated in (1) can be written as

$$\begin{aligned} & \text{Minimize } (-k, \theta), (x, \xi) \\ & \text{subject to } (A, I) \begin{pmatrix} x \\ \xi \end{pmatrix} = r, \quad x \geq \theta, \xi \geq \theta. \end{aligned}$$

That is, the problem LPI can be written as a canonical linear programming (CLP) problem

$$\begin{aligned} & \text{Minimize } (-k, \alpha) \\ & \text{subject to } Ax = r, \quad \alpha \geq \theta. \end{aligned}$$

The simplex method is applied to problems in canonical form.

Conversely, the problem CLP can be written as a problem LPI by writing the equality constraint $Ax = r$ as two inequality constraints, $Ax \leq r$ and $-Ax \leq -r$, and then incorporating the constraint $x \geq \theta$ in the matrix. That is, we write CLP as

$$\begin{aligned} & \text{Minimize } (-k, \alpha) \\ & \text{subject to } \begin{pmatrix} A \\ -A \\ -I \end{pmatrix} x \leq \begin{pmatrix} r \\ -r \\ \theta \end{pmatrix}. \end{aligned}$$

Remark 4.4. We warn the reader that the definitions "standard form" and "canonical form" are not uniform throughout the linear programming literature. Problems that we have designated as being in canonical form are designated as being in standard form by some authors who use the designation "canonical form" differently from ours. In consulting the literature the reader should always check the definitions used.

Exercise 4.1. Prove the following corollary to Theorem 4.1.

COROLLARY 1. Let x_0 be a solution of the standard linear programming problem SLP. Then there exists a vector $k, \alpha \geq \theta$ in \mathbb{R}^n such that (x_0, k) satisfies

- (i) $\langle A\mathbf{x}_0, \mathbf{c} \rangle, \mathbf{x}_0 \geq 0$
 (ii) $\sum_{j=1}^m \lambda_j a_{ij} \geq b_j, j = 1, \dots, m$, with equality holding if $\mathbf{x}_{0j} > 0$
 (iii) $\lambda \geq 0$
 (iv) $\langle A\mathbf{x}_0 - \mathbf{c}, \lambda \rangle = 0$, and
 (v) $(\lambda, \mathbf{x}_0) = (\mathbf{x}, \lambda)$.

If there exist an \mathbf{x}_0 in \mathbb{R}^n and a $\lambda, j \neq 0$ in \mathbb{R}^m such that (i)–(v) hold, then \mathbf{x}_0 is a solution of CLP.

COROLLARY 2. Let \mathbf{x}_0 be a solution of the canonical linear programming problem CLP. Then there exists a vector $\lambda, j \neq 0$ in \mathbb{R}^m such that (\mathbf{x}_0, λ) satisfies

- (i) $\langle A\mathbf{x}_0, \mathbf{c} \rangle, \mathbf{x}_0 \geq 0$
 (ii) $\sum_{j=1}^m \lambda_j a_{ij} \geq b_j, j = 1, \dots, m$, with equality holding if $\mathbf{x}_{0j} > 0$
 (iii) $\langle A\mathbf{x}_0 - \mathbf{c}, \lambda \rangle = 0$, and
 (iv) $(\lambda, \mathbf{x}_0) = (\mathbf{x}, \lambda)$.

If there exist an \mathbf{x}_0 in \mathbb{R}^n and a $\lambda, j \neq 0$ in \mathbb{R}^m such that (i)–(iv) hold, then \mathbf{x}_0 is a solution of CLP.

Note that in Corollary 2 the condition $\lambda \geq 0$ is absent.

Example 4.2. Show that the problem LP1 can be written as a problem in standard form.

5. FIRST-ORDER CONDITIONS FOR DIFFERENTIABLE NONLINEAR PROGRAMMING PROBLEMS

In this section we shall derive first-order necessary conditions for differentiable programming problems. As noted in the introduction of this chapter, the set of points that satisfy the necessary conditions need not furnish a solution to the programming problem. If a solution exists, then it belongs to this set. Satisfaction of the necessary conditions corresponds to "bringing in the usual suspects" in a criminal investigation. Further investigation is required to determine which, if any, of these points, or "suspects," furnish a solution. In this vein, we shall show that if the data of the problem are convex, then a point that satisfies the necessary conditions is a solution.

Problem P1

Let X_0 be an open convex set in \mathbb{R}^n . Let f, g , and \mathbf{h} be C^1 functions with domain X_0 and ranges in \mathbb{R} , \mathbb{R}^m , and \mathbb{R}^p , respectively. Let

$$K = \{x \mid x \in X_0, g(x) \leq 0, \mathbf{h}(x) = 0\}.$$

Minimize f over K .

The problem is often stated as

$$\begin{aligned} & \text{Minimize } f(\mathbf{x}) \\ & \text{subject to } \mathbf{g}(\mathbf{x}) \leq \mathbf{0}, \quad \mathbf{h}(\mathbf{x}) = \mathbf{0}. \end{aligned}$$

If the constraints $\mathbf{g}(\mathbf{x}) \leq \mathbf{0}$ and $\mathbf{h}(\mathbf{x}) = \mathbf{0}$ are absent, the problem becomes an unconstrained problem, which was treated in Section 2. We saw there that a necessary condition for a point \mathbf{x}_0 to be a solution is that all first-order partial derivatives of f vanish at \mathbf{x}_0 . For $n = 2$, l and $1 \leq k < n$, the problem with equality constraints and no inequality constraints is also usually treated in elementary calculus courses. It is stated there, with or without proof, that if \mathbf{x}_0 is a solution of the problem and if the l gradient vectors $\nabla h_j(\mathbf{x}_0)$ are linearly independent, then there exist scalars μ_j , $j = 1, \dots, l$, such that for each $j = 1, \dots, l$

$$\frac{\partial f}{\partial x_j}(\mathbf{x}_0) + \sum_{i=1}^l \mu_i \frac{\partial h_i(\mathbf{x}_0)}{\partial x_j} = 0$$

and

$$h_j(\mathbf{x}_0) = 0, \quad j = 1, \dots, l.$$

This is the Lagrange multiplier rule.

We now give a first-order necessary condition that is satisfied by a solution of problem PL. The condition includes a more general version of the Lagrange multiplier rule in the case that inequality constraints are absent.

We recall the notation established in Section 11 of Chapter 1. If f is a real-valued function and $\mathbf{g} = (g_1, \dots, g_m)$ is a vector-valued function that is differentiable at \mathbf{x}_0 , then

$$\nabla f(\mathbf{x}_0) = \left(\frac{\partial f}{\partial x_1}(\mathbf{x}_0), \frac{\partial f}{\partial x_2}(\mathbf{x}_0), \dots, \frac{\partial f}{\partial x_n}(\mathbf{x}_0) \right)$$

and

$$\nabla \mathbf{g}(\mathbf{x}_0) = \begin{pmatrix} \nabla g_1(\mathbf{x}_0) \\ \nabla g_2(\mathbf{x}_0) \\ \vdots \\ \nabla g_m(\mathbf{x}_0) \end{pmatrix} = \left(\frac{\partial g_i(\mathbf{x}_0)}{\partial x_j} \right).$$

THEOREM 5.1. Let \mathbf{x}_0 be a solution of problem PL. Then there exists a real number $\lambda_0 \geq 0$, a vector $\lambda \geq \mathbf{0} \in \mathbb{R}^m$, and a vector $\mu \in \mathbb{R}^l$ such that

- (i) $(\lambda_0, \lambda, \mu) \neq 0$,
 (ii) $\lambda_0 g(\mathbf{x}_0) = 0$, and
 (iii) $\lambda_0 \nabla f(\mathbf{x}_0) + \sum_{j=1}^m \lambda_j \nabla g_j(\mathbf{x}_0) + \sum_{j=1}^r \mu_j \nabla h_j(\mathbf{x}_0) = 0$.

Vectors $(\lambda_0, \lambda, \mu)$ having the properties stated in the theorem are called *multipliers*, as are the components of these vectors.

Definition 5.1. A point \mathbf{x}_0 at which the conclusion of Theorem 5.1 holds will be called a *critical point* for problem P1.

Thus, Theorem 5.1 says that a solution \mathbf{x}_0 of problem P1 is a critical point.

Remark 5.1. The necessary condition asserts that $\lambda_0 \geq 0$ and not the stronger statement $\lambda_0 > 0$. If $\lambda_0 > 0$, then we may divide through by λ_0 in (ii) and (iii) and relabel λ/λ_0 and μ/λ_0 as λ and μ , respectively, and then obtain statements (i)–(iii) with $\lambda_0 = 1$. In the absence of further conditions that would guarantee that $\lambda_0 > 0$, we cannot assume that $\lambda_0 = 1$. We shall return to this point later.

Remark 5.2. If $\mu = 0$, that is, the inequality constraints are absent, Theorem 5.1 becomes the Lagrange multiplier rule.

Remark 5.3. Since $\lambda \geq 0$ and $g_j(\mathbf{x}_0) \leq 0$, condition (ii) is equivalent to

$$\lambda_j g_j(\mathbf{x}_0) = 0, \quad j = 1, \dots, m. \quad (5')'$$

Condition (iii) is a system of s equations

$$\lambda_0 \frac{\partial f}{\partial x_j}(\mathbf{x}_0) + \sum_{i=1}^m \lambda_i \frac{\partial g_i(\mathbf{x}_0)}{\partial x_j} + \sum_{i=1}^r \mu_i \frac{\partial h_i(\mathbf{x}_0)}{\partial x_j} = 0, \quad j = 1, \dots, s. \quad (5'')$$

Remark 5.4. The necessary conditions are also necessary conditions for a local minimum. For if \mathbf{x}_0 is a local minimum, then there exists a $\delta > 0$ such that $f(\mathbf{x}_0) \leq f(\mathbf{x})$ for all \mathbf{x} that are in $\mathcal{R}(\mathbf{x}_0, \delta) \cap X_0$ and that satisfy the constraints $g(\mathbf{x}) \leq 0$ and $h(\mathbf{x}) = 0$. Thus, \mathbf{x}_0 is a global solution to the problem in which X_0 is replaced by $X_0 \cap \mathcal{R}(\mathbf{x}_0, \delta)$.

Remark 5.5. Theorem 5.1, under additional hypotheses guaranteeing that $\lambda_0 > 0$, was proved in Kuhn and Tucker [1951]. This paper had great impact since interest in nonlinear programming was developing at that time in economics and other areas of application. The theorem with $\lambda_0 \geq 0$ had been proved in Fiacri [1948], and prior to that, again under conditions ensuring that $\lambda_0 > 0$, in Karush, [1959]. The results of Fiacri and of Karush were unnoticed when they appeared, probably because interest in programming problems had not yet come about. These papers were looked upon only as mathematical

generalizations of the Lagrange multiplier rule. In the literature, Theorem 5.1, with additional hypotheses guaranteeing that $\lambda_0 = 0$, is often referred to as the Karush–Kuhn–Tucker Theorem. Theorem 5.1, as stated here, is referred to as the Fritz John theorem. We shall adopt this terminology and refer to Theorem 5.1 as the *FJ theorem* and call the multipliers (λ_0, λ, k) *FJ multipliers*.

The next example, due to Fiacco and McCormick [1968], illustrates the use of Theorem 5.1 and will be used to illustrate some subsequent theorems.

Example 4F. Let $X_n = \mathbb{R}^2$, let

$$f(x) = (x_1 - 1)^2 + x_1^2,$$

and let

$$g(x) = 2kx_1 - x_1^2 \leq 0, \quad k \geq 0.$$

For each value of $k \geq 0$ we obtain a programming problem π_k . The feasible set X_k for problem π_k is the set of points on and exterior to the graph of the parabola $x_1^2 = 2kx_1$, and the function to be minimized is the distance from \bar{x} to the point $(1, 0)$. In Figure 4.1 we have sketched some level curves of f and the parabolas for two values of k . Although the solution can be read from Figure 4.1, we will apply the appropriate theorems in this section and in the next section to solve this problem.

By straightforward calculation

$$\nabla f = (2x_1 - 2, 2x_1), \quad \nabla g = (2k, -2x_1).$$

From Theorem 5.1 we get that if a point \bar{x} is a solution to the programming problem π_k , then there exist a $\lambda_0 \geq 0$ and a $\lambda \geq 0$ with $(\lambda_0, \lambda) \neq (0, 0)$ such that (λ_0, λ) and \bar{x} satisfy

$$\lambda_0(2x_1 - 1, 2x_1) + \lambda(2k, -2x_1) = (0, 0) \quad (1)$$

$$\lambda(2kx_1 - x_1^2) = 0. \quad (2)$$

If $\lambda_0 = 0$, then since $\lambda \geq 0$, equations (1) would imply that $k = 0$, contradicting the assumption that $k \geq 0$. Hence we may take $\lambda_0 = 1$. If $\lambda = 0$, then (1) would imply that $x_1 = 1$ and $x_2 = 0$. The point $(1, 0)$ is not feasible for any value of $k \geq 0$, so we may assume that $\lambda \neq 0$ in (2). It then follows from (2) that $2kx_1 - x_1^2 = 0$; that is, the point \bar{x} must lie in the parabola. Thus we may replace (1) and (2) by

$$\lambda(x_1 - 1, x_1) + \lambda(k, -x_1) = (0, 0), \quad x_1^2 = 2kx_1. \quad (3)$$

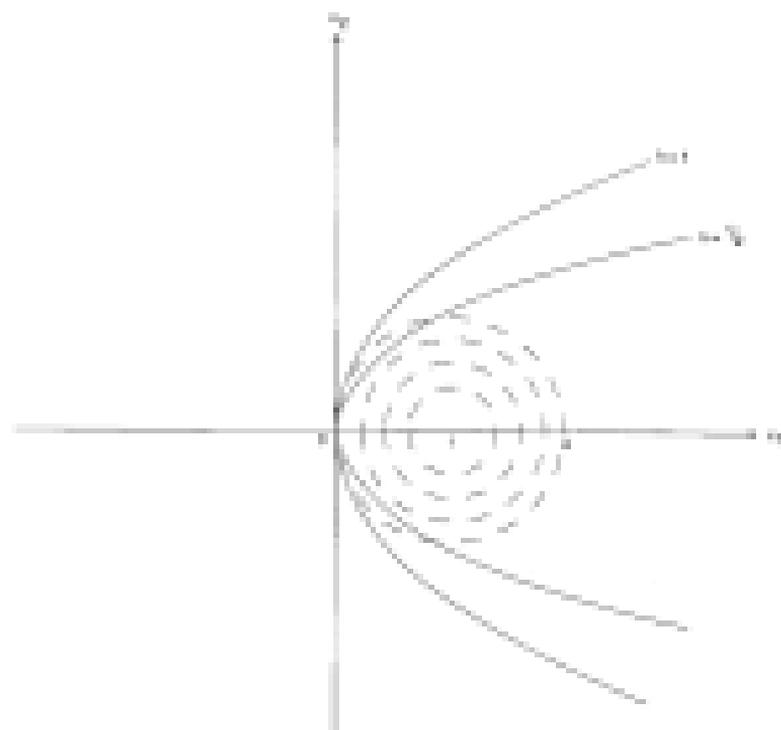


Figure 3.1.

If $x_2 \neq 0$, then we must have $k = 1$, and so

$$x_1 = 1 - k, \quad x_2 = \pm \sqrt{2k(1 - k)}, \quad 0 < k < 1.$$

If $x_2 = 0$, then from the second equation in (1) we get that $x_1 = 0$, and consequently $k = 1/2$.

In summary, we have shown that if $0 < k < 1$, then the only points $\mathbf{x} = (x_1, x_2)$ and multipliers (λ, μ) that satisfy the necessary conditions for problem n_k are

$$\lambda = k, \quad x_1 = 1 - k, \quad x_2 = \pm \sqrt{2k(1 - k)}, \quad (4)$$

$$\lambda = \frac{1}{2}, \quad x_1 = 0, \quad x_2 = 0. \quad (5)$$

If $k \geq 1$, the necessary conditions are only satisfied by (5). Note that if $k = 1$, then (4) and (5) both give $k = 1$ and $x_1 = x_2 = 0$.

We emphasize that Theorem 5.1 is a necessary condition, so that without further assumptions we cannot determine which, if any, of the points (4) and (5) furnish a minimum. In the problem n_0 , we are minimizing the distance from the point $(1, 0)$ to the closed set X_1 . By Lemma 4.1.3 a solution to this problem exists. Since for $k \geq 1$ the only point that satisfies the necessary conditions is the origin, this point is the solution to problem n_0 for $k \geq 1$. If we denote the two points in (4) by x_1 and x_2 , then

$$f(x_1) = f(x_2) = -k^2 + 2k.$$

The point in (5) is the origin, and $f(\bar{x}) = 1$. Since $-k^2 + 2k < 1$ for $k \neq 1$, a solution to problem n_0 for $k \neq 1$ is given by each of the points x_1 and x_2 .

Before we present our proof of Theorem 5.1, we shall present the principal idea of a proof which, with variations, is popular in the literature. We suppose that the equality constraints have been converted into inequality constraints, $h(x) \leq 0$ and $-h(x) \leq 0$, so that the only constraints are $g(x) \leq 0$. By a translation of coordinates, we can assume that $x_0 = 0$ and that $f(x_0) = f(\bar{x}) = 1$. We also assume that $\nabla f(\bar{x}) \neq 0$.

As in the linear programming problem, let

$$E = \{i \mid g_i(\bar{x}) = 0\}, \quad I = \{i \mid g_i(\bar{x}) < 0\}. \quad (6)$$

The constraints $g_i(x) = 0$ with $i \in E$ are said to be active or binding at \bar{x} ; those with $i \in I$ are said to be inactive at \bar{x} .

Since f and g are C^1 , they are differentiable by Theorem 4.1.2. Therefore, since $f(\bar{x}) = 1$ and $g_i(\bar{x}) = 0$,

$$\begin{aligned} f(x) &= \langle \nabla f(\bar{x}), x \rangle + o(\|x\|), \\ g_i(x) &= \nabla g_i(\bar{x})x + w_i(x), \end{aligned}$$

where $o(\|x\|)$ and $w_i(\|x\|)$ represent terms such that $o(\|x\|)/\|x\| \rightarrow 0$ and $w_i(\|x\|)/\|x\| \rightarrow 0$ as $\|x\| \rightarrow 0$. Thus, $o(\|x\|)$ and $w_i(\|x\|)$ represent higher order terms.

Since g is continuous and $g_i(\bar{x}) < 0$, there exists a $\delta > 0$ such that $g_i(x) < 0$ whenever $\|x\| < \delta$. The origin is therefore a solution of the local problem: Minimize $f(x)$ subject to $g_i(x) \leq 0$ and $\|x\| = \delta$. If we drop the terms $o(\|x\|)$ and $w_i(\|x\|)$, that is, if we linearize the problem about the solution \bar{x} and assume that the problem is the linearized problem, then $x_0 = \bar{x}$ is the solution of the linear programming problem:

$$\begin{aligned} \text{Minimize } \langle \nabla f(\bar{x}), x \rangle &= \langle -(-\nabla f(\bar{x})), x \rangle, \\ \text{subject to } \nabla g_i(\bar{x})x &\leq 0, \quad \|x\| \leq \delta. \end{aligned} \quad (7)$$

By Remark 4.3, we can now apply Theorem 4.1 with $J = \nabla g_0(\bar{x})$ and $\mathbf{h} = -\nabla f(\bar{x})$ to the linearized problem and obtain the existence of a vector $\mathbf{v}_1 \neq \mathbf{0}$ in \mathbb{R}^m such that $\mathbf{v}_1 \geq \mathbf{0}$ and

$$\mathbf{v}_1^T \nabla g_0(\bar{x}) = -\nabla f(\bar{x}).$$

If we now set $\lambda = (\mathbf{v}_1, \bar{\lambda}_1)$, then $\lambda \geq \mathbf{0}$ and we can rewrite this equation as

$$\nabla f(\bar{x}) + \mathbf{v}_1^T \nabla g_0(\bar{x}) + \bar{\lambda}_1^T \nabla g_0(\bar{x}) - \nabla f(\bar{x}) + \lambda^T \nabla g(\bar{x}) = \mathbf{0}.$$

The rightmost equation is statement (ii) of Theorem 3.1 with $\lambda_0 = 1$. (Recall that we are taking all the constraints to be inequality constraints.) Since $\lambda_0 = 1$, $(\lambda_0, \lambda) \neq \mathbf{0}$, also

$$\langle \lambda, \mathbf{g}(\bar{x}) \rangle = \langle (\mathbf{v}_1, \bar{\lambda}_1), \mathbf{g}_0(\bar{x}), \mathbf{g}_1(\bar{x}) \rangle = 0,$$

which is (i). Note that if the linearization is valid, then we can get a set of multipliers (λ_0, λ) with $\lambda_0 = 1$.

To complete the proof, one must show that the terms $\omega(\mathbf{x}_k)$ and $\omega(\mathbf{x}_k)$ can indeed be neglected. That is, it must be shown that if the terms represented by $\omega(\mathbf{x}_k)$ and $\omega(\mathbf{x}_k)$ are absent, then $\bar{\mathbf{x}}$ is a solution of (P). This part of the proof becomes quite technical and is carried out in different ways by different authors depending on the additional assumptions that are made. In Exercise 3.10 below the reader will be asked to carry out one such proof, the original proof by Kuhn and Tucker.

Our proof of Theorem 3.1, which we now present, will be a "penalty function" proof and is due to McShane [1977].

Proof. As noted in the preceding discussion, we may assume that $\mathbf{x}_0 = \bar{\mathbf{x}}$ and that $f(\bar{\mathbf{x}}_0) = 0$.

Let E and F be as in (6). To simplify the notation, we assume that $J = \{1, 2, \dots, r\}$ and that $I = \{r+1, \dots, m\}$. Since E or F can be empty, we have $0 \leq r \leq m$.

Let α be a function from $(-\infty, \infty)$ to \mathbb{R}^r such that (i) α is strictly increasing on $(0, \infty)$, (ii) $\alpha(u) = \mathbf{0}$ for $u \leq 0$, (iii) α is C^1 , and (iv) $\alpha(u) = \mathbf{0}$ for $u = 0$. Note that $\alpha(u) > \mathbf{0}$ for $u > 0$.

Since \mathbf{g} is continuous and N_{δ_0} is open, there exists an $\alpha_0 > 0$ such that $\mathbb{R}^r \alpha_0 \subset N_{\delta_0}$ and $\mathbf{g}_0(\mathbf{x}) = \mathbf{0}$ for $\mathbf{x} \in \mathbb{R}^r \alpha_0$.

Define a penalty function F as follows:

$$F(\mathbf{x}, \beta) = f(\mathbf{x}) + \beta \left\{ \sum_{i \in I} \alpha_i(\beta x_i) + \sum_{i \in J} \lambda_i \alpha_i(\beta x_i) \right\}. \quad (8)$$

where $\mathbf{x} \in X_p$ and p is a positive integer. We assert that for each ϵ satisfying $0 < \epsilon < \alpha_p$ there exists a positive integer $p(\epsilon)$ such that for \mathbf{x} with $\|\mathbf{x}\| = \epsilon$

$$F(\mathbf{x}, p(\epsilon)) > 0. \quad (9)$$

We prove the assertion by assuming it to be false and arriving at a contradiction. If the assertion were false, then there would exist an ϵ' with $0 < \epsilon' < \alpha_p$ such that for each positive integer p there exists a point \mathbf{x}_p with $\|\mathbf{x}_p\| = \epsilon'$ and $F(\mathbf{x}_p, p) \leq 0$. Hence, from (8)

$$F(\mathbf{x}_p) + \|\mathbf{x}_p\|^2 \leq -p \left\{ \sum_{i=1}^r \omega_i g_i(\mathbf{x}_p) + \sum_{i=1}^k \lambda_i h_i(\mathbf{x}_p)^2 \right\}. \quad (10)$$

Since $\|\mathbf{x}_p\| = \epsilon'$ and since $\overline{BB}(\epsilon') = \{\mathbf{y} \mid \|\mathbf{y}\| = \epsilon'\}$ is compact, there exist sub-sequences, which we relabel as \mathbf{x}_p and p , and a point \mathbf{x}_0 with $\|\mathbf{x}_0\| = \epsilon'$ such that $\mathbf{x}_p \rightarrow \mathbf{x}_0$. Since f , g , and h are continuous,

$$f(\mathbf{x}_p) \rightarrow f(\mathbf{x}_0), \quad g_i(\mathbf{x}_p) \rightarrow g_i(\mathbf{x}_0), \quad h_i(\mathbf{x}_p) \rightarrow h_i(\mathbf{x}_0).$$

Therefore, if in (10) we divide through by $-p$ and then let $p \rightarrow \infty$, we get

$$0 \geq \sum_{i=1}^r \omega_i g_i(\mathbf{x}_0) + \sum_{i=1}^k \lambda_i [h_i(\mathbf{x}_0)]^2 \geq 0.$$

Hence for each $i = 1, \dots, r$ we have $g_i(\mathbf{x}_0) \leq 0$ and for each $i = 1, \dots, k$ we have $h_i(\mathbf{x}_0) = 0$. Since $\|\mathbf{x}_0\| = \epsilon' < \alpha_p$, we have $g_i(\mathbf{x}_0) = 0$. Thus \mathbf{x}_0 is a feasible point, and so

$$f(\mathbf{x}_0) \geq f^* = 0. \quad (11)$$

On the other hand, from (10) and from $\|\mathbf{x}_p\| = \epsilon'$ we get that $f(\mathbf{x}_p) \leq -\epsilon'^2$, and so

$$f(\mathbf{x}_0) \leq -\epsilon'^2 < 0,$$

which contradicts (11). This proves the assertion.

For each $\epsilon \in (0, \alpha)$ the function $F(\cdot, p(\epsilon))$ is continuous on the closed ball $\overline{BB}(\epsilon)$. Since $\overline{BB}(\epsilon)$ is compact, $F(\cdot, p(\epsilon))$ attains its minimum on $\overline{BB}(\epsilon)$ at some point \mathbf{x}_ϵ with $\|\mathbf{x}_\epsilon\| \leq \epsilon$.

Since $F(\mathbf{x}, p(\epsilon)) = 0$ for \mathbf{x} with $\|\mathbf{x}\| = \alpha$, and since $F(\mathbf{x}, p(\epsilon)) = f(\mathbf{x}) = 0$, it follows that $F(\cdot, p(\epsilon))$ attains its minimum on $\overline{BB}(\epsilon)$ at an interior point \mathbf{x}_ϵ of $\overline{BB}(\epsilon)$. Hence,

$$\frac{\partial F}{\partial x_j}(\mathbf{x}_\epsilon, p(\epsilon)) = 0, \quad j = 1, \dots, n.$$

Calculating $\partial F/\partial x_j$ from (8) gives

$$\begin{aligned} \frac{\partial F}{\partial x_j}(x_j) + \lambda_{r+1} + \sum_{i=1}^r \mu(i) \lambda_i(x_j) \frac{\partial g_i}{\partial x_j}(x_j) \\ + \sum_{i=1}^k \lambda_i \mu(i) h_i(x_j) \frac{\partial h_i}{\partial x_j}(x_j) = 0 \quad \text{for } j = 1, \dots, n. \end{aligned} \quad (12)$$

Define

$$\begin{aligned} \lambda_i(x) &= 1 + \sum_{j=1}^n [\mu(i) \lambda_j(x_j) x_j]^2 + \sum_{j=1}^k [\lambda_j \mu(i) h_j(x_j)]^2, \\ \lambda_p(x) &= \frac{1}{\sqrt{\lambda(x)}}, \\ \lambda_q(x) &= \frac{\mu(i) \lambda_j(x_j) x_j}{\sqrt{\lambda(x)}}, \quad i = 1, \dots, r, \\ \lambda_r(x) &= \lambda_{r+1}, \quad i = r + 1, \dots, m_1, \\ \lambda_{k+i}(x) &= \frac{\lambda_i \mu(i) h_i(x_j)}{\sqrt{\lambda(x)}}, \quad i = 1, \dots, k. \end{aligned}$$

Note that

$$\begin{aligned} \text{(I)} \quad \lambda_p(x) &> 0, \\ \text{(II)} \quad \lambda_q(x) &\geq 0, \quad i = 1, \dots, r, \\ \text{(III)} \quad \lambda_r(x) &= \lambda_{r+1}, \quad i = r + 1, \dots, m_1, \\ \text{(IV)} \quad \|\lambda_p(x), \lambda_q(x), \mu(x)\| &= 1, \end{aligned} \quad (13)$$

where $\lambda(x) = (\lambda_1(x), \dots, \lambda_{m_1}(x))$ and $\mu(x) = (\mu_1(x), \dots, \mu_k(x))$.

If we divide through by $\sqrt{\lambda(x)}$ in (12), we get

$$\lambda_p(x) \frac{\partial F}{\partial x_j}(x_j) + \frac{\lambda_{r+1}}{\sqrt{\lambda(x)}} + \sum_{i=1}^r \lambda_q(x) \frac{\partial g_i(x_j)}{\partial x_j} + \sum_{i=1}^k \lambda_{k+i}(x) \frac{\partial h_i(x_j)}{\partial x_j} = 0. \quad (14)$$

Now let $x \rightarrow 0$ through a sequence of values x_n . Then since $\|x_n\| = \alpha_n$, we have that

$$\alpha_n \rightarrow 0. \quad (15)$$

Since the vectors $(\lambda_p(x), \lambda_q(x), \mu(x))$ are all unit vectors, [see (13)], there exists a subsequence, that we again denote as x_n , and a unit vector $(\lambda_p, \lambda_q, \mu)$ such that

$$(\lambda_p(x_n), \lambda_q(x_n), \mu(x_n)) \rightarrow (\lambda_p, \lambda_q, \mu). \quad (16)$$

Since $(\lambda_p, \lambda_q, \mu)$ is a unit vector, it is different from zero.

From (14), (15), (16), and the continuity of the partials of f , g , and h , we get

$$\lambda_0 \frac{\partial f}{\partial x_j}(\bar{x}) + \sum_{i=1}^r \lambda_i \frac{\partial g_i(\bar{x})}{\partial x_j} + \sum_{i=1}^m \mu_i \frac{\partial h_i(\bar{x})}{\partial x_j} = 0. \quad (17)$$

From (14) and (15) we see that $\lambda_i > 0$ for $i = 0, 1, \dots, r$ and $\lambda_i = 0$ for $i = r + 1, \dots, m$. Thus, $\lambda_0 > 0$ and $\lambda > \mathbf{0}$. Since $g_i(\bar{x}) = 0$ for $i = 1, \dots, r$ and $h_i = 0$ for $i = r + 1, \dots, m$, we have that $\lambda_i g_i(\bar{x}) = 0$ for $i = 1, \dots, m$. Also, we can take the upper limit in the second term in (17) to be m and write

$$\lambda_0 \frac{\partial f}{\partial x_j}(\bar{x}) + \sum_{i=1}^m \lambda_i \frac{\partial g_i(\bar{x})}{\partial x_j} + \sum_{i=1}^m \mu_i \frac{\partial h_i(\bar{x})}{\partial x_j} = 0.$$

This completes the proof of the theorem.

We now present conditions that ensure that $\lambda_0 = 0$ cannot occur in Theorem 3.1. Such conditions are called *constraint qualifications* and are statements about the geometry of the feasible set in a neighborhood of x_0 . There are many constraint qualifications in the literature. We have selected two that are relatively simple and easy to apply. The second is easier to apply, but the first, given in Definition 3.2 below, will expose more points at which $\lambda_0 = 0$ cannot occur. All of the constraint qualifications that are appropriate to problem P1 have the defect that, in general, they require knowledge of the point x_0 rather than requiring knowledge of global properties of the functions involved. The reader who is interested in other constraint qualifications and the relationships among them is referred to *Bonnans, Denari, and Sibiry [1993]* and *Mangasarian [1985]*.

Definition 3.2. The functions g and h satisfy the *constraint qualification (CQ)* at a feasible point x_0 if

- (i) the vectors $\nabla g_1(x_0), \dots, \nabla g_r(x_0)$ are linearly independent and
- (ii) the system

$$\nabla g_i(x_0)x < 0, \quad \nabla h_i(x_0)x = 0, \quad (18)$$

has a solution x in \mathbb{R}^n . Here, $\mathcal{I} = \{i \mid g_i(x_0) = 0\}$.

If the equality constraints are absent, then the statements involving $\nabla h_i(x_0)$ are to be deleted.

Condition (i) of the constraint qualification says that "to first order" we can move from x_0 into the feasible set by moving along the surface defined by the equality constraints $g(x) = 0$ and into the interior of the set defined by $g(x) < 0$, $h(x) = 0$. It turns that this implies that we can actually move from x_0 into the

feasible set in this fashion. We shall prove this in the corollary to Lemma 5.1 below.

Note that CQ does not impose any condition on the point x_0 other than that it be feasible. Also note that $z = 0$ cannot be a solution of (18).

The constrained qualification CQ of Definition 5.1 was introduced by Mangasarian and Fromowitz [1967].

THEOREM 5.2. Let x_0 be a solution of problem P1 and let CQ hold at x_0 . Then $\lambda_0 > 0$ and there exists a vector $\lambda > 0$ in \mathbb{R}^m and a vector g in \mathbb{R}^n such that

- (i) $\langle g, p(x_0) \rangle = 0$ and
 (ii) $\nabla f(x_0) + \lambda^T \nabla g(x_0) + p^T \nabla h(x_0) = 0$.

We shall refer to the necessary conditions of Theorem 5.2 as the Karush-Kuhn-Tucker (KKT) conditions and to the multipliers of Theorem 5.2 as the KKT multipliers.

We now prove the theorem. First we will show that if CQ holds at x_0 , then λ_0 in Theorem 5.1 cannot be zero. Hence if (λ_0, λ, g) is as in Theorem 5.1, so is $(\lambda, \lambda, g, p(x_0))$ and Theorem 5.2 follows. Note that, since we have $\lambda_0 > 0$, then (λ_0, λ, g) is in $\text{int } \Theta_{\lambda_0, \Theta_1}$.

We now show that $\lambda_0 = 0$ cannot occur. If $\lambda_0 = 0$ did occur, then by (ii) of Theorem 5.1

$$\lambda^T \nabla g(x_0) + p^T \nabla h(x_0) = 0. \quad (19)$$

If $\lambda = 0$, then $g = 0$, since $(\lambda_0, \lambda, g) = (0, 0, 0) \notin \Theta$. But then $p^T \nabla h(x_0) = 0$, contradicting the linear independence of $\nabla h_1(x_0), \dots, \nabla h_k(x_0)$. Thus, $\lambda \neq 0$.

Since $\lambda = (\lambda_1, \lambda_2)$ and $\lambda_2 = \Theta_1$, we have that $\lambda_1 \neq 0$. We may therefore rewrite (19) as

$$\lambda_1^T \nabla g_1(x_0) + p^T \nabla h(x_0) = 0. \quad (20)$$

Multiplying on the right by a vector x that is a solution of (18) gives

$$\lambda_1^T \nabla g_1(x_0)x = 0.$$

Since $\lambda_1 \neq 0$, $\lambda_1 \in \mathbb{R}$, and $\nabla g_1(x_0)x \neq 0$, we have a contradiction. Thus $\lambda_0 \neq 0$.

COROLLARY 1. Let x_0 be a solution of problem P1 such that the set J defined in (9) is $\{1, \dots, r\}$ and such that the vectors

$$\nabla h_1(x_0), \dots, \nabla h_r(x_0), \nabla h_{r+1}(x_0), \dots, \nabla h_k(x_0) \quad (21)$$

are linearly independent. Then the conclusion of Theorem 5.2 holds.

Proof. If $\lambda_0 = 0$ in Theorem 5.1, then (20) holds. Since the vectors in (21) are linearly independent, this implies that $\nabla g_0(\mathbf{x}_0) = \mathbf{0}$. Hence $\nabla g_0(\mathbf{x}_0) = \mathbf{0}$, which cannot occur, and so $\lambda_0 > 0$.

We leave to the reader the modifications of the arguments in the proofs of Theorem 5.2 and Corollary 1 if the equality constraints are absent.

Remark 5.6. We noted in Remark 3.3 that the necessary condition (i) of Theorem 5.2 is equivalent to the m equations

$$\lambda_i g_i(\mathbf{x}_0) = 0, \quad i = 1, \dots, m. \quad (27)$$

Condition (ii) in component form gives the system of n equations

$$\frac{\partial f}{\partial x_i}(\mathbf{x}_0) + \sum_{j=1}^m \lambda_j \frac{\partial g_j(\mathbf{x}_0)}{\partial x_i} + \sum_{k=1}^l \mu_k \frac{\partial h_k}{\partial x_i} = 0, \quad i = 1, \dots, n. \quad (28)$$

Since \mathbf{x}_0 is a solution, it is feasible, and so we obtain an additional k equations

$$h_j(\mathbf{x}_0) = 0, \quad j = 1, \dots, k.$$

Thus we have a system of $m + n + k$ equations in the $m + n + k$ unknowns

$$\lambda = (\lambda_1, \dots, \lambda_m), \quad \mu = (\mu_1, \dots, \mu_k), \quad \mathbf{p} = (p_1, \dots, p_n).$$

Remark 5.7. If the inequality constraints are absent and we only have the equality constraints $\mathbf{h}(\mathbf{x}) = \mathbf{0}$, then CQ_2 becomes " $\nabla h_1(\mathbf{x}_0), \dots, \nabla h_k(\mathbf{x}_0)$ are linearly independent." The constraint in Corollary 1 reduces to the same statement. Under these circumstances, the multiplier (λ, μ) is unique. This follows from the observation that if (λ', μ') with $\mu' \neq \mu$ is another multiplier, then $(\lambda, \mu - \mu')$ will be a multiplier. This, however, is impossible since $\lambda_0 = 0$ cannot be. A similar argument shows that if all the inequality constraints $g_j(\mathbf{x}) \leq 0$ are inactive at \mathbf{x}_0 , then the multiplier $(\lambda, \mu, \mathbf{p})$ is unique.

In Remark 5.7 we have established the form of the Lagrange multiplier rule that is usually given in multivariate calculus courses. We state it as a corollary to Theorem 5.1 and 5.2.

COROLLARY 2. Let \mathbf{x}_0 be a solution to the problem of minimizing $f(\mathbf{x})$ subject to the equality constraints $\mathbf{h}(\mathbf{x}) = \mathbf{0}$ and such that the vectors $\nabla h_1(\mathbf{x}_0), \dots, \nabla h_k(\mathbf{x}_0)$ are linearly independent. Then there exists a unique vector \mathbf{p} in \mathbb{R}^n such that

$$\nabla f(\mathbf{x}_0) + \mathbf{p}^T \nabla \mathbf{h}(\mathbf{x}_0) = \mathbf{0}$$

or, in component form,

$$\frac{\partial f}{\partial x_j}(x_0) + \sum_{i=1}^k \lambda_i \frac{\partial g_i(x_0)}{\partial x_j} = 0, \quad j = 1, \dots, n.$$

The following example illustrates CQ by showing what happens when it is not fulfilled. Also, the vectors in (2) are linearly dependent in this example.

Example 4.1. Let $N_0 = \mathbb{R}^2$, let $f(x) = -x_2$, and let

$$g(x) = \begin{pmatrix} -x_2 \\ -x_2 \\ x_1 + (x_2 - 1)^2 \end{pmatrix} \in \begin{pmatrix} \mathbb{R}^1 \\ \mathbb{R}^1 \\ \mathbb{R}^1 \end{pmatrix}.$$

The feasible set N is sketched in Figure 4.2.

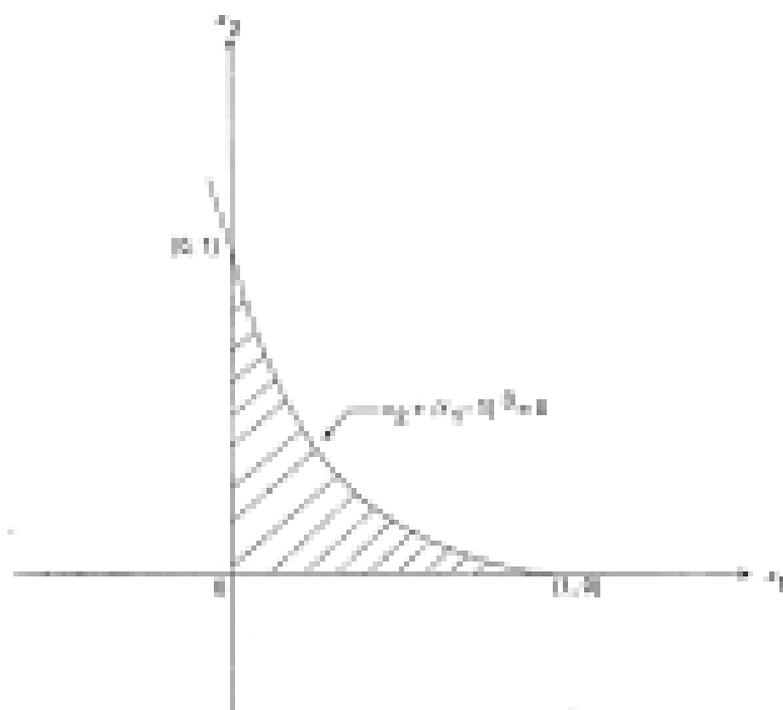


Figure 4.2.

The minimum is attained at $\mathbf{x}_0 = (1, 0)$, and so

$$\mathbf{g}_0(\mathbf{x}) = \begin{pmatrix} -x_2 \\ x_1 + (x_1 - 1)^2 \end{pmatrix}.$$

Hence

$$\nabla \mathbf{g}_0(\mathbf{x}_0)\mathbf{x} = \begin{pmatrix} 0 & -1 \\ 1 & 2x_1 \end{pmatrix} \begin{pmatrix} x_1 \\ x_2 \end{pmatrix} = \begin{pmatrix} -x_2 \\ x_1 + 2x_1^2 \end{pmatrix} < \begin{pmatrix} 0 \\ 0 \end{pmatrix}$$

has no solution in \mathbb{R}^2 . Thus neither CQ nor the linear independence of the vectors in (21) holds at $\mathbf{x}_0 = (1, 0)$.

The necessary condition with $\lambda_0 = 1$ would imply

$$\nabla f(\mathbf{x}_0) + \lambda_0 \lambda_1 \begin{pmatrix} 0 & -1 \\ 1 & 1 \end{pmatrix} = (0, 0)$$

or

$$(-1, 0) + (0, -\lambda_1 + \lambda_1) = (0, 0),$$

which is impossible. On the other hand, if we took $\lambda_0 = 0$, then any (λ_1, λ_2) with $\lambda_1 = \lambda_2 \neq 0$ would satisfy the conditions of Theorem 5.1.

Example 3.3 In this example we consider a problem with equality constraints only in which CQ does not hold at the minimum point. Let $\mathbf{X}_0 = \mathbb{R}^2$, let $f(\mathbf{x}) = 2x_1^2 - x_2^2$, and let $h(\mathbf{x}) = x_1^2(x_1 - x_2)$. The constraint set \mathcal{X} consists of the lines $x_2 = 0$ and $x_2 = \pm x_1$. The level curves of f are the two branches of the hyperbola $2x_1^2 - x_2^2 = c$ and the asymptotes of this family, the lines $x_2 = \pm x_1\sqrt{2}$. See Figure 4.5.

From the figure we see that the point $\mathbf{x}_0 = (0, 0)$ furnishes a minimum. Since $\nabla f(\mathbf{x}) = (4x_1, -2x_2)$ and $\nabla h(\mathbf{x}) = (2x_1x_2, x_1^2 - 2x_2^2)$, we get that

$$\nabla f(\mathbf{x}_0) = \nabla f(0, 0) = (0, 0), \quad \nabla h(\mathbf{x}_0) = (0, 0).$$

Since $\nabla h(\mathbf{x}_0) = \mathbf{0}$, the constraint qualification CQ does not hold at \mathbf{x}_0 . Any (λ_0, μ) of the form $(0, \mu)$, μ arbitrary, will serve as a Lagrange multiplier, as will any $(\lambda, 0)$, λ arbitrary.

Example 3.4 (Continued). In this example CQ holds at a point \mathbf{x}_0 , $\nabla_{\mathcal{P}} f(\mathbf{x}_0) = \mathbf{0}$ and there exists a solution \mathbf{u} to the equation

$$\langle \nabla f(\mathbf{x}_0), \mathbf{u} \rangle = (2\lambda_0 - 2x_{01}) \langle \mathbf{u}, \mathbf{e}_1 \rangle + \lambda_0 \langle \mathbf{u}, \mathbf{e}_2 \rangle = 0.$$

Clearly, $\mathbf{u} = (-1, 0)$ is a solution. Thus CQ holds at all boundary points of the constraint set. This is consistent with our being able to rule out $\lambda_0 = 0$ in the outset in our first treatment of this example.

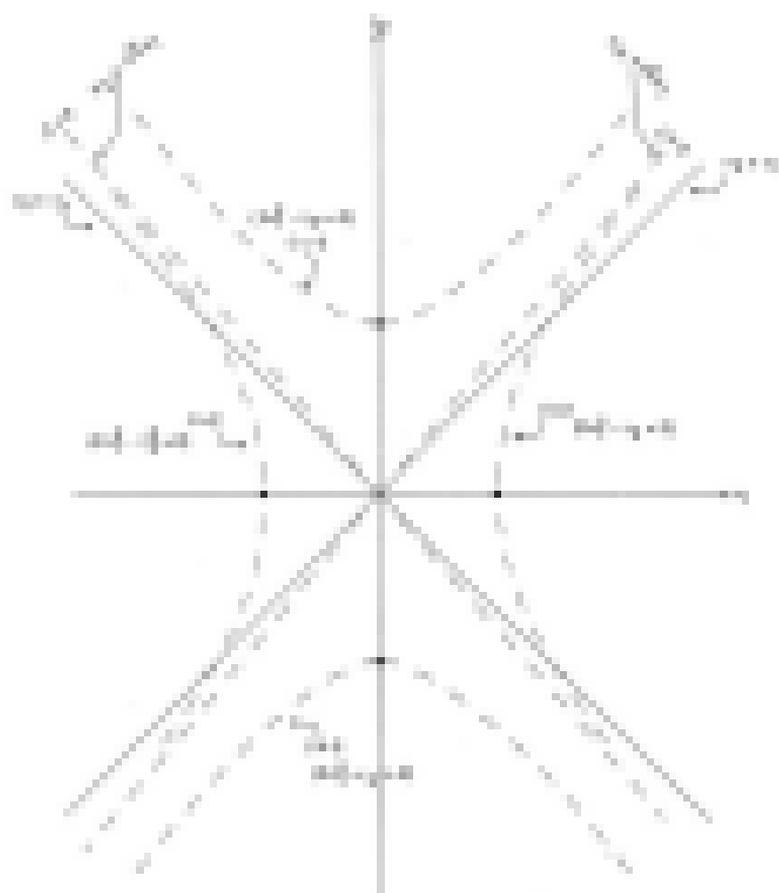


FIGURE 4.4.

The next theorem is a sufficiency theorem which states that if X_{α_0} , f , and g are convex and if the KKT necessary conditions hold, then α_0 is a solution.

Theorem 4.5. Let f and g be as in the statement of problem P1 and let X_{α_0} , f , and g be convex. Let $\alpha_0 \in \mathbb{R}^n$ and $\lambda \in \mathbb{R}^m$ be such that

$$(i) \quad g(\alpha_0) = \lambda^T b$$

$$(ii) \quad \lambda \geq 0.$$

- (iii) $\langle \lambda, g(x_0) \rangle = 0$, and
 (iv) $\nabla f(x_0) + \lambda^T \nabla g(x_0) = 0$.

Then x_0 is a solution of problem *PT*.

Remark 5.8. The sufficiency theorem can also be applied to problems with equality constraints $h(x) = 0$ provided the function h is affine, that is, has the form $h = Ax + b$. For then we can write the equality constraint $h(x) = 0$ as a pair of inequality constraints $h(x) \leq 0$ and $-h(x) \leq 0$, with both h and $-h$ convex. The function g in the theorem is to be interpreted as incorporating this pair of inequality constraints.

Remark 5.9. Under the assumptions of the theorem the feasible set X is convex. To see this, let $X_i = \{x \mid x \in X_0, g_i(x) \leq 0\}$. Since each g_i is convex, each set X_i is convex. Since $X = \bigcap_{i=1}^m X_i$, we get that X is convex.

Proof. Condition (i) states that $x_0 \in X$. Since f is convex,

$$f(x) - f(x_0) \geq \langle \nabla f(x_0), x - x_0 \rangle, \quad x \in X.$$

Substituting (iv) into this inequality gives

$$\begin{aligned} f(x) - f(x_0) &\geq \langle \lambda, -\nabla g(x_0), x - x_0 \rangle \\ &= -\langle \lambda, \nabla g(x_0)(x - x_0) \rangle, \quad x \in X. \end{aligned} \quad (21)$$

Since g is convex,

$$g(x) - g(x_0) \geq \nabla g(x_0)(x - x_0).$$

Therefore, since $\lambda \geq 0$,

$$\langle \lambda, g(x) - g(x_0) \rangle \geq \langle \lambda, \nabla g(x_0)(x - x_0) \rangle.$$

Substituting this inequality into (21) and then using (iii) give

$$f(x) - f(x_0) \geq \langle \lambda, g(x_0) - g(x) \rangle = -\langle \lambda, g(x) \rangle.$$

Since $\lambda \geq 0$ and since for x in X we have $g(x) \leq 0$, it follows that $-\langle \lambda, g(x) \rangle \geq 0$. Hence $f(x) \geq f(x_0)$ for all x in X , and so x_0 is a solution to problem *PT*.

The function g in Example 5.1 is not convex, as a calculation of the Hessian matrix of g shows. Therefore Theorem 5.3 cannot be used to analyze Example 5.1.

Exercise 5.1. State and prove the form of Theorem 5.3 that results when affine equality constraints $h(x) = 0$ are also present.

Exercise 5.2. Use the following to complete the exercise:

- Sketch the feasible set.
- Use an appropriate theorem to show that a solution exists.
- Find all critical points and associated multipliers.
- Without using the second-order criteria of the next section determine, if possible, whether a given critical point furnishes a solution.

(i) Minimize $f(x) = x_1^2 + (x_2 + 1)^2$ subject to

$$x_1^2 + x_2^2 \leq 4, \quad x_2^2 \leq x_1 - 1, \quad -2x_1 \leq x_1 - 1.$$

(ii) Minimize $f(x) = x_1^2 + x_2^2$ subject to

$$x_1^2 + x_2^2 \leq 4, \quad x_1^2 \geq x_2, \quad -2x_1 \leq x_2 - 1, \quad -2x_2 \geq x_2 - 2.$$

(iii) Minimize $(x_1 - 3)^2 + x_2^2$ subject to

$$x_1^2 + x_2^2 \leq 4, \quad x_2^2 = 2x_1 + x_2^2 \leq 4, \quad x_1 + x_1^2 \leq 2.$$

(iv) Minimize $f(x) = \|x\|^2$, where $x \in \mathbb{R}^2$, subject to

$$x_1 x_2 \geq 1, \quad x_1 \geq 0, \quad x_2 \geq 0.$$

(v) Maximize $f(x) = x_1^2 + x_2^2 + ax_1 x_2$, $a > 0$, subject to $x_1^2 + x_2^2 = 1$. What can you say about minimizing $f(x)$ subject to $x_1^2 + x_2^2 = 1$. [In the minimization problem you need not answer questions (i) and (ii).]

Exercise 5.3. Minimize $f(x) = -(x_1 - 2)^2 + (x_2 - 1)^2$ subject to

$$x_1^2 - x_1 \leq 0, \quad x_1 + x_2 - 4 \leq 0, \quad -x_1 \leq 0, \quad -x_2 \leq 0.$$

Give a full justification of your answer. You may use geometric considerations to assist you in finding critical points.

Exercise 5.4. Minimize $f(x) = x_1^2 + a_1 x_1 + x_2^2$ subject to

$$-2x_1 - 2x_2 + 4 \leq 0, \quad -x_1 + x_2 - 3 \leq 0, \quad 4x_1 - 2x_2^2 + 4a_1 - 1^2 \leq 25.$$

Hint: Compare with Exercise 5.2.

Exercise 5.4. Minimize $f(\mathbf{x}) = e^{-2x_1 + 3x_2}$ subject to $x_1^2 + x_2^2 \leq 20$.

Exercise 5.5. Minimize $f(\mathbf{x}) = \sum_{i=1}^n c_i x_i$ subject to

$$\sum_{i=1}^n a_i x_i = b, \quad x_i \geq 0, \quad i = 1, \dots, n,$$

where $c_i, a_i, i = 1, \dots, n$, and b are positive constants.

Exercise 5.6. Let A be a positive-definite $n \times n$ symmetric matrix and let \mathbf{y} be a fixed vector in \mathbb{R}^n . Show that the maximum value of $f(\mathbf{x}) = \langle \mathbf{y}, \mathbf{x} \rangle$ subject to the constraint $\langle \mathbf{x}, A\mathbf{x} \rangle \leq 1$ is $\langle \mathbf{y}, A^{-1}\mathbf{y} \rangle^{1/2}$. Use this result to establish the generalization of the Cauchy-Schwarz inequality

$$|\langle \mathbf{x}, \mathbf{y} \rangle|^2 \leq \langle \mathbf{x}, A\mathbf{x} \rangle \langle \mathbf{y}, A^{-1}\mathbf{y} \rangle.$$

Exercise 5.7. Use the Lagrange multiplier rule to show that the distance $\rho(\mathbf{x}_0)$ from the point \mathbf{x}_0 to the hyperplane $H_c = \{\mathbf{x} : \langle \mathbf{x}, \mathbf{n} \rangle = c\}$ is given by

$$\rho(\mathbf{x}_0) = \left| \frac{\langle \mathbf{x}_0, \mathbf{n}_0 \rangle - c}{\|\mathbf{n}\|} \right|.$$

(*Hint: In problems involving the minimization of a norm or of a distance it is often simpler to minimize the square of the norm or distance.*)

Exercise 5.8. In this problem we will use the Lagrange multiplier rule to discuss important properties of eigenvectors and eigenvalues of a real $n \times n$ symmetric matrix A .

(a) Show that the maximum of $f(\mathbf{x}) = \langle \mathbf{x}, A\mathbf{x} \rangle$ subject to $\|\mathbf{x}\| = 1$ is attained at a unit eigenvector \mathbf{x}_1 and that $f(\mathbf{x}_1)$ is the largest eigenvalue λ_1 of A .

(b) Suppose $m < n$ mutually orthogonal real eigenvectors $\mathbf{x}_1, \dots, \mathbf{x}_m$ of A have been determined. Show that the minimum of $f(\mathbf{x}) = \langle \mathbf{x}, A\mathbf{x} \rangle$ subject to $\|\mathbf{x}\| = 1$ and $\langle \mathbf{x}, \mathbf{x}_j \rangle = 0, j = 1, \dots, m$, is attained at a point \mathbf{x}_{m+1} that is an eigenvector of A and that $f(\mathbf{x}_{m+1})$ is the corresponding eigenvalue.

(c) Use (a) and (b) to show that a real $n \times n$ symmetric matrix has a set of n orthonormal eigenvectors $\mathbf{x}_1, \dots, \mathbf{x}_n$ with corresponding eigenvalues $\lambda_1 \geq \lambda_2 \geq \dots \geq \lambda_n$.

Exercise 5.9. In this exercise we will use Farkas's lemma and a constraint qualification introduced by Karh and Tucker [1991] that is different from CQ to establish the conclusion of Theorem 5.2 without using Theorem 5.1. The Karh-Tucker constraint qualification KTQ is said to hold at a point \mathbf{x}_0 if there

exists a vector α such that

$$\nabla_{\alpha} g_j(x) \leq 0, \quad \forall j, \alpha_j = 0$$

and if for each such α there exists a C^2 function $\eta(\cdot)$ defined on an interval $[0, \tau]$ such that

$$\begin{aligned} \eta(0) &= x_0, & \eta'(0) &= \alpha, \\ g_j(\eta(t)) &\leq 0, & h_j(\eta(t)) &= 0 \quad \text{for all } t \in [0, \tau]. \end{aligned}$$

The Karlin-Tucker constraint qualification is difficult to check. In Lemma 5.1 we give a verifiable condition that implies KTQ.

Let x_0 be a solution of problem PI and let KTQ hold at x_0 . Show that the conclusion of Theorem 5.2 holds without invoking Theorem 5.1. *Hint:* What can you say about the right-hand derivatives at $t=0$ of the mappings $t \rightarrow \lambda_j^*(t)$, $t \rightarrow g_j(\eta(t))$, and $t \rightarrow h_j(\eta(t))$ in relation to Farkas's theorem?

6. SECOND-ORDER CONDITIONS

In this section we first obtain second-order necessary conditions. Points that satisfy the first-order necessary conditions but fail to satisfy the second-order conditions cannot solve the programming problem. We conclude this section with a presentation of some second-order sufficiency conditions. Points that satisfy the first-order necessary conditions and these sufficiency conditions are solutions to the programming problem.

Since second-order conditions involve second derivatives, in this section we assume that the functions f , g , and h that occur in the statement of problem PI are of class C^2 on X_0 .

The derivation of our second-order necessary condition will require an application of the implicit function theorem, which we now discuss.

Consider the system of equations

$$\begin{aligned} \alpha_1 f(x_1, \dots, x_n, y_1, \dots, y_m) &= 0, \\ &\vdots \\ \alpha_m f(x_1, \dots, x_n, y_1, \dots, y_m) &= 0, \end{aligned} \tag{1}$$

where the real-valued functions α_i , $i = 1, \dots, m$, are defined on an open set $D = D_x \times D_y$ where D_x is open in the x -space \mathbb{R}^n and D_y is open in the y -space \mathbb{R}^m . Equation (1) asks us to find all (x, y) in D such that (1) is satisfied. For a given x in D_x , the system (1) can therefore be interpreted as a system of m equations for the m unknowns (y_1, \dots, y_m) . If for each x in some subset I of D_x we can find a unique solution $y(x) = (y_1(x), \dots, y_m(x))$, then in this fashion we obtain a function $y(\cdot)$ with domain I . In this event, we say that (1) defines y as a function of x implicitly.

Let (x_0, y_0) be a point in D that satisfies (1). The implicit function theorem gives conditions ensuring that (1) has a unique solution for points x sufficiently close to x_0 . The resulting function $y(x)$ will inherit the regularity properties of the functions g_1, \dots, g_m .

To illustrate these ideas, take $D_1 = \mathbb{R}^2$, $D_2 = \mathbb{R}^1$, and $D = \mathbb{R}^3 = \mathbb{R}^2 \times \mathbb{R}^1$. Consider

$$w(x, y) = x_1^2 + x_2^2 + y^2 - 1 = 0. \quad (2)$$

The set of points $(x, y) = (x_1, x_2, y)$ in \mathbb{R}^3 that satisfy (2) is the surface of the ball with center at the origin and radius 1.

Given any point $(x_0, y_0) = (x_{01}, x_{02}, y_0)$ with $y_0 > 0$ that satisfies (2), we can find an $\alpha > 0$ and a $\beta > 0$ such that

$$y = \sqrt{1 - x_1^2 - x_2^2}$$

is the unique solution of (2) for x in $B(x_0, \alpha)$ having range in $B(y_0, \beta)$. The point $(x_0, -y_0)$ also satisfies (2), and

$$y = -\sqrt{1 - x_1^2 - x_2^2}$$

is the unique solution of (2) for x in $B(x_0, \alpha)$ having range in $B(-y_0, \beta)$.

Although all points of $\mathcal{S} = \{(x, y) \mid x_1^2 + x_2^2 + y^2 = 1, y \neq 0\}$ are solutions of (2), equation (2) does not define y as a function of x on any ball about a point x_0 such that $\|x_0\| = 1$. Since $dy/dy = 2y$, we see that at any point (x_0, y_0) that satisfies (2) we have $dy/dy \neq 0$ if $y_0 \neq 0$ and $dy/dy = 0$ if $y_0 = 0$. Thus at all points of \mathcal{S} , $dy/dy \neq 0$. The reader should return to this paragraph after reading Theorem 6.1.

Equation (1) in vector form is

$$w(x, y) = 0.$$

THEOREM 6.1 (IMPLICIT FUNCTION THEOREM). Let D_1 be an open set in \mathbb{R}^n and let D_2 be an open set in \mathbb{R}^m . Let $D = D_1 \times D_2$. Let

$$w(x, y) = 0, \quad x \in D_1, y \in D_2$$

be a mapping that is of class C^k on D and with range in \mathbb{R}^m . Let (x_0, y_0) , $x_0 \in D_1$, $y_0 \in D_2$, be a solution of

$$w(x, y) = 0$$

and suppose that the $m \times m$ matrix

$$J^2(w_{ij}, x_0) = \left(\frac{\partial^2 w_{ij}}{\partial x_j^2} (x_0, F_{ij}) \right), \quad i = 1, \dots, m, \quad j = 1, \dots, m, \quad (3)$$

is nonsingular. Then there exist an $\alpha > 0$, a $\beta > 0$, and a function $g(x)$ defined on $B(x_0, \alpha) \cap D_1$ with range in $B(F_{ij}, \beta) \cap E_1$ such that

$$g(\cdot) \text{ is of class } C^2, \quad (4)$$

$$w(x, g(x)) = 0 \quad \text{for all } x \in B(x_0, \alpha). \quad (5)$$

Moreover, for x in $B(x_0, \alpha)$, $g(x)$ is the only solution of $w(x, y) = 0$ such that $y(x) \in B(F_{ij}, \beta)$.

For a proof of this theorem we refer the reader to Stadin [1976] or Fleming [1977].

Definition 4.1. Let x_0 be a point of X_0 . The vector $z \in \mathbb{R}^r$ is said to satisfy the tangential constraints at x_0 if

$$\nabla_{E_j}(x_0)z \leq 0, \quad \forall j(x_0)z = 0, \quad (6)$$

where $E = \{j(x_0)z = 0\}$.

For definiteness we shall henceforth suppose that $E = \{1, \dots, r\}$. Let $F = \{j \in \{1, \dots, m\} : E \cap j(x_0)z = 0\}$. Since either E or F may be empty, $0 \leq r \leq m$.

Let z be a vector that satisfies the tangential constraints. Let

$$E_1 = \{j \in E : \langle \nabla_{E_j}(x_0), z \rangle = 0\}$$

$$E_2 = \{j \in E : \langle \nabla_{E_j}(x_0), z \rangle < 0\}.$$

Then $E = E_1 \cup E_2$.

$$\nabla_{E_1}(x_0)z = 0, \quad \forall j(x_0)z = 0, \quad (7)$$

and

$$\nabla_{E_2}(x_0)z < 0. \quad (8)$$

For the sake of definiteness suppose that $E_1 = \{1, \dots, q\}$. Since E_1 may be empty, $0 \leq q \leq r$. If $q = 0$, statements involving E_1 in (7) in the ensuing discussion should be deleted. The reader should keep in mind that the sets E_1 and E_2 depend on z and x_0 .

The next lemma gives a condition under which, given a vector \mathbf{a} that satisfies the tangential constraints, we can construct a curve $\xi(\cdot)$ emanating from \mathbf{x}_0 and going into the feasible set in the direction given by \mathbf{a} .

Lemma 6.1. Let g and h be of class C^{2p} on \mathbb{R}^n , $p \geq 1$. Let \mathbf{x}_0 be feasible and let \mathbf{z} satisfy the tangential constraints at \mathbf{x}_0 . Suppose that the vectors

$$\nabla g(\mathbf{x}_0), \dots, \nabla g_j(\mathbf{x}_0), \nabla h_1(\mathbf{x}_0), \dots, \nabla h_k(\mathbf{x}_0) \quad (7)$$

are linearly independent. Then there exists a $\epsilon > 0$ and a C^{2p} mapping $\xi(\cdot)$ from $(-\epsilon, \epsilon)$ to \mathbb{R}^n such that

$$\begin{aligned} \xi(0) &= \mathbf{x}_0, & \xi'(0) &= \mathbf{z}, \\ \mathbf{g}_j(\xi(t)) &= 0, & \mathbf{h}(\xi(t)) &= 0, & \xi(t) \in \mathbb{R}^n \\ \mathbf{x}_0 + t\mathbf{z} &= 0, & \text{for } 0 < t < \epsilon. \end{aligned}$$

Proof. Let $p = q + k$. In view of (7), $p < n$. If $p = n$, then again in view of (7) the only solution of (5) is $\mathbf{z} = \mathbf{0}$. Hence, $\mathcal{R}_p = \{\mathbf{0}\}$ and $\mathcal{R}_k = \mathbb{R}^k$. The mapping $\xi(\cdot)$ is then the constant map $\xi(t) = \mathbf{x}_0$.

We now suppose that $p < n$. Let Γ denote the $p \times n$ matrix whose rows are the vectors in (7). Thus

$$\Gamma = \begin{pmatrix} \nabla g_1(\mathbf{x}_0) \\ \vdots \\ \nabla h_k(\mathbf{x}_0) \end{pmatrix}$$

and Γ has rank p . By relabeling the coordinates corresponding to p linearly independent columns of Γ , we may suppose without loss of generality that the first p columns of the matrix Γ are linearly independent. Thus, for $p \times p$ matrix

$$\Gamma_p = \begin{pmatrix} \frac{\partial g_1(\mathbf{x}_0)}{\partial x_1} \\ \vdots \\ \frac{\partial g_p(\mathbf{x}_0)}{\partial x_1} \\ \vdots \end{pmatrix} \quad i = 1, \dots, q; \quad i = 1, \dots, k; \quad j = 1, \dots, p,$$

is nonsingular.

The system of equations

$$\begin{aligned} \mathbf{g}_1(\xi) &= 0, \\ \mathbf{h}(\xi) &= 0, \\ x_{2p+1} - x_{2p+2} - \alpha_{2p+1} &= 0, \\ &\vdots \\ x_n - x_{2p} - \alpha_n &= 0, \end{aligned} \quad (8)$$

where x_{p+1}, \dots, x_n are the last $n-p$ components of the vector \mathbf{x} , has a solution $(\mathbf{x}, t) = (\mathbf{x}_0, 0)$. We shall use the implicit function theorem to show that the system (8) determines \mathbf{x} as a function of t near $t = 0$.

For the system (8), viewed as defining \mathbf{x} implicitly as a function of t , the matrix J defined in (1) becomes

$$J = \begin{pmatrix} \Gamma_p & \Gamma_{n-p} \\ 0_{n-p, p} & I_{n-p} \end{pmatrix}$$

where Γ_{n-p} is the $p \times (n-p)$ matrix whose columns are the last $n-p$ columns of Γ , the matrix $0_{n-p, p}$ is the $(n-p) \times p$ zero matrix, and I_{n-p} is the $(n-p) \times (n-p)$ identity matrix. Since Γ_p is nonsingular, so is J . Hence by Theorem 6.1, there exists a $\tau > 0$ and a C^2 function $\mathbf{g}(\cdot)$ defined on the interval $(-\tau, \tau)$ and with range in \mathbb{R}^p such that

$$\begin{aligned} \mathbf{g}(0) &= \mathbf{x}_0 \\ \mathbf{g}_p(\mathbf{g}(t)) &= \mathbf{0}, & \mathbf{h}(\mathbf{g}(t)) &= \mathbf{0}, \\ \mathbf{g}'(t) &= \mathbf{x}_{p+1} + \mathbf{a}_p, & t &= p + 1, \dots, n. \end{aligned} \quad (9)$$

Differentiation of the equations in (8) with respect to t and then setting $t = 0$ give

$$\begin{pmatrix} \Gamma \\ 0_{n-p, p} \end{pmatrix} \mathbf{g}'(0) = \begin{pmatrix} \mathbf{0} \\ \mathbf{d} \end{pmatrix}$$

where $\mathbf{d} = (d_{p+1}, \dots, d_n)$. The matrix on the left is the matrix J for the system (8) and is nonsingular. Therefore the system

$$\begin{pmatrix} \Gamma \\ 0_{n-p, p} \end{pmatrix} \mathbf{w} = \begin{pmatrix} \mathbf{0} \\ \mathbf{d} \end{pmatrix}$$

has a unique solution. Since $\mathbf{g}'(0)$ and \mathbf{w} are both solutions, we get that $\mathbf{g}'(0) = \mathbf{w}$.

Since $\mathbf{g}_p(\mathbf{x}_0) < 0$ and $\mathbf{g}(0) = \mathbf{x}_0$, it follows from the continuity of \mathbf{g} and \mathbf{g}' that for τ sufficiently small we have $\mathbf{g}_p(\mathbf{g}(t)) < 0$ for $|t| < \tau$.

Since

$$\frac{d\mathbf{g}_p(\mathbf{g}(t))}{dt} = \mathbf{g}'_{\mathbf{g}_p}(\mathbf{g}(t))\mathbf{g}'(t)$$

and since $\mathbf{g}'(0) = \mathbf{x}_p$ and $\mathbf{g}'(0) = \mathbf{w}$, it follows from (9) that $d\mathbf{g}_p(\mathbf{g}(t))/dt < 0$. It then follows from continuity that by taking τ to be smaller, if necessary, there is an interval $(0, \tau)$ on which all of the previous conclusions hold and on which

$d\mathbf{g}_{\mathbf{x}_0}(\mathbf{f})(\mathbf{v}) < 0$. Since $\mathbf{g}_{\mathbf{x}_0}(\mathbf{f})(\mathbf{v}) = \mathbf{g}_{\mathbf{x}_0}(\mathbf{x}_0) = 0$, we have $\mathbf{g}_{\mathbf{x}_0}(\mathbf{f})(\mathbf{v}) < 0$ on (δ, ϵ) . This completes the proof.

COROLLARY 1. Let CQ hold at a point \mathbf{x}_0 . Then for every \mathbf{v} satisfying relation (18) in Section 5 there exists a C^2 function $\mathbf{q}(\cdot)$ such that $\mathbf{q}(0) = \mathbf{x}_0$, $\mathbf{q}'(0) = \mathbf{v}$ and $\mathbf{q}(t)$ is feasible for all t in some interval $(0, \delta)$. Moreover, $\mathbf{g}(\mathbf{q}(t)) = 0$ on $(0, \delta)$.

Proof. A vector \mathbf{v} satisfying relation (18) in Section 5 satisfies the tangential constraints, and the set of indices E_1 corresponding to \mathbf{v} is empty.

Example 5.2 (Continued). At the point $\mathbf{x}_0 = (1, 0)$

$$\mathbf{V}_{\mathbf{g}_0}(\mathbf{x}_0)\mathbf{v} = \begin{pmatrix} 0 & -1 \\ 0 & 1 \end{pmatrix} \begin{pmatrix} v_1 \\ v_2 \end{pmatrix} = \begin{pmatrix} -v_2 \\ v_2 \end{pmatrix}.$$

Therefore, since there are no equality constraints, every vector of the form $(v_1, 0)$ satisfies the tangential constraints. Clearly, there exists no curve $\mathbf{q}(\cdot)$ emanating from $\mathbf{x}_0 = (1, 0)$ with tangent vector $(1, 0)$ at \mathbf{x}_0 such that $\mathbf{q}(t)$ is feasible for t in some interval $(0, \delta)$. Lemma 6.1 is not contradicted, since we have $E_1 = E$ and the rows of $\mathbf{V}_{\mathbf{g}_0}(\mathbf{x}_0)$ are linearly dependent. Also, the conclusion of the corollary is not contradicted since CQ does not hold at $(1, 0)$.

We now present a second-order necessary condition.

Definition 6.2. Let $(\mathbf{x}_0, \lambda, \mu)$ be as in Theorem 5.2. Let \mathbf{v} be a vector that satisfies the tangential constraints (9) at \mathbf{x}_0 . Let

$$E_1 = E_1(\mathbf{v}) = \{i \in E, \langle \nabla g_i(\mathbf{x}_0), \mathbf{v} \rangle = 0\},$$

$$E_2 = E_2(\mathbf{v}) = \{i \in E, \langle \nabla g_i(\mathbf{x}_0), \mathbf{v} \rangle < 0\}.$$

The vector \mathbf{v} will be called a *second-order test vector* if (i) $\lambda_i = 0$ for $i \in E_1$ and (ii) the rows of the matrix

$$\begin{pmatrix} \nabla^2 g_i(\mathbf{x}_0, \lambda) \\ \nabla h_j(\mathbf{x}_0) \end{pmatrix}$$

are linearly independent.

THEOREM 6.2. Let f, g , and h be of class C^2 on X_0 . Let \mathbf{x}_0 be a solution to problem P1. Let λ and μ be KKT multipliers as in Theorem 5.2 and let $H(\cdot, \lambda, \mu)$ be defined by

$$F(\mathbf{x}, \lambda, \mu) = f(\mathbf{x}) + \langle \lambda, \mathbf{g}(\mathbf{x}) \rangle + \langle \mu, \mathbf{h}(\mathbf{x}) \rangle.$$

Then for every second-order test vector λ ,

$$\langle \lambda, F_{xx}(x_{\alpha}, \lambda, p(x)) \rangle \geq 0, \quad (5.2)$$

where

$$F_{xx} = \left(\frac{\partial^2 \mathcal{L}}{\partial x_i \partial x_j} \right) \quad i = 1, \dots, n, \quad j = 1, \dots, n.$$

We first establish an elementary result.

LEMMA 5.2. Let ϕ be a C^2 function defined on an interval $(-\epsilon, \epsilon)$ such that $\phi'(0) = 0$ and $\phi(t) \geq \phi(0)$ for all t in $(0, \epsilon)$. Then $\phi''(0) \geq 0$.

Proof. By Taylor's theorem, for $0 < t < \epsilon$

$$\phi(t) - \phi(0) = \frac{1}{2} \phi''(\theta) t^2,$$

where $0 < \theta < t$. If $\phi''(0)$ were negative, then by continuity there would exist an interval $[0, \epsilon')$ on which ϕ'' would be negative. Hence $\phi(t) < \phi(0)$ on this interval, contradicting the hypothesis.

Proof (Theorem 5.1). Let x be a second-order test vector. Since $\lambda_i = 0$ for $i \in E_2$ and

$$\lambda = (\lambda_{E_1}, \lambda_{E_2}) = (\lambda_{E_1}, \lambda_{E_2} - \lambda_{E_2}) = (\lambda_{E_1}, \mathbf{0}),$$

we have

$$F(x, \lambda, p) = f(x) + \langle \lambda_{E_1}, g_{E_1}(x) \rangle + \langle \lambda_{E_2}, h_{E_2}(x) \rangle.$$

Let $\tilde{Q}(\cdot)$ be the function corresponding to x as in Lemma 5.1, with $x_0 = x_{\alpha}$. Since $g_{E_1}(\tilde{Q}(t)) = 0$ and $h_{E_2}(\tilde{Q}(t)) = 0$ for $|t| < \epsilon$, we have

$$F(\tilde{Q}(t)) = F(\tilde{Q}(t), \lambda, p) \quad \text{for } |t| < \epsilon.$$

Since for $0 < t < \epsilon$, all points $\tilde{Q}(t)$ are feasible, and since $\tilde{Q}(0) = x_{\alpha}$, the mapping $t \rightarrow F(\tilde{Q}(t))$, where $0 < t < \epsilon$ has a minimum at $t = 0$. Hence, so does the mapping $\phi(\cdot)$ defined on $(0, \epsilon)$ by

$$\phi(t) = F(\tilde{Q}(t), \lambda, p).$$

Weight-normalized calculation gives

$$\phi(t) = (\nabla F(x_0), \lambda, \mu) \psi(t)$$

$$\phi'(t) = (\lambda^T, \nabla_{x_0} \phi(x_0), \lambda, \mu \nabla G(x_0)) + (\nabla^2 F(x_0), \lambda, \mu) \psi'(t).$$

Setting $t = 0$ in the first equation and using the relations $\psi(0) = x_0$ and $\nabla F(x_0, \lambda, \mu) = 0$ give $\phi(0) = 0$. Hence by Lemma 6.2 we have $\phi'(0) \geq 0$. Setting $t = 0$ in the second equation and using $\psi(0) = x_0$, $\nabla F(x_0, \lambda, \mu) = 0$, and $\psi'(0) = \alpha$ give the conclusion of the theorem.

The following corollary requires that (10) holds for fewer α than the theorem does and thus is more restrictive than the theorem but easier to apply.

COROLLARY 2. Let the functions f , g , and h be as in Theorem 6.2 and let x_0 be a solution to problem P1. Suppose that the vectors

$$\nabla h_1(x_0), \dots, \nabla h_r(x_0), \nabla h_{r+1}(x_0), \dots, \nabla h_k(x_0)$$

are linearly independent. Then there exist KKT multipliers λ, μ as in Theorem 6.2, and (10) holds for every vector α satisfying

$$\nabla g_j(x_0)\alpha = 0, \quad \nabla h_{r+i}(x_0)\alpha = 0, \quad (11)$$

Proof. The existence of KKT multipliers follows from Corollary 1 of Theorem 6.2. The second conclusion follows from the theorem by taking $E_1 = E$, in which case $E = \{1, \dots, r\}$ and $E_2 = \emptyset$.

If the inequality constraints are absent, then the linear independence hypothesis of Theorem 6.2 and Corollary 2 becomes "the vectors $\nabla h_1(x_0), \dots, \nabla h_r(x_0)$ are linearly independent," which is also the constraint qualification CQ for this problem. Thus, we get the following corollary.

COROLLARY 3. Let f and h be as in Theorem 6.2, let the inequality constraints be absent, let x_0 be a solution of problem P1 and let CQ hold at x_0 . Then there exists a unique vector μ in \mathbb{R}^r such that the function $\tilde{M}(\cdot, \mu)$ defined by

$$\tilde{M}(x, \mu) = f(x) + \langle \mu, h(x) \rangle$$

satisfies

$$(i) \quad \nabla f(x_0) + \mu \nabla h(x_0) = 0$$

and

$$(10) \quad \langle \lambda, \nabla_{x_1} \pi(x) \rangle \geq 0$$

for all x in \mathbb{R}^2 satisfying

$$\nabla h(x) = 0.$$

Corollary 3 is the classical second-order necessary condition for equality-constrained problems.

Example 5.4 (Continued). The reader will recall that for $0 < k < 1$, the first-order necessary conditions were satisfied by three points and corresponding multipliers given by

$$\begin{array}{lll} x_1 = 1 - k, & x_2 = \pm \sqrt{2k(1-k)}, & \lambda = 1 \\ x_1 = k, & x_2 = 0, & \lambda = 1/k. \end{array}$$

We will use the second-order necessary condition as given in Corollary 2 to rule out the origin as a candidate for optimality when $0 < k < 1$.

Recall that for any x we have $\nabla g(x) = (2k, -2x_2)$. At $x_0 = (0, 0)$, $\nabla g(x_0) \neq 0$, so the linear independence requirement of Corollary 2 holds. Thus for λ to satisfy (11) we require that

$$\langle \nabla g(x_0), \lambda \rangle = \langle (2k, 0), \lambda \rangle = 2\lambda_1 = 0.$$

This is satisfied by all vectors (λ_1, λ_2) of the form

$$\lambda_1 = 0, \quad \lambda_2 \text{ arbitrary.} \quad (12)$$

The function F with $\lambda = (0, k)$ becomes

$$F = (x_1 - 1)^2 + x_2^2 + \frac{1}{k}(2kx_1 - x_2^2).$$

Condition (12) with $x_0 = (0, 0)$ becomes

$$\langle \lambda, \lambda_2 \rangle \begin{pmatrix} 2 & 0 \\ 0 & 2(1-k) \end{pmatrix} \begin{pmatrix} x_1 \\ x_2 \end{pmatrix} = 2\lambda_2^2 + 2\lambda_2^2 \left(1 - \frac{k}{k}\right) \geq 0$$

for all (x_1, x_2) satisfying (12). Since $x_1 = 0$ and for $0 < k < 1$ the coefficient of $2\lambda_2^2$ is negative, the last relation only holds for $\lambda_2 = 0$. Thus (12) fails to hold for all (λ_1, λ_2) satisfying (12), and so $x_0 = (0, 0)$ is not optimal if $0 < k < 1$.

For $\lambda \geq 1$ the only critical point is $x_\lambda = (0, 0)$ with $\lambda = 1/k$. We now have $(x, F_\lambda x) \geq 0$ for all x , so the second-order conditions of Theorem 5.2 and Corollary 2 both hold at $(0, 0)$.

We now show that for $0 < \lambda < 1$ the second-order condition (10) holds at the two points $(-k, \pm\sqrt{2k(1-\lambda)})$. At these points $\lambda = 1$. Thus, for a second-order test vector x we have $E_1 = 2I$ and $E_2 = 2I$. Hence (5) and (6) reduce to $\nabla_{\lambda\lambda}^2(x_\lambda)x = 0$ at these points. Thus a test vector x satisfies

$$(2k, \pm 2\sqrt{2k(1-\lambda)})(x_1, x_2) = 2\alpha x_1 \pm 2\alpha\sqrt{2k(1-\lambda)}x_2 = 0. \quad (11)$$

The function F with $\lambda = 1$ becomes

$$F = (x_1 - 1)^2 + x_2^2 + 2k(x_1 - x_2^2)$$

and condition (10) becomes

$$(x_1, x_2) \begin{pmatrix} 2 & 0 \\ 0 & 2 \end{pmatrix} \begin{pmatrix} x_1 \\ x_2 \end{pmatrix} - 2kx_2^2 \geq 0.$$

This inequality is satisfied for all x and hence for all x satisfying (11).

We now state and prove a second-order sufficiency theorem [Hiri and Mangasarian, 1979]. This theorem is more useful and easier to apply than the second-order necessary conditions.

Theorem 5.3. Let $(x_\lambda, \lambda, k, g)$ satisfy the *FF* necessary conditions of Theorem 5.1. Suppose that every $\lambda \neq 0$ that satisfies the tangential constraint (4) at $x_\lambda = x_\lambda$ and the inequality $(\nabla^2 f)(x_\lambda, \lambda) \ll 0$ also satisfies

$$(x, F_\lambda^2(x_\lambda, k, g)x) > 0, \quad (14)$$

where

$$F_\lambda^2(x, k, g) = k_\lambda f''(x) + \langle k, g \rangle(x) + \langle g, k \rangle(x).$$

Then f attains a strict local minimum for problem *PI* at x_λ .

Note that strict inequality is required in (14), while equality is allowed in the necessary condition (10).

Proof. To simplify notation, we assume that $x_\lambda = \Phi$ and that $f(\Phi) = 0$. For any real-valued function $\varphi: \mathbb{R}^n \rightarrow \mathbb{R}$, the symbol φ_{xx} will denote the Hessian matrix whose (i, j) th entry is $\partial^2 \varphi / \partial x_i \partial x_j$.

If f did not attain a strict local minimum at Φ , then for every positive integer q there would exist a limit-like point $x_q \neq \Phi$ such that $x_q \in \mathbb{R}^n$, $1/q \in \mathbb{R}$ and $f(x_q) < 0$.

Then there would exist a sequence of points $\{x_j\}$ in X such that

$$x_j \neq \bar{x}, \quad \lim_{j \rightarrow \infty} x_j = \bar{x}, \quad f(x_j) < \bar{v}, \quad g(x_j) \leq \bar{v}, \quad \lim_{j \rightarrow \infty} g(x_j) = \bar{v}.$$

Let $E = \{1, \dots, r\}$. It then follows from Taylor's theorem (Corollary 1, Theorem 1.1.2) and our assumption that $x_0 = \bar{x}$ and $f(\bar{x}) = \bar{v}$ that

$$f(x_j) - (F)(x_0, x_j, x_j) < 0 \quad (15)$$

and that for $i = 1, \dots, r$ and $j = 1, \dots, j$

$$\begin{aligned} a_i(x_j) &= (F_{a_i})(x_0, x_j, x_j) < 0, \\ b_j(x_j) &= (F_{b_j})(x_0, x_j, x_j) = 0 \end{aligned} \quad (16)$$

where the x_j are numbers in the open interval $(0, 1)$.

It also follows from Taylor's theorem that

$$f(x_j) = (N)(\bar{x}, x_j) + \frac{1}{2}(x_j, J_{nn}(\bar{x}, x_j)x_j) < 0 \quad (17)$$

and that for $i = 1, \dots, r$ and $j = 1, \dots, j$

$$\begin{aligned} a_i(x_j) &= (N_{a_i})(\bar{x}, x_j) + \frac{1}{2}(x_j, A_{a_i}(\bar{x}, x_j)x_j) < 0, \\ b_j(x_j) &= (N_{b_j})(\bar{x}, x_j) + \frac{1}{2}(x_j, A_{b_j}(\bar{x}, x_j)x_j) = 0, \end{aligned} \quad (18)$$

where the \bar{x}_j are real numbers in the open interval $(0, 1)$.

Since $x_0 \neq \bar{x}$, the sequence $\{x_j/\|x_0\|\}$ is a sequence of unit vectors. Hence there exists a unit vector v and a subsequence that we relabel as $\{x_j\}$ such that

$$x_j/\|x_0\| \rightarrow v.$$

Now suppose that the sequence $\{x_j\}$ in (15) and (16) is the subsequence. If we divide through by $\|x_0\|$ in the rightmost inequality and the equalities in (15) and then let $j \rightarrow +\infty$, we get that v satisfies

$$(F)(\bar{x}, v) < \bar{v}, \quad \nabla_{x_0}^2(\bar{x})v < \bar{v}, \quad \nabla b_j(\bar{x})v = 0.$$

Thus v satisfies the tangential constraints (4).

Next, multiply the first inequality in (16) by $\lambda_1 > 0$, the i th inequality in (16) by $\lambda_i > 0$ (the i th component of λ), and the j th equation in (16) by μ_j (the j th component of μ) and then add the resulting equations and inequality. Since

$\lambda = (\lambda_0, \lambda_1)$ and $k_2 = 0$, we get

$$(\lambda_0 F_0(\bar{x}) + \lambda^T F_0(\bar{x}) + \mu^T \nabla h(\bar{x}, \bar{y})) + \frac{1}{2} (\lambda_1 P_{\bar{x}}^0(\bar{y}_0, \bar{y}_1)) \leq 0,$$

where

$$P_{\bar{x}}^0(\bar{y}_0, \bar{y}_1) = \lambda_0 L_{\bar{x}}(\bar{y}_0, \bar{y}_1) + \sum_{j=1}^k \lambda_j L_{\bar{x}}(\bar{y}_j, \bar{y}_1) + \sum_{j=1}^k \mu_j L_{\bar{x}}(\bar{y}_0, \bar{y}_j).$$

Since $(\bar{x}, \bar{y}^0, \lambda, \mu)$ satisfies the *FJ* necessary conditions of Theorem 5.1, we get that

$$\frac{1}{2} (\lambda_1 P_{\bar{x}}^0(\bar{y}_0, \bar{y}_1)) \leq 0.$$

If we now divide both sides of this inequality by $|\lambda_1|^2$ and then let $q \rightarrow \infty$, we get that $(\lambda_1 P_{\bar{x}}^0(\bar{y}^q)) \leq 0$ for a nonzero vector that satisfies the tangential constraint (4) and $(\nabla F_0(\bar{x}), \bar{v}) \leq 0$. This contradicts the assumption that (14) holds for all such \bar{v} .

If there exists a set of multipliers at \bar{x}_0 with $\lambda_0 = 1$, we have a slightly different version of the sufficient conditions.

COROLLARY 4. Let $(\bar{x}_0, \lambda_0, \mu)$ satisfy the *KKT* necessary conditions of Theorem 5.2. Suppose that every $\bar{v} \neq 0$ in \mathbb{R}^n that satisfies the tangential constraint (4) at $\bar{x}_0 = \bar{x}_*$ and the equality $(\nabla F(\bar{x}_0), \bar{v}) = 0$ also satisfies

$$(\lambda_0 F_{\bar{x}}(\bar{x}_0, \lambda_0, \mu)) > 0, \quad (17)$$

where

$$F(\bar{x}, \lambda, \mu) = f(\bar{x}) + (\lambda_0 g(\bar{x})) + (\mu, h(\bar{x})). \quad (18)$$

Then f attains a *strict* local minimum for problem *PJ* at \bar{x}_0 .

The corollary differs from the theorem in that the condition $(F(\bar{x}_0), \bar{v}) \leq 0$ is replaced by $(\nabla F(\bar{x}_0), \bar{v}) = 0$.

Proof. If $(\bar{x}_0, \lambda, \mu)$ satisfies the conditions of Theorem 5.1, then

$$\nabla f(\bar{x}_0) + \lambda^T \nabla g(\bar{x}_0) + \mu^T \nabla h(\bar{x}_0) = 0.$$

Since $k_2 = 0$,

$$\nabla f(\bar{x}_0) = -\lambda_0^T \nabla g(\bar{x}_0) - \mu^T \nabla h(\bar{x}_0).$$

Hence, for any $x \in E^*$

$$\langle \nabla f(x_0), x \rangle = -\langle \lambda_0, \nabla g_0(x_0, x) \rangle - \langle \mu, \nabla h(x_0, x) \rangle.$$

Since $\lambda_0 \geq 0$, it follows that for any x satisfying the tangential constraints (8) we have $\langle \nabla f(x_0), x \rangle \geq 0$. Hence the condition $\langle \nabla f(x_0), x \rangle \leq 0$ reduces to $\langle \nabla h(x_0), x \rangle = 0$.

If the inequality constraints are absent, then a vector x satisfies the tangential constraints at a point x_0 if and only if $\nabla h(x_0)x = 0$. Thus, if at a point x_0 the necessary conditions of Corollary 2 of Theorem 5.2 are satisfied, the proof of Corollary 4 shows that for every vector x such that $\nabla h(x_0)x = 0$ we have $\langle \nabla f(x_0), x \rangle = 0$. Therefore the following corollary holds.

COROLLARY 5. Let (x_0, μ) satisfy the necessary condition of Corollary 2 of Theorem 5.2. Suppose that for every x such that $\nabla h(x_0)x = 0$

$$\langle \mu, F_{xx}(x_0, x) \rangle = 0,$$

where

$$F(x, \mu) = f(x) + \langle \mu, h(x) \rangle.$$

Then f attains a strict local minimum for the problem of minimizing f subject to $h(x) = 0$.

There is a sufficient condition due to Ponsot [1933] which is not as convenient to apply as Corollary 4 but is equivalent to it. We now state this condition and prove its equivalence to Corollary 4.

COROLLARY 6. Let (x_0, λ, μ) satisfy the KKT necessary conditions of Theorem 5.2. Let $E^* = \{x \mid x \in E \text{ and } \lambda_j > 0\}$. Suppose that for every $x \neq 0 \in E^*$ that satisfies

$$\begin{aligned} \langle \nabla g_0(x_0), x \rangle = 0, \quad x \in E^*, \quad \langle \nabla g_0(x_0), x \rangle < 0, \quad x \in E, \quad \forall E^*, \\ \nabla h(x_0)x = 0, \end{aligned} \quad (19)$$

the inequality

$$\langle \lambda, F_{xx}(x_0, \lambda, \mu) \rangle > 0$$

holds, where F is as in (8). Then f attains a strict local minimum for problem P1 at x_0 .

Proof. To prove the corollary it suffices to show that the set K_1 of vectors that satisfy (19) is the same as the set K_2 of vectors that satisfy the tangential

constraints (4) and the equality $\langle \nabla f(x_0), x \rangle = 0$, and then invoke Corollary 4.

We first show that $V_1 \subseteq V_2$. If $V_1 = \emptyset$, there is nothing to prove, so we assume that $V_1 \neq \emptyset$. From (6) of Theorem 5.2 we get, that for x in E^2

$$\langle \nabla f(x_0), x \rangle + \lambda_1 \nabla g_1(x_0) \cdot x + \langle \mu, \nabla h(x_0) \rangle = 0. \quad (20)$$

Now let $x \in V_1$. Then $\nabla g_1(x_0) \cdot x \leq 0$, and $\nabla h(x_0) \cdot x = 0$. Thus x satisfies the tangential constraints (9). Moreover, since $\lambda_1 = 0$ for $x \in E$, we have that $\lambda_1 = 0$ for $x \in E^2$. Thus (20) becomes

$$\langle \nabla f(x_0), x \rangle + \sum_{i=2}^m \lambda_i \langle \nabla g_i(x_0), x \rangle = 0$$

for all x in V_1 . For such x and for $x \in E^2$ we have $\langle \nabla g_2(x), x \rangle = 0$. Hence $\langle \nabla f(x_0), x \rangle = 0$, and so $V_1 \subseteq V_2$.

We now show that $V_2 \subseteq V_1$. We assume that $V_2 \neq \emptyset$. Since $\lambda_1 = 0$ for $x \in E^2$ and $\langle \nabla f(x_0), x \rangle = 0$, it follows that, for $x \in V_2$, (20) becomes

$$\sum_{i=2}^m \lambda_i \langle \nabla g_i(x_0), x \rangle = 0.$$

Since for any vector x that satisfies the tangential constraints (4) we have $\langle \nabla g_2(x_0), x \rangle \leq 0$, and since $\lambda_1 > 0$ for $x \in E^2$, the last equality implies that $\langle \nabla g_i(x_0), x \rangle = 0$ for $x \in E^2$. Since $\nabla g_i(x_0) \cdot x \leq 0$ for any vector satisfying the tangential constraints, we get that $V_2 \subseteq V_1$.

Example 5.7 (Continued). We now apply Corollary 4 to show that the points given in relation (4) in Section 5 furnish local minima for problems π_0 when $0 < k < 1$ and that the origin furnishes a local minimum for problems π_0 when $k > 1$.

We first consider the problem π_0 for $0 < k < 1$. At the point $x_{10} = (1 - k, \sqrt{2k(1 - k)})$, the condition $\langle \nabla f(x_{10}), x \rangle = 0$ becomes

$$(1 - 2k, \sqrt{2k(1 - k)}) \cdot (x_1, x_2) = 0.$$

At the point $x_{20} = (1 - k, -\sqrt{2k(1 - k)})$ the condition $\langle \nabla f(x_{20}), x \rangle = 0$ becomes

$$(1 - 2k, -\sqrt{2k(1 - k)}) \cdot (x_1, x_2) = 0.$$

Thus at x_{10} and x_{20} we require that

$$x_{10,1} - 2x_1 = 2x_2 \sqrt{\frac{2k(1 - k)}{k}}, \quad x_{20,1} - 2x_1 = -2x_2 \sqrt{\frac{2k(1 - k)}{k}}. \quad (21)$$

The tangential constraints at the points \mathbf{x}_{1a} and \mathbf{x}_{2a} become

$$\begin{aligned}\mathbf{x}_{1a}: (2k_1 - 2\alpha_1\sqrt{2k_1}) - k_1(\alpha_1 - \alpha_1^2) &\leq 0, \\ \mathbf{x}_{2a}: (2k_2 - 2\alpha_2\sqrt{2k_2(1-d)}) - k_2(\alpha_2 - \alpha_2^2) &\leq 0\end{aligned}$$

or

$$\mathbf{x}_{1a}: \alpha_1 \leq \alpha_1 \sqrt{\frac{2k_1 - k_1}{k_1}}, \quad \mathbf{x}_{2a}: \alpha_2 \leq -\alpha_2 \sqrt{\frac{2k_2 - k_2}{k_2}}.$$

Hence at \mathbf{x}_{1a} the vectors $\mathbf{z} = (z_1, z_2)$ that satisfy $(\nabla F)(\mathbf{x}_{1a}, \mathbf{z}) = 0$ and the tangential constraints are given in (21). Similarly, at \mathbf{x}_{2a} the vectors that satisfy $(\nabla F)(\mathbf{x}_{2a}, \mathbf{z}) = 0$ and the tangential constraints are given in (21). Since $d = 1$ at \mathbf{x}_{1a} and \mathbf{x}_{2a} , the function F in (16) becomes

$$F = (\alpha_1 - 1)^2 + \alpha_1^2 + (2k_1\alpha_1 - \alpha_1^2).$$

Thus

$$\langle \mathbf{z}, \nabla F(\mathbf{x}_{1a}, 1, \mathbf{z}) \rangle = (\alpha_1, \alpha_2) \begin{pmatrix} 2 \\ 0 \end{pmatrix} \begin{pmatrix} \alpha_1 \\ \alpha_2 \end{pmatrix} = 2\alpha_1^2. \quad (22)$$

At \mathbf{x}_{1a} and at \mathbf{x}_{2a} vectors $\mathbf{z} \neq \mathbf{0}$ at which (21) holds have $\alpha_1 \neq 0$. It therefore follows that the hypotheses of Corollary 4 hold at \mathbf{x}_{1a} and \mathbf{x}_{2a} , and so both of these points furnish local minima. We leave it as an exercise to show that these points furnish minima for problem π_k with $0 < k < 1$.

We now consider the problem π_k , $k > 1$. The origin is the only point that satisfies the conditions of Theorem 5.2. Since $\nabla f(\mathbf{0}) = (-2, 0)$, all vectors \mathbf{z} of the form (z_1, z_2) with z_1 arbitrary, satisfy $(\nabla f)(\mathbf{0}, \mathbf{z}) = 0$. Since $\nabla g(\mathbf{0}) = (2k, 0)$, all vectors \mathbf{z} of the form $\mathbf{z} = (z_1, z_2)$ with $z_1 \leq 0$, and z_2 arbitrary, satisfy the tangential constraint $(\nabla g)(\mathbf{0}, \mathbf{z}) \leq 0$. Hence the vectors that satisfy the tangential constraints and $(\nabla f)(\mathbf{0}, \mathbf{z}) = 0$ are vectors of the form (z_1, z_2) where z_1 is arbitrary.

The function F with $d = 1/k$ becomes

$$F = (\alpha_1 - 1)^2 + \alpha_1^2 + \frac{2k\alpha_1 - \alpha_1^2}{k}.$$

Thus (17) with $\mathbf{z} = (z_1, z_2)$ becomes

$$\langle \mathbf{z}, \mathbf{z} \rangle \begin{pmatrix} 2 \\ 0 \end{pmatrix} \begin{pmatrix} \alpha_1 \\ \alpha_2 \end{pmatrix} = 2 \left(1 - \frac{1}{k} \right) (\alpha_1)^2 > 0, \quad \alpha_1 \neq 0.$$

Hence the origin furnishes a strict local minimum for problem π_k , $k > 1$. Again, we leave it as an exercise for the reader to show that the origin furnishes a global minimum.

If $k = 1$, then $\mathbf{x}_a = (0, 0)$ and

$$\langle \nabla f(\mathbf{0}), \mathbf{x} \rangle = \langle (-2, 0), (x_1, x_2) \rangle = -2x_1,$$

$$\langle \nabla g(\mathbf{0}), \mathbf{x} \rangle = \langle (2, 0), (x_1, x_2) \rangle = 2x_1.$$

Hence a vector \mathbf{x} satisfies the tangential constraint $\langle \nabla g(\mathbf{0}), \mathbf{x} \rangle \leq 0$ and the constraint $\langle \nabla f(\mathbf{0}), \mathbf{x} \rangle = 0$ if and only if it has the form $\mathbf{x} = (0, x_2)$, x_2 arbitrary. If $\mathbf{x}_a = \mathbf{0}$, then $\lambda = 1$ and (22) gives $\langle \mathbf{x}, F_{xx}(\mathbf{0}, 1, \rho)\mathbf{x} \rangle = 0$. Hence the sufficient condition gives no information in this case.

Exercise 6.1. Show that the points \mathbf{x}_{1a} , \mathbf{x}_{2a} and $\mathbf{0}$ in Example 5.1 furnish minima for the relevant problems.

Exercise 6.2. In Exercises 5.1–5.5 use an appropriate second-order sufficiency condition to determine whether the critical points furnish a local minimum.

Exercise 6.3. Given m distinct points y_1, \dots, y_m in \mathbb{R}^n , find the smallest closed ball $\overline{B(\mathbf{x}, \rho)}$ that encloses these points. This problem can be formulated analytically as follows: Minimize $f(\mathbf{x}, \rho) = \rho^2$ subject to $\|\mathbf{x} - y_i\|^2 \leq \rho^2$.

- (a) Show that a solution (\mathbf{x}_0, ρ_0) exists.
 (b) Show that at a critical point (\mathbf{x}_0, ρ_0)

(i) $\lambda_0 = 1$,

(ii) $\lambda_i \geq 0$, $i = 1, \dots, m$,

(iii) $\sum_{i=1}^m \lambda_i = 1$,

(iv) $\mathbf{x}_0 = \sum_{i=1}^m \lambda_i y_i$, and

(v) $\rho_0^2 + \|\mathbf{x}_0\|^2 = \sum_{i=1}^m \lambda_i \|y_i\|^2$.

Note that (ii), (iii), and (iv) show that \mathbf{x}_0 is in the simplex determined by y_1, \dots, y_m .

- (c) Show that with the KKT multipliers as in (b) the function F can be written as

$$F(\rho, \mathbf{x}, \lambda) = \|\mathbf{x}\|^2 - 2 \left\langle \mathbf{x}, \sum_{i=1}^m \lambda_i y_i \right\rangle.$$

- (d) Does the problem have a unique solution?
 (e) Solve the special case $y_i = \mathbf{e}_i$, $i = 1, 2, 3$, in detail.

V

CONVEX PROGRAMMING AND DUALITY

1. PROBLEM STATEMENT

In the preceding chapter we studied programming problems with differentiable data. In this chapter we shall replace the assumption that the data are differentiable by the assumption that the data are convex. This assumption will enable us to obtain much more detailed information about the behavior of solutions. In Section 1 of Chapter IV we defined a convex programming problem to be a programming problem in which the underlying set X_0 is convex, the functions f and g are convex, and there are no equality constraints.

Convex Programming Problem 1 (CPI)

Let X_0 be a convex set and let $f: X_0 \rightarrow \mathbb{R}$ and $g: X_0 \rightarrow \mathbb{R}^m$ be convex.

$$\text{Minimize } [f(x)] \text{ over } X,$$

where

$$X = \{x \mid g(x) \in \theta\}.$$

The problem is often stated as

$$\text{Minimize } f(x) \quad \text{subject to } g(x) \in \theta.$$

Remark 1.1. A convex programming problem with the additional equality constraint $h(x) = \theta$ can be cast in the format CPI if the function h is affine. One replaces the constraint $h(x) = \theta$ by the two inequality constraints $h(x) \leq \theta$ and $-\mathbf{h}(x) \leq \theta$. Since h is affine, both h and $-\mathbf{h}$ are convex.

In Remark IV.5.3 we pointed out that if the feasible set X in CPI is not empty, then it is convex. Thus in CPI we are minimizing a convex function

defined on the convex set X . Therefore, the properties of solutions listed in Theorem IV.5.1 are valid for CPl.

If the functions f and g are differentiable on N_p , then CPl is a special case of the programming problem Pl. Hence Theorem IV.5.1 (the Fritz John necessary conditions), Theorem IV.5.2 (the KKT necessary conditions), and the sufficiency theorem IV.5.3 are all valid.

It turns out that in many problems involving affine equality constraints $\mathbf{h}(\mathbf{x}) = \mathbf{0}$, sharper results can be obtained if we retain the equality constraints as such rather than replace them by two sets of inequality constraints. It will also be convenient to take $N_p = R^n$ when affine constraints are present. We therefore formulate the convex programming problem with affine constraints as follows.

Convex Programming Problem II (CPII)

Let $f: R^n \rightarrow R$ be convex, let $g: R^n \rightarrow R^k$ be convex and let $\mathbf{h}: R^n \rightarrow R^l$ be affine. Minimize

$$f(\mathbf{x}) \quad (\mathbf{x} \in X),$$

where

$$X = \{\mathbf{x} : g(\mathbf{x}) \leq \mathbf{0}\} \cap \{\mathbf{x} : \mathbf{h}(\mathbf{x}) = \mathbf{0}\}.$$

If the feasible set X is not empty, the problem is said to be consistent.

Remark 1.1. The components of the affine equality constraints $\mathbf{h}(\mathbf{x}) = \mathbf{0}$ can be written as

$$(\mathbf{a}_i, \mathbf{x}) - b_i = 0, \quad i = 1, \dots, l. \quad (1)$$

Thus, if the problem is consistent, then the system

$$A\mathbf{x} = \mathbf{b}, \quad (2)$$

where A is the $l \times n$ matrix whose i th row is \mathbf{a}_i and $\mathbf{b} = (b_1, \dots, b_l)^T$ has solutions. We recall from linear algebra that for (1) to have solutions the rank of A must equal the rank of the augmented matrix (A, \mathbf{b}) . If the system (2) has solutions and if the rank of (A, \mathbf{b}) were less than l , then one of the equations in (1) could be written as a linear combination of the others and thus eliminated from the system. Since we are only interested in problems that are consistent, there is no loss of generality in assuming that A has maximum possible rank, namely l . In that event the ranks of A and (A, \mathbf{b}) both equal l . If the feasible set consists of only one point, then the programming problem is trivial. Therefore, we further assume that $l < n$. For if $l = n$ and A has rank

$k = n$, then (2) has a unique solution and the feasible set consists of a single point. The assumptions that A has rank k and that $k = n$ will be in force henceforth.

In the next section we shall develop necessary conditions and a sufficient condition for a point x_0 to be a solution of CPL without assuming differentiability.

2. NECESSARY CONDITIONS AND SUFFICIENT CONDITIONS

Our first necessary condition will involve a function A with domain $E^1 \times E^1 \times E^r \times E^r$ and range in E defined by

$$A(x, \lambda, k, p) = \lambda_0 f(x) + (x, g(x)) + (p, h(x)). \quad (1)$$

The function A is sometimes called a *generalized Lagrangian*.

THEOREM 2.1. Let f and g be convex on E^1 and let h be affine. Let x_0 be a solution of CPL. Then there exist a real number λ_0 , a vector k_0 in E^r , and a vector p_0 in E^r with $(\lambda_0, k_0) \geq (0, 0)$ and $(\lambda_0, k_0, p_0) \neq (0, 0, 0)$ such that

$$A(x_0, g(x_0)) = 0 \quad (2)$$

and such that for all x in E^1 , all $k \geq 0$ in E^r , and all p in E^r

$$A(x_0, \lambda_0, k, p) \leq A(x_0, \lambda_0, k_0, p_0) \leq A(x, \lambda_0, k_0, p_0). \quad (3)$$

Moreover,

$$A(x_0, \lambda_0, k_0, p_0) = \lambda_0 f(x_0). \quad (4)$$

Proof. Since x_0 is a solution of CPL, the system

$$\begin{aligned} f(x) - \beta &< 0, \\ g(x) &< 0, \quad h(x) = 0, \end{aligned}$$

has no solution in E^1 . Hence by Theorem III.4.1 there exists a vector $(\lambda_0, k_0, p_0) \neq (0, 0, 0)$ with $(\lambda_0, k_0) \geq 0$ such that

$$\lambda_0 f(x) - f(x_0) + (x_0, g(x)) + (p_0, h(x)) \geq 0$$

for all x in \mathbb{R}^n . Hence

$$\lambda_{10} f(x_0) \leq \lambda_{10} f(x) + \langle \lambda_{10}, g(x) \rangle + \langle \mu_{10}, h(x) \rangle \quad (5)$$

for all x in \mathbb{R}^n . In particular, (5) holds for $x = x_0$. Since $h(x_0) = 0$, we get that $\langle \lambda_{10}, g(x_0) \rangle \geq 0$. On the other hand, since x_0 is feasible, $g(x_0) \leq 0$. But $\lambda_{10} \geq 0$, so $\langle \lambda_{10}, g(x_0) \rangle \leq 0$. Hence, $\langle \lambda_{10}, g(x_0) \rangle = 0$, and (2) is established.

From (2), (5), and $h(x_0) = 0$ we now get

$$\begin{aligned} \lambda_{10} f(x_0) + \langle \lambda_{10}, g(x_0) \rangle + \langle \mu_{10}, h(x_0) \rangle \\ = \lambda_{10} f(x_0) \leq \lambda_{10} f(x) + \langle \lambda_{10}, g(x) \rangle + \langle \mu_{10}, h(x) \rangle \end{aligned}$$

for all x . This establishes (4) and the rightmost inequality in (3). From (2) and from $g(x_0) \leq 0$ it follows that for $\lambda \geq 0$

$$\langle \lambda, g(x_0) \rangle \leq 0 = -\langle \lambda_{10}, g(x_0) \rangle.$$

From this and from $h(x_0) = 0$ we get that

$$\lambda_{10} f(x_0) + \langle \lambda, g(x_0) \rangle + \langle \mu_{10}, h(x_0) \rangle \leq \lambda_{10} f(x_0) + \langle \lambda_{10}, g(x_0) \rangle + \langle \mu_{10}, h(x_0) \rangle,$$

which is the leftmost inequality in (3).

Remark 2.1. Since $\lambda_{10} \geq 0$ and $g(x_0) \leq 0$, condition (2) is equivalent to

$$\lambda_{10} g(x_0) = 0 \quad i = 1, \dots, m. \quad (7)$$

Remark 2.2. Although (3) is written as a saddle point condition for the function \mathcal{L} , the leftmost inequality does not involve g since $h(x_0) = 0$.

Remark 2.3. If the functions f and g are differentiable, then since x_0 minimizes the function $\mathcal{L}(\cdot, \lambda_{10}, \lambda_{20}, \mu_{10})$ on \mathbb{R}^n , where

$$\mathcal{L}(x, \lambda_{10}, \lambda_{20}, \mu_{10}) = \lambda_{10} f(x) + \langle \lambda_{20}, g(x) \rangle + \langle \mu_{10}, h(x) \rangle,$$

we have

$$\frac{\partial \mathcal{L}}{\partial x_i}(\lambda_{10}, \lambda_{20}, \lambda_{20}, \mu_{10}) = 0, \quad i = 1, \dots, n.$$

Thus, if A is the matrix in (1.2),

$$\lambda_{10} \nabla f(x_0) + \lambda_{20}^T \nabla g(x_0) + \mu_{10}^T A = 0,$$

which together with (2) and the condition $\langle \lambda_{10}, \lambda_{20}, \mu_{10} \rangle \neq (0, 0, 0)$ constitute the KJ necessary condition of Theorem IV.3.1.

Remark 1.4. In the proof of Theorem 2.1, the optimality of x_0 is shown to imply the existence of a vector $(\lambda_{00}, \lambda_0, \mu_0) \in (0, \infty) \times \mathbb{R}^m$ with $\lambda_{00} > 0$ and $\lambda_0 \geq \mathbf{0}$ such that relation (2) holds. All the conclusions of the theorem then follow from (2) and the feasibility of x_0 . We shall use this observation in the sequel.

As in the study of differentiable problems, we seek a constraint qualification that guarantees that $\lambda_{00} \neq 0$, for then we can take $\lambda_{00} = 1$. For the convex programming problem the following simple constraint qualification suffices.

Definition 1.1. The problem CP1 is said to be *strongly consistent*, or to satisfy the *Slater condition*, if there exists an x_0 such that $g(x_0) < \mathbf{0}$ and $h(x_0) = \mathbf{0}$.

Theorem 2.1. Let x_0 be a solution of CP1 and let the problem be strongly consistent. Then there exists a $\lambda_0 \geq \mathbf{0}$ in \mathbb{R}^m such that

$$\langle \lambda_0, g(x_0) \rangle = 0 \quad (9)$$

and a vector μ_0 in \mathbb{R}^l such that for all x , all $\lambda \geq \mathbf{0}$ in \mathbb{R}^m , and all μ in \mathbb{R}^l

$$L(x, \lambda, \mu) \leq L(x_0, \lambda_0, \mu_0) \leq L(x, \lambda_0, \mu_0), \quad (10)$$

where

$$L(x, \lambda, \mu) = A(x, \lambda, \mu) = f(x) + \langle \lambda, g(x) \rangle + \langle \mu, h(x) \rangle. \quad (11)$$

Moreover,

$$f(x_0) = A(x_0, \lambda_0, \mu_0). \quad (12)$$

Proof. Since x_0 is a solution of CP1, the conclusion of Theorem 1.1 holds. If $\lambda_{00} = 0$, then since (2) holds and $h(x_0) = \mathbf{0}$, the rightmost inequality in (3) becomes

$$0 \leq \langle \lambda_0, g(x) \rangle + \langle \mu_0, h(x) \rangle \quad (13)$$

for all x . If we also have $\lambda_0 = \mathbf{0}$, then the condition $(\lambda_{00}, \lambda_0, \mu_0) \neq (0, 0, 0)$ implies that $\mu_0 \neq \mathbf{0}$. Since $h(x) = Ax - b$, the inequality (13) and $\lambda_0 = \mathbf{0}$ imply that

$$\langle \mu_0^T A, x \rangle \geq \langle \mu_0, b \rangle \quad (14)$$

for all x . The standing hypothesis made in Section 1 that A has rank l and the condition $\mu_0 \neq \mathbf{0}$ imply that $\mu_0^T A \neq \mathbf{0}$. Hence if we take $x = \mu_0^T A$ and then let $y \rightarrow +\infty$, we see that (14) cannot hold for all x . Hence $\lambda_0 \neq \mathbf{0}$.

Since the problem is strongly consistent, there exists an x_0 such that $g(x_0) < \mathbf{0}$ and $h(x_0) = \mathbf{0}$. Thus, from (10), $\langle \lambda_0, g(x_0) \rangle > 0$. On the other hand,

$\lambda_{i_0} \geq 0$ and $\lambda_{i_0} \neq 0$ imply that $\langle \lambda_{i_0}, g(x_{i_0}) \rangle < 0$. This contradiction shows that $\lambda_{i_0} = 0$. Therefore, we can divide through by $\lambda_{i_0} = 0$ in (2), (3), and (4), relabel $\lambda_{i_0}/\lambda_{i_0}$ as λ_{i_0} and relabel μ_{i_0}/λ_{i_0} as μ_{i_0} to get the conclusion of Theorem 2.2.

Remark 2.5. The rightmost inequality in (7) and equation (8) give

$$f(x_{i_0}) = f(x_{i_0}, \lambda_{i_0}, \mu_{i_0}) \leq f(x) + \langle \lambda_{i_0}, g(x) \rangle + \langle \mu_{i_0}, h(x) \rangle$$

for all x . Therefore, if we assume that CFI has a solution x_{i_0} and if we know the multipliers λ_{i_0} and μ_{i_0} , then the solution x_{i_0} of the constrained problem

$$\min\{f(x) : x \in X\}, \quad X = \{x : g(x) \leq 0, h(x) = 0\}$$

will be a solution of the unconstrained problem

$$\min\{f(x) + \langle \lambda_{i_0}, g(x) \rangle + \langle \mu_{i_0}, h(x) \rangle\}. \quad (12)$$

The set S of solutions of the unconstrained problem (12) may contain points that do not satisfy the constraints. To find solutions of the constrained problem, one must determine those points of S that satisfy the constraints. It follows from the rightmost inequality in (7) that if (12) has a unique solution, then it must be a solution of the constrained problem.

One sometimes finds the statement that we may replace the constrained problem by the unconstrained problem of minimizing a Lagrangian. This statement is only true in the sense described in the preceding paragraph.

It is conceivable that there exist vectors (λ, μ) with $\lambda \geq 0$ such that the unconstrained problem (12) with $(\lambda_{i_0}, \mu_{i_0}) = (\lambda, \mu)$ has the same infimum as constrained problem. Such a vector will be called a *Kuhn-Tucker vector*. The multiplier vectors $(\lambda_{i_0}, \mu_{i_0})$ of Theorem 2.2 are Kuhn-Tucker vectors. In the next theorem we show that if CFI has a solution, then a Kuhn-Tucker vector for CFI is a multiplier vector as in Theorem 2.2.

Theorem 2.3. Let CFI have a solution x_{i_0} . Then a Kuhn-Tucker vector $(\bar{\lambda}, \bar{\mu})$ for CFI is a multiplier vector $(\lambda_{i_0}, \mu_{i_0})$ as in Theorem 2.2.

Proof. Since x_{i_0} is a solution of CFI and since $(\bar{\lambda}, \bar{\mu})$ is a Kuhn-Tucker vector,

$$f(x_{i_0}) = \inf\{f(x) + \langle \bar{\lambda}, g(x) \rangle + \langle \bar{\mu}, h(x) \rangle : x \in \mathbb{R}^n\}.$$

Thus (7) holds with $\lambda_{i_0} = 1$. Hence by the indicated statement in Remark 2.4, the vector $(\bar{\lambda}, \bar{\mu})$ is a multiplier vector as in Theorem 2.2.

Theorem 2.1 and 2.3 imply the following:

COROLLARY 1. Let CPL have a solution \mathbf{x}_0 . Let \mathcal{M} denote the set of multiplets associated with \mathbf{x}_0 , as in Theorem 2.2 and let \mathcal{K} denote the set of Kuhn-Tucker vectors for CPL. Then $\mathcal{M} = \mathcal{K}$.

Remark 2.6. The hypothesis of strong consistency in Theorem 2.2 guaranteed that the set \mathcal{M} is not empty. Thus if CPL is strongly consistent and has a solution, then \mathcal{K} is also nonempty.

We now give an example of a problem, taken from Rockafellar [1970], for which no Kuhn-Tucker vector, and hence no vector in \mathcal{M} , exists.

Example 2.7. Consider \mathbb{R}^2 , and let

$$f(x_1, x_2) = x_2, \quad g_1(x_1, x_2) = x_2, \quad g_2(x_1, x_2) = x_1^2 - x_2.$$

Then the only point in \mathbb{R}^2 satisfying $\mathbf{g}(\mathbf{x}_1, \mathbf{x}_2) \leq \mathbf{0}$ is $(0, 0)$, and the value of the minimum is zero. Thus, a Kuhn-Tucker vector $\mathbf{k} = (k_1, k_2)$ would have to satisfy $k_1 > 0, k_2 > 0$ and

$$0 \leq k_1 + k_2 x_1 + k_2(x_1^2 - x_2) = x_1 + (k_1 - k_2)x_2 + k_2 x_1^2$$

for all \mathbf{x} . We leave it to the reader to check that this is impossible.

Note that the problem in this example is not strongly consistent.

Condition (I) of Theorem 2.1 and condition (7) of Theorem 2.2 are called *saddle point conditions*. The points $(\mathbf{x}_0, \lambda_0, \mu_0)$ in (6) and $(\mathbf{x}_0, \lambda_0, \mu_0)$ in (7) are called *saddle points*. We previously encountered saddle point conditions and saddle points in our discussion of game theory in Section 3 of Chapter III. If the convex programming problem CPL has a solution and if the problem is strongly consistent, then the saddle point condition (7) gives us a method of computing $f(\mathbf{x}_0)$, the value of the minimum, without finding the optimal point \mathbf{x}_0 , as is shown in the following theorem. We shall return to this point in our study of duality in Section 4.

THEOREM 2.8. Let CPL be strongly consistent, let \mathbf{x}_0 be a solution, let v denote the value of the minimum, and let λ_0, μ_0 be the multiplets associated with \mathbf{x}_0 as in Theorem 2.2. Then

$$\begin{aligned} v &= f(\mathbf{x}_0, \lambda_0, \mu_0) = \min_{\mathbf{x}} L(\mathbf{x}, \lambda_0, \mu_0) \\ &= \sup_{\mathbf{k}, \mathbf{g}} \inf_{\mathbf{x}} L(\mathbf{x}, \mathbf{k}, \mathbf{g}) = \inf_{\mathbf{k}, \mathbf{g}} \sup_{\mathbf{x}} L(\mathbf{x}, \mathbf{k}, \mathbf{g}), \end{aligned} \quad (11)$$

where $\mathbf{k} \in \mathbb{R}^m, \mathbf{k} \geq \mathbf{0}, \mathbf{g} \in \mathbb{R}^l, \mathbf{x} \in \mathbb{R}^n$.

Proof. By a slight modification of the argument in Lemma III.5.1 to take into account the possible value of $+\infty$ for a supremum and $-\infty$ for an infimum, we have that

$$\sup_{\lambda \in \Lambda} \inf_{x \in X} L(x, \lambda, \mu) \leq \inf_{\lambda \in \Lambda} \sup_{x \in X} L(x, \lambda, \mu). \quad (14)$$

From (7) of Theorem 2.2 we get

$$\sup_{\lambda \in \Lambda} L(x_{\lambda^*}, \lambda, \mu) = L(x_{\lambda^*}, \lambda_{\lambda^*}, \mu_{\lambda^*}) = \inf_{\lambda} L(x_{\lambda^*}, \lambda, \mu_{\lambda^*}). \quad (15)$$

Hence

$$\inf_{\lambda} \sup_{x \in X} L(x, \lambda, \mu) \leq L(x_{\lambda^*}, \lambda_{\lambda^*}, \mu_{\lambda^*}) \leq \sup_{\lambda} \inf_{x \in X} L(x, \lambda, \mu). \quad (16)$$

We also note that by (7) of Theorem 2.2 we can replace the infimum in (15) by a minimum. This observation and (14) and (16) establish all the equalities in (11) except the first, which follows from (5) of Theorem 2.2.

Remark 2.7. It follows from Remark 2.6 and the proof of Theorem 2.4 that the hypothesis of strong consistency in the theorem can be replaced by the hypothesis that \mathcal{R} is not empty.

Theorem 2.2 states that if CPT is strongly consistent and x_{λ^*} is a solution, then we may take $\lambda_0 = 1$. Hence, if the functions f and g are differentiable, we may take $\lambda_0 = 1$ in Remark 2.3. Theorem IV.5.5 states that, for differentiable convex problems, if the KKT conditions hold at a feasible point x_{λ^*} , then x_{λ^*} minimizes. Thus, we have the following theorem.

Theorem 2.8. Let f and g be convex and differentiable and let h be affine. Let CPT be strongly consistent. A necessary and sufficient condition that a feasible point x_{λ^*} be a solution of CPT is that there exist multipliers $\lambda_{\lambda^*} \in \mathbb{R}^r$, $\lambda_{\mu^*} \geq 0$, and $\mu_{\lambda^*} \in \mathbb{R}^m$ such that

$$\begin{aligned} \nabla f(x_{\lambda^*}) + \lambda_{\lambda^*} \nabla g(x_{\lambda^*}) + \mu_{\lambda^*} h &= 0, \\ \langle \lambda_{\mu^*}, g(x_{\lambda^*}) \rangle &= 0. \end{aligned}$$

In the next theorem we shall show that, with no restrictions placed on f and g , the saddle point condition (7) is a sufficient condition that x_{λ^*} be a solution to the programming problem P: Minimize f over the set X , where $X = \{x : h(x) \in K, g(x) \leq \theta\}$.

Note that since we place no conditions on the constraints, equality constraints can be replaced by inequality constraints. We assume that this has been done, and so $h = 0$.

THEOREM 2.6. Let X_0 be a set in \mathbb{R}^n , let f be a real-valued function defined on X_0 , and let g be a function from X_0 to \mathbb{R}^m . Let x_0 be a point in X_0 , and let $\lambda_0 \geq \mathbf{0}$ be a vector in \mathbb{R}^m such that the saddle point condition (7) holds for all x in X_0 and all $\lambda \in \mathbf{0}$ in \mathbb{R}^m . Then x_0 is a solution to problem P .

Proof. We first show that x_0 is feasible. The leftmost inequality in (7) is

$$f(x_0) + \langle \lambda, g(x_0) \rangle \leq f(x_0) + \langle \lambda_0, g(x_0) \rangle, \quad \lambda \geq \mathbf{0}.$$

Hence

$$\langle \lambda - \lambda_0, g(x_0) \rangle \leq 0. \quad (17)$$

Now fix a value j in the set $\{1, \dots, m\}$. Let

$$\lambda_j = (\lambda_{1j}, \dots, \lambda_{(j-1)j}, \lambda_{jj} + 1, \lambda_{(j+1)j}, \dots, \lambda_{mj}).$$

If for each $j = 1, \dots, m$ we take $\lambda = \lambda_j$ in (17), we get

$$g_j(x_0) \leq 0, \quad j = 1, \dots, m.$$

Thus x_0 is feasible.

We next show that (6) holds. If we take $\lambda = \mathbf{0}$ in (17), we get that $\langle \lambda_0, g(x_0) \rangle \geq 0$. On the other hand, $g(x_0) \leq \mathbf{0}$ and $\lambda_0 \geq \mathbf{0}$ imply that $\langle \lambda_0, g(x_0) \rangle \leq 0$. Hence $\langle \lambda_0, g(x_0) \rangle = 0$, which is (6).

From the rightmost inequality in (7) and from (6) we get that

$$f(x_0) \leq f(x) + \langle \lambda_0, g(x) \rangle, \quad x \in X_0.$$

For x feasible, $g(x) \leq \mathbf{0}$. Hence for x feasible

$$f(x_0) \leq f(x),$$

and so $f(x_0) = \min\{f(x) : x \in X\}$.

Remark 2.8. Note that the proof of Theorem 2.6 only involves very elementary arguments.

The next theorem emphasizes some of the points implicit in the material presented in this section.

THEOREM 2.7. A necessary and sufficient condition for a point (x_0, λ_0, μ_0) to be a saddle point for the Lagrangian L defined in (8) on the domain

$$D = \{x, \lambda, g : x \in \mathbb{R}^n, \lambda \in \mathbb{R}^m, \lambda \geq \mathbf{0}, \mu \in \mathbb{R}^r\}$$

is that

- (i) $f(x_{\mu}, \lambda_{\mu}, \mu_{\mu}) = \min\{f(x, \lambda, \mu) : x \in R^*\}$
 (ii) $g(x_{\mu}) \leq 0, \quad h(x_{\mu}) = 0$
 (iii) $G_{\mu} g(x_{\mu}, \bar{c}) = 0.$
- (18)

Moreover, if (18) and (19) hold, then

$$f(x_{\mu}) = L(x_{\mu}, \lambda_{\mu}, \mu_{\mu}). \quad (19)$$

Proof. It follows from the definition of L that if (18) and (19) hold, then (19) holds.

Let $(x_{\mu}, \lambda_{\mu}, \mu_{\mu})$ be a saddle point for L . Then (18) follows from the definition of saddle point. The inequality in (18) and the equation (18ii) were established in the course of proving Theorem 2.8. The equality in (18i) also follows from Theorem 2.8, since there we converted the equality constraints to two inequality constraints, $h(x) \leq 0$ and $-h(x) \leq 0$, and then $h(x_{\mu}) \leq 0$ and $-h(x_{\mu}) \leq 0$. Now suppose that (18) holds. Condition (18ii) states that x_{μ} is feasible. Conditions (18i) and (18) imply that (5) holds. It now follows from Remark 2.4 that $(x_{\mu}, \lambda_{\mu}, \mu_{\mu})$ is a saddle point for $L(x, \lambda, \mu)$.

5. PERTURBATION THEORY

In problems arising in applications the values of various parameters are usually not known precisely but are known to lie in certain ranges of values. Thus one might assign a value to the parameters and then study the effect of perturbing the parameters. We shall assume that this is modeled by perturbing the constraints and shall study the effect of perturbations on the value of the infimum of the problem. Given the uncertainty in the parameter values, one would hope for a certain stability of the problem in the sense that small perturbations in the constraints will not lead to large changes in the value of the infimum. We shall show that if the unperturbed convex programming problem is strongly consistent and has a solution, then the change in the value of the infimum is bounded below by a linear function of the perturbations. In the preceding section we showed that under reasonable assumptions the set \mathcal{M} of multipliers and the set \mathcal{N} of Kuhn-Tucker vectors are nonempty and equal. In this section we shall show that the vectors in \mathcal{N} and \mathcal{M} are the coefficient vectors of the aforementioned linear function of the perturbations. Moreover, the sets \mathcal{N} and \mathcal{M} are the negative subdifferential of the infimum as a function of the perturbation, evaluated at the unperturbed value. Thus, in a sense, the multiplier measures the rate of change of the value of the infimum as we perturb the constraints. We shall also give an economic interpretation of the multiplier. These statements will be made precise in this section.

Throughout this section we assume that the problem CPH is consistent, that is, it has a nonempty feasible set. For each $\alpha = (\alpha_1, \alpha_2)$ with $\alpha_1 \in \mathbb{R}^m$ and $\alpha_2 \in \mathbb{R}^l$ we define a perturbed programming problem CPH(α) as follows.

Problem CPH(α)

Minimize $f(x)$ subject to x in $X(\alpha)$, where

$$X(\alpha) = \{x \mid g(x) \leq \alpha_1, h(x) = \alpha_2\},$$

the function f is a convex function from \mathbb{R}^n to \mathbb{R} , the function g is a convex function from \mathbb{R}^n to \mathbb{R}^m , and the function h is an affine function from \mathbb{R}^n to \mathbb{R}^l .

The constraint set $X(\alpha)$ can also be described by

$$X(\alpha) = \{x \mid g(x) - \alpha_1 \leq 0, h(x) - \alpha_2 = 0\}.$$

For each fixed α the function $g(x) - \alpha_1$ is convex and the function $h(x) - \alpha_2$ is affine. Thus, each problem CPH (α) is a convex programming problem. Note that the problem CPH is the problem CPH(θ).

For each α in \mathbb{R}^{m+l} the feasible set $X(\alpha)$ is either empty or convex. For each α such that $X(\alpha) \neq \emptyset$, let

$$v(\alpha) = \inf\{f(x) \mid x \in X(\alpha)\}.$$

Thus we have a function v whose domain is the set of α such that $X(\alpha)$ is not empty and which can assume the value $-\infty$. We shall designate the domain of v by $\text{dom}(v)$ and shall call v the value function of the family of perturbed problems CPH(α), or simply the value function. For α in $\text{dom}(v)$ we shall call $v(\alpha)$ the value of CPH(α). We shall denote the infimum of the unperturbed problem CPH by v . Thus

$$v = \inf\{f(x) \mid g(x) \leq 0, h(x) = 0\},$$

so $v = v(\theta)$.

Since $v(\alpha) = -\infty$ is possible, we shall adjoin $-\infty$ to the real \mathbb{R} . Although we shall not need to do so in this section, we shall also adjoin $+\infty$ to \mathbb{R} . To work with $-\infty$ and $+\infty$ in a consistent way, we define the following operations:

$$\begin{aligned} (+\infty) + (+\infty) &= +\infty, \\ (-\infty) + (-\infty) &= -\infty, \\ (-\infty) + \alpha &= -\infty, & \alpha \text{ real}, \\ \alpha - (-\infty) &= -\infty, & \alpha \text{ real } \alpha > 0, \\ \alpha - (+\infty) &= +\infty, & \alpha \text{ real } \alpha < 0, \\ -\infty < \alpha < +\infty, & \alpha \text{ real}. \end{aligned}$$

The operation $(-\infty) + (+\infty)$ is not defined.

For functions f defined on a convex set C and taking on either real values or the value of $-\infty$, the definition of convexity is unchanged. Namely, f is said to be convex on C if for every x_1, x_2 in C and every (α, β) in P_1

$$\beta f(x_1) + \alpha f(x_2) \leq f(\beta x_1 + \alpha x_2)$$

where the values of $-\infty$ are to be operated on as in (1). As in Chapter 3, an equivalent definition of f convex is that the set

$$\text{epi}(f) = \{(x, y) \mid x \in C, y \in \mathbb{R}, y \geq f(x)\}$$

is convex.

Lemma 3.1. *The domain of the value function v is a convex set and v is a convex function on $\text{dom}(v)$. If the problem CPH is strongly consistent, then Φ is an inverse point of $\text{dom}(v)$.*

Proof. We first show that $\text{dom}(v)$ is a convex set. We shall suppose that the affine equality constraints have been written as pairs of inequality constraints. Since we assume that CPH is consistent, Φ is in $\text{dom}(v)$ and thus $\text{dom}(v) \neq \emptyset$. We must show that if w_1 and w_2 are two points in $\text{dom}(v)$, then so is $\alpha w_1 + \beta w_2$ for all (α, β) in P_1 . Since w_1 and w_2 are in $\text{dom}(v)$, there exist points x_1 and x_2 in R^n such that

$$g(x_1) \leq w_1 \quad \text{and} \quad g(x_2) \leq w_2.$$

Since g is convex, for any (α, β) in P_1

$$g(\alpha x_1 + \beta x_2) \leq \alpha g(x_1) + \beta g(x_2) \leq \alpha w_1 + \beta w_2.$$

Thus, $\alpha w_1 + \beta w_2$ is in $N(w_1 + \beta w_2)$ and $\alpha w_1 + \beta w_2$ is in $\text{dom}(v)$.

To show that v is convex, again let w_1 and w_2 be two points in $\text{dom}(v)$ and let (α, β) be in P_1 . Let $x_1 \in N(w_1)$ and let $x_2 \in N(w_2)$. Then as in the preceding paragraph $\alpha x_1 + \beta x_2$ is in $N(\alpha w_1 + \beta w_2)$. Hence

$$\begin{aligned} v(\alpha w_1 + \beta w_2) &= \inf\{f(x) \mid g(x) \leq \alpha w_1 + \beta w_2\} \\ &\leq \inf\{f(x) \mid x \in N(\alpha w_1 + \beta w_2), g(x) \leq w_1, g(x) \leq w_2\} \\ &= \inf\{f(\alpha x_1 + \beta x_2) \mid \alpha x_1 \in N(w_1), \alpha x_2 \in N(w_2)\} \\ &\leq \inf\{v(\alpha x_1) + \beta f(x_2) \mid \alpha x_1 \in N(w_1), \alpha x_2 \in N(w_2)\}. \end{aligned} \quad (1)$$

If $v(w_1)$ and $v(w_2)$ are both finite, then by Exercise 1.6.4 the rightmost expression in (1) is equal to

$$\inf\{v(\alpha x_1) \mid \alpha x_1 \in N(w_1)\} + \inf\{f(x_2) \mid \alpha x_2 \in N(w_2)\} = \alpha v(w_1) + \beta v(w_2).$$

If either $\alpha(\mathbf{x}_1) = -\infty$ or $\alpha(\mathbf{x}_2) = -\infty$, then the rightmost expression in (1) is equal to $-\infty$. Thus,

$$\alpha(\mathbf{x}_1 + \beta\mathbf{x}_2) \leq \alpha(\mathbf{x}_2) + \beta\alpha(\mathbf{x}_1), \quad (7)$$

and the convexity of α is established.

We now suppose that the problem CPH is strongly consistent and shall show that \mathbf{B} is an interior point of $\text{dom}(\alpha)$. We shall not incorporate the affine equality constraints into the inequality constraints. Recall that we are writing the vector \mathbf{a} which is the perturbation of the constraints as $\mathbf{a} = \mathbf{a}_1, \mathbf{a}_2$, where $\mathbf{a}_1 \in \mathbb{R}^m$ is the perturbation of the inequality constraints and $\mathbf{a}_2 \in \mathbb{R}^k$ is the perturbation of the equality constraints.

Since CPH is strongly consistent, there exists an ϵ_0 such that

$$g(\mathbf{x}_0) < 0 \quad \text{and} \quad \mathbf{h}(\mathbf{x}_0) = \mathbf{0}. \quad (8)$$

Let $\mathbf{x}_0 = (x_{01}, \dots, x_{0n})$. For each $j = 1, \dots, m$, the function G_j defined on $\mathbb{R}^n \times \mathbb{R}$ by

$$G_j(\mathbf{x}, \alpha_j) = g_j(\mathbf{x}) - \alpha_j$$

is continuous. This holds because G_j is convex and a convex function is continuous on the interior of its domain of definition, which in this case is $\mathbb{R}^n \times \mathbb{R}$. From (8) we get that $G_j(\mathbf{x}_0, 0) < 0$. Hence, by the continuity of G_j , there exist $\alpha_j > 0$ and a $\delta_j > 0$ such that whenever $|\alpha_j| < \delta_j$ and $\|\mathbf{x} - \mathbf{x}_0\| < \delta_j$

$$g_j(\mathbf{x}) - \alpha_j < 0.$$

Let $\delta_0 = \min\{\delta_j : j = 1, \dots, m\}$ and let $\alpha_0 = \min\{\alpha_j : j = 1, \dots, m\}$. Then

$$g(\mathbf{x}) < \alpha_0 \quad (9)$$

when $\|\mathbf{x}\| < \delta_0$ and $\|\mathbf{x} - \mathbf{x}_0\| < \delta_0$.

We now consider the perturbed equality constraints, which in accordance with (2) we write as

$$\mathbf{A}\mathbf{x} = \mathbf{b} + \mathbf{a}_2, \quad (10)$$

Recall that in Remark 1.1 we assumed that A has rank k . We assume that the square matrix A_1 consisting of the first k columns of A is nonsingular. There is no loss of generality in this assumption since we can relabel the components to achieve this. Let $A_{2,2}$ be the matrix consisting of the last $n - k$ columns of A . Let $\mathbf{x} = (\mathbf{x}_1, \mathbf{x}_{2,2})$, where \mathbf{x}_1 is the vector consisting of the first k components of \mathbf{x} and $\mathbf{x}_{2,2}$ is the vector consisting of the last $n - k$ components of \mathbf{x} . Then

we may write (3) as

$$A_1x_1 + A_{2+1}x_{2+1} = b + d_1.$$

From this equation and the equation $A_1x_1 = b$ we get that

$$(x_1 - x_{1j}) = A_1^{-1}[-d_1 + d_1 + (x_{2+1} - x_{2+1j}) + d_1]. \quad (5)$$

where $x_{1j} = (x_{1j1}, \dots, x_{1j1})$ and $x_{2+1j} = (x_{2+1j1}, \dots, x_{2+1j1})$. It follows from (5) that there exist a $P > 0$ and a $Q < p < p_0$ such that if $\|x_{2+1} - x_{2+1j}\| < \frac{1}{2}p$ and $\|x_1\| < P$, then $\|x_1 - x_{1j}\| < \frac{1}{2}p_0$. Since for any x we have $\|x\|^2 = \|x_1\|^2 + \|x_{2+1}\|^2$, it follows that for every vector x_{2+1} satisfying $\|x_{2+1} - x_{2+1j}\| < \frac{1}{2}p$ and for every vector x_1 satisfying $\|x_1\| < P$ the vector $x = (x_1, x_{2+1})$, where x_1 is given by (5), is such that $\|x - x_{1j}\| < p_0/2$ and (3) is satisfied. The last statement together with (4) show that for every $x = (x_1, x_2)$ with $\|x\| < \min\{Q, P\}$ there exists a vector x_1 that is feasible. Thus Θ is an interior point of $\text{dom}(f)$.

LEMMA 3.1 guarantees that $\text{dom}(f)$ is a convex set and f is a convex function. It also guarantees that if C/P is strongly consistent, $\text{dom}(f)$ has nonempty interior. There is no guarantee, however, that in general $\text{dom}(f)$ will have a nonempty interior. To state two important corollaries to Lemma 3.1, it will therefore be necessary to consider the notions of relative interior and relative boundary, which were introduced in Section 7 of Chapter II.

COROLLARY 1. If $f(x_{1j}) = -\infty$ at some point x_{1j} in $\text{dom}(f)$, then $f(x) = -\infty$ at all points of $\text{dom}(f)$ with the possible exception of relative boundary points.

Proof. It follows from (2) that if $f(x_{1j}) = -\infty$ at some point x_{1j} in $\text{dom}(f)$, then $f(x) = -\infty$ at all points of the half open line segment $[x_{1j}, x_2]$, where $x_2 \neq x_{1j}$ is another point of $\text{dom}(f)$.

Let x_1 be a point of $\text{dom}(f)$ that is not a relative boundary point. Then x_2 is a relative interior point of $\text{dom}(f)$. Hence the line segment $[x_1, x_2]$, which lies in $\text{dom}(f)$, can be extended to a line segment $[x_1, x_2]$ that lies in $\text{dom}(f)$. The point x_1 is interior to $[x_1, x_2]$, and so $f(x_1) = -\infty$.

COROLLARY 2. If f is finite at a relative interior point of its domain, then f is finite over its entire domain.

This follows from Corollary 1, for if f were not finite at some point, it would have to be infinite at all relative interior points of its domain.

Remark 3.1. The only property of f that was used in the proof of Corollary 1 was the convexity of f . Thus Corollaries 1 and 2 are valid for all convex functions admitting $-\infty$ as a value.

Lemma 3.2. Let problem CPH be strongly convex and let $v = \alpha(\theta)$ be finite. Then α is finite on $\text{dom}(\alpha)$, the subdifferential $\partial\alpha(\theta)$ of α is nonempty and each η in $-\partial\alpha(\theta)$ can be written as $\eta = \lambda_1 \eta_1$, where $\lambda_1 \in \mathbb{R}^+$, $\lambda_1 \geq 0$, $\eta_1 \in \mathbb{R}^n$ and such that for all $x = (x_1, x_2)$ in $\text{dom}(\alpha)$

$$\alpha(x) \geq \alpha(\theta) + \langle \eta, x \rangle = \alpha(\theta) + \langle \lambda_1 \eta_1, x \rangle = \langle \lambda_1, x_1 \rangle. \quad (7)$$

Proof. It follows from Lemma 3.1 that θ is an interior point of $\text{dom}(\alpha)$ and that α is a convex function on the convex set $\text{dom}(\alpha)$. Since $\alpha(\theta)$ is finite, it follows from Corollary 2 to Lemma 3.1 that α is finite at all points of $\text{dom}(\alpha)$. Therefore, by Theorem III.2.1, the set of subgradients of α at θ is not empty. Relation (7) then follows by taking $\eta = -\lambda_1$, where λ_1 is any element in $\partial\alpha(\theta)$.

We now show that $\lambda_1 \geq 0$. Since the problem is convex, $\{x : g(x) \leq 0, h(x) = 0\}$ is nonempty. Hence the set $X(\lambda)$ is nonempty for any vector $\lambda = (\lambda_1, \lambda_2)$ with $\lambda_1 \geq 0$ and $\lambda_2 = 0$. Thus, all such vectors are in $\text{dom}(\alpha)$.

If λ_1 were not nonnegative, then at least one component of λ_1 , say λ_{1j} , would be negative. Let $\lambda_1 = (\lambda_{1j}, \lambda_{1j'})$ where $\lambda_{1j} = \lambda_{1j}$, the vector in \mathbb{R}^n all of whose components except the j th are zero and whose j th component is 1. Then λ_1 is in $\text{dom}(\alpha)$ and

$$\alpha(\lambda_1) \geq \alpha(\theta) + \langle \lambda_1, \lambda_1 \rangle = \langle \lambda_1, \lambda_1 \rangle = \alpha(\theta) + \lambda_{1j} = \alpha(\theta), \quad (8)$$

the last inequality holding because $\lambda_{1j} < 0$.

Since $\lambda_{1j} \geq 0$ and $\lambda_{1j} = 0$, we have $X(\theta) \subseteq X(\lambda_1)$. Therefore,

$$\alpha(\lambda_1) = \inf\{f(x) : x \in X(\lambda_1)\} \leq \inf\{f(x) : x \in X(\theta)\} = \alpha(\theta),$$

which contradicts (8). Thus $\lambda_1 \geq 0$.

Remark 3.1. The vector $-\eta$ is any vector in $\partial\alpha(\theta)$ and we need not be unique. By Theorem III.5.1 the vector $-\eta$ is unique if and only if α is differentiable at θ , in which case $-\eta = \nabla\alpha(\theta)$. Thus, $-\eta$ gives the direction of maximum rate of increase at θ of the value function α .

If α is not differentiable, then by Exercise III.2.3,

$$\alpha(\theta, v) \geq \langle -\eta, v \rangle \quad \text{for all } v \text{ in } \mathbb{R}^n,$$

where $\alpha(\theta, v)$ is the directional derivative of α at θ in the direction v . Thus, $-\eta$ again in a sense measures the changes in α as we move away from the origin. The vector η is therefore called a *sensitivity vector*.

We noted in Remark 2.5 that if we know the multipliers (λ_1, λ_2) of Theorem 2.2, then we may replace the constrained problem by an unconstrained problem. In other words, the multipliers (λ_1, λ_2) of Theorem 2.2 are Kuhn-Tucker vectors. By Corollary 1 to Theorem 2.5, if CPH has a solution, then the set X^* of Kuhn-Tucker vectors is equal to the set \mathcal{M} of multiplier vectors.

In the next lemma we shall show that if $\alpha(\mathbb{R})$ is finite, then every vector in $-\text{int}(\mathbb{R})$ is a Karush-Tucker vector.

Lemma 3.1. Let

$$\alpha(\mathbb{R}) = \inf\{f(x) : x \in X\},$$

where $X = \{x \in \mathbb{R}^n : h(x) = 0\}$, is finite and let $\beta \in \mathbb{R}$. Then for each $q = (\lambda, \mu)$ such that $-\eta \in \text{int}(\mathbb{R})$ the vector k is nonnegative and

$$-\alpha(\mathbb{R}) = \inf\{f(x) + \langle \lambda, g(x) \rangle + \langle \mu, h(x) \rangle : x \in \mathbb{R}^n\}.$$

Proof. Since $-\eta \in \text{int}(\mathbb{R})$, for each x in $\text{dom}(f)$

$$\alpha(\mathbb{R}) \geq \alpha(\mathbb{R}) - \langle \eta, x \rangle. \quad (9)$$

From the proof of Lemma 3.2 we get that $\lambda \geq 0$. For all x in \mathbb{R}^n , $g(x) \leq g(x)$ and $h(x) = h(x)$. Therefore $x - \langle g(x), h(x) \rangle$ is in $\text{dom}(f)$. Taking $x = \langle g(x), h(x) \rangle$ in (9) gives

$$-\langle g(x), h(x) \rangle + \langle \lambda, g(x) \rangle + \langle \mu, h(x) \rangle \geq -\alpha(\mathbb{R}) \quad (10)$$

for all x in \mathbb{R}^n . Now

$$-\langle g(x), h(x) \rangle = \inf\{f(x) : x \in N(g(x), h(x))\},$$

where $N(g(x), h(x)) = \{y \in \mathbb{R}^n : y \leq g(x), h(y) = h(x)\}$. But $x \in N(g(x), h(x))$. Hence

$$f(x) \geq -\langle g(x), h(x) \rangle.$$

Substituting this into (10) gives

$$f(x) + \langle \lambda, g(x) \rangle + \langle \mu, h(x) \rangle \geq -\alpha(\mathbb{R})$$

for all x . Hence

$$\alpha(\mathbb{R}) \leq \inf\{f(x) + \langle \lambda, g(x) \rangle + \langle \mu, h(x) \rangle\}.$$

We now establish the reverse inequality and thus prove the lemma. We have

$$\begin{aligned} \inf\{f(x) + \langle \lambda, g(x) \rangle + \langle \mu, h(x) \rangle : x \in \mathbb{R}^n\} \\ \leq \inf\{f(x) + \langle \lambda, g(x) \rangle + \langle \mu, h(x) \rangle : g(x) \leq 0, h(x) = 0\}. \end{aligned}$$

Since $\lambda \geq 0$, $g(x) \leq 0$, and $h(x) = 0$, the right-hand side of this inequality is less than or equal to

$$\inf_{x \in N} (g(x) + \lambda h(x) - \lambda) = \alpha(\mathbb{R}).$$

Remark 3.1. Because of relationships (7) and (9), which bound the change in α from below by a linear function, problems with $\text{int}(\mathbb{R}) \neq \emptyset$ are said to be *stable*.

The next theorem is an immediate consequence of Lemmas 3.2 and 3.3.

Theorem 3.1. Let the programming problem CPM be strongly consistent and let $\alpha(\mathbb{R})$ be finite. Then $\text{int}(\mathbb{R})$ is nonempty and for every $q = (q_1, q_2)$ such that $-q \in \text{int}(\mathbb{R})$ the vector k is nonnegative and

$$\alpha(\mathbb{R}) = \inf_{x \in N} (g(x) + (q_1, q_2)x) + (q_1, k)(x) : x \in \mathbb{R}^n.$$

The next theorem summarizes the relationships among the set $\text{int}(\mathbb{R})$, the multiplier vectors, and the Karush-Tucker vectors.

Theorem 3.2. Let the programming problem CPM be strongly consistent and let it have a solution. Then the set \mathbb{M} of multiplier vectors, the set \mathcal{K} of Karush-Tucker vectors, and the set $-\text{int}(\mathbb{R})$ are all nonempty and equal.

Proof. That \mathbb{M} and \mathcal{K} are nonempty and equal was noted in Corollary 1 of Theorem 2.3 and Remark 2.6. Theorem 3.1 states that if CPM is strongly consistent and $\alpha(\mathbb{R})$ is finite, then $\text{int}(\mathbb{R})$ is nonempty and every vector in $-\text{int}(\mathbb{R})$ is a Karush-Tucker vector.

To complete the proof of the theorem, we must show that each Karush-Tucker vector is an element of $-\text{int}(\mathbb{R})$. Let $q_0 = (q_1, q_2)$ be a Karush-Tucker vector. Since CPM is consistent, $\mathbb{R} \cap \text{dom}(c)$ and hence $\text{dom}(c) \neq \emptyset$. For any x in $\text{dom}(c)$

$$\begin{aligned} \alpha(\mathbb{R}) &= \inf_{x \in N} (g(x) + (q_1, q_2)x) + (q_1, k)(x) : x \in \mathbb{R}^n \\ &\leq \inf_{x \in N} (g(x) + (q_1, q_2)x) + (q_1, k)(x) : x \in \text{dom}(c). \end{aligned}$$

Since $g(x) \leq c_1$ and $h(x) = c_2$ for x in $N(x)$ and since $\lambda_0 \geq 0$, we get that

$$\begin{aligned} \alpha(\mathbb{R}) &\leq \inf_{x \in N} (g(x) + (q_1, q_2)x) + (q_1, k_0)x : x \in \text{dom}(c) \\ &= \alpha(\mathbb{R}) + (q_1, k_1) + (q_2, k_2). \end{aligned}$$

Hence

$$\alpha(\mathbb{R}) \geq \alpha(\mathbb{R}) - (q_1, k_1) - (q_2, k_2)$$

for all x in $\text{dom}(c)$. Thus $q_0 \in -\text{int}(\mathbb{R})$.

The assumption of strong consistency in Theorem 3.2, as already noted, guarantees that the sets \mathbb{W} and $-\text{int}\mathbb{W}$ are nonempty. In any specific problem, the validity of the strong consistency assumption is easily verified a priori from the data of the problem. In the next theorem we show that if we drop the assumption of strong consistency but assume that one of the sets \mathcal{J} , \mathbb{W} , or $-\text{int}\mathbb{W}$ is nonempty, then all of the sets are nonempty and are equal. This assumption, however, is not verifiable a priori.

Theorem 3.3. Let the programming problem (P1) have a solution. Let one of the sets $-\text{int}\mathbb{W}$, the set \mathcal{J} of Karush–Kuhn multipliers, or the set \mathbb{W} of multipliers as in Theorem 3.2 be nonempty. Then all of the sets \mathbb{W} , \mathcal{J} , and $-\text{int}\mathbb{W}$ are nonempty and

$$\mathbb{W} = \mathcal{J} = -\text{int}\mathbb{W}.$$

Proof. By Corollary 1 to Theorem 3.2, the sets \mathbb{W} and \mathcal{J} are equal. By Lemma 3.1, if $\text{int}\mathbb{W}$ is not empty, then $-\text{int}\mathbb{W} \subset \mathcal{J}$. In reverse, if $\text{int}\mathbb{W} = \emptyset$, this is also true. In proving Theorem 3.2, we showed that $\mathcal{J} \subset -\text{int}\mathbb{W}$. Thus, $\mathcal{J} = -\text{int}\mathbb{W}$, and the theorem is proved. \square

Remark 3.4. The relationship between perturbations and multiplier vectors has an interesting economic interpretation. We can consider the function f as giving the cost of an activity described by a vector x , where x is subject to constraints $g(x) \leq b$ and $h(x) = c$. Our objective is to minimize f subject to the constraints. We perturb the constraints by a vector $\alpha = (\alpha_1, \alpha_2)$ in the hope of reducing the optimal cost $\omega(\alpha)$. If we must pay a price $\beta = (\beta_1, \beta_2)$ per unit change in the constraints, then the total cost of the perturbed process with constraints $g(x) \leq b_1$ and $h(x) = c_1$ is

$$\omega(\alpha) + \langle \beta, \alpha \rangle.$$

Thus it will be advantageous to purchase a change in the constraints if and only if

$$\omega(\beta) + \langle \beta, \alpha \rangle < \omega(\alpha).$$

But

$$\omega(\alpha) \geq \omega(\beta) + \langle \beta, \alpha_1 \rangle + \langle \beta, \alpha_2 \rangle, \quad \langle \beta, \beta \rangle \in \mathbb{W}, \quad (\alpha_1, \alpha_2) \in \mathbb{R}^{m+p},$$

and so

$$\omega(\alpha) + \langle \beta, \alpha_1 \rangle + \langle \beta, \alpha_2 \rangle \geq \omega(\beta).$$

Also, any $\langle \beta, \beta \rangle$ for which the last inequality holds is in $-\text{int}\mathbb{W}$, and so is in

If, then we can consider a multiplier vector (λ, μ) to be a price such that the purchase of a perturbation at that price will not result in a lowering of the cost. Since at these prices no perturbation is worth buying, we may consider the cost to be at equilibrium at these prices. Therefore, such prices are called equilibrium prices.

We now illustrate the results of this section by several examples.

Example 1.1. Minimize

$$f(\mathbf{x}) = e^{-2x_1 + x_2}$$

subject to

$$e^{2x_1} + e^{x_2} - 20 \leq \alpha.$$

This problem is a perturbation of the problem in Example IV.5.5. Since f^* and the constraint function $g(\mathbf{x}) = e^{2x_1} + e^{x_2} - 20$ are convex, the unperturbed problem is a convex programming problem, as are the perturbed problems. The problem is strongly consistent since $g^*(0) < 20$.

Let

$$\tilde{L}(\mathbf{x}, \lambda, \alpha) = e^{-2x_1 + x_2} + \lambda(e^{2x_1} + e^{x_2} - 20 - \alpha).$$

It follows from Theorem 2.5 that a solution $\xi = (\xi_1, \xi_2)$ is characterized by the existence of a positive number $\tilde{\lambda}$ such that $\tilde{\lambda}(g(\xi) - \alpha) = 0$ and $\partial L_0/\partial x_1 = \partial L_0/\partial x_2 = 0$, where the partial derivatives are evaluated at $(\xi, \tilde{\lambda})$. Thus

$$-e^{-2\xi_1 + \xi_2} + \tilde{\lambda}e^{2\xi_1} = 0, \quad -e^{-2\xi_1 + \xi_2} + \tilde{\lambda}e^{\xi_2} = 0. \quad (11)$$

From this it follows that $\tilde{\lambda} \neq 0$ and that

$$\tilde{\lambda}e^{2\xi_1} = \tilde{\lambda}e^{\xi_2}.$$

Hence $\xi_1 = \xi_2$. From $\tilde{\lambda} \neq 0$ and $\tilde{\lambda}(g(\xi) - \alpha) = 0$ it follows that

$$e^{2\xi} + e^{2\xi} = 20 + \alpha.$$

Let ζ denote the common value of ξ_1 and ξ_2 . Then

$$2e^{2\zeta} = 20 + \alpha \quad (12)$$

and so $\zeta = \ln(10 + \frac{1}{2}\alpha)$. It follows from (11) and (12) that

$$\tilde{\lambda} = e^{-2\zeta} = (10 + \frac{1}{2}\alpha)^{-2}. \quad (13)$$

It also follows from (12) and from the relation $f(\bar{x}) = z^*$ that

$$w(\bar{x}) = z^* - 10 = (10 + \frac{1}{2}z^*)^{-1}.$$

Hence

$$w(\bar{x}) = -(10 + \frac{1}{2}z^*)^{-2}.$$

From (13) we get that the value of the multiplier λ_0 for the unperturbed problem is 10^{-2} . Also, $w(\bar{x}) = -10^{-2}$ and $w(\bar{x})\bar{x} = -\lambda_0$, in agreement with the theory.

We now show that the solution of the constrained problem can be obtained by solving the unconstrained problem of minimizing $F(x)$, where

$$F(x) = e^{-2(x+10)} + 2\lambda_0 x^2 + x^2 - 20x.$$

Substituting the value of λ_0 in the preceding expression gives

$$F(x) = e^{-2(x+10)} + 10^{-2}x^2 + x^2 - 20x.$$

Since F is convex, any critical point furnishes a minimum. Setting $\partial F/\partial x_1$ and $\partial F/\partial x_2$ equal to zero gives

$$-e^{-2(x+10)} + 2(10^{-2}x) = 0, \quad -e^{-2(x+10)} + (2^{-2}x^2) = 0.$$

Hence $x_1 = x_2$. Solving $x = x_1 = x_2$ gives

$$x^* = 10^2 e^{-20},$$

whence $x = 10$. This is in agreement with the solution of the constrained problem given by (12) with $z = 0$.

Example 5.3. Find the point in the half-space defined by $x_1 + x_2 \leq 1$ that is closest to the origin. This problem has a solution. (Why?)

We may formulate this problem analytically as follows: Minimize $\|x\|^2$ subject to $x_1 + x_2 \leq 1$. Thus the problem is a problem of type CPH(1). If we take $z = 0$, we see that CPH is strongly consistent. On observing the feasible sets for $a \geq 0$, it becomes clear that for $a > 0$ the solution is the origin and that $w(\bar{x}) = 0$. For $a = 0$ the minimum is the point of intersection of the line $x_1 = x_2$ and the line $x_1 + x_2 = 1$. Thus the minimum is achieved at $(\frac{1}{2}, \frac{1}{2})$ and $w(\bar{x}) = \frac{1}{2}a^2$ for $a < 0$. The function w is differentiable at all points, and $w(\bar{x}) = 0$. Therefore 0 is the unique Kuhn-Tucker vector, and the constrained problem CPH can be replaced by the problem of minimizing $\|x\|^2$. The conditions of Theorem 2.5 are clearly satisfied if we take $x_0 = (0, 0)$ and $\lambda_0 = 0$.

We may also formulate the problem analytically as follows: Minimize $|x|$ subject to $x_1 + x_2 \leq \alpha$. Again, the solution is the origin if $\alpha \geq 0$ and the point $(\frac{1}{2}\alpha, \frac{1}{2}\alpha)$ if $\alpha < 0$. The function v is now given by $v(x) = 0$ for $x \geq 0$ and $v(x) = -\alpha/\sqrt{2}$ for $x < 0$. Thus, v is not differentiable at the origin. It is a straightforward calculation to show that $\text{dom}(v)$ is the interval $[-1/\sqrt{2}, 0]$. Hence by Theorem 3.2 the Kuhn-Tucker vectors and multiplier vectors are all points in the interval $(0, 1/\sqrt{2}]$. We now verify this assertion independently.

For λ to be a Kuhn-Tucker vector, we must have

$$\text{inf}\{|\mathbf{x}| + \lambda(x_1 + x_2)\} = 0.$$

When $\alpha = 0$, the quantity in brackets is clearly zero. Therefore, for λ to be a Kuhn-Tucker vector, we require that $\lambda \geq 0$ and

$$\sqrt{x_1^2 + x_2^2} + \lambda(x_1 + x_2) \geq 0 \quad \text{for all } \mathbf{x} \in \mathbb{R}^2.$$

To see what requirement this places on λ , we introduce polar coordinates. The preceding inequality is then satisfied at the origin and whenever

$$\lambda(\cos \theta + \sin \theta) \geq -1, \quad 0 \leq \theta \leq 2\pi. \quad (14)$$

The function $y(\theta) = \cos \theta + \sin \theta$ achieves its minimum on the interval $[0, 2\pi]$ at $\theta = 3\pi/4$ and the value of the minimum is $-\sqrt{2}$. Hence we must have $\lambda \leq 1/\sqrt{2}$ as well as $\lambda \geq 0$.

We can replace the constrained problem by the unconstrained problem

$$\text{min}\{|\mathbf{x}| + \lambda(x_1 + x_2)\}$$

where λ is any point in $(0, 1/\sqrt{2}]$. The preferred choice of λ is clearly $\lambda = 0$. We leave it to the reader to verify that any λ in $(0, 1/\sqrt{2}]$ will serve as a multiplier λ_0 as in Theorem 3.2.

Example 3.2 furnishes another illustration of the advantage of minimizing the square of the distance rather than the distance itself.

Example 3.3. Minimize $f(\mathbf{x}) = x_1$ subject to

$$g_1(\mathbf{x}) = x_1 \leq \alpha, \quad g_2(\mathbf{x}) = x_1^2 - x_2 \leq \alpha_2.$$

When $\alpha = 0$, the problem reduces to the problem considered in Example 2.1. We see from that when $\alpha = 0$, there exists no Kuhn-Tucker vector. By drawing an appropriate figure, the reader can see that $\text{dom}(v) = \{x_1, x_2 \mid x_1 \geq 0\}$ and that $v(x) = -\sqrt{x_1 + x_2}$. We leave the proof of this assertion as an exercise. Thus neither θ nor any point on the line $x_1 + x_2 = 0$ are interior to $\text{dom}(v)$. It is easy to verify that $\text{sup}(v) = \frac{1}{2}$ at each θ on the boundary of $\text{dom}(v)$.

The next example has discontinuous value function at the origin.

Example 3.4. Minimize e^{-x} subject to $\sqrt{x_1^2 + x_2^2} - x_1 \leq z$. For $z < 0$, the feasible set is $\emptyset = \{(x_1, x_2) \mid x_1 \geq 0, x_2 = 0\}$. The minimum is attained at each point of the feasible set since $e^{-x} = e^{-x_1} = 1$. Hence $\text{val}(z) = 1$. If $z = 0$, the feasible set $K(z)$ is empty. If $z > 0$, points x of the feasible set satisfy

$$x_1^2 \leq 2zx + z^2.$$

Thus the feasible set is the set of points on and "inside" the parabola $x_1^2 = 2zx + z^2$. This parabola is symmetric with respect to the x_2 -axis, has vertex at $(-z/2, 0)$, and opens to the right. Hence $\inf\{e^{-x} \mid x \in K(z)\} = 0$, and $\text{val}(z) = 0$ for $z > 0$. Thus, $\text{dom}(\text{val}) = [0, \infty)$, and val is discontinuous at $z = 0$. Also, $\text{epi}(\text{val}) = \mathbb{R}^2$. Note that if $z > 0$, then the infimum is never attained at any point of $K(z)$. Thus the problem $\text{CPI}(z)$ has no solution if $z > 0$.

Exercise 3.1. Prove the assertions in Example 3.3 that

- $\text{dom}(\text{val}) = \{z \mid x_1 + x_2 \geq 0\}$,
- $\text{val}(z) = -\sqrt{z^2 + 2z}$, and
- $\text{val}(z) = 0$ at each z on the boundary of $\text{dom}(\text{val})$.

Exercise 3.2. In each of the following problems find val , $\text{dom}(\text{val})$, and $\text{epi}(\text{val})$ if it is not empty:

- Minimize e^{-x} subject to $-x \leq 0$.
- Minimize x_1 subject to $x_1^2 + x_2^2 \leq 1$.
- Minimize $x_1^2 + x_2^2 + 2x_1 + 2x_2$ subject to $x_1 + x_2 \geq 2$.
- Minimize x_1 subject to $e^{-x_1} = x_2 \leq 0$.

Exercise 3.3. Show directly that there are no Karush–Tucker vectors for the problem with $z = 0$ in Example 3.4.

4. LAGRANGIAN DUALITY

If CPI has a solution and is strongly consistent, then Theorem 2.4 suggests a method for computing the value of the minimum. Let

$$X_0 = \{\eta \mid \eta = 0, \mu \text{ (} \lambda \in \mathbb{R}^m, \lambda \geq 0, \mu \in \mathbb{R}^n)\}. \quad (1)$$

and for each $q \in F_q$ let

$$\theta(q) = \inf\{f(x) + (q, g(x)) + (q_0, h(x)) : x \in \mathbb{R}^n\}. \quad (1)$$

Let

$$\delta = \sup\{\theta(q) : q \in \Sigma_q\}. \quad (2)$$

Then by Theorem 2.4, $\delta = v$, where

$$v = \inf\{f(x) : g(x) \leq 0, h(x) = 0\}. \quad (3)$$

A possible procedure for finding x_{opt} , a point at which the minimum is achieved, is to first find all solutions to the equation $f(x_{opt}) = v$ and then retain those solutions that satisfy the constraints.

This discussion suggests that it might be fruitful to consider the following problem, which we call the *problem dual to CPE*, or simply the *dual problem*, and denote by DCPE.

Problem DCPE

Let

$$F = \{q : q \in F_q, \theta(q) > -\infty\}, \quad (4)$$

where F_q is defined in (1) and $\theta(q)$ in (2). Then

$$\text{Maximize } \{\theta(q) : q \in F\}.$$

Note that in formulating the dual problem it is not assumed that CPE has a solution or is strongly consistent. The set F is the feasible set for DCPE. If $F \neq \emptyset$, then the dual problem is said to be consistent. The original problem CPE is often referred to as the *primal problem*.

The next theorem indicates how the dual problem gives information about the primal problem and vice-versa.

THEOREM 4.1 (LAGRANGIAN DUALITY THEOREM FOR CONVEX PROGRAMMING).

- (i) If $X \neq \emptyset$ and $F \neq \emptyset$, then for each $x \in X$ and $q \in F$

$$\theta(q) \leq f(x). \quad (5)$$

Moreover, if δ and v are as in (2) and (3), then both are finite and $\delta \leq v$.

- (ii) Let $F \neq \emptyset$. If $\theta(q)$ is unbounded above on F , then $X = \emptyset$.
 (iii) Let $X \neq \emptyset$. If $f(x)$ is unbounded below on X , then $F = \emptyset$.
 (iv) If there exists an $q_0 \in F$ and an $x_0 \in X$ such that $f(x_0) = \theta(q_0)$, then

$$\delta = \theta(q_0) = f(x_0) = v.$$

Thus x_0 is a solution of CPE and q_0 is a solution of DCPE.

Proof. The key to the proof of the theorem is the observation that since $\mathbf{A} \geq \mathbf{0}$ for $\mathbf{a}_i \in F$ and since $g(\mathbf{x}) \leq \theta$ and $h(\mathbf{x}) = \theta$ for $\mathbf{x} \in X$, it follows that for $\mathbf{x} \in X$ and $\mathbf{a}_i \in F$

$$f(\mathbf{x}) + \langle \mathbf{a}_i, g(\mathbf{x}) \rangle + \langle \mathbf{a}_i, h(\mathbf{x}) \rangle \leq f(\mathbf{x}). \quad (7)$$

In (7) we take the infimum over all \mathbf{x} in E . This infimum is less than the infimum in (2) taken over all $\mathbf{x} \in X$. From these observations and from (7) the inequalities (i) and (i') follow. Assertion (ii) follows from (i), for if K were not empty, then $\theta(g)$ would be bounded above. Also, (iii) follows from (ii) by a similar argument. Statement (iv) follows from the chain

$$\theta(\mathbf{a}_i) \leq \theta \leq \gamma \leq f(\mathbf{x}_i) = \theta(\mathbf{a}_i).$$

We now see how the dual problem gives useful information about the primal problem and vice-versa. Statement (i) shows how a feasible point for the dual problem gives a lower bound for the value of the primal problem and how a feasible point for the primal problem gives an upper bound for the value of the dual problem. Statement (ii) gives a way of checking whether $K = \emptyset$ in the absence of more direct checks. Statement (iii) gives a condition under which the dual problem cannot be formulated.

Definition 4.1. Problems CPE and DCPH are said to exhibit a duality gap if $\theta < \gamma$.

The next theorem gives a sufficient condition for the absence of a duality gap and relates the multipliers $\mathbf{a}_i = (\mathbf{a}_i, \mu_i)$ to the solutions of the dual problem.

Theorem 4.2. Let the primal problem CPE be strongly consistent and have a solution. Then

- (i) There is no duality gap.
- (ii) If $\mathbf{a}_i = (\mathbf{a}_i, \mu_i)$ is a multiplier for CPH, then \mathbf{a}_i is a solution of the dual problem.
- (iii) If $\mathbf{a}_i = (\mathbf{a}_i, \mu_i)$ is a solution of the dual problem, then \mathbf{a}_i is a multiplier vector for the primal problem.

Proof. Let \mathbf{x}_i denote the point at which CPE achieves its minimum. Since the problem is strongly consistent by Theorem 3.2, there exists a multiplier vector $\mathbf{a}_i = (\mathbf{a}_i, \mu_i)$ for CPE. Thus,

$$f(\mathbf{x}_i) \leq f(\mathbf{x}) + \langle \mathbf{a}_i, g(\mathbf{x}) \rangle + \langle \mu_i, h(\mathbf{x}) \rangle \quad (8)$$

for all \mathbf{x} in E . Hence, from (2), $\theta(\mathbf{a}_i) \geq f(\mathbf{x}_i)$. By Theorem 4.1, $\theta(\mathbf{a}_i) \leq f(\mathbf{x}_i)$, so $\theta(\mathbf{a}_i) = f(\mathbf{x}_i)$. Conclusions (i) and (ii) now follow from (iv) of Theorem 4.1.

Now let $\mathbf{q}_0 = (\lambda_0, \mathbf{p}_0)$ be a solution of the dual problem. From (2) and from the absence of a duality gap we get that

$$\inf\{f(x) + (\lambda_0, \mathbf{p}_0)(x) : x \in \mathbb{R}^n\} = \theta(\mathbf{q}_0) = f(\mathbf{x}_0).$$

Thus $(\lambda_0, \mathbf{p}_0)$ is a Kuhn-Tucker vector, and so by Theorem 2.1, $(\lambda_0, \mathbf{p}_0)$ is a multiplier vector for the primal problem.

We now present several examples to illustrate the theory developed in this section.

Example 4.1. We take the primal problem to be the problem in Exercise IV.5.9 or the problem in Example 3.1 with $\varepsilon = 0$. The primal problem is

$$\begin{aligned} \text{Minimize } f(x) &= e^{-2x+2x^2} \\ \text{subject to } x^2 + x^4 &\leq 20. \end{aligned}$$

Let

$$L(x, \lambda) = e^{-2x+2x^2} + \lambda(x^2 + x^4 - 20).$$

Then for each $\lambda \geq 0$ define

$$\theta(\lambda) = \inf\{L(x, \lambda) : x \in \mathbb{R}^1\}.$$

The dual problem is to maximize $\theta(\lambda)$.

We saw in Example 3.1 that the primal problem is strongly consistent and that it has a solution at $\mathbf{x}_0 = \frac{1}{2}, \frac{1}{2}$, where $\frac{1}{2} = \ln 10$. The value of the minimum is 10^{-2} and the value of the multiplier λ_0 is 10^{-2} . Thus we would expect the value of the maximum of θ to be 10^{-2} and for θ to assume this value at $\lambda_0 = 10^{-2}$. We now show that this is indeed so.

We first determine the function θ . For each $\lambda \geq 0$, the function $L(\cdot, \lambda)$ is readily seen to be a strictly convex function of x . Hence a necessary and sufficient condition that $L(\cdot, \lambda)$ attain a unique minimum at $\frac{1}{2} = (\frac{1}{2}, \frac{1}{2})$ is that $\partial L/\partial x$, and $\partial L/\partial x$, both equal zero at $\frac{1}{2}$. As in Example 3.1 we see that for $\lambda \neq 0$ this implies that $\frac{1}{2} = \frac{1}{2}$. Denoting this common value by $\frac{1}{2}$, we get that

$$\theta(\lambda) = \inf\{g(\lambda, \frac{1}{2}) : \lambda \in \mathbb{R}\}, \quad (9)$$

where

$$g(\lambda, \frac{1}{2}) = e^{-1+2\lambda} + 2\lambda e^{\lambda} - 10.$$

For fixed positive λ , $g(\cdot, \frac{1}{2})$ is a strictly convex function of λ . A necessary and sufficient condition that $g(\cdot, \frac{1}{2})$ attain a unique minimum at a point λ_0 is that

$\lambda \phi(\lambda)$ vanish at λ_0 . Now

$$\frac{d\phi}{d\lambda} = -2e^{-2\lambda} + 2\lambda e^{\lambda}.$$

Setting $\lambda \phi(\lambda) = 0$, we get that $\exp(-\lambda_0) = \lambda_0^2$. Substituting this into (9) gives

$$\theta(\lambda) = 2\lambda^{2\lambda} - 2\lambda\lambda, \quad \lambda > 0.$$

Since $\theta(\lambda) = 0$ when $\lambda = 0$, the preceding formula is valid for all $\lambda \geq 0$. An elementary calculus argument shows that θ is maximized at the unique point at which $\theta'(\lambda) = 0$. This value of λ is readily seen to be $1/2$ and the corresponding value of θ is $1/2$. These values are as required.

Example 4.2. Let the primal problem be as follows: Minimize $\|x\|$ subject to $x_1 + x_2 \leq 0$. This problem was analyzed in Example 3.2 with $\gamma = 0$. This problem is strongly consistent and has $x = 0$ as its solution. The value of the minimum is clearly zero. Any λ in the interval $[0, 1/\sqrt{2}]$ is a multiplier vector.

The dual problem is as follows: Maximize $\{\theta(\lambda) : \lambda \geq 0, \theta(\lambda) = -\infty\}$, where $\theta(\lambda) = \inf\{\|x\| + \lambda(x_1 + x_2) : x \in \mathbb{R}^2\}$. We introduce polar coordinates (r, ϕ) and obtain

$$\theta(\lambda) = \inf\{r\} + \lambda(r \cos \phi + \sin \phi) : r \geq 0, 0 \leq \phi \leq 2\pi\}.$$

Let

$$\gamma(\phi, \lambda) = 1 + \lambda \cos \phi + \sin \phi.$$

Then $\theta(\lambda) = 0$ if $\gamma(\phi, \lambda) \geq 0$ for all $0 \leq \phi \leq 2\pi$, and $\theta(\lambda) = -\infty$ if this is not so and $r \neq 0$. The argument used in Example 3.2 shows that $\gamma(\phi, \lambda) \geq 0$ for all ϕ in $[0, 2\pi]$ if and only if $0 \leq \lambda \leq 1/\sqrt{2}$. Thus $\theta(\lambda) = 0$ for λ in $[0, 1/\sqrt{2}]$ and $\theta(\lambda) = -\infty$ otherwise. Hence $\max\{\theta(\lambda) : \lambda \in F\} = 0$ and the maximum is attained at each $\lambda \in [0, 1/\sqrt{2}]$. Comparison with the solution to the primal problem shows that there is no duality gap and the points at which the dual problem obtains its maximum are the multipliers for the primal problem.

Example 4.3. Let the primal problem be the problem treated in Example 2.1 and in Example 3.1 with $\gamma = 0$. The primal problem has the origin as the only feasible point. This point furnishes the solution and the value of the minimum is zero. This problem is clearly not strongly consistent.

The dual problem is as follows: Maximize $\{\theta(\lambda) : \lambda \in Y\}$, where

$$Y = \{\lambda, \lambda \geq 0, \theta(\lambda) > -\infty\} \quad \text{and} \quad \theta(\lambda) = \inf\{x_1 + \lambda_1 x_1 + \lambda_2(x_1^2 - x_2) : x \in \mathbb{R}^2\}.$$

If we write

$$x_1 + \lambda_1 x_2 + \lambda_2(x_1^2 - x_2^2) - x_1 = x_1 + \lambda_2 x_1^2 + (\lambda_1 - \lambda_2)x_2,$$

we see that if $\lambda_1 \neq \lambda_2$, then $\theta(\lambda) = -\infty$. For $\lambda \in \mathbb{R}$ with $\lambda_1 = \lambda_2$, we have $\lambda_2 \geq 0$ and

$$\theta(\lambda) = \inf\{x_1 + \lambda_2 x_1^2 : x_1 \in \mathbb{R}\}.$$

The function inside the curly brackets attains its minimum at $x_1 = -1/(2\lambda_2)$ and

$$\theta(\lambda) = -1/(4\lambda_2). \quad (F)$$

For $\lambda = 0$ we have

$$\theta(\mathbf{0}) = -\infty.$$

Thus $F = \{\lambda = (\lambda_1, \lambda_2) : \lambda_1 = \lambda_2, \lambda_2 \geq 0\}$. Since $\lambda_2 \geq 0$, it follows from (F) that $\sup\{\theta(\lambda) : \lambda \in F\} = 0$. The supremum is not attained in the dual problem, even though there is no duality gap. Thus we have an example in which there is no duality gap, the primal problem has a solution but the dual problem has no solution. We also recall from Example 3.3 that $\text{int}(\mathbf{0}) = \emptyset$ and there are no Kuhn-Tucker vectors.

Example 3.4. In this example we will exhibit a duality gap. Let the primal problem be the problem in Example 3.4 with $\alpha = 1$. We saw there that the minimum is attained at every point of the feasible set and that the value of the minimum is 1.

The dual problem is to maximize $\theta(\lambda) : \lambda \geq 0, \theta(\lambda) > -\infty$, where

$$\theta(\lambda) = \inf\{e^{-x_1} + \lambda_1\sqrt{x_1^2 + x_2^2} - x_1 : x_1, x_2 \in \mathbb{R}\}.$$

Let

$$f(x, y) = e^{-x} + \lambda_1\sqrt{x^2 + y^2} - x.$$

Then $f(x, y) \geq 0$ for all x in \mathbb{R}^2 . We shall show that $\theta(\lambda) = 0$.

Let $\lambda_1 > 0$ and let $x_2 > 0$. Then

$$\begin{aligned} \theta(\lambda) &= e^{-x_1} + \frac{\lambda_1 x_2^2}{\sqrt{x_1^2 + x_2^2} + x_1} \\ &= e^{-x_1} + \lambda_1 \frac{x_2^2}{x_1 \left[\sqrt{1 + \left(\frac{x_2}{x_1}\right)^2} + 1 \right]}. \end{aligned}$$

Now let $x_1 = (x_2)^2$ and then let $x_2 \rightarrow \infty$. We see that $f(x, \lambda) \rightarrow 0$ along this curve. Thus $\theta(\lambda) = 0$ for all $\lambda \geq 0$ and $\max\{\theta(\lambda) : \lambda \geq 0\} = 0$. Thus $\delta = 0$, and since $v = 1$, we have a duality gap.

Theorem 4.2 states that δ a solution to the primal problem exists, then strong consistency is a sufficient condition for the absence of a duality gap and the existence of a solution to the dual problem. The following example shows that strong consistency is not a necessary condition for the absence of a duality gap and for the existence of a solution to the dual problem.

Example 4.3. Minimize $f(x) = x_1$ subject to

$$g_1(x) = (x_1 + 1)^2 + (x_2)^2 - 1 \leq 0, \quad g_2(x) = -x_1 \leq 0.$$

The problem is clearly a convex programming problem. The feasible set consists of the origin \bar{x} in \mathbb{R}^2 , and the problem is not strongly consistent. The solution to the primal problem is $x_1 = \bar{x}$ and $v = 0$.

The dual objective function is given by

$$\theta(\lambda) = \inf\{x_1 + \lambda_1 [(x_1 + 1)^2 + (x_2)^2 - 1] - \lambda_2 x_1 : (x, \lambda) \in \mathbb{R}^2\},$$

where $\lambda \geq 0$. If we take $\lambda = \lambda_0 = (0, 1)$ we get that $\theta(\lambda_0) = 0$. It then follows from (a) of Theorem 4.1 that there is no duality gap and that λ_0 is a solution of the dual problem.

We will now obtain two necessary and sufficient conditions for the absence of a duality gap and the existence of a solution to the dual problem.

Theorem 4.1. Let problem CPM have a solution x_0 . Then a necessary and sufficient condition for the absence of a duality gap and for $\mu_0 = (\lambda_0, \mu_0)$, $\lambda_0 \geq 0$, to be a solution of the dual problem is that (x_0, λ_0, μ_0) is a saddle point for L .

Proof. Let (x_0, λ_0, μ_0) be a saddle point for L . Then (2.16) and (2.18) of Theorem 2.7 hold. Thus (b) of Theorem 4.2 holds. That there is no duality gap and that (λ_0, μ_0) is a solution to the dual problem follows by the same arguments used to establish (i) and (ii) in Theorem 4.2.

Now let $\mu_0 = (\lambda_0, \mu_0)$ be a solution of the dual problem and let there be no duality gap. Then, as in the proof of statement (ii) in Theorem 4.2 with $\lambda_0 = \lambda_0$ and $\mu_0 = \mu_0$, we get that (λ_0, μ_0) is a multiplier (as in Theorem 2.2 for the primal problem). But (7) of Theorem 2.2 states that (x_0, λ_0, μ_0) is a saddle point for L .

Theorem 4.4. Let CPM have a solution x_0 . Then a necessary and sufficient condition for the absence of a duality gap and the existence of a solution to the dual problem is that $\text{int}W \neq \emptyset$. In this case every vector q in $-\text{int}W$ is a solution of the dual problem.

Proof. We first prove the sufficiency of the condition $\text{dim}(\mathcal{K}) = \mathcal{G}$. Since the primal problem has a solution, $\text{inf}(\mathcal{K})$ is finite. By hypothesis, $\text{dim}(\mathcal{K}) = \mathcal{G}$. Hence by Lemma 3.3 each vector $\eta = (\lambda, \mu)$ such that $-\eta \in \text{int}(\mathcal{K})$ is a Kuhn-Tucker vector, and such vectors exist. By Theorem 3.3 each such vector is in \mathcal{R} , the set of multiplier vectors, as in Theorem 2.5. Thus (x_0, λ, μ) is a saddle point for L , and so by Theorem 4.3 there is no duality gap and (λ, μ) is a solution to the dual problem. Now suppose that there is no duality gap and that the dual problem has a solution $\lambda_0 = (\lambda_0, \mu_0)$. Then by Theorem 4.3 the point (x_0, λ_0, μ_0) is a saddle point for L . By Theorem 2.7, $(\lambda_0, \mu_0) \in \text{int}(\mathcal{K})$ is a Kuhn-Tucker vector. It then follows from Theorem 3.3 that $(\lambda_0, \mu_0) \in -\text{int}(\mathcal{K})$.

COROLLARY 1. Let CPM have a solution x_0 . Then a necessary and sufficient condition for the absence of a duality gap and the existence of a solution to the dual problem is that one of the sets $\text{int}(\mathcal{K})$, \mathcal{K} , or \mathcal{R} be nonempty.

The corollary is an immediate consequence of Theorem 4.4 and 3.3.

Remark 4.1. In Remark 3.4 we indicated an economic interpretation of the multiplier vectors as equilibrium prices. In view of the preceding corollary, we see that the solution of the dual problem gives the equilibrium prices.

Exercise 4.1. Formulate, and solve whenever possible, the problems dual to the problems in Exercise 3.5. Discuss your results as they relate to Theorem 4.1.

Exercise 4.2. The problem posed in Exercise II.5.8 is a convex programming problem. Solve this problem by formulating the dual and solving the dual. Hint: To simplify the algebra involved, first solve the problem for $\lambda_0 = \mathbf{0}$. Then use a translation of coordinates to obtain the result for general λ_0 .

Exercise 4.3. Show that the function f defined by (2.7) is convex.

8. GEOMETRIC INTERPRETATION

In this section we present a geometric interpretation of the duality and perturbation results. We shall assume that all the constraints are inequality constraints and that we are considering CPM: Minimize $f(x)$ subject to $g(x) \leq \mathbf{0}$.

The mapping

$$[\cdot] = f(x), \quad x = g(x),$$

is a mapping from \mathbb{R}^n to \mathbb{R}^{n+1} . We have indicated two image sets,

$$I = \{(\mathbf{0}, \zeta), \quad \zeta = f(x), \quad \mathbf{0} = g(x), x \in \mathbb{R}^n\}, \quad (1)$$

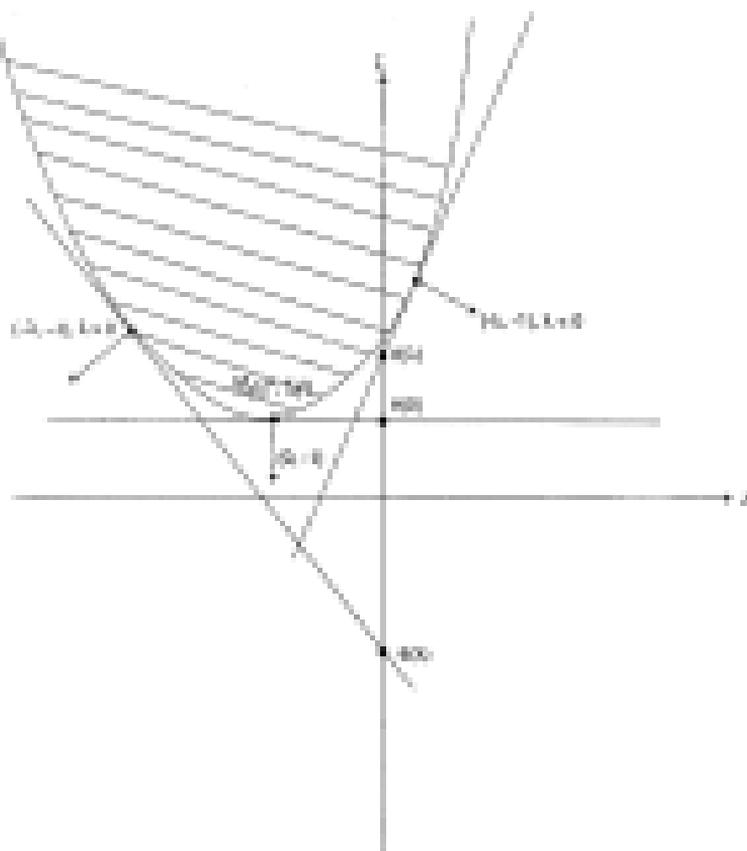


Figure 5.1.

schematically for $m=1$ as the hatched area in Figures 5.1 and 5.2. In geometric terms, to solve CPl, we must find a point in the intersection of F and the closed half space $x \leq \theta$ with smallest C coordinate if such a point exists. In Figure 5.1 this is the point M_1 with coordinates (x^*, z^*) . In Figure 5.2 this is the point M_2 with coordinates (θ, z^*) .

Let $\lambda \geq 0$. The dual problem is to maximize $f(\lambda)$, where f is defined in relation (2) of Section 4. We may therefore also write

$$\begin{aligned} f(\lambda) &= \inf\{z \mid (x, z) \in F\} = f(x), \quad F = g(x) \\ &= \inf\{z \mid z \in J + (x, x) \cap (0, J) + \delta\}. \end{aligned}$$

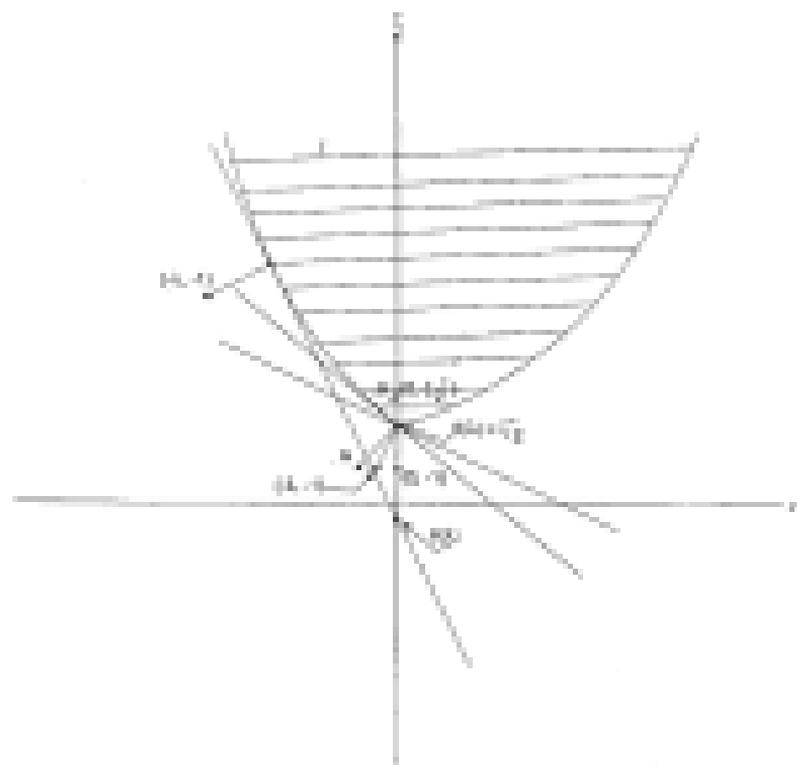


Figure 3.1.

Thus, for all (z, x) in F

$$(z - z_0 - 1) \text{ or } (z) \leq -\theta(z, x)$$

and so the hyperplane

$$(z - z_0 - 1) \text{ or } (z) = -\theta(z, x) \quad (2)$$

is a supporting hyperplane for the set F . The equation of this hyperplane can also be written as

$$z = (z_0 + 1) + \theta(z, x) \quad (3)$$

Hence $\theta(z, x)$ is the (z) -intercept of the hyperplane (2). See Figures 3.1 and 3.2. The

geometric interpretation of the dual problem is to find, among supporting hyperplanes (C) to f with $\lambda \geq 0$, those whose (λ) -intercept is maximized provided such hyperplanes exist.

In Figure 5.1 the supporting hyperplane to f at any point P with x -coordinate greater than x^* has $\lambda = 0$. Thus these hyperplanes do not enter into the competition. The unique hyperplane through M_1 with normal $(0, -1)$ is the required hyperplane. The (λ) -intercept of this horizontal hyperplane is $\frac{c^2}{4}$. Thus there is no duality gap and the dual problem has a solution. The point N_1 corresponds to a point x_1 such that $g(x_1) = 0$. Therefore, from the complementary slackness condition $(\lambda_1, g(x_1, 0)) = 0$, we have that $\lambda_1 = 0$. This is in agreement with the geometry.

In Figure 5.2 any hyperplane through M_2 whose normal lies in the angle formed by \mathbf{n} and the vector $(0, -1)$ is a supporting hyperplane with $\lambda \geq 0$ having maximal (λ) -intercept equal to $\frac{c^2}{4}$. Thus, again there is no duality gap and the dual problem has a solution.

From Figures 5.1 and the definitions of the function ω we see that the graph of ω is the lower boundary curve of J for $x < x_1^*$ and then the straight line $C = \frac{c^2}{4}$ for $x \geq x_1^*$. At M_1 the function ω is differentiable and $\omega'(x) = 0$. The dual problem has the unique solution $\lambda_1 = 0$, and so $\omega'(x) = \lambda_1$.

In Figure 5.2 the graph of ω is the lower boundary curve of f for $x < 0$ and then the straight line $C = \frac{c^2}{4}$ for $x \geq 0$. At N_2 the function ω is not differentiable. The set of vectors $\lambda_2 \geq 0$ such that $(-\lambda_2, -1)$ lies in the angle formed by \mathbf{n} and $(0, -1)$ is the set $\text{co}(N_2)$. The dual problem attains its solution at any vector in $-\text{co}(N_2)$.

Exercise 5.1. In Exercises 3.13, 3.15, and 3.16 sketch the sets J and carry out the geometric arguments of this section in these specific problems, whenever possible. For problems in which the arguments made in connection with Figures 5.1 and 5.2 seem to fail, explain why.

6. QUADRATIC PROGRAMMING

Quadratic Programming Problem (QP)

Let Q be a positive semidefinite $n \times n$ matrix, let \mathbf{h} be a vector in \mathbb{R}^n , let A be an $m \times n$ matrix, and let \mathbf{r} be a vector in \mathbb{R}^m . Minimize

$$f(\mathbf{x}) = \frac{1}{2}(\mathbf{x}, Q\mathbf{x}) - (\mathbf{h}, \mathbf{x})$$

subject to

$$A\mathbf{x} \leq \mathbf{r}.$$

The formulation just given includes equality constraints of the form $A_i\mathbf{x} = d_i$. This is achieved by the device already used of writing such equality constraint $(\mathbf{k}_i, \mathbf{x}) = d_i$, where \mathbf{k}_i is the i th row of A as two equality constraints.

Because \tilde{Q} is positive semidefinite, the function f is convex. The constraints are affine. Hence the problem is a convex programming problem. We leave it as an exercise to show that if the matrix \tilde{Q} is positive definite, the problem does have a solution, while the problem need not have a solution if \tilde{Q} is only positive semidefinite. If \tilde{Q} is positive definite, then f is strictly convex. Since in this case we are minimizing a strictly convex function over a convex set, the minimum is achieved at a unique point x_0 .

For the quadratic programming problem the Lagrangian L is defined by

$$L(x, \lambda) = \frac{1}{2}\langle x, \tilde{Q}x \rangle - \langle b, x \rangle + \langle \lambda, Ax - c \rangle. \quad (1)$$

THEOREM 4.1. A necessary and sufficient condition for a point x_0 to be a solution of QP is that

$$Ax_0 \leq c \quad (2)$$

and there exist a vector $\lambda_0 \geq 0$ in \mathbb{R}^m such that

$$\tilde{Q}x_0 - b + A^T\lambda_0 = 0, \quad (3)$$

$$\langle \lambda_0, c - Ax_0 \rangle = 0. \quad (4)$$

Proof. We first prove the necessity of the conditions. If x_0 is a solution of QP, then x_0 is feasible, so (2) holds. By Theorem IV.3.1 or by Remark 2.3, there exists a real number $\lambda_0 \geq 0$ and a vector $\lambda_0 \geq \mathbf{0}$ in \mathbb{R}^m such that $(\lambda_0, b_0) \in F(\lambda, \mathbf{0})$.

$$\lambda_0(\tilde{Q}x_0 - b) + A^T\lambda_0 = 0 \quad (5)$$

and (4) holds. We shall show that $\lambda_0 = 0$ is not possible. Hence we may take $\lambda_0 = 1$ and get (3).

If $\lambda_0 = 0$, then $\lambda_0 \neq \mathbf{0}$ and from (5) we get that $A^T\lambda_0 = 0$. It then follows from (4) that $\langle \lambda_0, c \rangle = 0$. Thus the system

$$\begin{pmatrix} A^T \\ c^T \end{pmatrix} y = 0, \quad y \geq 0, \quad y \neq 0, \quad y \in \mathbb{R}^n,$$

has a solution $y = \lambda_0 \neq 0$. Then, follows from Gordan's theorem (Theorem II.3.2) that

$$(A, -c) \begin{pmatrix} x \\ 1 \end{pmatrix} \geq 0$$

for all $x \in \mathbb{R}^n$, $\zeta \in \mathbb{R}$. Hence $Ax \geq \zeta c$ for all x in \mathbb{R}^n and $\zeta \in \mathbb{R}$. In particular,

this inequality holds for $\lambda = 0$. Hence $Ax \geq b$ and $A(-x) \geq b$ for all x in \mathbb{R}^n . Thus $Ax = b$ for all x in \mathbb{R}^n , which is not possible since A is not the zero matrix.

The sufficiency of the conditions follows from Theorem IV.3.1 with $f(x) = \frac{1}{2}(x, Qx)$ and $g(x) = Ax - b$.

Remark 6.1. In Theorem 6.1 we showed that in the quadratic programming problem the KKT condition is a necessary condition for a minimum without assuming a priori that a constraint qualification holds. The linearity of the constraints, however, played a crucial role.

Theorem 6.1 suggests the following procedure for solving QP. Solve the linear system (2) and then reject all solutions of (2) that do not satisfy (3) and (4). In actual practice, however, various computational algorithms related to the simplex method have been devised to solve quadratic programming problems. For a discussion of these see Rosen, Solov'yev, and Shorukh [1981].

An alternate, and possibly more fruitful, approach is through the dual problem, which is

$$\max\{\theta(\lambda) : \lambda \geq 0, \theta(\lambda) \geq -\alpha\}, \quad (5)$$

where

$$\theta(\lambda) = \inf\{L(x, \lambda) : x \in \mathbb{R}^n\} \quad (6)$$

and L is defined in (1). A justification for this procedure is provided by the following theorem.

Theorem 6.2. Let Problem QP have a solution x_0 . Then there is no duality gap between QP and the dual problem, the dual problem has a solution, and the solution λ given by the multipliers λ_0 of Theorem 6.1.

Proof. For fixed $\lambda \geq 0$, the function $L(\cdot, \lambda)$ defined in (1) is convex on \mathbb{R}^n . Therefore, by the corollary to Theorem IV.3.2, a necessary and sufficient condition for a point x_0 to minimize L is that $\nabla_x L(x_0, \lambda) = 0$, where ∇_x denotes the gradient with respect to x . Since x_0 is a solution to QP, it follows from (4) of Theorem 6.1 that $\nabla_x L(x_0, \lambda_0) = 0$. Hence,

$$L(x_0, \lambda_0) = \min\{L(x, \lambda_0) : x \in \mathbb{R}^n\}.$$

From this and from (2) and (4) we see that (2.10) of Theorem 2.7 holds, where \mathbf{h} is the zero function. Thus, (x_0, λ_0) is a saddle point for L . The conclusion of the theorem now follows from Theorem 4.1.

Remark 6.2. Note that we have not assumed strong consistency, but we have assumed that QP has a solution. We have already remarked that a sufficient condition for the existence of a solution is that Q be positive definite.

In determining the explicit form of the dual problem, it is useful to consider the unconstrained quadratic programming problem: Minimize

$$f(\mathbf{x}) = \frac{1}{2}(\mathbf{x}, \mathbf{Q}\mathbf{x}) - (\mathbf{b}, \mathbf{x}), \quad \mathbf{x} \in \mathbb{R}^n. \quad (8)$$

Since \mathbf{Q} is positive semidefinite, there exists an orthogonal matrix P such that $P\mathbf{Q}P^T$ is a diagonal matrix. Under the transformation of coordinates $\mathbf{y} = P\mathbf{x}$, (8) becomes

$$F(\mathbf{y}) = f(P^T\mathbf{y}) = \frac{1}{2}\mathbf{y}^T P\mathbf{Q}P^T\mathbf{y} - \mathbf{b}^T P\mathbf{y} = \frac{1}{2}(\alpha_1 y_1^2 + \cdots + \alpha_k y_k^2) + \sum_{i=k+1}^n \alpha_i y_i,$$

where the α_i are the positive eigenvalues of \mathbf{Q} . If \mathbf{b} is positive definite, then $k = n$; if \mathbf{b} is positive semidefinite, then $k < n$. If a coefficient α_i with $i > k$ is different from zero, then by considering vectors $\rho\mathbf{e}_i$, where ρ is a scalar and \mathbf{e}_i is the vector whose i th component is 1 and whose other components are zero, we see that $\inf\{F(\mathbf{y}) : \mathbf{y} \in \mathbb{R}^n\} = -\infty$. If $\alpha_i = 0$ for all $i > k$, then $\nabla F(\mathbf{y}_0) = \mathbf{0}$ at any point $\mathbf{y}_0 = (y_{01}, \dots, y_{0n})$ with

$$y_{0i} = \begin{cases} \frac{b_i}{\alpha_i}, & i = 1, \dots, k, \\ \text{arbitrary}, & i = k+1, \dots, n. \end{cases}$$

Since F is convex, it follows from the corollary to Theorem IV.1.1 that F attains a minimum at any such point \mathbf{y}_0 . If \mathbf{Q} is positive definite, then $k = n$ and the point \mathbf{y}_0 is unique. We summarize this discussion in the following lemma.

LEMMA 4.1. The unconstrained programming problem either has infimum equal to $-\infty$ or attains a minimum at a finite point.

We return to the determination of the explicit form of the dual problem. For fixed $\lambda \in \mathbb{R}$, the problem of minimizing $L(\mathbf{x}, \lambda)$ over all \mathbf{x} in \mathbb{R}^n is an unconstrained quadratic programming problem, since

$$\begin{aligned} L(\mathbf{x}, \lambda) &= \frac{1}{2}(\mathbf{x}, \mathbf{Q}\mathbf{x}) - (\mathbf{b}, \mathbf{x}) + (\mathbf{b}, \lambda\mathbf{x} - \mathbf{c}) \\ &= \frac{1}{2}(\mathbf{x}, \mathbf{Q}\mathbf{x}) - (\mathbf{b} - \lambda\mathbf{A}\mathbf{x}, \mathbf{x}) - (\mathbf{b}, \mathbf{c}) \\ &= \frac{1}{2}(\mathbf{x}, \mathbf{Q}\mathbf{x}) - (\mathbf{d}, \mathbf{x}) - (\mathbf{b}, \mathbf{c}), \end{aligned}$$

where $\mathbf{d} = \mathbf{b} - \lambda\mathbf{c}$. It then follows from Lemma 4.1 that if $\mathbf{Q}\mathbf{x}$ is finite, then we may replace the infimum in (7) by a minimum. Thus, for fixed λ such that

$\theta(\lambda) \geq -\alpha$, if we denote a point at which $\min \{L(\cdot, \lambda)\}$ is attained by \mathbf{x}_λ , then

$$\theta(\lambda) = L(\mathbf{x}_\lambda, \lambda) = \frac{1}{2}(\mathbf{x}_\lambda, Q\mathbf{x}_\lambda) - (\mathbf{b}, \mathbf{x}_\lambda) + (\lambda, A\mathbf{x}_\lambda - \mathbf{c}). \quad (9)$$

Since for fixed λ , the function $L(\cdot, \lambda)$ is convex on \mathbb{R}^n , it follows from the corollary to Theorem IV.3.2 that the minimum is attained at a point \mathbf{x}_λ if and only if $\nabla_{\mathbf{x}} L(\mathbf{x}_\lambda, \lambda) = \mathbf{0}$. A straightforward calculation shows that $\nabla_{\mathbf{x}} L(\mathbf{x}_\lambda, \lambda) = \mathbf{0}$ if and only if

$$Q\mathbf{x}_\lambda = \mathbf{b} + A^T\lambda. \quad (10)$$

If we take the inner product of each side of (10) with \mathbf{x}_λ , we get

$$(\mathbf{x}_\lambda, Q\mathbf{x}_\lambda) - (\mathbf{b}, \mathbf{x}_\lambda) + (\lambda, A\mathbf{x}_\lambda) = 0.$$

Substitution of this equation into (9) gives

$$\theta(\lambda) = -\frac{1}{2}(\mathbf{x}_\lambda, Q\mathbf{x}_\lambda) - (\lambda, \mathbf{c}).$$

Thus the explicit formulation of the dual problem is as follows.

Dual Quadratic Programming Problem (DQP)

Maximize

$$-\frac{1}{2}(\mathbf{x}, Q\mathbf{x}) - (\lambda, \mathbf{c}) \quad (11)$$

subject to

$$Q\mathbf{x} + A^T\lambda = \mathbf{b}, \quad \lambda \geq \mathbf{0}. \quad (12)$$

If Q is positive definite, then the dual problem can be simplified. From (12) we get that

$$\mathbf{x} = Q^{-1}(\mathbf{b} - A^T\lambda). \quad (13)$$

If we substitute this into (11), then a straightforward calculation, using the fact that Q^{-1} is symmetric, gives

$$\theta(\lambda) = -\frac{1}{2}(\lambda, R\lambda) + (\mathbf{d}, \lambda) + \alpha, \quad (14)$$

where

$$R = AQ^{-1}A^T, \quad \mathbf{d} = \mathbf{b}Q^{-1}A^T - \mathbf{c}, \quad \alpha = -\frac{1}{2}(\mathbf{b}, Q^{-1}\mathbf{b}).$$

Thus the dual problem becomes the following: Maximize $\theta(\lambda)$ subject to $\lambda \geq 0$, where $\theta(\lambda)$ is given by (14).

This problem is a quadratic programming problem whose constraint set has a simple structure, namely $\{\lambda \in \mathbb{R}^n, \lambda \geq 0\}$. By Theorem 6.1 there is no duality gap and this problem has a solution λ_0 , where λ_0 is a multiplier as in Theorem 6.1. Thus, from (2) we get that the solution of QP* is

$$x_0 = Q^{-1}(\bar{b} - A^T \lambda_0).$$

Exercise 6.1. Show that if Q is positive definite, then the quadratic programming problem has a unique solution. Does your proof apply to the unconstrained problem?

Exercise 6.2. The problems posed in Exercise 10.5.8 and Exercise 5.2.10) are quadratic programming problems. Use the duality developed in this section to solve these problems. Compare with the duality in Exercises 4.1 and 4.2.

7. DUALITY IN LINEAR PROGRAMMING

The linear programming problem LPL formulated in Section 1 in Chapter IV can be viewed as a convex programming problem with $f(x) = c \cdot b_0(x)$ and $g(x) = Ax - c$. Thus all of the results of this chapter hold for the linear programming problem. The linear nature of LPL, however, results in a much stronger duality than that of Theorem 4.1. We proceed to determine this duality.

To maintain a symmetry in the duality, we shall not incorporate the constraint $x \geq 0$ in the matrix A and the vector c , as was done in Chapter IV, but shall write our primal problem (P) as

$$\text{Minimize } (c - b_0, x) \quad \text{subject to } \begin{pmatrix} A \\ -I_n \end{pmatrix} x \leq \begin{pmatrix} a \\ b_0 \end{pmatrix}, \quad (P)$$

where I_n is the $n \times n$ identity matrix and b_0 is the zero vector in \mathbb{R}^n .

We now formulate the dual problem. Let $\eta = (k, \mu)$, where $k \in \mathbb{R}^m$, $k \geq 0$, $\mu \in \mathbb{R}^n$, $\mu \geq 0$. Let

$$\begin{aligned} \theta(\eta) &= \inf \{ (c - b_0, x) + (k, Ax - c) - (k, x) : x \in \mathbb{R}^n \} \\ &= \inf \{ (c - b_0 - \mu + A^T k, x) - (k, c) : x \in \mathbb{R}^n \}. \end{aligned}$$

The expression inside the curly braces is linear in x . Therefore, for it to have a finite infimum over all x in \mathbb{R}^n , the coefficient of each component of x must be zero, and so

$$A^T k - b_0 - \mu = 0. \quad (1)$$

Since $\mathbf{p} \geq \mathbf{0}$, equation (1) is equivalent to

$$A^T \mathbf{x} - \mathbf{b} \geq \mathbf{0}.$$

If (1) holds, then $\theta(\mathbf{q}) = \langle -\mathbf{a}, \mathbf{x} \rangle$. Thus the problem dual to **P** is

$$\text{Maximize } \langle -\mathbf{a}, \mathbf{x} \rangle \quad \text{subject to } \begin{pmatrix} A^T \\ -I_m \end{pmatrix} \mathbf{x} \geq \begin{pmatrix} \mathbf{b} \\ \mathbf{0}_m \end{pmatrix}, \quad (2)$$

where I_m is the $m \times m$ identity matrix, and $\mathbf{0}_m$ is the zero vector in \mathbb{R}^m .

The dual problem (2) can be written as a linear programming problem:

$$\text{Minimize } \langle -\langle -\mathbf{a}, \mathbf{x} \rangle, \mathbf{k} \rangle \quad \text{subject to } \begin{pmatrix} -A^T \\ -I_m \end{pmatrix} \mathbf{k} \leq \begin{pmatrix} -\mathbf{b} \\ \mathbf{0}_m \end{pmatrix}.$$

The dual of this problem is

$$\text{Maximize } \langle \mathbf{b}, \mathbf{y} \rangle \quad \text{subject to } \begin{pmatrix} -A \\ I_m \end{pmatrix} \mathbf{y} \geq \begin{pmatrix} -\mathbf{a} \\ \mathbf{0}_m \end{pmatrix}.$$

If we write this problem as a minimization problem and relabel the vector $\mathbf{y} \in \mathbb{R}^m$ as \mathbf{x} , we see that the dual of the dual problem is the primal problem. We summarize this discussion in the following lemma.

LEMMA 3.1. *The dual of the linear programming problem*

$$\text{Minimize } \langle -\mathbf{b}, \mathbf{x} \rangle \quad \text{subject to } \begin{pmatrix} A \\ -I_m \end{pmatrix} \mathbf{x} \leq \begin{pmatrix} \mathbf{a} \\ \mathbf{0}_m \end{pmatrix} \quad (3)$$

is

$$\text{Maximize } \langle -\mathbf{a}, \mathbf{k} \rangle \quad \text{subject to } \begin{pmatrix} A^T \\ I_m \end{pmatrix} \mathbf{k} \geq \begin{pmatrix} \mathbf{b} \\ \mathbf{0}_m \end{pmatrix}. \quad (4)$$

The dual of the dual problem (4) is the primal problem (3).

THEOREM 3.1 (DUALITY THEOREM OF LINEAR PROGRAMMING). *Let X denote the feasible set for the primal problem and let Y denote the feasible set for the dual problem. Then:*

- (i) For each \mathbf{x} in X and \mathbf{k} in Y

$$\langle -\mathbf{a}, \mathbf{k} \rangle \leq \langle -\mathbf{b}, \mathbf{x} \rangle \quad (\text{weak duality})$$

- (ii) Let F be nonempty. Then X is empty if and only if $\zeta(-c, \lambda)$ is unbounded above on X .
- (iii) Let X be nonempty. Then F is empty if and only if $\zeta(-k, \lambda)$ is bounded below on X .
- (iv) (Strong duality). Let the primal problem (P) have a solution x_0 . Then the dual problem (D) has a solution λ_0 and $\zeta(\lambda_0) = (k, x_0)$.

Remark 7.1. Since problem P is also a convex programming problem, all of the conclusions of Theorem 4.1 apply to the linear programming problem. Thus, (i) is a special case of (i) of Theorem 4.1. In (ii) the statement that if $\zeta(-c, \lambda)$ is unbounded above then X is empty is a special case of (ii) of Theorem 4.1. In the linear programming problem the stronger "if and only if" statement holds. Analogous comments hold for (iii).

Also note that the phenomenon exhibited in Example 4.3 in which there is no duality gap, yet the dual problem has no solution, cannot occur.

We now prove the theorem. We pointed out in Remark 7.1 that (i) follows from (i) of Theorem 4.1.

From the discussion of (ii) in Remark 7.1 it follows that to establish (ii) we must show that if X is empty then $\zeta(-c, \lambda)$ is unbounded above. We first note that if X is empty, then we must have $c \neq 0$. For if $c = 0$, then $x = 0$ would be feasible and X would be nonempty.

If X is empty, then the system

$$\begin{pmatrix} A \\ -I \end{pmatrix} x \leq \begin{pmatrix} b \\ r \end{pmatrix}$$

has no solution. It then follows from Theorem II.5.4 that if X is empty, then there exists a vector $y = (y_1, y_2)$, where $y_1 \in \mathbb{R}^m$, $y_1 \geq 0$, $y_2 \in \mathbb{R}^n$, $y_2 \geq 0$, such that

$$(y_1, y_2) (b, r) = -1, \quad (y_1, -Iy_2) \begin{pmatrix} b_0 \\ r_0 \end{pmatrix} = (R_0, \theta_0). \quad (2)$$

Hence (2) becomes

$$(y_1, y_2) = -1, \quad Ay_1 = R_0, \quad y_2 \geq \theta_0.$$

Since $F \neq \emptyset$, there exists a $\lambda_0 \geq 0$ such that $A'\lambda_0 \geq b$. Then for any real $\alpha > 0$, the vector $\lambda_0 + \alpha y_1$ is in F and

$$\zeta(-c, \lambda_0 + \alpha y_1) = \zeta(-c, \lambda_0) + \alpha.$$

Thus, since α can be taken to be arbitrarily large, $\zeta(-c, \lambda)$ is unbounded above.

We view the primal problem as the dual to the dual and view the dual problem as primal. Then from (ii) we conclude that if X is nonempty, then F is empty if and only if $\langle \mathbf{b}, \mathbf{x} \rangle$ is unbounded above, or equivalently that $\langle -\mathbf{b}, \mathbf{x} \rangle$ is unbounded below. Thus, (iii) is proved.

We now take up the proof of (iv). Since the primal problem has a solution, there exists a vector \mathbf{x} in E^n as in Theorem IV.4.1. It is readily seen that this vector is a multiplier as in the KKT conditions of Theorem IV.5.2. The set \mathcal{B} is thus nonempty. The absence of a duality gap and the existence of a solution follows from Corollary 1 to Theorem 4.4. Since there is no duality gap and both the primal and dual have a solution, $\langle -\mathbf{c}, \mathbf{x}_0 \rangle = \langle -\mathbf{b}, \mathbf{x}_0 \rangle$.

COROLLARY 1. A necessary and sufficient condition that a point \mathbf{x}_0 in X and a point λ_0 in F be solutions of the primal and dual problems respectively is that

$$\langle \mathbf{b}, \mathbf{x}_0 \rangle = \langle \mathbf{c}, \mathbf{x}_0 \rangle.$$

Proof. The necessity of the condition follows from (iv) of Theorem 3.1 and the sufficiency from (ii) of Theorem 4.1.

COROLLARY 2. Necessary and sufficient conditions that a point \mathbf{x}_0 in X and a point λ_0 in F be solutions of the primal and dual problems respectively are

- (i) $\langle \mathbf{A}\mathbf{x}_0 - \mathbf{c}, \lambda_0 \rangle = 0$ and
- (ii) $\langle \mathbf{A}^T\lambda_0 - \mathbf{b}, \mathbf{x}_0 \rangle = 0$.

Proof. Let \mathbf{x}_0 be optimal for the primal problem and let λ_0 be optimal for the dual problem. Then $\mathbf{A}\mathbf{x}_0 \leq \mathbf{c}$ and $\lambda_0 \geq 0$. Hence

$$\langle \mathbf{A}\mathbf{x}_0 - \mathbf{c}, \lambda_0 \rangle \leq 0. \quad (3)$$

Since λ_0 is feasible for the dual problem, $\mathbf{A}^T\lambda_0 \geq \mathbf{b}$. Since $\mathbf{x}_0 \geq 0$, we have

$$\langle \mathbf{A}^T\lambda_0 - \mathbf{b}, \mathbf{x}_0 \rangle \geq 0. \quad (4)$$

By Corollary 1, $\langle \mathbf{b}, \mathbf{x}_0 \rangle = \langle \mathbf{c}, \mathbf{x}_0 \rangle$, so $\langle \mathbf{A}^T\lambda_0, \mathbf{x}_0 \rangle = \langle \mathbf{c}, \mathbf{x}_0 \rangle \geq 0$. Therefore $\langle \mathbf{A}\mathbf{x}_0 - \mathbf{c}, \lambda_0 \rangle \geq 0$. Comparing this with (3) gives (i). Also,

$$\langle \mathbf{A}\mathbf{x}_0 - \mathbf{c}, \lambda_0 \rangle = \langle \mathbf{A}\mathbf{x}_0, \lambda_0 \rangle - \langle \mathbf{c}, \lambda_0 \rangle = \langle \mathbf{A}^T\lambda_0, \mathbf{x}_0 \rangle;$$

and so from (4) we get that $\langle \mathbf{A}^T\lambda_0 - \mathbf{b}, \mathbf{x}_0 \rangle \leq 0$. This inequality and (4) give (ii).

To prove the sufficiency, note that (i) and (ii) give

$$\langle \mathbf{c}, \mathbf{x}_0 \rangle = \langle \mathbf{A}\mathbf{x}_0, \lambda_0 \rangle = \langle \mathbf{A}^T\lambda_0, \mathbf{x}_0 \rangle = \langle \mathbf{b}, \mathbf{x}_0 \rangle.$$

It then follows from the equality of the extremum quantities and Corollary 3 that x_p is a solution of the primal problem and that λ_p is a solution of the dual.

The next theorem is a stronger version of the strong duality statement (iv) of Theorem 7.1) in that we obtain strong duality by merely assuming the feasible sets X and Y to be nonempty.

THEOREM 7.2. *The feasible sets X and Y for the primal and dual problems respectively are nonempty if and only if there exists a solution x_p to the primal problem and a solution λ_p to the dual problem. In this event*

$$\langle c, \lambda_p \rangle = \langle b, x_p \rangle. \quad (5)$$

Proof. If the primal and dual problems have solutions, then the sets X and Y are clearly nonempty.

To show that $X \neq \emptyset$ and $Y \neq \emptyset$ imply that both the dual and primal problems have solutions and that (5) holds, it suffices to show that the primal problem has a solution. For from (v) of Theorem 7.1, the dual problem has a solution and (5) holds.

In Remark IV.4.1 we noted that relation (2v) of Theorem IV.4.1 is a consequence of (i), (ii), (iii), and (v). Thus in Corollary IV.4.3 the same observation holds. Hence a necessary and sufficient condition for a vector x_p to be a solution of P is that there exist a vector λ_p in \mathbb{R}^m such that

$$\begin{pmatrix} A \\ -I_n \end{pmatrix} x_p = \begin{pmatrix} c \\ b \end{pmatrix}, \quad \begin{pmatrix} A' \\ I_m \end{pmatrix} \lambda_p = \begin{pmatrix} b \\ 0 \end{pmatrix}, \quad (6)$$

$$\langle b, \lambda_p \rangle = \langle c, x_p \rangle.$$

Since the existence of vectors x_p in \mathbb{R}^n and vector λ_p in \mathbb{R}^m that satisfy (6) is a necessary and sufficient condition that x_p be a solution of P , it follows that if P has no solution, then the system

$$\begin{aligned} \text{(i)} \quad & \begin{pmatrix} A \\ -I_n \end{pmatrix} x = \begin{pmatrix} c \\ b \end{pmatrix}, \\ \text{(ii)} \quad & \begin{pmatrix} -A' \\ -I_m \end{pmatrix} u = \begin{pmatrix} -b \\ 0 \end{pmatrix}, \\ \text{(iii)} \quad & \langle -b, u \rangle + \langle c, x \rangle = 0 \end{aligned} \quad (7)$$

has no solution (x, u) , $x \in \mathbb{R}^n$, $u \in \mathbb{R}^m$. Condition (i) is the condition that x be feasible for P and condition (ii) is the condition that u be feasible for the dual problem. By assumption, there exist feasible vectors $x \in X$ and $u \in Y$. By (i) of Theorem 7.1, for such vectors we have $\langle -b, u \rangle \geq \langle -c, x \rangle$, and thus $\langle -b, u \rangle + \langle c, x \rangle \geq 0$. Hence for a vector (x, u) that satisfies (i) and (ii) to

fail to satisfy (7b), it must fail to satisfy

$$c(-\mathbf{h}, \mathbf{x}) + C_0(\mathbf{x}) < 0, \quad (7c)$$

so we may replace (7b) by (7c) in (7).

Now consider the system

$$\begin{aligned} & \lambda > 0, \quad \mu \in \mathbb{R}, \\ & \begin{pmatrix} A \\ -I_m \end{pmatrix} \mathbf{x} \leq \begin{pmatrix} \mathbf{b} \\ \mathbf{a} \end{pmatrix}, \quad \begin{pmatrix} -A' \\ -I_m \end{pmatrix} \mathbf{x} \leq \begin{pmatrix} -\mathbf{b} \\ \mathbf{a} \end{pmatrix}, \\ & c(-\mathbf{h}, \mathbf{x}) + C_0(\mathbf{x}) < 0. \end{aligned} \quad (8)$$

If (7) has a solution $(\mathbf{x}_0, \mathbf{u}_0)$, then (8) has a solution $(\lambda_0, \mathbf{x}_0, \mathbf{u}_0) = (1, \mathbf{x}_0, \mathbf{u}_0)$. Conversely, if (8) has a solution $(\lambda', \mathbf{x}_0, \mathbf{u}_0)$, then $(\mathbf{x}_0, \mathbf{u}_0) = (\lambda')^{-1}(\lambda', \mathbf{x}_0, \mathbf{u}_0)$ is a solution of (7). Therefore (7) has no solution if and only if (8) has no solution.

System (8) can be written as

$$\begin{aligned} & c(\mathbf{x}_0, \mathbf{u}_0, 1), \quad c_0(\mathbf{x}_0, \mathbf{u}_0) < 0, \\ & \begin{pmatrix} A & O & -c' \\ -I_m & O & \mathbf{b}_0 \\ O & -A' & \mathbf{b} \\ O & -I_m & \mathbf{u}_0 \\ -\lambda' & c' & 0 \end{pmatrix} \begin{pmatrix} \mathbf{x}_0 \\ \mathbf{u}_0 \\ \zeta \end{pmatrix} \leq \begin{pmatrix} \mathbf{b}_0 \\ \mathbf{b} \\ \mathbf{u}_0 \\ 0 \end{pmatrix}, \end{aligned} \quad (9)$$

where O represents a zero matrix of appropriate dimension and \mathbf{b}_0 is the zero vector in \mathbb{R}^m . Since this system has no solution, it follows from Theorem B.5.1 (Farkas's lemma) that there exists a vector $(\mathbf{p}_1, \mathbf{p}_2, \mathbf{q}_1, \mathbf{q}_2, \sigma) \geq \mathbf{0}_{m+m+m+m}$, with $\mathbf{p}_1 \in \mathbb{R}^m$, $\mathbf{p}_2 \in \mathbb{R}^m$, $\mathbf{q}_1 \in \mathbb{R}^m$, $\mathbf{q}_2 \in \mathbb{R}^m$, and $\sigma \in \mathbb{R}$, such that

$$\begin{pmatrix} A' & -I_m & O & O & -\mathbf{b}' \\ O & O & -A & -I_m & c \\ -c' & \mathbf{b}_0 & \mathbf{b}' & \mathbf{u}_0 & 0 \end{pmatrix} \begin{pmatrix} \mathbf{p}_1 \\ \mathbf{p}_2 \\ \mathbf{q}_1 \\ \mathbf{q}_2 \\ \sigma \end{pmatrix} = \begin{pmatrix} \mathbf{b}_0 \\ \mathbf{b} \\ 0 \end{pmatrix}.$$

Thus

$$\begin{aligned} & A'\mathbf{p}_1 - \mathbf{p}_2 - \sigma\mathbf{b} = \mathbf{b}_0, \\ & -A\mathbf{q}_1 - \mathbf{q}_2 + \sigma c = \mathbf{u}_0, \\ & c(-\mathbf{p}_1, c) + C_0(\mathbf{b}) = 1. \end{aligned} \quad (10)$$

We shall show that this system cannot have a solution $(\mathbf{p}_1, \mathbf{p}_2, \mathbf{q}_1, \mathbf{q}_2) \in \mathcal{D}_{\mathbf{p}_1, \mathbf{p}_2, \mathbf{q}_1, \mathbf{q}_2}$. This contradiction will prove that the primal problem has a solution, since we arrived at (10) by assuming that it did not.

Let $\mathbf{x}_1 \in \mathcal{X}$ and let $\mathbf{y}_2 \in \mathcal{Y}$. Such points exist since $\mathcal{X} \neq \emptyset$ and $\mathcal{Y} \neq \emptyset$. Suppose first that $\alpha = 0$. Then $\lambda \mathbf{p}_1 = \mathbf{p}_1$ and $\lambda \mathbf{q}_1 = -\mathbf{q}_1$. Since $\mathbf{x}_1 \in \mathcal{X}$, we have $\lambda \mathbf{p}_1 \leq \mathbf{c}$. Thus, since $\mathbf{p}_1 \geq \mathbf{0}$,

$$\langle -\mathbf{p}_1, \mathbf{c} \rangle \leq \langle -\mathbf{p}_1, \lambda \mathbf{p}_1 \rangle = -\langle \lambda \mathbf{p}_1, \mathbf{p}_1 \rangle = -\langle \mathbf{p}_1, \mathbf{p}_1 \rangle \leq 0, \quad (11)$$

where the last inequality follows from $\mathbf{p}_1 \geq \mathbf{0}$, $\mathbf{p}_1 \geq \mathbf{0}$. Similarly,

$$\langle \mathbf{q}_2, \mathbf{b} \rangle \leq \langle \mathbf{q}_2, \lambda \mathbf{p}_2 \rangle = \langle \lambda \mathbf{q}_2, \mathbf{p}_2 \rangle = -\langle \mathbf{q}_2, \mathbf{p}_2 \rangle \leq 0. \quad (12)$$

From (11) and (12) we get

$$\langle -\mathbf{p}_1, \mathbf{c} \rangle + \langle \mathbf{q}_2, \mathbf{b} \rangle \leq 0, \quad (13)$$

which contradicts the last equation in (10).

Suppose now that $\alpha > 0$. The first equation in (10) implies that \mathbf{p}_1/α is feasible for the dual problem. The second equation in (10) implies that \mathbf{q}_2/α is feasible for the primal problem. By the weak duality, $\langle -\mathbf{b}, \mathbf{q}_2/\alpha \rangle \geq \langle -\mathbf{c}, \mathbf{p}_1/\alpha \rangle$. Hence we again arrive at (13), which contradicts the last equation in (10).

VI

SIMPLEX METHOD

1. INTRODUCTION

In Section 4 of Chapter IV we formalized the linear programming problem and derived a necessary and sufficient condition that a point x_0 be a solution to a linear programming problem. In Section 7 of Chapter V we presented the duality theorem for linear programming. We remarked in Section 4 of Chapter IV that in practice the dimension of the space \mathbb{R}^n precluded the use of the necessary and sufficient conditions to solve linear programming problems and that instead a method known as the simplex method is used. In this chapter we shall present the essentials of the simplex method, which is an algorithm for solving linear programming problems.

To motivate the rationale of the simplex method, we consider several examples in \mathbb{R}^2 .

Example 1.1. Minimize $f(x) = x_1 - 2x_2$ subject to

$$\begin{aligned} -x_1 - x_2 &\leq -1, \\ -x_1 + x_2 &\leq 0, \\ x_1 + 2x_2 &\leq 4, \\ x_1 &\geq 0, \quad x_2 \geq 0. \end{aligned} \tag{1}$$

The feasible set is indicated by the hatched area in Figure 6.1. The level curves $f(x) = 0$ and $f(x) = -1$ are indicated by dashed lines; the arrows indicate the directions of the gradient vectors of f . Since the gradient vectors indicate the direction of increasing f , we see from the figure that f attains a minimum at the extreme point $\left(\frac{1}{3}, \frac{1}{3}\right)$ of the feasible set.

Example 1.2. This example differs from Example 1.1 in that we take the function f to be $f(x) = x_1 + x_2$. The level curves are lines parallel to the line $x_1 + x_2 = 1$, a segment of which is part of the boundary of the feasible set. Since the gradient vector of f is $(1, 1)$ and this is the direction of increasing f ,

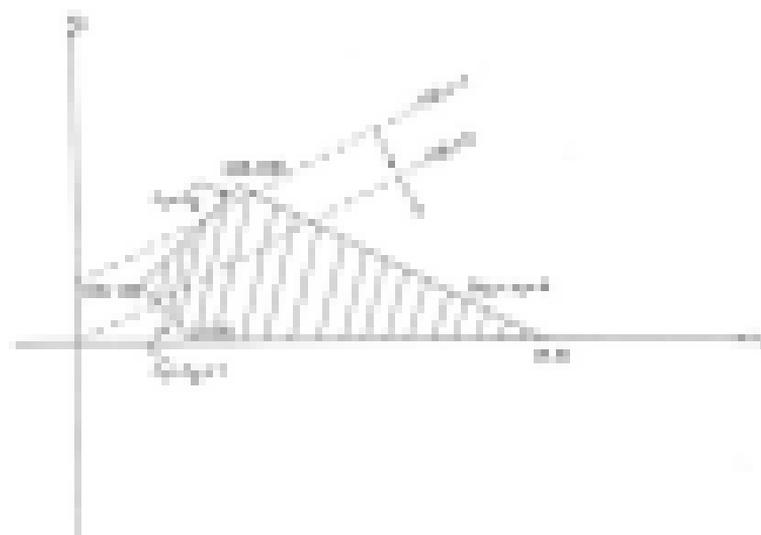


Figure 4.1.

it follows that F attains its minimum on all boundary points of the feasible set that lie on the line $x_1 + x_2 = 6$. In particular, the minimum is attained at the two extreme points $(2, 4)$ and $(3, 3)$.

Example 2.1. Minimize $F(x) = -x_1 - x_2$ subject to

$$-x_1 + x_2 \leq 0,$$

$$x_1 - 2x_2 \leq -2,$$

$$x_1 \geq 0, \quad x_2 \geq 0.$$

The feasible set is the shaded area in Figure 4.2. The level curve $F(x) = -10$ is indicated by a dashed line. Since the gradient of F is $(-1, -1)$, it follows from the figure that this problem has no solution. The feasible set in this example is unbounded, but unbounded feasible sets, however, does not imply that a linear programming problem has no solution. One can still, for example, by taking $F(x) = x_1 + x_2$, find the problem has a solution at the extreme point $(2, 0)$.

In the examples that we considered either the problem had no solution or the problem had a solution and there was an extreme point of the feasible set at which the solution was attained. There is, of course, a third possibility,

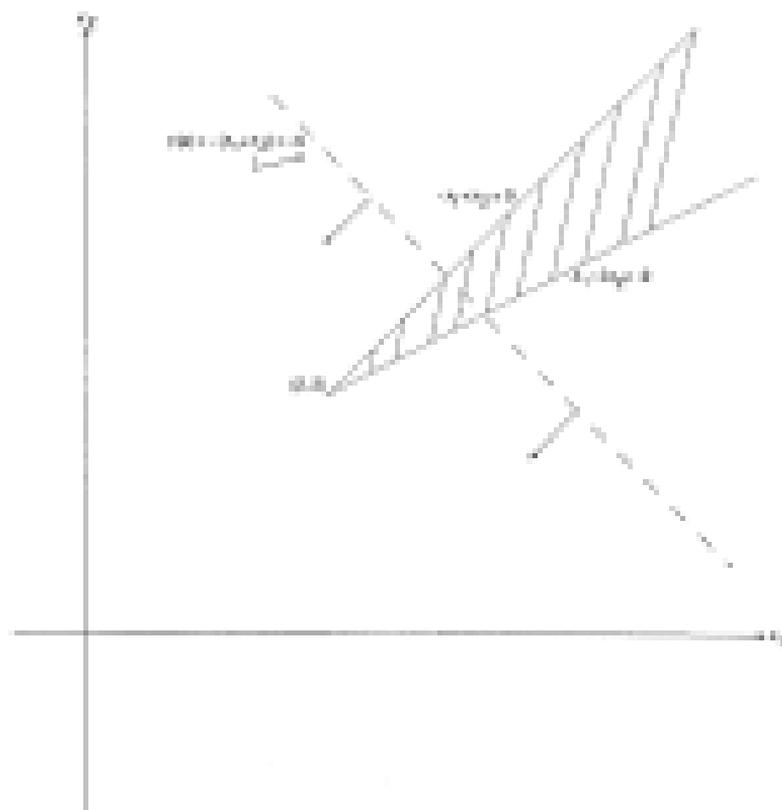


Figure 5.1

surely that the feasible set is empty. In the next section we shall show that in a linear programming problem there are the only possibilities.

The simplex method is applied to problems in canonical form, CLP (see Section 4 of Chapter IV for definitions). To put the problem in Example 1.1 in canonical form we introduce additional slack variables x_3, x_4, x_5 and write the problem as

$$\begin{aligned}
 & \text{Minimize } f(x) = x_1 - 2x_2 \\
 & \text{subject to } -x_1 - x_2 + x_3 = -4 \\
 & \quad \quad -x_1 + x_2 + x_4 = 0, \\
 & \quad \quad x_1 + 2x_2 + x_5 = 4.
 \end{aligned} \tag{2}$$

Since the variables x_4, x_5, x_6 do not appear in the formula for f , the problem with feasible set defined by (2) has its minimum attained at

$$x_1 = \frac{1}{2}, \quad x_2 = \frac{1}{2}, \quad x_3 = \frac{1}{2}, \quad x_4 = 0, \quad x_5 = 0.$$

It is readily verified that this point is an extreme point of the set defined by (2). We shall leave it to the reader to show that the extreme points of the constraint set in problem SLP corresponds to the extreme points in the problem CLP obtained by introducing slack variables.

Exercise 1.1. Solve the following linear programming problem graphically. Minimize $f(x) = x_1 + 2x_2$ subject to

$$\begin{aligned} x_1 + x_2 &\leq 1, \\ -x_1 + x_2 &\leq 1, \\ x_2 - x_1 &\leq 0, \\ -2x_1 - x_2 &\leq -1, \\ x_1 &\geq 0, \quad x_2 \geq 0. \end{aligned}$$

Exercise 1.2. Let x_0 be a feasible point for problem SLP and let (x_0, \bar{z}_0) be the corresponding feasible point for the related problem CLP; then

$$\begin{aligned} Ax_0 &\leq b, & Ax_0 + \bar{z}_0 &= b, \\ x_0 &\geq 0, & \bar{z}_0 &\geq 0. \end{aligned}$$

Show that x_0 is an extreme point of the feasible set for problem SLP if and only if (x_0, \bar{z}_0) is an extreme point for problem CLP.

Exercise 1.3. In problem LP, which is Minimize $(-k, x)$ subject to $ax \leq c$, show that if the feasible set is compact, then the problem has a solution and there exists an extreme point of the feasible set at which the minimum is attained.

2. EXTREME POINTS OF FEASIBLE SET

In this section we shall show that if a linear programming problem has a solution, then there exists an extreme point at which the minimum is attained. We shall prove this for problems in canonical form without assuming that the feasible set is compact, as was done in Exercise 1.3. By virtue of Exercise 1.2 this also holds for problem SLP. We shall also give an analytic characterization of the extreme points of the feasible set for CLP problems. The significance of

these results is that in looking for a solution to a linear programming problem we can confine ourselves to the extreme points of the feasible set, points for which we have an analytic characterization. The simplex method is an algorithm for efficiently determining whether a solution exists and, if so, for finding an extreme point which solves the problem.

We now consider the linear programming problem in canonical form:

$$\begin{aligned} & \text{Maximize } \langle \mathbf{b}, \mathbf{x} \rangle \\ & \text{subject to } A\mathbf{x} = \mathbf{a}, \quad \mathbf{x} \geq \mathbf{0}. \end{aligned} \quad (1)$$

In keeping with much of the linear programming literature we have replaced the objective of minimizing $\langle -\mathbf{b}, \mathbf{x} \rangle$ by the equivalent objective of maximizing $\langle \mathbf{b}, \mathbf{x} \rangle$. We assume that \mathbf{a} and $\mathbf{0}$ are in \mathbb{R}^m , that \mathbf{c} is in \mathbb{R}^n , that A is in $m \times n$, and that

$$(i) \quad m < n, \quad (ii) \quad \text{rank } A = m. \quad (2)$$

The rationale for the assumption (2) is precisely the same as given in Remark 5.1.2 for the assumption that the $k \times n$ matrix A defining the affine constraints in the convex programming problem CPH has rank k and that $k < n$.

Before presenting the principal results of this section, we emphasize that the reader should keep in mind that the constraint $\mathbf{x} \geq \mathbf{0}$ implies that the components x_j of \mathbf{x} are either positive or zero.

THEOREM 2.1. *A feasible point \mathbf{x}_0 is an extreme point of the feasible set if and only if the columns \mathbf{A}_j of the matrix A that correspond to positive components of \mathbf{x}_0 are linearly independent.*

Proof. Suppose that the columns \mathbf{A}_j of A that correspond to positive components of \mathbf{x}_0 are linearly independent. We shall show that \mathbf{x}_0 is an extreme point of the feasible set.

If \mathbf{x}_0 were not an extreme point of the feasible set, there would exist two distinct feasible points \mathbf{x}_1 and \mathbf{x}_2 and a scalar λ such that $0 < \lambda < 1$ and

$$\mathbf{x}_0 = \lambda \mathbf{x}_1 + (1 - \lambda) \mathbf{x}_2. \quad (3)$$

Let B denote the submatrix of A formed by the columns of A that correspond to positive components of \mathbf{x} and let N denote the submatrix of A formed by the columns of A corresponding to zero components of \mathbf{x} . To simplify notation, let us assume that B consists of the first k columns of A and that N consist of the remaining $n - k$ columns. There is no loss of generality in making this assumption since it can be achieved by a relabeling of the coordinates. Since the columns of B are linearly independent, it follows from (2) that $k \leq m$. It then further follows from (2) that $n - k > 0$. For \mathbf{x} in \mathbb{R}^n , let

$$\mathbf{x}_B = (x_1, \dots, x_k), \quad \mathbf{x}_N = (x_{k+1}, \dots, x_n).$$

Then

$$\begin{aligned}x_{1i} &= (x_{1i0}, x_{1i1}) = (x_{1i0}, \mathbf{0}_k) \\x_i &= (x_{ij0}, x_{ij1}), \quad i = 1, 2,\end{aligned}$$

where x_{10} , x_1 , and x_2 are as in (1) and $\mathbf{0}_k$ is the zero vector in \mathbb{R}^k .

Since $d > 0$ and $1 - d > 0$, it follows from (1) that a component of x_{1i} is equal to zero if and only if the corresponding components of x_1 and x_2 are zero. Hence $x_{1i0} = x_{ij0} = x_{ij1} = \mathbf{0}_k$ and we may write

$$x_i = (x_{ij0}, \mathbf{0}_k), \quad i = 0, 1, 2.$$

Since x_1 is feasible,

$$c = Ax_1 = (B, N) \begin{pmatrix} x_{1j0} \\ \mathbf{0}_k \end{pmatrix} = Bx_{1j0}.$$

Similarly, we get that $c = Bx_{2j0}$. Hence

$$B(x_{1j0} - x_{2j0}) = \mathbf{0}. \quad (4)$$

Since $x_1 \neq x_2$ and $x_{1j0} = x_{2j0}$, we have that $x_{1j0} - x_{2j0} \neq \mathbf{0}$. But this cannot be, since the columns of B are linearly independent. This contradiction shows that x_{10} is an extreme point.

Now let us suppose that x_{10} is an extreme point of the feasible set. To simplify notation, we suppose that the first k components of x_{10} are positive and the remaining components are zero. We shall show that the first k columns A_1, \dots, A_k of A are linearly independent. If these columns were not linearly independent, there would exist constants $\alpha_1, \dots, \alpha_k$ not all zero such that

$$\alpha_1 A_1 + \dots + \alpha_k A_k = \mathbf{0}.$$

Let the α -vector α be defined by

$$\alpha = (\alpha_1, \dots, \alpha_k, \mathbf{0}, \dots, \mathbf{0}).$$

Then for any real ϵ ,

$$\alpha \epsilon \alpha = \epsilon \sum_{j=1}^k \alpha_j A_j = \mathbf{0}. \quad (5)$$

For $\epsilon > 0$ and sufficiently small the vectors $x_{10} + \epsilon \alpha$ and $x_{10} - \epsilon \alpha$ have their first k components positive and the remaining $n - k$ components equal to zero. Also

$$B(x_{10} \pm \epsilon \alpha) = Bx_{10} \pm \epsilon \alpha B = c,$$

where the last equality follows from (5) and the feasibility of x_0 . Thus for sufficiently small $\epsilon > 0$ the vectors $x_0 + \epsilon u$ and $x_0 - \epsilon u$ are feasible and feasible. But then the relation

$$x_0 = \frac{1}{2}(x_0 + \epsilon u) + \frac{1}{2}(x_0 - \epsilon u)$$

contradicts the hypothesis that x_0 is an extreme point. Hence the vectors A_1, \dots, A_k are linearly independent.

The following corollary is an immediate consequence of the theorem and the fact that A has rank m .

COROLLARY 1. *An extreme point of the feasible set has at most m positive components.*

COROLLARY 2. *The number of extreme points of the feasible set is less than or equal to*

$$\binom{n}{m} = \frac{n!}{m!(n-m)!}.$$

Proof. If x is an extreme point of the feasible set, then x has $k \leq m$ positive components. These k positive components correspond to k linearly independent columns of A . We assert that if x' is another extreme point with the same nonzero components as x , then $x = x'$. To see this, let B denote the submatrix of A corresponding to the nonzero components of x and x' , let $x = (x_B, 0)$, and let $x' = (x'_B, 0)$. Then by the argument used to establish (4), we get that $B(x_B - x'_B) = 0$. Since the columns of B are linearly independent, $x_B = x'_B$, and so $x = x'$.

If $k < m$, then since $\text{rank } A = m$, we can adjoin an additional $m - k$ columns of A to obtain a set of m linearly independent columns associated with x . If $k = m$, then this certainly holds. Hence the number of extreme points cannot exceed the number of ways in which m columns can be selected from among n columns.

THEOREM 2.1. *Let the canonical linear programming problem as stated in (1) have a solution x_0 . Then there exists an extreme point x of the feasible set that is also a solution.*

Proof. By Theorem 2.1, if the columns of A that correspond to positive components of x_0 are linearly independent, then x_0 is an extreme point of the feasible set and we let $x = x_0$.

Let us now suppose that the columns of A that correspond to the positive components of x_0 are linearly dependent. To simplify notation, we suppose that the last p components of x_0 are positive and that the remaining $n - p$ com-

ponents are equal to zero. Thus the first p columns of A are linearly dependent and there exist scalars $\gamma_1, \dots, \gamma_p$ not all zero such that

$$\sum_{j=1}^p \gamma_j \mathbf{A}_j = \mathbf{0}. \quad (6)$$

For each real number σ define a vector $\mathbf{x}(\sigma)$ as follows:

$$x_j(\sigma) = \begin{cases} b_{0j} - \sigma \gamma_j & j = 1, \dots, p \\ b_{0j} & j = p+1, \dots, n \end{cases} \quad (7)$$

Thus, from the fact that $x_{0j} = 0$ for $j = p+1, \dots, n$, from the feasibility of \mathbf{x}_0 and from (6) we get that

$$A\mathbf{x}(\sigma) = \mathbf{b}\mathbf{x}_0 - \sigma \sum_{j=1}^p \gamma_j \mathbf{A}_j = \mathbf{c}.$$

Let

$$\sigma_1 = \min \left\{ \frac{b_{0j}}{|\gamma_j|} : j = 1, \dots, p \text{ and } \gamma_j > 0 \right\}. \quad (8)$$

It follows from (7) that $\mathbf{x}(\sigma) \geq \mathbf{0}$ whenever $|\sigma| \leq \sigma_1$. Hence $\mathbf{x}(\sigma)$ is feasible whenever $|\sigma| \leq \sigma_1$. Thus,

$$\langle \mathbf{b}, \mathbf{x}_0 \rangle \geq \langle \mathbf{b}, \mathbf{x}(\sigma) \rangle = \langle \mathbf{b}, \mathbf{x}_0 \rangle - \sigma \sum_{j=1}^p b_{0j} \gamma_j$$

for all σ such that $|\sigma| \leq \sigma_1$. Hence

$$\sum_{j=1}^p b_{0j} \gamma_j = 0 \quad \text{and} \quad \langle \mathbf{b}, \mathbf{x}_0 \rangle = \langle \mathbf{b}, \mathbf{x}(\sigma) \rangle.$$

Thus, $\mathbf{x}(\sigma)$ is a solution whenever $|\sigma| \leq \sigma_1$. Let the minimum in (8) be attained at the index r . Let $\sigma_1 = b_{0r} / \gamma_r$. Then $\sigma_1 = \sigma_2$ if $\gamma_r > 0$ and $\sigma_1 = -\sigma_2$ if $\gamma_r < 0$. Hence $\mathbf{x}(\sigma_1)$ is feasible and is a solution. Moreover $\mathbf{x}(\sigma_1)$ has fewer than p positive components. If the corresponding columns of A are linearly independent, then $\mathbf{x}(\sigma_1)$ is an extreme point. If the columns are not linearly independent, we repeat the process, taking \mathbf{x}_0 to be $\mathbf{x}(\sigma_1)$. In a finite number of iterations we shall obtain a solution such that the column vectors corresponding to the positive components are linearly independent. This solution will be an extreme point of the feasible set.

3. PRELIMINARIES TO SIMPLEX METHOD

In the last section we showed that if a solution to the canonical linear programming problem exists, then it exists at an extreme point of the feasible set. We also showed that there are at most $n!(n-1)!$ such points. For a moderate system of 6 constraint equations and 30 variables, there can be as many as

$$\binom{30}{6} = \frac{30!}{6!24!} = 591,779$$

extreme points. Thus the direct-minded scheme of evaluating the objective function at all possible extreme points and then selecting the point at which the largest value of the objective function is attained is not practical. The simplex method is an algorithm for finding a solution by going from one extreme point of the feasible set to another in such a way that the objective function is not decreased.

We now introduce the notation and definitions that will be used. Consider the system

$$Ax = r \quad (1)$$

that defines the feasible set for CLP, the canonical linear programming problem. Thus, A is $m \times n$ and of rank m . Let B denote any $m \times m$ submatrix of A obtained by taking m linearly independent columns of A and let N denote the submatrix of A consisting of the remaining $n - m$ columns. We can write a vector x in \mathbb{R}^n as

$$x = (\mathbf{x}_B, \mathbf{x}_N),$$

where \mathbf{x}_B consists of those components of x corresponding to the columns in B and \mathbf{x}_N consists of those components of x corresponding to the columns in N .

Definition 1.1. The columns of B are called a *basis*. The components of \mathbf{x}_B are called *basis variables*. The components of \mathbf{x}_N are called *nonbasic variables* or *free variables*.

We may write the system (1) as

$$(B, N) \begin{pmatrix} \mathbf{x}_B \\ \mathbf{x}_N \end{pmatrix} = r. \quad (2)$$

Any vector of the form $(\mathbf{x}_B, \mathbf{x}_N)$ where \mathbf{x}_N is arbitrary and

$$\mathbf{x}_B = B^{-1}r - B^{-1}N\mathbf{x}_N$$

will be a solution of (2). (This is why the components of x_B are called *basic variables*.) There is a unique solution of (2) with $x_N = \theta_N$, namely

$$\xi = (\theta_B, \theta_N), \quad \theta_B = B^{-1}b. \quad (3)$$

Definition 2.1. (i) A solution of (1) of the form (3) is called a *basic solution*. (ii) If $\xi_B \geq \theta_B$, then ξ is feasible and is called a *basic feasible solution*. (iii) If some of the components of θ_B are zero, then ξ is called a *degenerate basic solution*.

We emphasize the following aspects of our notation. By x_B we mean the vector consisting of those components of x that correspond to the columns of B . By x_N we mean the vector consisting of those components of x that correspond to the columns of N . By θ_B we mean the value of the basic variables in a basic solution. By ξ we mean the value of a solution of (1) and by θ_B we mean the value of a particular x_B .

We also call attention to a change in terminology. Up to now we have used the term *solution* to mean a solution of the linear programming problem. We are now using it to mean a solution of (3). Henceforth we shall use the term *optimal solution* to designate a solution of the linear programming problem, that is, a *feasible point* at which the maximum is attained.

Remark 2.1. It follows from the preceding and from Theorem 2.1 that a basic feasible solution is an extreme point and conversely. It follows from Theorem 2.2 that if an optimal solution exists, then there exists a basic feasible solution that is an optimal solution. Therefore, we need only consider basic feasible solutions in our search for the optimal solution.

To apply the simplex method, we reformulate the problem stated in (1) of Section 2 by introducing a scalar variable z and writing the problem as

$$\begin{aligned} \text{Maximize:} \\ \text{subject to } Ax = b, \\ (b, x) - z = 0, \\ x \geq 0. \end{aligned} \quad (4)$$

The simplex method consists of two phases. In phase I the method either determines a basic feasible solution or determines that the feasible set is empty. In phase II the method starts with a basic feasible solution and either determines that no optimal solution exists or determines an optimal solution. In the latter case it does so by moving from one basic feasible solution to another in such a way that the value of z is not decreased.

We conclude this section with a review of some aspects of linear algebra involving the solution of systems of linear equations and the calculation of inverses of matrices. These aspects of linear algebra arise in the simplex method.

Consider a system of linear equations

$$M\mathbf{y} = \mathbf{d}, \quad (5)$$

where M is a $k \times n$ matrix of rank k , \mathbf{y} is a vector in \mathbb{R}^n , and \mathbf{d} is a vector in \mathbb{R}^k . A system of linear equations $M'\mathbf{y} = \mathbf{d}'$ is said to be equivalent to (5) if the solution sets of both equations are equal. We define an elementary operation on the system (5) to be any one of the following operations.

- (i) Interchange two equations.
- (ii) Multiply an equation by a nonzero constant.
- (iii) Add a nonzero multiple of an equation to another equation.

Any sequence of elementary operations applied to (5) will yield a system equivalent to (5).

Let R denote a $k \times k$ submatrix of M that has rank k and let F denote the submatrix of M formed by the remaining $n - k$ columns. Then we may write (5) as

$$(R, F) \begin{pmatrix} \mathbf{y}_k \\ \mathbf{y}_n \end{pmatrix} = \mathbf{d}, \quad (6)$$

where \mathbf{y}_k is the vector obtained from \mathbf{y} by taking the components of \mathbf{y} that correspond to the columns in R and \mathbf{y}_n has an analogous meaning. By an appropriate sequence of elementary operations, we can obtain a system equivalent to (6) having the form

$$(U, F) \begin{pmatrix} \mathbf{y}_k \\ \mathbf{y}_n \end{pmatrix} = \mathbf{d}, \quad (7)$$

where U is a $k \times k$ upper triangular matrix with diagonal entries equal to 1 and F is a $k \times (n - k)$ matrix. To simplify notation, assume that R consists of the first k columns of M . Then (7) in component form reads

$$\begin{aligned} y_1 + u_{12}y_2 + \cdots + u_{1k}y_k + u_{1,k+1}y_{k+1} + \cdots + u_{1n}y_n &= d'_1, \\ y_2 + \cdots + u_{2k}y_k + u_{2,k+1}y_{k+1} + \cdots + u_{2n}y_n &= d'_2, \\ &\vdots \\ y_k + u_{k,k+1}y_{k+1} + \cdots + u_{kn}y_n &= d'_k. \end{aligned} \quad (8)$$

The last equation in (8) expresses y_k in terms of y_{k+1}, \dots, y_n . Proceeding

solved through the system (9) by backward substitution, we can express x_1, \dots, x_n in terms of x_{k+1}, \dots, x_n , and thus obtain the general solution of (7) and hence of (5). The method of solving (5) just described is called *Gauss elimination*.

In solving the system (5) or (9) we need not write down the equations and the variables. We need only to write down the augmented matrix (A, \mathbf{b}) or $(R, \mathbf{z}, \mathbf{d})$ and apply the elementary operations to the rows of the augmented matrix. In so doing we speak of elementary row operations. The reader will recall, or can verify, that an elementary row operation on a $k \times n$ matrix M can be effected in the following fashions. Perform the elementary operation on the $k \times k$ identity matrix to obtain a matrix E . Then premultiply the matrix M by E . The matrices E are called elementary matrices. In computer programs for solving systems of linear equations by Gauss elimination, elementary row operations are executed by premultiplications by elementary matrices. Such computer programs also incorporate rules for choosing the components \mathbf{p} , of \mathbf{p} so as to increase accuracy and speed-of-computation. We shall not take up these questions here.

By an additional sequence of elementary row operations, we can obtain a system equivalent to (7) having the form

$$(I_k, 0) \begin{pmatrix} \mathbf{y}_1 \\ \mathbf{y}_2 \end{pmatrix} = \mathbf{d}, \quad (9)$$

where I_k is the $k \times k$ identity matrix. We can rewrite (9) as

$$I_k \mathbf{y}_1 = \mathbf{d} - 0 \mathbf{y}_2, \quad (10)$$

and thus obtain \mathbf{y}_1 in terms of \mathbf{y}_2 directly, without backward substitution. The method of solving (5) by reducing it to the equivalent system (9) is known as *Gauss-Jordan elimination*. For most large systems, computer implementations of Gauss elimination are superior to those of Gauss-Jordan elimination.

We next recall an effective way of calculating the inverse of a nonsingular matrix \mathbf{R} . We first note that an elementary matrix E is nonsingular. Since \mathbf{R} is nonsingular, we can reduce \mathbf{R} to the identity matrix by applying a sequence of elementary row operations. But each elementary row operation is effected by a premultiplication by an elementary matrix. Thus, we have a sequence E_p, E_{p-1}, \dots, E_1 of elementary matrices such that

$$E_p \dots E_1 E_1 \mathbf{R} = \mathbf{I}.$$

Multiplying this equation through on the right by \mathbf{R}^{-1} gives

$$E_p \dots E_1 \mathbf{I} = \mathbf{R}^{-1}. \quad (11)$$

This is the justification for the following method of calculating R^{-1} for matrices of low dimension. Reduce the matrix (R, I) to the form (I, T) by a sequence of elementary operations. Then $T = R^{-1}$. The so-called product form of the inverse will also be used in the simplex method.

Remark 1.2. In reducing (8) to (9), we performed a sequence of elementary operations, which in effect was a sequence of premultiplications of (8) by elementary matrices. Thus,

$$E_p \cdots E_1 R, \Delta, \mathbf{d} = (I_p, R(\mathbf{d}^T)). \quad (12)$$

Since the sequence of elementary operations reduced R to I_p , we have that $R^{-1} = E_p \cdots E_1 I$. Hence, from (12) we have

$$\mathbf{W} = R^{-1}\mathbf{b} \quad \text{and} \quad \mathbf{d}^T = R^{-1}\mathbf{d}.$$

Substituting this into (10) gives

$$\mathbf{x}_1 = R^{-1}\mathbf{d} - R^{-1}\mathbf{d}_2.$$

4. PHASE II OF SIMPLEX METHOD

The aim in this and the next section is to explain the essentials of the simplex method; it is not to present the most current computer implementations of the simplex method. In Section 5, where we discuss the revised simplex method, we shall go into greater detail concerning the computations in current computer codes. For more complete discussions of computational questions see Dantzig [1962], Chvátal [1982], and Gass [1982].

We suppose that the problem is formulated as in relation (4) of Section 1. Phase I of the simplex, which we will describe in Section 8, either determines that the feasible set is empty or determines a basic feasible solution. We begin our discussion of phase II with the assumption that by means of phase I, or otherwise, we have obtained a basic feasible solution $\bar{\mathbf{q}} = (\bar{\mathbf{q}}_B, \bar{\mathbf{q}}_N)$, where $\bar{\mathbf{q}}_B$ is the zero vector in \mathbb{R}^{m-p} . We denote the value of the objective function at $\bar{\mathbf{q}}$ by ζ . Thus, $\zeta = (\bar{\mathbf{c}}_B, \bar{\mathbf{c}}_N)$. We call $\bar{\mathbf{q}} = (\bar{\mathbf{q}}_B, \bar{\mathbf{q}}_N)$ the *current basic feasible solution*, the matrix B associated with $\bar{\mathbf{q}}_B$ the *current basis*, and ζ the *current value* of the objective function. We abuse the phrase for good terminology and also call the components of $\bar{\mathbf{q}}_B$ the *current basis*. The simplex method consists of a sequence of iterations. In a given iteration we determine either that the current basic feasible solution is optimal or that there is no optimal solution or we determine a new basic feasible solution that does not decrease the value of the objective function. In the last event, we regard the iteration, with the new basic feasible solution as the current one. Each iteration consists of four steps. We shall illustrate the process using Example 1.1.

Step 1. Express x and the basic variables x_B in terms of the nonbasic variables x_N .

The equation in (1) of Section 3 can be written as

$$\begin{pmatrix} B & N & b_0 \\ K_B & K_N & -1 \end{pmatrix} \begin{pmatrix} x_B \\ x_N \\ x \end{pmatrix} = \begin{pmatrix} b \\ 0 \end{pmatrix}, \quad (1)$$

where b_0 is the zero vector in \mathbb{R}^m . If we multiply the augmented matrix corresponding to the first m equations in (1) by B^{-1} on the left, we see that (1) is equivalent to

$$\begin{pmatrix} I_m & B^{-1}N & b_0 \\ K_B & K_N & -1 \end{pmatrix} \begin{pmatrix} x_B \\ x_N \\ x \end{pmatrix} = \begin{pmatrix} B^{-1}b \\ 0 \end{pmatrix}, \quad (2)$$

where I_m is the $m \times m$ identity. Setting $x_B = \xi_B$ and $x_N = \xi_N$ in (2) gives

$$\xi_B = B^{-1}b. \quad (3)$$

From (2) and (3) we get that

$$x_B = \xi_B - B^{-1}N\xi_N. \quad (4)$$

From (2) and (4) we get that

$$x = (K_B \xi_B) + (K_N - K_B B^{-1}N) \xi_N.$$

Setting $x_N = \xi_N$ gives

$$z = (b_0, \xi_B) \quad (5)$$

and so

$$z = z + (K_N - K_B B^{-1}N) \xi_N. \quad (6)$$

Equations (5) and (6) express x_B and z in terms of x_N .

Let

$$d_j = \xi_B - K_B B^{-1}N = (d_{j_1}, \dots, d_{j_m}), \quad (7)$$

where j_1, \dots, j_m are the indices of the components of x that constitute x_B . Then we may write (6) as

$$z = z + (d_N, x_N) = z + \sum_{j=1}^k d_{j_i} x_{j_i}. \quad (8)$$

Remark 4.1. Note that by virtue of (3) the right-hand side of (2) can be written as \mathbf{b}_0/\mathcal{C} .

The difficult calculation in (4) and (6) is the calculation of \mathbf{B}^{-1} . The other calculations involve straightforward matrix and vector-matrix multiplications. We now present a second, seemingly different, method of obtaining \mathbf{x}_0 and z in terms of \mathbf{x}_0 . Computer implementation of the second method turns out in essence to be the same as the computer implementation of the previous method. Our use of the second method is didactic. For small-scale problems solved by hand this method is advantageous and will be the one used in our illustrative example. The tableau method is a method that systematically records the sequence of steps and calculations that we shall describe.

We again start with the augmented matrix of (1). By a sequence of elementary operations applied to the first m rows of (1), we can transform this matrix to an equivalent matrix

$$\left(\begin{array}{ccc|c} \mathbf{I}_m & \mathbf{B} & \mathbf{b}_0 & r' \\ \mathbf{K}_0 & \mathbf{K}_0 & -1 & 0 \end{array} \right). \quad (9)$$

Each of these operations can be effected by a premultiplication by an elementary matrix \mathcal{E}_i . Since the sequence of operations transforms \mathbf{B} to \mathbf{I} , the product of the elementary matrices associated with these transformations is equal to \mathbf{B}^{-1} , as was shown in the preceding section. Thus, the matrix (9) can also be obtained by a premultiplication of the first m rows of the augmented coefficient matrix (1) by the matrix \mathbf{B}^{-1} . Hence, $\mathbf{B} = \mathbf{B}^{-1}\mathbf{B}$ and $r' = \mathbf{B}^{-1}\mathbf{r}$, and we have, as before, that (1) is equivalent to a system with augmented matrix

$$\left(\begin{array}{ccc|c} \mathbf{I}_m & \mathbf{B}^{-1}\mathbf{B} & \mathbf{b}_0 & r' \\ \mathbf{K}_0 & \mathbf{K}_0 & -1 & 0 \end{array} \right), \quad (10)$$

where $r' = \mathbf{B}^{-1}\mathbf{r}$. From (10) we again get (3) and (4). Let us now suppose that the columns of \mathbf{B} are the columns $\mathbf{j}_1, \dots, \mathbf{j}_m$ of \mathbf{A} . If we successively multiply row 1 of (10) by $-\mathbf{j}_1$ and add the result to the last row, multiply row 2 of (10) by $-\mathbf{j}_2$ and add the result to the last row, and so on, then we transform (10) into

$$\left(\begin{array}{ccc|c} \mathbf{I}_m & \mathbf{B}^{-1}\mathbf{B} & \mathbf{b}_0 & r' \\ \mathbf{K}_0 & \mathbf{K}_0 - \mathbf{K}_0\mathbf{B}^{-1}\mathbf{B} & -1 & -(\mathbf{K}_0\mathbf{b}_0 + r') \end{array} \right). \quad (11)$$

From the last row of this matrix and from the relation $r' = \mathbf{B}^{-1}\mathbf{r}$, we immediately obtain

$$z = (\mathbf{K}_0\mathbf{b}_0 + r') + (\mathbf{K}_0 - \mathbf{K}_0\mathbf{B}^{-1}\mathbf{B})\mathbf{x}_0.$$

Remark 4.1. From (11) and $x' = B^{-1}b$ we again get (3) and (5). Consequently, the right-hand side of (11) may be written as $b_0 - z$. From this relation and from (3) and (5) we again get (8).

We illustrate the preceding description of Step 1 using the second method.

Example 4.1. We first introduce slack variables x_4, x_5, x_6 and write the problem as in (2) of Section 1. We then write the problem as a maximization problem in the form of (4) of Section 3.

Maximize z subject to

$$\begin{aligned} -x_2 - x_3 + x_4 &= -1, \\ -x_2 + x_3 + x_5 &= 0, \\ x_1 + 2x_2 + x_3 + x_6 &= 4, \\ -x_1 + 2x_2 &- z = 0. \end{aligned} \quad (12)$$

It is readily verified that $\bar{x} = (1, 0, 0, 1, 0, 0)$, which corresponds to the point (1, 0) in Figure 5.1, is a basic feasible solution. The value of z at this point is -1 . Thus, x_4, x_5, x_6 are the basic variables and x_1, x_2 are the nonbasic variables. The augmented coefficient matrix of (12) is

$$\begin{array}{c} x_4 \\ x_5 \\ x_6 \end{array} \left\{ \begin{array}{cccccc|c} x_1 & x_2 & x_3 & x_4 & x_5 & z & c \\ -1 & -1 & 1 & 0 & 0 & 0 & -1 \\ -1 & 1 & 0 & 1 & 0 & 0 & 0 \\ 1 & 2 & 0 & 0 & 1 & 0 & 4 \\ -1 & 2 & 0 & 0 & 0 & -1 & 0 \end{array} \right\} \quad (13)$$

Above each column of the matrix we have indicated the variable corresponding to the column. At the left of the matrix we have listed the current basic variables. The last column is the column of right-hand sides of (12). In the tabular description of the simplex method the matrix (13) is the initial tableau. The matrix B in (13) is the submatrix formed by the first three rows and columns 4, 5, and 6. Thus

$$B = \begin{pmatrix} -1 & 0 & 0 \\ -1 & 1 & 0 \\ 1 & 0 & 1 \end{pmatrix}$$

To obtain the matrix corresponding to (13), we must first transform B to the identity matrix by a sequence of elementary row operations on (13). From the form of B it is clear that the most effective sequence of operations is the following. (i) Multiply row 1 by -1 . (ii) Add the transformed row 1 to row 2.

(iii) Add the negative of the transformed row 1 to row 3. To obtain the row corresponding to the last row of (11), add the transformed row 1 to row 4. We thus obtain

$$\begin{array}{l} x_1 \\ x_2 \\ x_3 \\ x_4 \\ z \end{array} \begin{pmatrix} x_1 & x_2 & x_3 & x_4 & x_5 & z & | \\ 1 & 1 & -1 & 0 & 0 & 0 & 1 \\ 0 & 2 & -1 & 1 & 0 & 0 & 1 \\ 0 & 1 & 1 & 0 & 1 & 0 & 2 \\ 0 & 2 & -1 & 0 & 0 & -1 & 1 \end{pmatrix}. \quad (14)$$

From this we read off

$$\begin{aligned} x_1 &= 1 - x_2 + x_{3p} \\ x_2 &= 1 - 2x_{3p} + x_{4p} \\ x_3 &= 2 - x_1 - x_{4p} \end{aligned} \quad (15)$$

and

$$z = -1 + 2x_{3p} - x_{4p} \quad (16)$$

which are (4) and (5) for this example. Note that the last column of (14) gives the current values of ζ_{1p} , ζ_{2p} , ζ_{3p} , $-\zeta^*$ in agreement with (4) and (5) and Remark 4.2.

Step 2. Check the current basic feasible solution for optimality.

If every component of the vector d_p defined in (7) is nonpositive (i.e., $d_p \leq 0$), then the current basic feasible solution is optimal and the process terminates.

To see that this is so, recall that by Remark 3.1 we need only consider basic feasible solutions in our search for the optimal solution. The current basic feasible solution is the unique basic feasible solution with $x_p = \bar{x}_p$. Thus any other basic feasible solution will have at least one of the components of x_p positive. But since $d_p \leq 0$, we have from (8) that the value of the objective function will either decrease or be unchanged. Thus the current basic feasible solution is optimal.

In our example, we see from (16) that the current solution is not optimal.

Step 3. If the current basic feasible solution is not optimal, either determine a new basic and corresponding basic feasible solution that does not decrease the current value of z or determine that z is unbounded and that therefore the problem has no solution.

We first motivate Step 3 using Example 4.1.

From (18) we see that if we keep the value of x_2 at zero and increase the value of x_1 , then we shall increase the value of z . The variables x_1, x_2, x_3, x_4, x_5 must satisfy (15) and the condition $x_i \geq 0, i = 1, \dots, 5$, for feasibility. Setting $x_2 = 0$ in (15) and taking into account the requirement that x_1, x_3 , and x_5 must be nonnegative give the following inequalities that x_1 must satisfy:

$$1 - x_1 \geq 0, \quad 1 - 2x_1 \geq 0, \quad 3 - x_1 \geq 0.$$

For all three of these inequalities to hold we can only increase the value of x_1 to $\frac{1}{2}$. If $x_1 = \frac{1}{2}$ and $x_2 = 0$, then from (15) we get that $x_3 = \frac{1}{2}, x_4 = 0$, and $x_5 = \frac{3}{2}$. We now have another feasible solution $(\frac{1}{2}, \frac{1}{2}, 0, 0, \frac{3}{2})$. This corresponds to the point $(\frac{1}{2}, \frac{1}{2})$ in Figure 6-1. The value of z is increased from -1 to $\frac{1}{2}$.

We now show that this solution is basic. The matrix corresponding to the system (15) is the submatrix of (14) consisting of the first three rows of (14). The matrix B corresponding to $(\frac{1}{2}, \frac{1}{2}, 0, 0, \frac{3}{2})$ is the matrix formed by taking columns 1, 2, and 5 of this submatrix. Thus

$$B = \begin{pmatrix} 1 & 1 & 0 \\ 0 & 2 & 0 \\ 3 & 1 & 1 \end{pmatrix}.$$

Since the columns of B are linearly independent, the point $(\frac{1}{2}, \frac{1}{2}, 0, 0, \frac{3}{2})$ is a basic feasible solution.

In summary, the original basic variable x_1 , whose value was 1 is now a nonbasic variable and the original nonbasic variable x_2 is now a basic variable with value $\frac{1}{2}$. The variable x_1 is said to be the leaving variable or to leave the basis and the variable x_2 is said to be the entering variable or to enter the basis.

Before presenting Step 3 in general, we call attention to another possible situation. Suppose that we have an example such that instead of (15) we arrived at a system

$$x_1 = 1 + \alpha x_2 + x_3,$$

$$x_4 = 1 + \beta x_2 + x_5,$$

$$x_2 = 2 + \gamma x_1 + x_6,$$

with α as in (18) and where α, β , and γ are all nonnegative. If we set $x_2 = 0$ and let $x_1 = t$, then we obtain feasible solutions for all $t > 0$. If we let $t \rightarrow \infty$, we get that $z \rightarrow \infty$. Thus, the objective function is unbounded and the problem has no solution.

We now describe Step 3, which we break up into two substeps, for the general problem.

Step 3(i). Choose the nonbasic variable that will enter the new basis.

Let \mathbf{d}_q be as in (7). Since the current basic feasible solution is not optimal, at least one of the components of \mathbf{d}_q is positive. Choose the component of \mathbf{d}_q that is the largest. If there is more than one such component, choose the one with the smallest index. Suppose that the component of \mathbf{d}_q that is chosen is the q th component of $\mathbf{x} = (x_1, \dots, x_n)$. The variable x_q will enter the basis. The column of A corresponding to x_q is \mathbf{a}_q .

Step 3(i). Determine that the problem has no solution or choose the basic variable that will leave the basis.

Let \mathbf{x}_p be the column of $B^{-1}M$ corresponding to x_q . Then $\mathbf{x}_p = B^{-1}\mathbf{a}_q$. Let

$$\mathbf{x}_p = (x_{1p}, \dots, x_{mp}) \quad \mathbf{d}_p = (d_{1p}, \dots, d_{mp}) \quad (17)$$

and let $\mathbf{x}_q = (x_{1q}, \dots, x_{mq})$. Set $x_q = t$, where $t \geq 0$ and let the other components of \mathbf{x}_q remain at zero. For the resulting vector to be a feasible solution, \mathbf{x}_p must satisfy the conditions (4) and $x_q \geq 0$. Equation (4) becomes

$$\mathbf{x}_p = \mathbf{d}_p - t\mathbf{x}_q \quad (18)$$

Thus, the condition $\mathbf{x}_p \geq 0$ imposes the restrictions

$$d_{rp} \leq t x_{rq}, \quad r = 1, \dots, m. \quad (19)$$

From (8) we get that

$$z = \bar{z} + d_q t \quad (20)$$

Since \bar{z} is feasible, $\bar{z}_0 \geq 0$ and so $d_q \geq 0$ for all $r = 1, \dots, m$. Hence if $x_{rq} \leq 0$ for all $r = 1, \dots, m$, then (18) and (19) hold for all $t \geq 0$. Since $d_q > 0$, it follows that z is unbounded above and the problem has no solution.

If not all of the x_{rq} are nonpositive, let

$$t^* = \min \left\{ \frac{d_{rp}}{x_{rq}}; r = 1, \dots, m \text{ and } x_{rq} > 0 \right\}. \quad (21)$$

Let p denote the smallest index r at which the minimum is achieved. Let $x_{pq} = t^*$ and let

$$\begin{aligned} \bar{z}_p &= \bar{z}_0 + t^* d_q & \bar{z} &= \bar{z}_0 \\ \bar{z}_q &= t^* & & \\ \bar{z}_r &= 0 & & \\ \bar{z}_i &= 0 & \text{all other indices.} & \end{aligned} \quad (22)$$

Let ξ^q denote the vector whose coordinates are given by (22). Then

$$\xi^q = k_{1,p}^q \dots k_{j,p}^q \delta + k_{j,p}^q \dots k_{j,p}^q \delta + r^q \delta_1 \dots \delta_j$$

where r^q is the q th coordinate of ξ^q . It follows from the way ξ^q was constructed that ξ^q is a feasible solution. We shall show that ξ^q is also basic.

We can write ξ^q as $\xi^q = (\xi_{1,p}^q, \xi_{j,p}^q)$ with

$$\xi_{1,p}^q = k_{1,p}^q \dots k_{j,p}^q r^q, \quad \xi_{j,p}^q = \delta + k_{j,p}^q \delta$$

We shall show that the coordinate variables of

$$x_p = (x_{1,p}, \dots, x_{j,p}, x_p, x_{j,p}, \dots, x_{n,p}) \quad (23)$$

are a basis. Since x_p is obtained from $x_{j,p}$ by replacing $x_{j,p}$ by x_p , this will also justify our saying that $x_{j,p}$ enters the basis and $x_{j,p}$ leaves the basis.

From (17) we get that the matrix J_p in (2) and (5) is

$$(x_1 \dots x_{j-1}, x_j, x_{j+1}, \dots, x_n)$$

where x_j is the vector whose j th component is 1 and whose other components are zero. Let

$$K = (x_1 \dots x_{j-1}, x_p, x_{j+1}, \dots, x_n)$$

That is, K is the matrix obtained from J_p by replacing the j th column by x_p . Let N be the matrix obtained from $K^{-1}K$ by replacing column x_p by $x_{j,p}$. Let x_p be as in (23) and let $x_{j,p}$ be the vector obtained from x_p by replacing x_p by $x_{j,p}$. Let d_p denote the q th coordinate of the vector x_p defined in (7) and let $d_{j,p}$ denote the vector obtained from d_p by replacing the q th coordinate of d_p by 0. The system of equations

$$\begin{pmatrix} x_1 & \dots & x_{j-1} & x_j & x_{j+1} & \dots & x_n & N & x_{j,p} \\ 0 & \dots & 0 & d_p & 0 & \dots & 0 & d_{j,p} & -1 \end{pmatrix} \begin{pmatrix} x_1 \\ \vdots \\ x_{j-1} \\ x_j \\ \vdots \\ x_n \\ x_{j,p} \end{pmatrix} = \begin{pmatrix} 0 \\ \vdots \\ 0 \\ -d_p \\ \vdots \\ 0 \\ -d_{j,p} \end{pmatrix} \quad (24)$$

is equivalent to the system whose augmented matrix is (11), since (24) is obtained from (11) by a reordering of the variables.

We assert that the matrix K is a basis and hence that the variables x_p are basic. To prove this, we must show that the columns of K are linearly independent. By (21) and the definition of the index p , the j th component of x_j is positive. The j th component of each of the other columns of K is zero. Hence x_j cannot be written as a linear combination of the other columns of K . Since the remaining columns of K are linearly independent, it follows that all the columns of K are linearly independent. Hence the coordinate variables of

\mathbf{a}_p are a basis, and the vector ζ is a basic feasible solution. From (20) we get that

$$\zeta = \zeta + d_p t^2.$$

Since $d_p > 0$ and $t > 0$, it follows that the value of the objective function is not decreased.

We say that in Step 3 we have updated the unpriced quantities to the priced quantities. The priced, or updated, quantities are the "new" or "updated current" quantities.

Remark 4.1. In Step 3(i) we choose the entering variable x_p to be the variable corresponding to the largest component of \mathbf{d}_p . Upon returning to Step 1, it should be clear that the choice of entering variable could be any variable corresponding to a positive component of \mathbf{d}_p . The choice of largest component was motivated by the fact that at a given iteration this choice would appear to result in the largest increase in the value of z .

Step 4. Return to Step 1 using the updated current quantities and start the next iteration.

We now illustrate Steps 3 and 4 using Example 4.1.

Step 3(i). From (14) we see that $\mathbf{d}_p = (d_1, d_2)$ and that $d_1 > 0$ and $d_2 < 0$. Thus, we take x_1 as the entering variable. The entering column \mathbf{v}_1 is the second column of (14), and $d_p = d_1 = 3$.

Step 3(ii). By Remark 4.2, the entries in the last column of (14) give

$$(\zeta_1, \zeta_2, \zeta_3) = (\zeta_1, \zeta_2, \zeta_3) = (1, 1, 3), \quad \zeta = -1.$$

From this and the fact that the first three entries in column 2 are positive, we get that

$$\frac{\zeta_1}{v_{11}} = \frac{\zeta_1}{v_{11}} = 1, \quad \frac{\zeta_2}{v_{21}} = \frac{\zeta_2}{v_{21}} = \frac{1}{2}, \quad \frac{\zeta_3}{v_{31}} = \frac{\zeta_3}{v_{31}} = 3.$$

Hence $t^* = \frac{1}{2}$ and x_2 is the leaving variable. The new, or updated, basic variables are (x_1, x_2, x_3) and the updated nonbasic variables are x_4 and x_5 . From (21) we get the updated basic feasible solution

$$\begin{aligned} \zeta'_1 = \zeta_1 - t^*(1) &= \frac{1}{2}, & \zeta'_2 = \zeta_2 - t^* &= \frac{1}{2} \\ \zeta'_3 = \zeta_3 &= 3, & \zeta'_4 = \zeta_4 = \zeta_5 &= 3 - \frac{1}{2}(1) = \frac{5}{2}. \end{aligned}$$

From (20) we get

$$z^* = (-1) + 3\left(\frac{1}{3}\right) = \frac{2}{3}.$$

These calculations, however, are unnecessary as we will get them automatically in Step 1 of the next iteration.

Step 4. We go to Step 1 and start the next iteration using the updated data. We do so dropping the primes that indicated updated data. The variables x_1, \dots, x_6, z satisfy the system of equations represented by (14). We rewrite (14) with the labeling of the new basic variables:

$$\begin{array}{l} x_1 \\ x_2 \\ x_3 \end{array} \begin{pmatrix} x_1 & x_2 & x_3 & x_4 & x_5 & z & 1 \\ 1 & 1 & -1 & 0 & 0 & 0 & 1 \\ 0 & 2 & -1 & 1 & 0 & 0 & 1 \\ 0 & 1 & 1 & 0 & 1 & 0 & 2 \\ 0 & 3 & -1 & 0 & 0 & -1 & 1 \end{pmatrix}. \quad (21)$$

Step 1 requires us to transform (21) into a form corresponding to (11). The matrix B is now

$$B = \begin{pmatrix} 1 & 1 & 0 \\ 0 & 2 & 0 \\ 0 & 1 & 1 \end{pmatrix}.$$

Thus, to achieve the desired result, we apply the following sequence of transformations to (21).

- (i) Multiply row 2 by $\frac{1}{2}$.
- (ii) Add the negative of the new row 2 to row 1.
- (iii) Add the negative of the new row 2 to row 3.
- (iv) Multiply the new row 2 by -3 and add the resulting row to row 4.

We get

$$\begin{array}{l} x_1 \\ x_2 \\ x_3 \end{array} \begin{pmatrix} x_1 & x_2 & x_3 & x_4 & x_5 & z & 1 \\ 1 & 0 & -\frac{1}{2} & -\frac{1}{2} & 0 & 0 & \frac{3}{2} \\ 0 & 1 & -\frac{1}{2} & \frac{1}{2} & 0 & 0 & \frac{1}{2} \\ 0 & 0 & \frac{1}{2} & -\frac{1}{2} & 1 & 0 & \frac{5}{2} \\ 0 & 0 & \frac{1}{2} & -\frac{1}{2} & 0 & -1 & -\frac{1}{2} \end{pmatrix}. \quad (22)$$

From (22) we can read off, if we desire, the basic variables $(x_1, x_2, x_3) = (x_1, x_2, x_3)$ and z in terms of the nonbasic variables (x_4, x_5) . By Remark 4.3 the last column of (22) gives

$$z_1 = \frac{3}{2}, \quad z_2 = \frac{1}{2}, \quad z_3 = \frac{5}{2}, \quad z = -\frac{1}{2}.$$

for \mathbf{x}_p , the current basic feasible solution, and ζ , the current value of the objective function. Note that we have agreement with the previous calculation of these quantities. The current basic feasible solution corresponds to the point $(\frac{1}{2}, \frac{1}{2})$ in Figure 6.1. The current value $\zeta = \frac{1}{2}$ is greater than the previous-current value of -1 .

Before proceeding to Step 2, we introduce some terminology and comment on the carrying-out of Step 1 in general. The entry a_{rp} in row r , column p in this example is called the *pivot* and the position in the matrix is called the *pivot position*. In performing the sequence of elementary operations used to obtain (26) from (25), we say that we have *pivoted on the element* a_{rp} .

In our example the pivot was the entry in the row and column of the entering variable. We now show that in each iteration after the initial iteration we always pivot on the element in the row and column of the entering variable. The system (26) gives the relationship among the variables x_1, x_2 , and s at the end of Step 4 of each iteration. It expresses s and each component of

$$\mathbf{x}_p = (x_1, \dots, x_{p-1}, x_p, x_{p+1}, \dots, x_m)$$

except x_p in terms of the components of \mathbf{x}_p , and the variable x_p . Step 1 of the current iteration requires that \mathbf{x}_p and s be expressed in terms of the components of \mathbf{x}_p . This will be achieved if we express x_p in terms of the components of \mathbf{x}_p and do so in a way that does not destroy the already established expressions of s and the a_{ij} , $j = 1, \dots, m$, $j \neq p$, in terms of x_1 and x_2 . We assert that pivoting on row i_p for entering row, and on column i_p (column \mathbf{x}_p) the entering column, achieves this.

By the definition of ρ and (21), we have that $a_{i_p} > 0$. Hence in (24) we may multiply row i_p by $1/a_{i_p}$ and obtain a row with entry 1 in the (i_p, i_p) position. If we now multiply row i_p by $-a_{ij}$ and add to row i , the resulting first row will have a zero in the i_p th position; that is, the column vector \mathbf{x}_p now has a 1 in the i_p entry and a zero in the first position. Since each of the vectors \mathbf{a}_j , $j = 1, \dots, m$, $j \neq p$, has a zero in the i_p th entry, the first entries of these vectors are not changed. Similarly, if for each $j = 2, \dots, m$, $j \neq p$, we multiply row i_p by $-a_{ij}$ and add the resulting row to row i , then the resulting i th row will have a zero in the i_p th position. Each of the vectors \mathbf{a}_j , $j = 1, \dots, m$, $j \neq p$, will have its i th entries unchanged. Finally, if we multiply row i_p by $-b_{i_p}$ and add the resulting row to the last row of (26), then we shall get a system

$$\begin{pmatrix} \mathbf{A}_p & \mathbf{N} & \mathbf{b}_p \\ \mathbf{a}_p & \mathbf{d}_p & -1 \end{pmatrix} \begin{pmatrix} \mathbf{x}_p \\ \mathbf{x}_n \\ s \end{pmatrix} = \begin{pmatrix} \mathbf{c}' \\ -r \end{pmatrix}$$

This system expresses \mathbf{x}_p and s in terms of \mathbf{x}_n .

We now return to Step 2 and (26) of the current iteration. Since $\mathbf{d}_p = \mathbf{d}_p - \mathbf{d}_p$, the current basic feasible solution is not optimal. The current value $\zeta = \frac{1}{2}$ is an improvement, however.

Step 3(i). Since $d_1 > 0$ and $d_2 < 0$, we take x_2 to be the entering variable.

Step 3(ii). The column corresponding to x_2 is the third column of (26). Since the only positive entry in this column is the entry in row 3 corresponding to x_3 , the variable x_2 will be leaving variable. We rewrite the matrix (26) with the updated basic variables indicated on the left:

$$\begin{array}{l} x_1 \\ x_2 \\ x_3 \end{array} \left\{ \begin{array}{cccccc|c} x_1 & x_2 & x_3 & x_4 & x_5 & r & \\ 1 & 0 & -\frac{1}{2} & -\frac{1}{2} & 0 & 0 & \frac{3}{2} \\ 0 & 1 & -\frac{1}{2} & \frac{1}{2} & 0 & 0 & \frac{5}{2} \\ 0 & 0 & \frac{1}{2} & -\frac{1}{2} & 1 & 0 & \frac{7}{2} \\ 0 & 0 & \frac{1}{2} & -\frac{1}{2} & 0 & -1 & -\frac{5}{2} \end{array} \right\}$$

We now follow Step 4 and begin another iteration with Step 1. We follow our pivot rule and pivot on $\frac{1}{2}$ in row 3, column 3. This will result in a matrix

$$\begin{array}{l} x_1 \\ x_2 \\ x_3 \end{array} \left\{ \begin{array}{cccccc|c} x_1 & x_2 & x_3 & x_4 & x_5 & r & \\ 1 & 0 & 0 & -\frac{1}{2} & \frac{1}{2} & 0 & \frac{4}{3} \\ 0 & 1 & 0 & \frac{1}{2} & \frac{1}{2} & 0 & \frac{7}{3} \\ 0 & 0 & 1 & -\frac{1}{2} & \frac{1}{2} & 0 & \frac{2}{3} \\ 0 & 0 & 0 & -\frac{1}{2} & -\frac{1}{2} & -1 & -\frac{11}{3} \end{array} \right\}$$

Step 2. $\bar{d}_3 = (-\frac{1}{2}, -\frac{1}{2})$. All components of \bar{d}_3 are negative, so the current basic feasible solution is optimal. From the last column we get that

$$z_{\max} = 6x_1 + 4x_2 + 2x_3 = \frac{2}{3} \left(\frac{4}{3} + \frac{7}{3} \right)$$

and that the value of the objective function is $z = \frac{20}{3}$. This solution corresponds to the point $(\frac{4}{3}, \frac{7}{3})$ in Figure 6.1.

8. TERMINATION AND CYCLING

In Section 4 we showed that each iteration of the simplex method results in a value of the objective function that is greater than or equal to the value of the objective function at the preceding iteration. If we could show that each iteration produces a value of the objective function that is strictly greater than the value at the preceding iteration, then the simplex method would terminate at an optimal solution in a finite number of steps, for there are only a finite number of basic feasible solutions and there are a finite number of basic feasible solutions. Unfortunately, it is not true that each iteration

produces a value of the objective function that is strictly greater than the preceding value, as the following simple example shows.

Let us suppose that at the end of Step 2 of an iteration we had arrived at a system given by

$$\begin{array}{l} x_3 \\ x_4 \\ x_2 \end{array} \begin{pmatrix} x_1 & x_2 & x_3 & x_4 & x_5 & b & \theta \\ 1 & 1 & 1 & 0 & 0 & 0 & 1 \\ 0 & 2 & -1 & 1 & 0 & 0 & 0 \\ 0 & 1 & 1 & 0 & 1 & 0 & 2 \\ 0 & 2 & -1 & 0 & 0 & -1 & 1 \end{pmatrix}.$$

Then the entering variable is x_1 and the entering column \mathbf{v}_1 is the second column. Recalling that the last column gives $(\mathbf{b}_B, -\zeta)$, we get that

$$\frac{\zeta_1}{v_1} = \frac{\zeta_2}{v_2} = \frac{1}{2}, \quad \frac{\zeta_3}{v_3} = \frac{\zeta_4}{v_4} = \frac{0}{2} = 0, \quad \frac{\zeta_5}{v_5} = \frac{\zeta_6}{v_6} = \frac{1}{2}. \quad (1)$$

Thus $r^* = 0$ and the leaving variable is x_4 . It then follows from (4.20) that $\zeta^* = \zeta$ and the updated value of \mathbf{a} is not changed. From (4.22) we get that the updated basic feasible solution is

$$\mathbf{b}_B = H_{21}^{-1} (\mathbf{b}_B, \zeta_2^*) = H_{21}^{-1} (\mathbf{b}_B, \zeta_2) = (1, 0, 2).$$

Looking at (1), we see that $r^* = 0$ because $\zeta_4 = 0$. In other words, $r^* = 0$ is a consequence of the current basic feasible solution being degenerate. In this example, the updated basic solution is also degenerate. In the next example we show that a degenerate basic solution need not lead to another degenerate solution with no increase in the value of \mathbf{a} .

Let us now suppose that at the end of Step 2 we arrived at the system described by

$$\begin{array}{l} x_3 \\ x_4 \\ x_2 \end{array} \begin{pmatrix} x_1 & x_2 & x_3 & x_4 & x_5 & b & \theta \\ 1 & 1 & 0 & 0 & 0 & 0 & 1 \\ 0 & -2 & -1 & 1 & 0 & 0 & 0 \\ 0 & 1 & 1 & 0 & 1 & 0 & 2 \\ 0 & 2 & -1 & 0 & 0 & -1 & 1 \end{pmatrix}.$$

The current basic feasible solution $(\mathbf{b}_B, \zeta_2, \zeta_3) = (1, 0, 2)$ is again degenerate. The entering variable is again x_1 . Now, however, the ratios used to determine r^* are

$$\frac{\zeta_1}{v_1} = \frac{\zeta_2}{v_2} = 1, \quad \frac{\zeta_3}{v_3} = \frac{\zeta_4}{v_4} = 2.$$

Thus $\theta^* = 1$ and the leaving variable is x_1 . Since $d_1^* = 1$, $\bar{c}^* = -1$, and $\theta^* = 1$, from (4.20) we get

$$\bar{c} = -1 + M(1) = 2.$$

Thus the value of s is increased. From (4.22) we calculate the new basic feasible solution to be

$$\bar{b}_0 = (4, 1, 0, 0) = (b_0^*, b_1^*, b_2^*, b_3^*),$$

which is nondegenerate.

We now consider the general situation. Let b_{0j} be a basic degenerate feasible solution with zero component b_{0j} . If the corresponding component a_{ij} of the leaving column v is positive, then $\theta^* = 0$, the updated basic feasible solution will be degenerate, and the updated value of the objective function will be unchanged. This follows from (4.21), (4.22), and (4.20).

The following question now arises: Can there exist a sequence of iterations k_1, k_2, \dots, k_{p-1} such that the basic feasible solutions $b_{0k_i}, i = 0, 1, \dots, p$, are all degenerate and such that $b_{0k_{i+1}} = b_{0k_i}$? If this occurs, then all current values for $b_{0k_1}, \dots, b_{0k_p}$ of the objective will be equal and the sequence of degenerate basic solutions $b_{0k_1}, b_{0k_2}, \dots, b_{0k_p}$ will repeat infinitely often. Thus, the simplex method will not terminate. The phenomenon just described is called *cycling* and the simplex method is said to *cycle*.

The bad news is that cycling can occur. The good news is that cycling appears to be extremely rare in problems arising in applications, and there are modifications to the simplex method that prevent cycling.

The following example due to Brink [1953] cycles in six iterations.

Example 3.1. Minimize $z = 22x_1 - 90x_2 - 21x_3 + 24x_4$ subject to

$$\begin{aligned} -6x_1 + 8x_2 + 2x_3 &+ x_4 - 9x_5 &= 0, \\ \frac{1}{2}x_1 - \frac{1}{2}x_2 - \frac{1}{2}x_3 + x_4 &+ x_5 + x_6 &= 0, \\ x_1 & &+ x_6 &= 1, \\ x_j &\geq 0, & j = 1, \dots, 6. \end{aligned}$$

From the preceding discussion and from Section 4 we have the following:

Lemma 3.1. Phase *M* of the simplex method does one of the following:

- (i) determines an optimal solution in a finite number of steps,
- (ii) determines that no solution exists in a finite number of steps, or
- (iii) cycles.

We shall present a very simple rule for the selection of entering and leaving variables that makes cycling impossible. This rule is the *smallest index rule* or *Bland's rule*, named after its originator R. G. Bland [1977]. Other methods for preventing cycling are the perturbation method and the related lexicographic entering rule. For a discussion of these methods see Chvátal [1983], Danzig [1963], and Gass [1985].

Before stating the smallest index rule, we recall that we pointed out in Remark 4.3 that the choice of entering variable could be any variable corresponding to a positive component of \mathbf{d}_k . Also, it is possible to choose any variable x_i at which r^* in (21) of Section 4 is attained.

Definition 5.1 (Smallest Index Rule). From among all possible entering variables, choose the one with smallest index. From among all possible leaving variables, choose the one with smallest index.

Remark 5.1. According to the smallest index rule, the choice of entering variable will, in general, be different from the choice of index such that d_k is a maximum. The choice of leaving variable is the same as that previously used.

Theorem 5.1. If at each iteration of the simplex method the smallest index rule is used, then cycling will not occur.

Proof. We shall prove the theorem by showing that the use of the smallest index rule and the assumption of cycling lead to a contradiction. The argument that we shall give follows the one by Chvátal [1983].

Let us assume that cycling occurs and that $\mathbf{B}_0, \mathbf{B}_1, \dots, \mathbf{B}_{p-1}$ is a sequence of bases corresponding to iterations I_0, I_1, \dots, I_{p-1} and such that $\mathbf{B}_{p-1} = \mathbf{B}_0$. Since cycling occurs, all of these bases are degenerate.

Let J denote the set of variables that are basic in some iterations and nonbasic in others. Let x_i denote the variable in J with the greatest index. If x_i is basic in \mathbf{B}_p , then there must exist an iteration \mathbf{B}_{c+1} with $0 < c < p$ such that x_i leaves \mathbf{B}_{c+1} and is nonbasic at the next iteration. If x_i is nonbasic in \mathbf{B}_p , then there must exist a basis \mathbf{B}_{l+1} with $0 < l < p$ such that x_i enters \mathbf{B}_{l+1} . Since $\mathbf{B}_{p-1} = \mathbf{B}_0$, there exists a basis \mathbf{B}_{j+1} with $l < j < p$ such that x_i leaves \mathbf{B}_{j+1} and is nonbasic at the next iteration. In conclusion, we have shown that there is a basis \mathbf{B} in the sequence $\mathbf{B}_0, \dots, \mathbf{B}_{p-1}$ such that x_i leaves \mathbf{B} at the iteration J corresponding to \mathbf{B} and becomes nonbasic in the next basis. Let x_r denote the entering variable at the iteration J . Thus, $x_r \in \mathbf{B}$ and x_i is in the next basis. Since cycling occurs, there is basis \mathbf{B}^* at an iteration I^* beyond J and in the sequence $\mathbf{B}_0, \mathbf{B}_1, \dots, \mathbf{B}_{p-1}, \mathbf{B}_0, \mathbf{B}_1, \dots, \mathbf{B}_{p-1}$ such that x_i enters the basis at I^* . That is, x_i is not in \mathbf{B}^* but is in the succeeding basis.

We shall write $i \in \mathbf{B}$ to mean that x_i is a basic variable in the basis \mathbf{B} and $i \notin \mathbf{B}$ to mean that x_i is nonbasic. If we denote the elements of the matrix $\mathbf{B}^{-1}\mathbf{b}$

in relation (4) of Section 4 by x_j , then we can write (4) and (5) of Section 4 as

$$\begin{aligned}x_j &= \zeta_j - \sum_{i \in B} a_{ij} x_i, & i \in B, \\z &= \zeta + \sum_{i \in B} d_i x_i.\end{aligned}\quad (3)$$

Since x_i enters at B , $i \notin B$. A solution of (3) is then given by

$$\begin{aligned}x_i &= y, & y \text{ arbitrary,} \\x_i &= 0, & i \in B, \quad i \neq k, \\x_i &= \zeta_i - a_{ik} y, & i \in B, \\z &= \zeta + d_k y.\end{aligned}\quad (4)$$

For the basic B^* equations (4) and (5) in Section 4 become

$$\begin{aligned}x_{B^*} &= \zeta_{B^*} - (B^*)^{-1} B^* x_B, \\z &= \zeta^* + \sum_{j \in B^*} d_j^* x_j,\end{aligned}\quad (5)$$

where $d_j^* = 0$ for $j \in B^*$.

The systems (3) and (4) are equivalent; that is, they have the same solution sets. In particular, (3) is a solution of (4). Since all of the bases from B to B^* are degenerate, the value of the objective function does not change from iteration to iteration. Hence, $\zeta^* = \zeta$. Equating the expressions for z in (3) and (4) and using (5) and $\zeta^* = \zeta$ give

$$\zeta + d_k y = \zeta + d_k^* y + \sum_{i \in B^*} d_i^* (\zeta_i - a_{ik} y).$$

Hence

$$\left(d_k - d_k^* + \sum_{i \in B^*} d_i^* a_{ik} \right) y = \sum_{i \in B^*} d_i^* \zeta_i \quad \text{for all } y.$$

Since the right-hand side is constant and the left-hand side holds for all y , we must have

$$d_k - d_k^* + \sum_{i \in B^*} d_i^* a_{ik} = 0. \quad (6)$$

Since x_k enters B at iteration L , $d_k = 0$. Also, since x_k is leaving, $i \neq k$. Hence $a_{ik} < 0$. Since x_k is not entering at B^* and $a_{ik} < 0$, it follows that $d_i^* < 0$. For

otherwise, by the smallest index rule, x_1 would not be the entering variable at P^r . It then follows from (7) that there exists an index r in B such that

$$d_j^r x_{10} < 0. \quad (8)$$

Since $d_j^r \neq 0$, the variable x_1 is not in the basis B^r . But $r \in B$, so $x_1 \in \mathcal{F}$. Hence $r \in \mathcal{L}_1$.

It follows from (11) in Section 4 that since x_1 is leaving in B and x_1 is entering, $x_{10} > 0$. Since x_1 is entering at B^r , it follows that $d_j^r > 0$. Hence $d_j^r x_{10} > 0$. Comparing this with (8) shows that $r \notin \mathcal{L}_1$. Hence $r \in \mathcal{L}_1$ since the variable x_1 is not in B^r , and since x_1 enters B^r , we conclude from the smallest index rule that $d_j^r \leq 0$. From (8) we conclude that $d_j^r < 0$ and

$$x_{10} > 0.$$

Since $x_1 \notin B^r$, the basic feasible solution Q^r at B^r has component $d_j^r = 0$. Let ζ_j denote the j th component of the basic feasible solution ξ corresponding to the basis B . We assert that $\zeta_j = 0$. To prove this assertion, we note that because all the bases B^1, \dots, B^r are degenerate, it follows from (21) of Section 4 that $d^j = 0$ at all the iterations $j \dots P^r$. Hence if $\zeta_j > 0$, then x_1 would not leave any of the bases B^1, \dots, B^r . This contradicts the fact that $x_1 \notin B^r$, and this proves the assertion.

Since $x_{10} > 0$ and $\zeta_j = 0$, it follows from (21) of Section 4 that x_1 is a candidate for leaving the basis B . Yet, we chose x_1 with $j > r$ to leave the basis, in contradiction to the smallest index rule.

Remark 5.2. Computational experience indicates that using the smallest subscript rule or the lexicographic rule instead of the usual d_j rule considerably increases the time required for the simplex method or the revised simplex method to find an optimal solution. Since very few instances of cycling have been reported, some computer codes do not incorporate any anticycling routines. Other computer codes introduce an anticycling routine only after a sequence of specified length of degenerate bases occurs. The anticycling routine will then lead to a nondegenerate basis, at which point the use of the largest coefficient rule can be resumed.

Exercise 5.1. Carry out the iterations of the simplex method for Example 5.1 and show that cycling occurs at the sixth iteration. Since $h_0 = Q_{h_0} = (p_1, 1) = (0, 0, 1)$ is a basic feasible solution.

Exercise 5.2. Solve the problem in Example 5.1 using the simplex method and Bland's rule.

6. PHASE I OF SIMPLEX METHOD

As noted in Section 3, phase I of the simplex method either determines a basic feasible solution or determines that none exists, in which case the problem has no optimal solution.

We begin with the linear programming problem in standard form (SLP):

$$\begin{aligned} & \text{Maximize } (\mathbf{k}, \mathbf{x}) \\ & \text{subject to } A\mathbf{x} \leq \mathbf{c}, \quad \mathbf{x} \geq \mathbf{0}. \end{aligned} \quad (1)$$

Let us first suppose that the vector \mathbf{c} is nonnegative. We then put the problem in canonical form (CLF) by adjoining a vector of slack variables,

$$\mathbf{x}_p = (x_{p_1}, \dots, x_{p_{m+n}}), \quad x_{p_j} \geq 0, \quad j = 1, \dots, m,$$

and write the problem as

$$\begin{aligned} & \text{Maximize } (\mathbf{k}, \mathbf{x}) \\ & \text{subject to } (A, I) \begin{pmatrix} \mathbf{x} \\ \mathbf{x}_p \end{pmatrix} = \mathbf{c}, \quad \mathbf{x} \geq \mathbf{0}, \quad \mathbf{x}_p \geq \mathbf{0}. \end{aligned} \quad (2)$$

Since $\mathbf{c} \geq \mathbf{0}$, the vector $(\mathbf{0}, \mathbf{x}_p) = (\mathbf{0}_m, \mathbf{c})$, where $\mathbf{0}_m$ is the zero vector in \mathbb{R}^m , is a basic feasible solution for the problem in (2). We then immediately pass to phase II with initial basic feasible solution $(\mathbf{0}_m, \mathbf{x}_p) = (\mathbf{0}_m, \mathbf{c})$.

If some of the components of \mathbf{c} are negative, the vector $(\mathbf{0}_m, \mathbf{c})$ is still a basic solution of the system of equations in (2), but it is not feasible. We therefore resort to another device and introduce an auxiliary problem.

Auxiliary Problem

Minimize x_0 subject to

$$\begin{aligned} & (A, I, -\mathbf{e}) \begin{pmatrix} \mathbf{x} \\ \mathbf{x}_p \\ x_0 \end{pmatrix} = \mathbf{c} \\ & \mathbf{x} \geq \mathbf{0}, \quad \mathbf{x}_p \geq \mathbf{0}, \quad x_0 \geq 0, \end{aligned} \quad (3)$$

where $\mathbf{e} = (1, \dots, 1)$ and I is the $m \times m$ identity matrix.

In component form the system in (3) is

$$\begin{aligned} & \sum_{j=1}^n a_{ij}x_j + x_{p_i} - x_0 = c_i, \quad i = 1, \dots, m, \\ & x_j \geq 0, \quad j = 0, 1, \dots, n+m. \end{aligned}$$

A vector (\bar{b}_1, \bar{b}_2) is feasible for relation (7) if and only if $(\bar{b}_1, \bar{b}_2, 0)$ is feasible for the auxiliary problem. Therefore, (\bar{b}_1, \bar{b}_2) is feasible for relation (7) if and only if $(\bar{b}_1, \bar{b}_2, 0)$ is optimal for the auxiliary problem.

If we write the objective function of the auxiliary problem as "Maximize $w = -s_0$," then the auxiliary problem becomes a linear programming problem in canonical form to which we can apply phase II of the simplex method, provided we can get an initial basic feasible solution. A basic solution to the system (1) is $(\bar{b}_1, \bar{b}_2, \bar{b}_3) = (\bar{b}, \bar{a}, 0)$. This solution, however, is not feasible since some components of \bar{a} are negative. We can obtain a basic feasible solution from this solution by replacing an appropriate basis variable by s_0 . Before we do this in general, we shall illustrate using Example 4.1.

In solving Example 4.1, we obtained an initial feasible solution from Figure 6.1. We now ignore the figure and write the auxiliary problem in canonical form: Maximize $w = -s_0$ subject to

$$\begin{aligned} -s_0 - x_1 + x_2 & & -s_0 & = -1, \\ -s_0 + x_1 & + x_2 & -s_0 & = 0, \\ x_1 + 2x_2 & & + x_3 - s_0 & = 4, \\ x_i & \geq 0, & i & = 0, 1, \dots, 3. \end{aligned} \quad (8)$$

The vector $(\bar{b}_0, \bar{b}_1, \bar{b}_2, \bar{b}_3, \bar{b}_4, \bar{b}_5) = (0, 0, 0, -1, 0, 4)$ is a basic solution but is not feasible. The only negative entry on the right-hand side is -1 and occurs in the first equation. Therefore, if in the augmented coefficient matrix of (8) we pivot on s_0 in row 1, then the right-hand side of the resulting system should be positive. Performing the indicated pivot gives

$$\begin{pmatrix} s_0 & x_1 & x_2 & x_3 & s_0 & s_0 & w \\ 1 & 1 & -1 & 0 & 0 & 1 & 1 \\ 0 & 2 & -1 & 1 & 0 & 0 & 1 \\ 1 & 3 & -1 & 0 & 1 & 0 & 5 \end{pmatrix} \quad (9)$$

Thus, $(\bar{b}_0, \bar{b}_1, \bar{b}_2, \bar{b}_3, \bar{b}_4, \bar{b}_5) = (1, 0, 0, 0, 1, 5)$ is a basic feasible solution. From (9) we also get that

$$w = -s_0 = -1 + x_1 + x_2 - s_0. \quad (10)$$

We can now start phase II for the problem of minimizing (6) subject to the constraint equations given by (7). We leave this to the reader and return to the general situation.

We now obtain a system equivalent to (8) and having a basic feasible solution. Let s_0 denote the most negative component of \bar{a} . Thus,

$$s_0 = \min\{a_i \mid i = 1, \dots, m\}.$$

If there are more than one such components, choose the one with smallest index. Next, pivot on a_{ip} in the p th row of the system (7). The result will be an equivalent system

$$(d', a_{1p}, \dots, a_{p-1p} = 0, a_{p+1p}, \dots, a_{mp}, a_p) \begin{pmatrix} x_1 \\ \vdots \\ x_p \\ \vdots \\ x_n \end{pmatrix} = d', \quad (7')$$

where $d' = (d, \dots, 1)$ and

$$d'_i = -d_i \geq 0, \quad d'_i = d_i - d_p \geq 0, \quad i = 1, \dots, p-1, p+1, \dots, m.$$

The system (7') in component form is

$$\begin{aligned} \sum_{j=1}^n a'_{ij} x_j - x_{n+p} + x_{n+1} &= d'_i, & i = 1, \dots, p, \\ \sum_{j=1}^n a'_{ij} x_j - x_{n+p} + x_i &= d'_i, & i = p+1, \dots, m. \end{aligned}$$

Thus a basic feasible solution to (7') is

$$\begin{aligned} x_{n+i} &= d'_i, & x_{n+p} &= d'_p, & i = 1, \dots, p-1, p+1, \dots, m, \\ x_i &= 0, & i &= 1, \dots, n, n+p. \end{aligned} \quad (8)$$

and

$$w = -d_p = -d'_p = -x_{n+p} + \sum_{j=1}^n a'_{pj} x_j. \quad (9)$$

The basic variables are

$$x_{n-p+1}, \dots, x_{n-p+2}, x_p, x_{n+p+1}, \dots, x_{n+m}. \quad (10)$$

We have transformed the auxiliary problem (5) to the following problem. Minimize w as given in (9) subject to (8) and $x_i \geq 0, i = 1, \dots, n+m$. For this problem we have determined a basic feasible solution given by (8) with basic variables given by (10). We can now apply phase II to the problem in this form. If we incorporate an anticycling sub-routine in our implementation of Phase II, then according to Lemma 5.1, we shall, in a finite number of steps, determine that the auxiliary problem either has no solution or determine an optimal solution to the auxiliary problem. If the auxiliary problem has no solution, then the original linear programming problem (1) has no feasible solution. For, as we saw, a vector (\bar{x}, \bar{x}_p) is feasible for the original problem if and only if $(\bar{x}, \bar{x}_p, 0)$ is optimal for the auxiliary problem.

The remaining alternative is that the auxiliary problem has a solution. We now explore the possibilities for this alternative. The variable x_p is basic initially. In applying phase II of the simplex method to the auxiliary problem,

we note that our rule for choosing the leaving variable is such that if at some iteration x_{ij} is a candidate for leaving the basis, then x_{ij} will be chosen to be the leaving variable. If at some iteration k , the variable x_{ij} leaves the basis, then at the iteration T_{k+1} , the variable x_{ij} will have the value zero. The solution $(\bar{b}_k, \bar{b}_{k+1}, \bar{c}_k)$ of the constrainting equations at T_{k+1} will also be a solution of (3). Hence if $\bar{c}_{ij} = 0$, then $w = 0$ and $(\bar{b}_k, \bar{b}_{k+1})$ is optimal for relation (1). Thus, if x_{ij} ever leaves the basis at an iteration T_k , at the next iteration T_{k+1} , we will have a solution to the auxiliary problem and hence a basic feasible solution for relation (2). We can then begin phase II with this basic feasible solution.

An optimal solution could also conceivably occur under the following circumstances:

- (i) x_{ij} basic, $w = 0$, and
- (ii) x_{ij} basic, $w \neq 0$.

If (ii) occurs, then the original problem has no feasible solution. We now show that (i) is impossible. If (i) did occur, then $x_{ij} = -w = 0$. Since the previous iteration did not end in an optimal solution, $x_{ij} < 0$ at the start of the previous iteration and became zero in the previous iteration. If this happened, then x_{ij} was a candidate for leaving the basis but did not do so. This contradicts the rule for determining whether x_{ij} leaves the basis.

The next lemma follows from the discussion in this section.

LEMMA 5.1. *If an unrestricting resource is included, then phase I of the simplex method either determines that no feasible solution exists or determines a basic feasible solution.*

Lemmas 5.1 and 5.2 imply the following.

THEOREM 5.1. *The simplex method incorporating an unrestricting resource applied to a linear programming problem determines one of the following in a finite number of iterations:*

- (i) No feasible solution exists.
- (ii) The problem does not have an optimal solution.
- (iii) An optimal solution.

Exercise 5.1. Complete the solution of the illustrative example.

Exercise 5.2. Let the linear programming problem be given originally in canonical form

$$\begin{aligned} & \text{Maximize } (\mathbf{b}, \mathbf{x}) \\ & \text{subject to } A\mathbf{x} = \mathbf{d}, \quad \mathbf{x} \geq \mathbf{0}. \end{aligned}$$

Show that an appropriate auxiliary problem for phase I is

$$\begin{aligned} \text{Maximize } w &= -\sum_{i=1}^m x_i \\ \text{subject to } (A, I) \begin{pmatrix} x \\ y \end{pmatrix} &= b, \quad x \geq 0, \quad y \geq 0. \end{aligned}$$

Show how one obtains an initial basic feasible solution for the auxiliary problem.

Exercise 4.3 Solve the following linear programming problems using the simplex method:

(a) Minimize $x_1 + 3x_2$ subject to

$$\begin{aligned} x_1 + 2x_2 + 3x_3 &\leq 4, \\ x_1 + 3x_2 + x_3 &\leq 5, \\ x_i &\geq 0, \quad i = 1, 2, 3. \end{aligned}$$

(b) Minimize $3x_1 - x_2$ subject to

$$\begin{aligned} x_1 + 2x_2 &\geq 2, \\ 3x_1 + x_2 &\leq 3, \\ x_j &\leq 4, \quad x_j \geq 0, \quad j = 1, 2, 3, 4. \end{aligned}$$

(c) Minimize $-2x_1 - x_2 + x_3 + x_4$ subject to

$$\begin{aligned} x_1 - x_2 + 2x_3 - x_4 &= 2, \\ 2x_1 + x_2 - 3x_3 + x_4 &= 8, \\ x_1 + x_2 + x_3 + x_4 &= 7. \end{aligned}$$

7. REVISED SIMPLEX METHOD

The revised simplex method does not differ from the simplex method in principle, that is, in going from one basic solution to another in such a way that the objective function is not decreased but differs in the manner in which some of the computations are carried out. In describing the revised simplex method, we assume, as we did in our description of the simplex method in Section 4, that the problem has been formulated as in (4) of Section 3. The reader will recall from Section 4 that at each iteration a basis B is determined

and the calculation of $B^{-1}N$ is used to determine the entering and leaving variable. The revised simplex method differs from the version of the simplex method given in Section 4 in that a system (1) in Section 4 with current B is not reduced to (11) in Section 4 and the current basis not given by (22) in Section 4. In the revised simplex method the entering variable x_k is chosen as before. If A_k is the column of A corresponding to x_k , then the updated basis B is obtained from B by replacing the column A_{j_0} of the leaving variable x_{j_0} by A_k . This procedure enables us to calculate $(B')^{-1}$ from B^{-1} very easily and expedite other calculations, as we shall show.

Before we present the steps in an iteration of the revised simplex method, we take up the calculation of $(B')^{-1}$, the updated inverse. At the current iteration l let the basis B be given by

$$B = (A_{j_1} \dots A_{j_m}), \quad (11)$$

where the A_{j_i} are columns of B . We then write

$$N = (A_{k_1} \dots A_{k_p}).$$

Let x_k be the entering variable at l , determined as in Section 4, and let A_k denote the column of N corresponding to x_k . Let the leaving variable x_{j_0} be determined as in Section 4. Let B' denote the matrix obtained from B by replacing the column A_{j_0} by A_k and let N' denote the matrix obtained from N by replacing column A_k by A_k . Thus

$$\begin{aligned} B' &= (A_{j_1} \dots A_{j_{p-1}} A_k A_{j_{p+1}} \dots A_{j_m}), \\ N' &= (A_{k_1} \dots A_{k_{p-1}} A_k A_{k_{p+1}} \dots A_{k_p}). \end{aligned} \quad (12)$$

Hence

$$\begin{aligned} B^{-1}B' &= B^{-1}(A_{j_1} \dots A_{j_{p-1}} A_k A_{j_{p+1}} \dots A_{j_m}) \\ &= \begin{pmatrix} 1 & 0 & \dots & 0 & v_{1k} & 0 & \dots & 0 \\ 0 & 1 & \dots & 0 & v_{2k} & 0 & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & 0 & v_{pk} & 0 & \dots & 1 \end{pmatrix} \\ &= (v_1 \dots v_{p-1} \quad v_k \quad v_{p+1} \dots v_m), \end{aligned} \quad (13)$$

where $v_j = B^{-1}A_{j_0}$. Note that v_k is the k th column of $B^{-1}N$ associated with x_k and this is as in Step 5 of an iteration in the simplex method. From the way p is determined we have that $v_{pk} > 0$, and hence the matrix on the right in (13) is nonsingular. If we multiply both sides of (13) by B on the left, we get that

$$B' = B(v_1 \dots v_{p-1} \quad v_k \quad v_{p+1} \dots v_m).$$

Since B is the product of nonsingular matrices, it is nonsingular and so is a basis. We also get

$$(B^{-1})^{-1} = (B_1 \cdots B_{i-1} (A_{ij} B_{i+1}) \cdots A_m)^{-1} B^{-1}. \quad (8)$$

It is readily verified that multiplying the matrix on the right-hand side of the following by the matrix on the right-hand side of (8) gives the $m \times m$ identity matrix. Hence

$$\begin{pmatrix} 1 & 0 & \cdots & 0 & v_{1j} & 0 & \cdots & 0 \\ 0 & 1 & \cdots & 0 & v_{2j} & 0 & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & 0 & v_{mj} & 0 & \cdots & 0 \\ & & & & \text{---} & & & \end{pmatrix}^{-1} = \begin{pmatrix} 1 & 0 & \cdots & 0 & v_{1j} & 0 & \cdots & 0 \\ 0 & 1 & \cdots & 0 & v_{2j} & 0 & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & 0 & v_{mj} & 0 & \cdots & 1 \\ & & & & \text{---} & & & \end{pmatrix}^{-1} \quad (9)$$

where

$$v_{ij} = -\frac{a_{ij}}{v_{jj}}, \quad i = 1, \dots, m, \quad i \neq j, \quad v_{jj} = \frac{1}{v_{jj}}.$$

The matrix on the right in (9) is called an *eta* matrix and will be denoted by E_j . Thus we may write (9) as

$$(B^{-1})^{-1} = E_j B^{-1}, \quad (10)$$

which updates B^{-1} for the next iteration.

We summarize this discussion in the following lemma:

Lemma 7.1. Let B be a basis matrix as in (1) and let B^{-1} be a matrix as in (2). Then B^{-1} is basic and (9) holds, where E_j is the matrix on the right in (9).

Remark 7.1. In the various calculations in the simplex method the essential matrix is B^{-1} , not B . The matrix B is needed, in principle, only to calculate B^{-1} . Equation (9) enables us to calculate $(B^{-1})^{-1}$ from B^{-1} without knowing B or B^{-1} . It is this fact that underlies the revised simplex method. Also note that an *eta* matrix differs from an identity matrix in one column.

We now present a generic revised simplex method. Most computer codes for solving linear programming problems use the revised simplex method. Codes used in practice incorporate various modifications to improve accuracy, reduce the number of iterations needed, decrease the running time, and so on. For details the reader is referred to Chvatal [1983], Dantzig [1983], Gass [1983], and Marjath [1981].

As in our presentation of the simplex method we assume that the problem is in canonical form: Maximize z subject to

$$Ax = b, \quad \langle b, e \rangle = z = 0, \quad x \geq 0. \quad (7)$$

Let the initial stage or iteration of phase II be called the *zeroth iteration*. At the k th iteration, $k = 0, 1, 2, \dots$ let $B(k)$ denote the basis, let $N(k)$ denote the submatrix of A complementary to $B(k)$, let $x_{B(k)}$ denote the vector of basic variables, let $x_{N(k)}$ denote the vector of nonbasic variables, let $q(k) = (q_{B(k)}, q_{N(k)})$ be the corresponding basic feasible solution, and let $z(k)$ be the current value of the objective function. Equations (4)–(5) from Section 4 yield, for $k = 0, 1, 2, \dots$

$$\begin{aligned} x_{B(k)} &= \bar{x}_{B(k)} - [B(k)]^{-1} N(k) x_{N(k)}, \\ z(k) &= \langle \bar{q}_{B(k)}, x_{B(k)} \rangle, \\ z &= z(k) - \langle \bar{q}_{B(k)} - \bar{q}_{N(k)} [B(k)]^{-1} N(k), x_{N(k)} \rangle. \end{aligned} \quad (8)$$

From Lemma 7.1 we get that

$$[B(k+1)]^{-1} = E_{k+1} [B(k)]^{-1}, \quad k = 0, 1, 2, \dots$$

where E_{k+1} is an approximate *eta* matrix. Proceeding inductively, we get

$$[B(k+1)]^{-1} = E_{k+1} E_k \dots E_1 [B(0)]^{-1} \dots b. \quad (9)$$

A sequential file in the computer storing the matrices $E_1, \dots, E_k, \dots, E_l$ is called an *eta file*. The calculation as indicated in (9) (from right to left) is called a *forward transformation*. The matrix $[B(k)]^{-1}$ is calculated using elementary matrices as in (11) from Section 3.

We now present the steps in the k th iteration, $0 \leq k$, of our genetic revised simplex algorithm. In essence, the steps mirror the steps of the simplex method as presented in Section 4. Some of the calculations are different and Step 2 of the simplex method is deleted. For typographical convenience we will replace the designations $B(k)$, $N(k)$, $x_{B(k)}$, $x_{N(k)}$, \dots of the current quantities by B , N , x_B , x_N, \dots . The updated quantities $B(k+1)$, $N(k+1)$, $x_{B(k+1)}$, $x_{N(k+1)}, \dots$ will be designated as B' , N' , $x_{B'}$, $x_{N'}$, \dots . As in (7) from Section 4 let

$$\bar{x}_k = \bar{x}_k - \bar{q}_k B^{-1} N = (\bar{x}_{k+1}, \dots, \bar{x}_l), \quad (10)$$

where $\bar{x}_{k+1}, \dots, \bar{x}_l$ are the components of \bar{x} that constitute $x_{N'}$. We may then write the last equation in (8) as

$$z = \zeta + \langle \bar{q}_k, x_k \rangle = \zeta + \sum_{j=k+1}^l \bar{q}_j x_{j'}$$

Step 1. Calculate \bar{d}_j . This is done in two steps.

(i) Calculate

$$w = \bar{b}_j^0 E_1 E_2 \dots E_k [B^{-1}]^{-1}$$

from left to right, that is, as

$$\{ \dots (\bar{b}_j^0 E_1 E_2 \dots E_k) \} \dots E_k [B^{-1}]^{-1}.$$

This backward computation involves repeated multiplications of a vector by a matrix, whereas the forward calculation (from right to left) involves repeated multiplications of matrices. Clearly, the backward transformation involves fewer operations.

(ii) Calculate

$$\bar{d}_j = \bar{c}_j - (c_j, \bar{A}_j), \quad j = m + 1, \dots, n,$$

where the \bar{c}_j are the indices of the cost coefficients at the k -th stage and the \bar{A}_j are the corresponding columns of N .

Step 2. If all the \bar{d}_j are less than or equal to zero, the current basic feasible solution is optimal and the algorithm terminates. If the algorithm terminates, then we use (C) to calculate ζ , the optimal value.

If the algorithm does not terminate at Step 2, we go to Step 3.

Step 3. Choose an entering variable.

Choose an index j_0 such that $\bar{d}_{j_0} > 0$. This may be done according to the smallest index rule, the ratio rule, or any other rule, as long as $\bar{d}_{j_0} > 0$. We denote the index j_0 by q and take x_q to be the entering variable. The column \bar{A}_q will be the entering column for the basis at the $(k + 1)$ st iteration.

Step 4. Either determine that the problem is unbounded or determine the leaving variable.

(i) Calculate

$$v_j = E_1 E_2 \dots E_k [B^{-1}]^{-1} \bar{A}_q$$

from right to left. Note that v_j is as in Section 4.

(ii) If all components of v_j are nonpositive, then the problem is unbounded and no solution exists. The justification for this conclusion is the same as the one given in Step 3(i) of Section 4.

- (iii) If $\theta = (v_{k_1}, \dots, v_{k_r})$ has at least one positive component, then as in (21) of Section 4 let ρ denote the smallest index r of which

$$\rho = \min \left\{ \frac{v_{k_r}}{v_{k_r}}; v_{k_r} > 0 \right\}$$

is attained. Choose s_{k_ρ} to be the leaving variable.

Step 5. Calculate K_{k+1} as given by the matrix on the right in (2).

Step 6. Update the data for the $(k + 1)$ st iteration.

- (i) Update \mathbf{x}_k to \mathbf{x}_{k+1} by replacing x_{k_ρ} by x_{k_ρ} .
- (ii) Update \mathbf{x}_k to \mathbf{x}_{k+1} by replacing x_{k_ρ} by x_{k_ρ} .
- (iii) Update \mathbf{B} to \mathbf{B}' by replacing \mathbf{B}_{k_ρ} by \mathbf{B}_{k_ρ} .
- (iv) Update \mathbf{N} to \mathbf{N}' by replacing \mathbf{B}_{k_ρ} by \mathbf{B}_{k_ρ} .
- (v) Update (\bar{b}_k, \bar{c}_k) to $(\bar{b}_{k+1}, \bar{c}_{k+1})$ using (22) from Section 4.
- (vi) Update \mathbf{B}^{-1} to $(\mathbf{B}')^{-1}$ using (8).
- (vii) Update \mathbf{h}_k and \mathbf{h}_k to \mathbf{h}_{k+1} and \mathbf{h}_{k+1} .

Step 7. Begin iteration $k + 1$ using the updated quantities.

Remark 7.2. If the initial basis consists of the slack variables, then $\mathbf{B}(0) = \mathbf{I}$ and (9) becomes

$$[\mathbf{B}(k) + \mathbf{I}]^{-1} = \mathbf{K}_{k+1} \mathbf{B}_1 \dots \mathbf{B}_k \mathbf{A}$$

Remark 7.3. The via file can become very long. In some computer codes, once the via file reaches a predetermined length, say $k + 1$, then $\mathbf{B}(k + 1)$ is set as the initial basis $\mathbf{B}(0)$, the matrix $[\mathbf{B}(k + 1)]^{-1}$ is set as $[\mathbf{B}(0)]^{-1}$, and a new via file is constructed.

Exercise 7A. Solve the linear programming problem in Example 3.1 using the revised simplex method. Start with phase I.

Exercise 7B. Solve the linear programming problem in Exercises 5.5(a), (b), and (c) using the revised simplex method.

BIBLIOGRAPHY

- Burke, K. G. and D. B. Stewart, *The Elements of Real Analysis*, 3rd ed., John Wiley & Sons, New York, 1999.
- Burton, M. E., H. D. Garfink, and C. M. Swamy, *Nonlinear Programming, Theory and Algorithms*, 2nd ed., John Wiley & Sons, New York, 1993.
- Burke, K. G. L., "Cycling in the Dual Simplex Algorithm," *Naval Research Logistics Quarterly* 4 (1955), 200-201.
- Burke, K. G., "New Finite Proving Rules for the Simplex Method," *Mathematics of Operations Research* 2 (1977), 103-107.
- Charnif, Y., *Linear Programming*, W. H. Freeman, New York, 1964.
- Dantzig, G. B., *Linear Programming and Extensions*, Education University Press, Englewood, N.J., 1963.
- Dantzig, M., *Games of Strategy: Theory and Applications*, Prentice-Hall, Englewood Cliffs, N.J., 1961.
- Farkas, I., "Über die Theorie der linearen Ungleichungen," *Journal für die reine und angewandte Mathematik* 124 (1903), 1-27.
- Fiacco, A. V. and G. P. McCormick, *Nonlinear Programming—Sequential Unconstrained Minimization Techniques*, John Wiley & Sons, New York, 1968.
- Fleming, W., *Functions of Several Variables*, 2nd ed., Springer, New York, 1971.
- Gale, D., *The Theory of Linear Economic Models*, McGraw-Hill, New York, 1960.
- Gomory, R. L., *Linear Programming, Methods and Applications*, 3rd ed., McGraw-Hill, New York, 1969.
- Golub, G. H. and C. F. Van Loan, *Matrix Computations*, Johns Hopkins University Press, Baltimore, 1980.
- Gordan, P., "Über die Auflösungen linearer Gleichungen mit reellen coefficients," *Mathematische Annalen* 6 (1873), 13-28.
- Hahn, S.-P. and W. L. Womersley, "Exact Penalty Functions in Nonlinear Programming," *Mathematical Programming* 17 (1979), 231-265.
- John, F., *Extremum Problems with Inequalities as Subsidiary Conditions, Studies and Essays: General Anniversary Volume* (R. G. Fetscherle, G. B. Neugebauer, and I. L. Steiner, Eds.), Interscience, New York, 1948, pp. 87-109.
- Karush, W., "Minima of Functions of Several Variables with Inequalities as Side Conditions," *Master's Dissertation*, University of Chicago, Chicago, IL, Dec. 1959.
- Kuhn, R. W. and A. W. Tucker, *Nonlinear Programming*, Proceedings of the Second Berkeley Symposium on Mathematical Statistics and Probability (J. Neyman, Ed.), University of California Press, Berkeley, CA, 1961, pp. 481-482.

- Mangasarian O. L. *Nonlinear Programming*, McGraw-Hill, New York, 1969.
- Moore, G. J., "The Lagrange Multiplier Rule," *American Mathematical Monthly* **68** (1961), 822-824.
- Morikita, T. K., "Übertrag zur Theorie Linearer Ungleichungen," *Inangewandte Mathematik, Band, Inventionen*, 1936.
- Marquardt, D. A., *Algorithms for Linear Programming: Computation and Practice*, McGraw-Hill, New York, 1969.
- Penin, L. L., "An Indirect (Necessary) Proof for the Problem of Lagrange with Differential Inequalities on Side Conditions," *Transactions of the American Mathematical Society* **74** (1953), 171-176.
- Rockafellar, R. T., *Convex Analysis*, Princeton University Press, Princeton, N.J., 1970.
- Rudin, W., *Principles of Mathematical Analysis*, 3rd ed., McGraw-Hill, New York, 1953.

INDEX

- Active set, 74
- Active transformations, 74–114
- Algebra of sets, 4–7
- Alternative theorems
 - convex functions, 113–117
 - convex sets, linear inequalities, 11–18
 - Farkas's lemma, 43–45, 67–69
 - Gordan's theorem, 48
 - Gerber's theorem, 45–47
 - Motzkin's theorem, 49–47
- Anticipating cycles, simplex method
 - phase I procedure, 214
 - termination and cycling, 210, 214
- Barzilai transformations, revised simplex method, 26
- Barzilai's coefficients, 14
- Barzilai's rule, simplex method, 24
- Basic variables
 - phase II procedure, 211–213
 - simplex method, 20
- Bland's rule, simplex method, termination and cycling, 248–250
- Bland–Wolfe's property, 13
- Boundary points
 - convex function optimization, 147
 - supporting hyperplanes, 8–17
- Convex linear programming problem,
 - 24–26
 - optimization, 149
 - simplex method, 23, 25
 - extreme points, feasible sets, 125–129
 - phase I, 211, 213
 - revised simplex method, 258–260
- Combinatorial functions, 4–6
- Conical products, 6–8
- Condy–Schwarz inequality, 3–4
- Closed ball space, 21–23
- Closed line segments, 31
- Closed sets, 4–8
- Closed point characterization, 23–24
- Convex functions, 20
- Convex Minkowski game, 120, 127
- Computation, 14–15
 - convex sets, separation theorem, 14–15
 - supporting hyperplanes, extreme points, 58–61
- Complementary slackness condition, linear programming, 148–149
- Completeness property, real numbers, 12–13
- Conver functions
 - defined, 49
 - game theory, 117
- Convexity
 - convex programming, (extremal theory), 182, 208
 - convex programming, (problem 1) (CPL), 180–181
 - strong is (CPL), 184–187
- Convex qualifications
 - convex programming
 - feasibility conditions, 183–188
 - quadratic programming, 211
 - differentiable convex programming, 188, 193
 - Kuhn–Tucker convex qualifications, 182–183
 - regularity conditions, 187–188
- Continuity, 9–11
- Convex cones, 44, 45
- Convex functions
 - alternative theorems, 113–117
 - definition and properties, 67–114
 - differentiable functions, 100–101
 - differentiable convex programming, 189–193

- Convex functions (Continued)
 - convex theory applications, 117–121
 - Duval's inequality, 91–94
 - optimization problems, 117–121
 - partial derivatives, 95, 106–109
 - subgradients, 101–104
- Convex hulls, 41–43
- Convex points, 41–43
- Convex programming, duality
 - geometric interpretation, 207–210
 - Lagrangian duality, 201–205
 - linear programming, 210–221
 - quadratic programming, 212
- Convex programming problem I (CPI), problem statement, 179–180
- Convex programming problem II (CPII), problem dual to (CPI), 200–207
- problem statement, 180–184
- Convex sets
 - affine geometry, 65–68
 - in parametric coordinates, 74
 - hyperplanes, 71–72
 - half-spaces, 74
 - independence, 72–75
 - linear manifold, 65–71
 - linear transformations, 78
 - non-void interiors, 76–77
 - parallel subspaces, 78
 - simplex dimension, 74
 - extreme points, 76–81
 - hyperplane support, 80–81
 - linear dependence, various (L), 41–48
 - Farkas's lemma, 43–45, 47
 - Gale's lemma, 48
 - Carathéodory's lemma, 45–47
 - Minkowski's lemma, 46, 47
 - minimally relative interiors, 81–82
 - properties of, 31–41
 - Carathéodory's theorem, 41–42
 - convex sets, 44
 - convex hull, 41–43
 - half-spaces, 71–74
 - planes, convex combinations, 41–43
 - relative combinations, 87
 - relative interior points, 81–82
 - separation theorems, 45–54
 - hyperplane support, 43–44, 47
 - proper separation, 46, 51
 - strict separation, 46
 - strong separation, 47–48, 44
- Convex form, simplex method, 284–285
- Convex value simplex method, 284–285
- Cycling, simplex method, 285–286
- Dependent basis relation, simplex method, 281, 283–287
- Derivatives
 - convex functions, 94–96, 97
 - Hessian matrix (H),
 - diagonalization, 111–113
 - linear transformations, 21–23
 - second derivative test, differentiable
 - unconstrained problems, 139–144
- Differentiability
 - convex functions, 95, 104–111
 - Hessian matrix (H), linear transformations, 24
- Differentiable nonlinear programming
 - Karush-Kuhn-Tucker conditions, 147–151
 - second-order conditions, 151–179
- Differentiable unconstrained problem, (CP), 134
 - linear supports, 134–135
 - local minimum, 131
 - refined non-optimal approximations, 133, 134
 - positive semidefinite functions, 131–132
 - second derivative test, 134
 - strict minimum, 131–134
- Differentiable, Hessian matrix (H), 111–113
- Directional derivative, convex functions, 91–93
- Direct convex sets, hyperplane separation, 43
- Duality theorems
 - convex programming
 - dual quadratic programming (DQP), 214–221
 - geometric interpretation, 207–210
 - Lagrangian duality, 201–207
 - linear programming, 211–221
 - weak, convex programming, 201–207
 - quadratic programming, 211
- Elementary operations, simplex method, 281, 283
- Empty set, 4–5
- Entering variable
 - revised simplex method, 284–285
 - simplex method, 289–291
- Epigraphs, convex functions, 46–49
- Equality constraints
 - convex programming problem (CPI), 180
 - differentiable nonlinear programming
 - Karush-Kuhn-Tucker conditions, 147–151, 148–171
 - second-order conditions, 151–169, 171–177
 - Equilibrium points, 178

- Equivariant norms, Euclidean norm (EF), 34–35
- Euclid's revised simplex method, 298–299
- Euclidean norm, revised simplex method, 297–300
- Euclidean norm, 2–3
- Euclidean norm (EF)
 - defined, 1
 - metric topology, 1–3
 - updates, 1–4
- Extreme points
 - convex sets, 56–61, 65
 - simplex method, feasible sets, 223–226
- Farkas's lemma
 - differentiable nonlinear programming, 145–148
 - linear inequalities, 61–66, 67–68
 - linear programming, 142
 - linear programming duality, 220
- Feasible sets
 - optimization problems, 12
 - simplex method
 - optimal problem, 211–220
 - subproblems, 221–226
- Finite-iterative property, 15
- First-order conditions, differentiable nonlinear programming, 145–154
- Forward transformation, revised simplex method, 298, 299
- Frechet differentiable, 24
- Free variables, simplex method, 290
- Fritz-Kuhn theorem, differentiable nonlinear programming, 149–163
 - second-order conditions, 171–174
- Functional constraint theorem, 18–20
- Gale's theorem, linear inequalities, 68
- Game theory, 127–127
 - optimal mixed strategies, 126–127
 - optimal pure strategies, 122–126
- Group functions, 66
- Group structures, 200–204
- Gauss–Jordan elimination, 331–334
- Gomory's inequalities, convex programming, duality, 205–208
- Gordan's theorem
 - generalization to convex functions, 210
 - linear inequalities, 61–67
 - quadratic programming, 211–213
- Halmos's five lectures, 6
- Half space, 14–16
 - intersection theorem, 51–54
- Rankin matrix
 - differentiable convex functions, 110–111
 - differentiable nonconvex problems, 111–120
- Rayleigh's theorem
 - convex functions, subgradients, 102–106
 - convex programming, duality, geometric interpretation, 208–210
 - convex set properties, 25–26
 - Farkas's lemma, 63
 - definition, 13–14
 - supporting hyperplanes, 56–59, 61, 65
- Revised simplex method, 291–294
- Revised, 12–14
- linear problem, 2
- Interior points
 - continuity of convex functions, 96–98
 - convex sets
 - n -dimensional, 76
 - relative interior points, 80–81
- Jordan matrix, 26, 101
- Jordan's inequality, convex functions, 96–98
- Karush–Kuhn–Tucker theorem, differentiable nonlinear programming, 150–164
 - constraint qualifications, 154–160
 - second-order conditions, 169–172
 - sublinear conditions, 159–160
- Krein–Milman theorem, 68
- Kuhn–Tucker theorem
 - definition, 164
 - examples, 167–168
 - in multiplicity, 164–165
 - in subgradients of value function, 164–166
- Lagrange multipliers
 - convex programming
 - in Kuhn–Tucker theorem, 164–168
 - necessary conditions, 161–164
 - optimality conditions, 167–168
 - in subgradients of value function, 164–166
 - sufficient conditions, 166–167
 - differentiable nonlinear programming, first-order conditions, 149–153
- Lagrangian duality, convex programming, 200–207
 - quadratic programming, 116, 201
- Lifting method
 - revised simplex method, 296–298
 - simplex method, 296–297
- Left-hand limit, real numbers, 14

- Local minima**
 convexity, 5–11
 nondiscrepancy, 6
 limits, 5–11
- Linear functional**
 Euclidean n -space (\mathbb{R}^n), 11–14
 hyperplane definition, 13–14
- Linear inequalities**
 convex sets, existence of the alternative,
 61–68
 Farkas's lemma, 43–45, 67
 Gale's theorem, 69
 Gordan's theorem, 65–67
 Motzkin's theorem, 66–67
- Linear inequalities, one- T_1**
 alternative, 58–71
- Linear programming, linear**
 nonlinear
 canonical linear programming problem, 144
 complementary slackness condition,
 146–149
 duality in, 141–151
 iteration, 135–140
 necessary and sufficient conditions,
 146–149
 simplex method
 auxiliary problem, 111–115
 examples, 117–123
 extreme points, feasible set, 126–128
 phase I procedure, 124–125
 phase II procedure, 124–128
 preliminary, 120–124
 revised method, 128–130
 termination and cycling, 141–143
 started in convex linear forms, 148–149
- Linear regression, 154, 175**
- Linear transformation**
 as derivative, 23–26
 Euclidean n -space (\mathbb{R}^n), 21
- Line segments, 8**
- Linear-convexity functions, 98**
- Local extrema**
 convex functions, 117
 differentiable nonlinear programming,
 second-order sufficient theorem,
 175–178
 differentiable unconstrained problems,
 178–179
- Logarithmic convexity, 114**
- Lower bounds, real numbers, 11, 13**
- Matrix game, 123–127**
- Maximization of convex functions, 117–119**
- Mean-square approximation by orthogonal functions, 105**
- Metric topology. Euclidean n -space (\mathbb{R}^n), 5–8**
- Minimization**
 convex function optimization, 117–119
 convex programming,
 necessary and sufficient conditions,
 165, 168
 quadratic programming, 213–215
 differentiable nonlinear programming,
 145–161
 fundamental sufficient theorem, 161–163
 linear programming, 146–147
- Minkowski distance function, 101**
- Mixed strategies, 126, 127**
- Morgan's law, 4–5**
- Motzkin's theorem, linear inequalities, 66–67**
- Multiplicity, as Kuhn–Tucker vectors,
 Lagrange multipliers**
- Necessary conditions**
 convex programming, 161–168
 differentiable nonlinear programming,
 first-order conditions, 161–163
 second-order conditions, 163–178
 linear programming duality, 141–143
 quadratic programming, 178–181
- Negative half space, convex set properties,
 11, 34**
- Nonbasic variables, simplex method, 128**
- Nonlinear programming, differentiable
 problems**
 first-order conditions, 161–163
 second-order conditions, 163–178
- Nonlinear supporting hyperplane, 56–57, 60**
- Not set, 1–1**
- Open sets, 10–16**
- Open line segments, 8**
- Open sets, metric topology, 4–8**
- Optimal solution, simplex method, 211,
 230–240**
- Optimal strategies, see Game theory**
- Orthogonality, 4**
- Parallel hyperplanes, 14**
- Parallel subspaces, 70**
- Parseval matrix, 105–107**
- "Pseudo-linear" point, differentiable
 nonlinear programming, 171, 173**
- Perturbation theory, 189–190**
- First problem, simplex method, 240–248**
- secondary problems, 252–254**

- Positive definite matrix, 111–112, 115–116
- Positive half-space, 13–14
- Primal versus programming problem (PVP), 180–181
 - Lagrangian duality, 200–207
 - Primal separation, correct sets, 46
- Quadratic programming, 208–210
 - first quadratic program (EQP), 209–211
- Real numbers, fundamental properties, 11–14
- Relative fundamental, correct sets, support and separation theorems, 81–86
- Relative interior, correct sets, support and separation theorems, 86–88
- Revised simplex method, 231–260
- Right-hand limit, 13–14
- Saddle point
 - convex programming, Lagrangian duality, 206–207
 - locally and sufficient conditions, 182–188
 - game theory applications, 125–127
- Second derivative test, differentiable unconstrained problems, 100–111
- Second-order conditions, differentiable nonlinear programming, 113–178
- Second-order test vector, differential nonlinear programming, 167–173
- Sensitivity vector, convex programming, perturbation theory, 193
- Separation theorems, 42–54
 - affiliated theorems, 44–48
 - best possible in SP, 83–85
 - disjoint convex sets, 33–36, 82
 - point and convex sets, 81–83
 - primal separation, 46–53–54
 - total separation, 46–47
 - strong separation, 47–49, 54
- Support cones, Hachian's support (S^c)
 - limits, 10–11
- Simplex method
 - examples, 221–226
 - convex points, feasible set, 105–110
 - phase I procedure, 231–240
 - phase II procedure, 244–245
 - prerequisites, 240–244
 - revised method, 253–260
 - termination and cycling, 245–250
- Slack variables
 - linear programming, 144
 - simplex method, phase I, 231–233
- Smallest index, into, simplex method, termination and cycling, 249–250
- Start level minimizers
 - defined, 128
 - differentiable nonlinear programming, second-order sufficient conditions, 170–178
 - differentiable unconstrained problems, 128–136
- Strictly convex function
 - defined, 87–88
 - differential criteria, 108–111
- Strong separation, correct sets, 46–47
- Strong separation, correct sets, 47–48
- Subdifferential, convex functions, 102–116
- Subgradients, convex functions, 102–108
 - differentiable functions, 105–107
- Submanifolds, games, 129–127
- Subspaces, convexity, 5–11
- Sufficient conditions
 - convex programming, 181–187
 - KKT criteria of duality gap, 206–207
 - quadratic programming, 211
 - differentiable nonlinear programming, 109–148
 - second-order conditions, 170–178
 - linear programming theory, 219–220
- Support function, 80
- Support function properties, 84–85, 89
- Supports, 127–131
- Tallies method, 126
- Tangential conditions, differentiable nonlinear programming, 149–178
 - second-order sufficiency theorems, 170–178
- Termination, simplex method, 245–250
- Theorems of the alternative, linear inequalities, 61–69
 - Farkas's lemma, 61–65, 67–68
 - Gale's theorem, 66
 - Carathéodory's theorem, 64–67
 - Minkowski's theorem, 66–67
- Three-start property, convex functions, 61–64
- Through inequality, 1
- Two-person zero-sum games, 122–127
- Unbounded set
 - revised simplex method, 260
 - simplex method, 225, 249–249
- Unconstrained problems
 - differentiable problems, 128–136

Unconstrained problems (continued)

linear regression, 184–185

orthogonal least-squares approximations,
155–156

reduced gradient test, 190–191

ridge minimizers, 186, 190

Euler's constant, 117

Upper bounds, real numbers, 11–14

Value function, perturbation theory, 189–200

Vector-valued convex functions, alternative
theorems, 105–117

Van Neumann minimax theorem, 117–121

PURE AND APPLIED MATHEMATICS

A Wiley-Interscience Series of Texts, Monographs, and Tracts

Founded by **RICHARD COURANT**

Editors: **MYRON B. ALLEN**, **DAVID A. COX**, **PETER LAX**

Editors Emeriti: **PETER HILTON**, **HARRY HOCHSTADT**, **JOHN TOLAND**

A complete list of the titles in this series appears at the end of this volume.

PURE AND APPLIED MATHEMATICS

A Wiley-Interscience Series of Texts, Monographs, and Tracts

Founded by RICHARD COURANT

Editors: MYRON H. ALLEN, ELI DAVIDSON, COO, PETER LAX

Editors Emeriti: PETER HILTON, HARRY HOCHSTADT, JOHN TOLAND

EDMUND HEWLETT, and STRICKER—Abstract and Concrete Categories
EDWARDS and DEBROSS—Logic of Mathematics
EDWARDS and DEBROSS—A Probabilistic Error Estimation in Finite Element Analysis
EKPIS and GOLDBERG—Continuum Differential Geometry and Its Generalizations
ELLEN and HALLON—Statistical Analysis for Applied Science

*EYTES—Geometric Algebra

FEEN—Applied Functional Analysis, Second Edition

FEYFO and GOLDFEVY—Linear Operators in Spaces with Indefinite Metric

FIBO—The Fourier-analytic Theory of Quadratic Reciprocity

FISMAN, FISMAN, and FIBO—Nonnegative Matrices in Dynamic Systems

FIRAPYTY—Convexity and Continuation in R^n

FOURDENTON—Methods of Solving Stiff Systems of Ordinary Differential Equations

FURE—Lagrange Multipliers and Interpolation: An Introduction

*GABER—Finite Groups of Lie Type

CASTELLO, CURIA, JUBATE, and POUNELA—Orthogonal Rats and Polar Methods in Linear Algebra: Applications to Matrix Calculations, Systems of Equations, Inequalities, and Linear Programming

CASTELLO, CORIO, FREESTAL, GABCLA, and ALGUAACE—Building and Solving Mathematical Programming Models in Engineering and Science

CHATELAIN—Eigenvalues of Matrices

CLARK—Mathematical Bioeconomics: The Optimal Management of Renewable Resources, Second Edition

COX—Primes of the Form $4x^2 + y^2$, Fermat, Class Field Theory, and Complex Multiplication

*COSTE and BEBOS—Representation Theory of Finite Groups and Associative Algebras

*COSTE and BEBOS—Methods of Representation Theory: With Applications to Finite Groups and Orders, Volume I

COSTE and BEBOS—Methods of Representation Theory: With Applications to Finite Groups and Orders, Volume II

DANTE RANJ—Vector Integration and Stochastic Integration in Banach Spaces

*DENFORD and SOWARTZ—Linear Operators

Part I—General Theory

Part 2—Spectral Theory, Self-Adjoint Operators in Hilbert Space

Part 3—Spectral Operators

FARNS and KIPALIS—Positive Linear Systems: Theory and Applications

FELDERS—Real Analysis: Modern Techniques and Their Applications

FELICHER and ERICL—Linear Spaces and Differentiation Theory

GRIBER—Tribonacci Theory and Quasiperiodic Differentials

GRINE and KLANZT—Function Theory of One Complex Variable

*Now available in a lower priced paperback edition in the Wiley Classic Library.
*Now available in paperback.

- KOEFITZ and HARRIS—Principles of Algebraic Geometry**
OSLETT—Algebra
OSOBY—Groups and Characters
QUIST and PIRON, KRÜSS and OLGER—Time-Dependent Problems and Difference Methods
RANGA and SCHWANE—Fourier Series, Transformations, and Boundary Value Problems, Second Edition
REINICK—Applied and Computational Complex Analysis
 Volume 1. *Power Series—Integration—Conformal Mapping—Location of Zeros*
 Volume 2. *Special Functions—Integral Transforms—Asymptotics—Continued Fractions*
 Volume 3. *Discrete Fourier Analysis, Cauchy Integrals, Construction of Conformal Maps, Uniform Functions*
SELTON and WU—A Course in Modern Algebra
SCHEIDT—Integral Equations
SCOTT—Two-Dimensional Geometries, Variational Problems
SHANN and CHEN—An Introduction to Metric Spaces and Fixed Point Theory
SOHAYBI and SCHULZ—Foundations of Differential Geometry, Volume I
SOHAYBI and SCHULZ—Foundations of Differential Geometry, Volume II
SOHET—Fibonacci and Lucas Numbers with Applications
LAG—Linear Algebra
LOUIS—An Introduction to Nonlinear Partial Differential Equations
MCCONNELL and ROBINSON—Noncommutative Noetherian Rings
MEISENER—Functional Analysis: An Introduction to Banach Space Theory
NATHAN—Perturbation Methods
NAYFER and BECK—Nonlinear Oscillations
PANJABI—The Hilbert Transform of Schwartz Distributions and Applications
PETCOV—Geometry of Reflecting Rays and Inverse Spectral Problems
PRESTER—Splines and Fractional Methods
RAG—Measure Theory and Integration
RABIER and SERA—Finite Sums Decompositions in Mathematical Analysis
RECHT—Matrix Systems and Quasiconformal Mappings
REYD—Chebyshev Polynomials: From Approximation Theory to Algebra and Number Theory, Second Edition
ROCKAFELLAR—Network Flows and Monotropic Optimization
ROTMAN—Introduction to Modern GP Theory
RUDIN—Functional Analysis on Groups
SENOK—The Averaged Method of Iterations: Applications in Numerical Methods and Approximation
SENOK and POPOV—The Averaged Method of Iterations
SHOGL—Topics in Complex Function Theory
 Volume 1—*Elliptic Functions and Uniformization Theory*
 Volume 2—*Automorphic Functions and Hilbert Integrals*
 Volume 3—*Abelian Functions and Modular Functions of Level 2*
SMITH and BROMBERG—Post-Modern Algebra
STARKER—Dirac's Functions and Boundary Value Problems, Second Edition
STURM—Differential Geometry
STURM—Nonlinear Vibrations in Mechanical and Electrical Systems
STURM—Wave Waves: The Mathematical Theory with Applications
WISSLING—An Introduction to Multigrid Methods
WEBBAM—Linear and Nonlinear Maps
WALTERS—Partial Differential Equations of Applied Mathematics, Second Edition

*Now available in paperback (paperback edition in the Wiley Classic Library).
 †Now available in paperback.