Mathematics for Economics and Finance

Jianfei Shen

School of Economics, The University of New South Wales

Sydney, Australia

[This version: March 2, 2012]

Jianfei Shen
School of Economics
The University of New South Wales
Sydney 2052
Australia
jianfei.shen@unsw.edu.au

Besides language and music, mathematics is one of the primary manifestations of the free creative power of the human mind.

— Hermann Weyl

Contents

1

MULTIVARIABLE CALCULUS

In this chapter we consider functions mapping $\mathbb{R}^m$ into $\mathbb{R}^n$, and we define what we mean by the derivative of such a function. It is important to be familiar with the idea that the derivative at a point a of a map between open sets of (normed) vector spaces is a *linear transformation* between the vector spaces (in this chapter the linear transformation is represented as a $n \times m$ matrix).

This chapter is based on Spivak (1965, Chapters 1 & 2) and Munkres (1991, Chapter 2)—one could do no better than to study theses two excellent books for multivariable calculus.

Notation

We use standard notation:

$\mathbb{N}$: The set of natural numbers: $\mathbb{N} = \{1, 2, 3, \ldots\}$.

$\mathbb{Z}$: The set of integers: $\mathbb{Z} = \{\ldots, -2, -1, 0, 1, 2, \ldots\}$.

$\mathbb{R}$: The set of real numbers.

$\mathbb{Q}$: The set of rational numbers: $\mathbb{Q} := \{x \in \mathbb{R} : x = p/q, p, q \in \mathbb{Z}, q \neq 0\}$.

We also define

$$\mathbb{R}_+ = \{x \in \mathbb{R} : x \geqslant 0\} \quad \text{and} \quad \mathbb{R}_{++} := \{x \in \mathbb{R} : x > 0\}.$$

1.1 Functions on Euclidean Space

Norm, Inner Product and Metric

Definition 1.1 (Euclidean n-space). *Euclidean n-space* $\mathbb{R}^n$ is defined as the set of all n-tuples $(x_1, \ldots, x_n)$ of real numbers x_i:

$$\mathbb{R}^n := \{(x_1, \ldots, x_n) : x_i \in \mathbb{R}, i = 1, \ldots, n\}.$$

An element of $\mathbb{R}^n$ is often called a point in $\mathbb{R}^n$, and $\mathbb{R}^1$, $\mathbb{R}^2$, $\mathbb{R}^3$ are often called the *line*, the *plane*, and *space*, respectively.

If $\boldsymbol{x}$ denotes an element of $\mathbb{R}^n$, then $\boldsymbol{x}$ is an n-tuple of numbers, the i^{th} one of which is denoted x_i; thus, we can write

$$\boldsymbol{x} = (x_1, \ldots, x_n).$$

A point in $\mathbb{R}^n$ is frequently also called a *vector in* $\mathbb{R}^n$, because $\mathbb{R}^n$, with

$$\boldsymbol{x} + \boldsymbol{y} = (x_1 + y_1, \ldots, x_n + y_n) \quad \boldsymbol{x}, \boldsymbol{y} \in \mathbb{R}^n$$

and

$$\alpha\boldsymbol{x} = (\alpha x_1, \ldots, \alpha x_n), \quad \alpha \in \mathbb{R} \text{ and } \boldsymbol{x} \in \mathbb{R}^n,$$

as operations, is a *vector space*.

We now introduce three structures on $\mathbb{R}^n$: the *Euclidean norm*, *inner product* and *metric*.

DEFINITION 1.2 (Norm). In $\mathbb{R}^n$, the length of a vector $\boldsymbol{x} \in \mathbb{R}^n$, usually called the *norm* $\|\boldsymbol{x}\|$ of $\boldsymbol{x}$, is defined by

$$\|\boldsymbol{x}\| = \sqrt{x_1^2 + \cdots x_n^2}.$$

REMARK 1.3. The norm $\|\cdot\|$ satisfies the following properties: for all $\boldsymbol{x}, \boldsymbol{y} \in \mathbb{R}^n$ and $\alpha \in \mathbb{R}$,

- $\|\boldsymbol{x}\| \geqslant 0$,
- $\|\boldsymbol{x}\| = 0$ iff[1] $\boldsymbol{x} = \boldsymbol{0}$,
- $\|\alpha\boldsymbol{x}\| = |\alpha| \cdot \|\boldsymbol{x}\|$,
- $\|\boldsymbol{x} + \boldsymbol{y}\| \leqslant \|\boldsymbol{x}\| + \|\boldsymbol{y}\|$ (Triangle inequality).

► EXERCISE 1.4. *Prove that* $\|\boldsymbol{x}\| - \|\boldsymbol{y}\| \leqslant \|\boldsymbol{x} - \boldsymbol{y}\|$ *for any two vectors* $\boldsymbol{x}, \boldsymbol{y} \in \mathbb{R}^n$ *(use the triangle inequality).*

DEFINITION 1.5 (Inner Product). Given $\boldsymbol{x}, \boldsymbol{y} \in \mathbb{R}^n$, the *inner product* of the vectors $\boldsymbol{x}$ and $\boldsymbol{y}$, denoted $\boldsymbol{x} \cdot \boldsymbol{y}$ or $\langle \boldsymbol{x}, \boldsymbol{y} \rangle$, is defined as

$$\boldsymbol{x} \cdot \boldsymbol{y} = \sum_{i=1}^{n} x_i y_i.$$

REMARK 1.6. The norm and the inner product are related through the following identity:

$$\|\boldsymbol{x}\| = \sqrt{\boldsymbol{x} \cdot \boldsymbol{x}}.$$

[1] "iff" is the abbreviation of "if and only if".

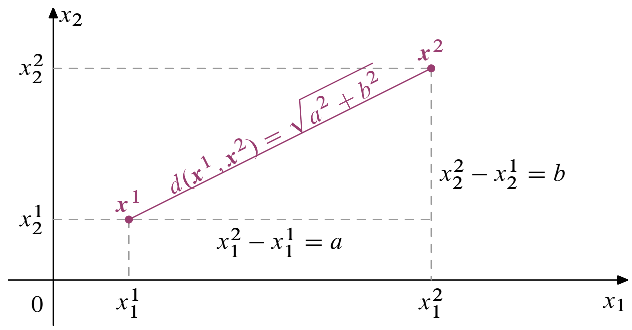

FIGURE 1.1. Distance in the plane.

THEOREM 1.7 (Cauchy-Schwartz Inequality). *For any $\boldsymbol{x}, \boldsymbol{y} \in \mathbb{R}^n$ we have*

$$|\boldsymbol{x} \cdot \boldsymbol{y}| \leqslant \|\boldsymbol{x}\| \, \|\boldsymbol{y}\|.$$

PROOF. We assume that $\boldsymbol{x} \neq \boldsymbol{0}$; for otherwise the proof is trivial. For every $a \in \mathbb{R}$, we have

$$0 \leqslant \|a\boldsymbol{x} + \boldsymbol{y}\|^2 = a^2\|\boldsymbol{x}\|^2 + 2a(\boldsymbol{x} \cdot \boldsymbol{y}) + \|\boldsymbol{y}\|^2.$$

In particular, let $a = -(\boldsymbol{x} \cdot \boldsymbol{y})/\|\boldsymbol{x}\|^2$. Then, from the above display, we get the desired result. □

► EXERCISE 1.8. *Prove the triangle inequality (use the Cauchy-Schwartz Inequality). Show it holds with equality iff one of the vector is a nonnegative scalar multiple of the other.*

DEFINITION 1.9 (Metric). The *distance* $d(\boldsymbol{x}, \boldsymbol{y})$ between two vectors $\boldsymbol{x}, \boldsymbol{y} \in \mathbb{R}^n$ is given by

$$d(\boldsymbol{x}, \boldsymbol{y}) = \sqrt{\sum_{i=1}^{n}(x_i - y_i)^2}.$$

The distance function d is called a *metric*.

EXAMPLE 1.10. In $\mathbb{R}^2$, choose two points $\boldsymbol{x}^1 = (x_1^1, x_2^1)$ and $\boldsymbol{x}^2 = (x_1^2, x_2^2)$ with $x_1^2 - x_1^1 = a$ and $x_2^2 - x_2^1 = b$. Then Pythagoras tells us that (Figure 1.1)

$$d(\boldsymbol{x}^1, \boldsymbol{x}^2) = \sqrt{a^2 + b^2} = \sqrt{\left(x_1^2 - x_1^1\right)^2 + \left(x_2^2 - x_2^1\right)^2}.$$

REMARK 1.11. The metric is related to the norm $\|\cdot\|$ through the identity

$$d(\boldsymbol{x}, \boldsymbol{y}) = \|\boldsymbol{x} - \boldsymbol{y}\|.$$

Subsets of $\mathbb{R}^n$

DEFINITION 1.12 (Open Ball). Let $\boldsymbol{x} \in \mathbb{R}^n$ and $r > 0$. The *open ball* $\mathcal{B}(\boldsymbol{x}; r)$ with center $\boldsymbol{x}$ and radius r is given by

$$\mathcal{B}(\boldsymbol{x}; r) := \left\{\boldsymbol{y} \in \mathbb{R}^n : d(\boldsymbol{x}, \boldsymbol{y}) < r\right\}.$$

DEFINITION 1.13 (Interior). Let $S \subset \mathbb{R}^n$. A point $\boldsymbol{x} \in S$ is called an *interior point* of S if there is some $r > 0$ such that $\mathcal{B}(\boldsymbol{x}; r) \subset S$. The set of all interior points of S is called its *interior* and is denoted S°.

DEFINITION 1.14. Let $S \subset \mathbb{R}^n$.

- S is *open* if for every $\boldsymbol{x} \in S$ there exists $r > 0$ such that $\mathcal{B}(\boldsymbol{x}; r) \subset S$.
- S is *closed* if its complement $\mathbb{R}^n \setminus S$ is open.
- S is *bounded* if there exists $r > 0$ such that $S \subset \mathcal{B}(\boldsymbol{0}; r)$.
- S is *compact* if (and only if) it is closed and bounded (Heine-Borel Theorem).[2]

EXAMPLE 1.15. On $\mathbb{R}$, the interval $(0, 1)$ is open, the interval $[0, 1]$ is closed. Both $(0, 1)$ and $[0, 1]$ are bounded, and $[0, 1]$ is compact. However, the interval $(0, 1]$ is neither open nor closed. But $\mathbb{R}$ is both open and closed.

Limit and Continuity

Functions

A *function* from $\mathbb{R}^m$ to $\mathbb{R}^n$ (sometimes called a vector-valued function of m variables) is a rule which associates to each point in $\mathbb{R}^m$ some point in $\mathbb{R}^n$. We write

$$f : \mathbb{R}^m \to \mathbb{R}^n$$

to indicate that $f(\boldsymbol{x}) \in \mathbb{R}^n$ is defined for $\boldsymbol{x} \in \mathbb{R}^m$.

The notation $f : A \to \mathbb{R}^n$ indicates that $f(\boldsymbol{x})$ is defined only for $\boldsymbol{x}$ in the set A, which is called the *domain* of f. If $B \subset A$, we define $f(B)$ as the set of all $f(\boldsymbol{x})$ for $\boldsymbol{x} \in B$:

$$f(B) := \{f(\boldsymbol{x}) : \boldsymbol{x} \in B\}.$$

If $C \subset \mathbb{R}^n$ we define

$$f^{-1}(C) := \{\boldsymbol{x} \in A : f(\boldsymbol{x}) \in C\}.$$

The notation $f : A \to B$ indicates that $f(A) \subset B$.

[2] This definition does not work for more general metric spaces. See Willard (2004) for details.

A function $f\colon A \to \mathbb{R}^n$ determines n *component functions* $f_1, \ldots, f_n\colon A \to \mathbb{R}$ by

$$f(\boldsymbol{x}) = \big(f_1(\boldsymbol{x}), \ldots, f_n(\boldsymbol{x})\big)\,.$$

Sequences

A *sequence* is a function that assigns to each natural number $n \in \mathbb{N}$ a vector or point $\boldsymbol{x}_n \in \mathbb{R}^n$. We usually write the sequences as $\{\boldsymbol{x}_n\}_{n=1}^{\infty}$ or $\{\boldsymbol{x}_n\}$.

EXAMPLE 1.16. Examples of sequences in $\mathbb{R}^2$ are

a. $\{\boldsymbol{x}_n\} = \{(n, n)\}$.

b. $\{\boldsymbol{x}_n\} = \{(\cos\frac{n\pi}{2}, \sin\frac{n\pi}{2})\}$.

c. $\{\boldsymbol{x}_n\} = \{((-1)^n/2^n, 1/2^n)\}$.

d. $\{\boldsymbol{x}_n\} = \{((-1)^n - 1/n, (-1)^n - 1/n)\}$.

See Figure 1.2.

DEFINITION 1.17 (Limit). A sequence $\{\boldsymbol{x}_n\}$ is said to have a *limit* $\boldsymbol{x}$ or to *converge* to $\boldsymbol{x}$ if for every $\varepsilon > 0$ there is $N_\varepsilon \in \mathbb{N}$ such that whenever $n > N_\varepsilon$, we have $\boldsymbol{x}_n \in \mathcal{B}(\boldsymbol{x}; \varepsilon)$. We write

$$\lim_{n\to\infty} \boldsymbol{x}_n = \boldsymbol{x} \quad \text{or} \quad \boldsymbol{x}_n \to \boldsymbol{x}\,.$$

EXAMPLE 1.18. In Example 1.16, the sequences (a), (b) and (d) do not converge, while the sequence (c) converges to $(0, 0)$.

Naive Continuity

Perhaps the simplest way to say that a function $f\colon A \to \mathbb{R}$ is *continuous* would be to say that one can draw its graph without taking the pencil off the paper. For example, a function whose graph looks like in Figure 1.3 would be continuous in this sense.[3]

But if we look at the function $f(x) = 1/x$, then we see that things are not so simple. The graph of this function has two parts—one part corresponding to negative x values, and the other to positive x values. The function is not defined at 0, so we certainly cannot draw both parts of this graph without taking our pencil off the paper; see Figure 1.4. Of course, $f(x) = 1/x$ is continuous near every point in its domain. Such a function deserves to be called continuous. So this characterization of continuity in terms of graph-sketching is too simplistic.

[3] I thank Prof. Wolfgang Buehler for mentioning this intuitive explanation to me. The current expression is from Crossley (2005, Chapter 2)

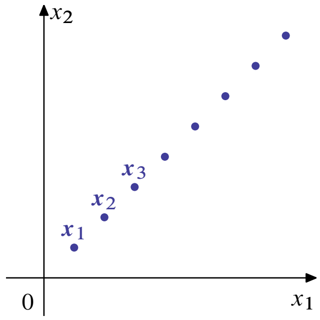

(a) $\{\boldsymbol{x}_n\} = \{(n, n)\}$.

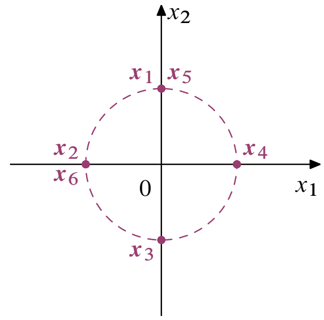

(b) $\{\boldsymbol{x}_n\} = \{(\cos \frac{n\pi}{2}, \sin \frac{n\pi}{2})\}$.

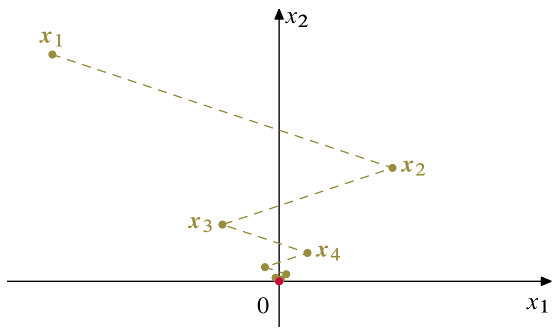

(c) $\{\boldsymbol{x}_n\} = \{((-1)^n/2^n, 1/2^n)\}$, which is convergent.

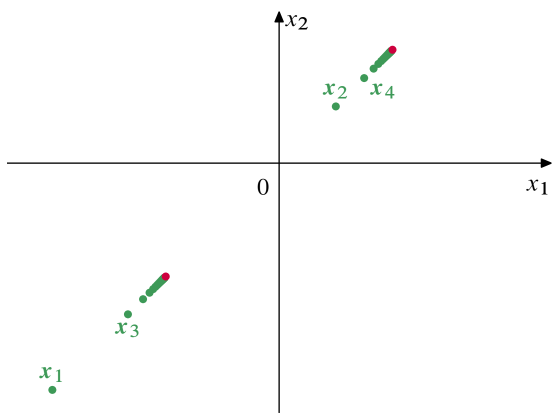

(d) $\{\boldsymbol{x}_n\} = \{((-1)^n - 1/n, (-1)^n - 1/n)\}$.

FIGURE 1.2. Examples of sequences.

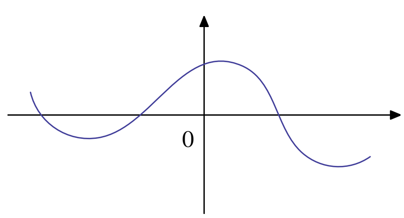

FIGURE 1.3. A continuous function.

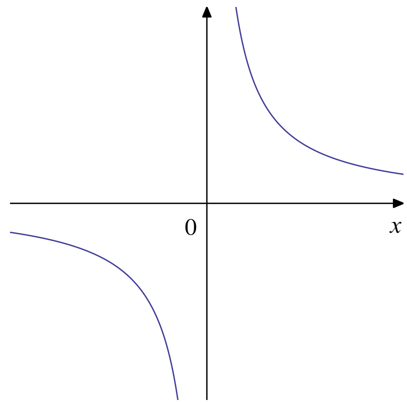

FIGURE 1.4. We cannot draw the graph of $1/x$ without taking our pencil off the paper.

Rigorous Continuity

The notation $\lim_{\boldsymbol{x}\to\boldsymbol{a}} f(\boldsymbol{x}) = \boldsymbol{b}$ means, as in the one-variable case, that we get $f(\boldsymbol{x})$ as close to $\boldsymbol{b}$ as desired, by choosing $\boldsymbol{x}$ sufficiently close to, but not equal to, $\boldsymbol{a}$. In mathematical terms this means that for every number $\varepsilon > 0$ there is a number $\delta > 0$ such that $\|f(\boldsymbol{x}) - \boldsymbol{b}\| < \varepsilon$ for all $\boldsymbol{x}$ in the domain of f which satisfy $0 < \|\boldsymbol{x} - \boldsymbol{a}\| < \delta$.

A function $f : A \to \mathbb{R}^n$ is called *continuous at* $\boldsymbol{a} \in A$ if $\lim_{\boldsymbol{x}\to\boldsymbol{a}} f(\boldsymbol{x}) = f(\boldsymbol{a})$, and f is continuous if it is continuous at each $\boldsymbol{a} \in A$.

► EXERCISE 1.19. *Let*

$$f(x) = \begin{cases} x & \text{if } x \neq 1 \\ 3/2 & \text{if } x = 1. \end{cases}$$

Show that $f(x)$ *is not continuous at* $a = 1$.

1.2 Directional Derivative and Derivative

Let us first recall how the derivative of a real-valued function of a real variable is defined. Let $A \subset \mathbb{R}$; let $f : A \to \mathbb{R}$. Suppose A contains a neighborhood of the point a, that is, there is an open ball $\mathcal{B}(a; r)$ such that $\mathcal{B}(a; r) \subset A$. We define the *derivative of* f *at* a by the equation

$$f'(a) = \lim_{t\to 0} \frac{f(a+t) - f(a)}{t}, \tag{1.1}$$

provided the limit exists. In this case, we say that f is *differentiable at* a. Geometrically, $f'(a)$ is the slope of the tangent line to the graph of f at the point $(a, f(a))$.

DEFINITION 1.20. For a function $f : (a, b) \to \mathbb{R}$, and point $x_0 \in (a, b)$, if

$$\lim_{t \uparrow 0} \frac{f(x_0 + t) - f(x_0)}{t}$$

exists and is finite, we denote this limit by $f'_-(x_0)$ and call it the *left-hand derivative* of f at x_0. Similarly, we define $f'_+(x_0)$ and call it the *right-hand derivative* of g at x_0. Of course, f is differentiable at x_0 iff it has left-hand and right-hand derivatives at x_0 that are equal.

Now let $A \subset \mathbb{R}^m$, where $m > 1$; let $f : A \to \mathbb{R}^n$. Can we define the derivative of f by replacing a and t in the definition just given by points of $\mathbb{R}^m$? Certainly we cannot since we cannot divide a point of $\mathbb{R}^n$ by a point of $\mathbb{R}^m$ if $m > 1$.

Directional Derivative

The following is our first attempt at a definition of "derivative".

DEFINITION 1.21 (Directional Derivative). Let $A \subset \mathbb{R}^m$; let $f : A \to \mathbb{R}^n$. Suppose A contains a neighborhood of $\boldsymbol{a}$. Given $\boldsymbol{u} \in \mathbb{R}^m$ with $\boldsymbol{u} \neq \boldsymbol{0}$, define

$$f'(\boldsymbol{a}; \boldsymbol{u}) = \lim_{t \to 0} \frac{f(\boldsymbol{a} + t\boldsymbol{u}) - f(\boldsymbol{a})}{t},$$

provided the limit exists. This limit is called the *directional derivative of f at $\boldsymbol{a}$ with respect to the vector $\boldsymbol{u}$*.[4]

EXAMPLE 1.22. Let $f : \mathbb{R}^2 \to \mathbb{R}$ be given by the equation

$$f(x_1, x_2) = x_1 x_2.$$

The directional derivative of f at $\boldsymbol{a} = (a_1, a_2)$ with respect to the vector $\boldsymbol{u} = (1, 0)$ is

$$\begin{aligned} f'(\boldsymbol{a}; \boldsymbol{u}) &= \lim_{t \to 0} \frac{f(\boldsymbol{a} + t\boldsymbol{u}) - f(\boldsymbol{a})}{t} \\ &= \lim_{t \to 0} \frac{(a_1 + t)a_2 - a_1 a_2}{t} \\ &= a_2. \end{aligned}$$

With respect to the vector $\boldsymbol{v} = (1, 2)$, the directional derivative is

$$f'(\boldsymbol{a}; \boldsymbol{v}) = \lim_{t \to 0} \frac{(a_1 + t)(a_2 + 2t) - a_1 a_2}{t} = 2a_1 + a_2.$$

[4] In calculus, one usually requires $\boldsymbol{u}$ to be a unit vector, i.e., $\|\boldsymbol{u}\| = 1$, but that is not necessary.

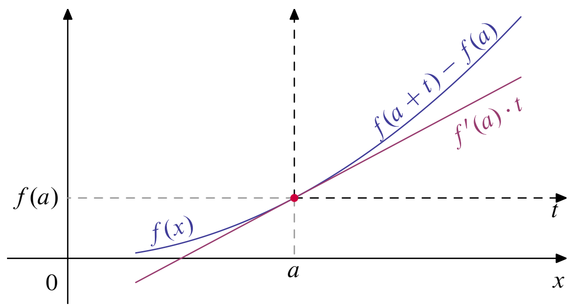

FIGURE 1.5. $f'(a)t$ is the linear approximation to $f(a+t)-f(a)$ at $\boldsymbol{a}$.

However, directional derivative is NOT the appropriate generalization of the notion of "derivative". The main problems are:

- Continuity does not follow from this definition of "differentiability" (There exists functions such that $f'(\boldsymbol{a};\boldsymbol{u})$ exists for all $\boldsymbol{u}\neq\mathbf{0}$ but are not continuous);
- Composites of "differentiable" functions may not differentiable.

Derivative

To give the right generalization of the notion of "derivative", let us rewrite (1.1) as

$$\lim_{t\to 0}\frac{f(a+t)-f(a)-f'(a)\cdot t}{t}=0,$$

or equivalently

$$\lim_{t\to 0}\frac{|f(a+t)-f(a)-f'(a)\cdot t|}{|t|}=0,$$

which makes precise the sense in which we approximate the *increment function* $f(a+t)-f(a)$ by the linear function $f'(a)\cdot t$. We often call $f'(a)\cdot t$ the "first-order approximation" or the "linear approximation" to the increment function.[5] See Figure 1.5. It is this idea leads to the following definition:

DEFINITION 1.23 (Differentiability). Let $A\subset\mathbb{R}^m$; let $f\colon A\to\mathbb{R}^n$. Suppose A contains a neighborhood of $\boldsymbol{a}$. We say that f is *differentiable at* $\boldsymbol{a}$ if there is an $n\times m$ matrix $\mathbf{B}_{\boldsymbol{a}}$ such that

$$\lim_{\boldsymbol{h}\to\mathbf{0}}\frac{\|f(\boldsymbol{a}+\boldsymbol{h})-f(\boldsymbol{a})-\mathbf{B}_{\boldsymbol{a}}\cdot\boldsymbol{h}\|}{\|\boldsymbol{h}\|}=0.$$

The matrix $\mathbf{B}_{\boldsymbol{a}}$, which is unique, is called the *derivative of f at $\boldsymbol{a}$*; it is denoted $\mathbb{D}f(\boldsymbol{a})$.

[5] "All science is dominated by the idea of approximation."—Bertrand Russell.

REMARK 1.24. Notice that $\boldsymbol{h}$ is a point of $\mathbb{R}^m$ and $f(\boldsymbol{a}+\boldsymbol{h}) - f(\boldsymbol{a}) - \mathbf{B}_{\boldsymbol{a}} \cdot \boldsymbol{h}$ is a point of $\mathbb{R}^n$, so the norm signs are essential. (Actually, it is enough to only take the norm of $\boldsymbol{h}$.)

REMARK 1.25. The derivative $\mathbb{D}f(\boldsymbol{a})$ depends on the point $\boldsymbol{a}$ as well as the function f. We are not saying that there exists a $\mathbf{B}$ which works for all $\boldsymbol{a}$, but that for a fixed $\boldsymbol{a}$ such a $\mathbf{B}$ exists.

EXAMPLE 1.26. Let $f\colon \mathbb{R}^m \to \mathbb{R}^n$ be defined by the equation

$$f(\boldsymbol{x}) = \mathbf{A} \cdot \boldsymbol{x} + \boldsymbol{b},$$

where $\mathbf{A}$ is an $n \times m$ matrix, and $\boldsymbol{a} \in \mathbb{R}^n$. Then

$$\lim_{\boldsymbol{h} \to \mathbf{0}} \frac{\| f(\boldsymbol{a}+\boldsymbol{h}) - f(\boldsymbol{a}) - \mathbf{A} \cdot \boldsymbol{h} \|}{\|\boldsymbol{h}\|} = 0;$$

that is, $\mathbb{D}f(\boldsymbol{a}) = \mathbf{A}$.

We now show that the definition of derivative is stronger than directional derivative. In particular, we have:

THEOREM 1.27. *Let $A \subset \mathbb{R}^m$; let $f\colon A \to \mathbb{R}^n$. If f is differentiable at $\boldsymbol{a}$, then f is continuous at $\boldsymbol{a}$.*

PROOF. Let $\mathbf{B}_{\boldsymbol{a}} = \mathbb{D}f(\boldsymbol{a})$. For $\boldsymbol{h}$ near $\mathbf{0}$ but different from $\mathbf{0}$, write

$$f(\boldsymbol{a}+\boldsymbol{h}) - f(\boldsymbol{a}) = \|\boldsymbol{h}\| \cdot \left[\frac{f(\boldsymbol{a}+\boldsymbol{h}) - f(\boldsymbol{a}) - \mathbf{B}_{\boldsymbol{a}} \cdot \boldsymbol{h}}{\|\boldsymbol{h}\|} \right] + \mathbf{B}_{\boldsymbol{a}} \cdot \boldsymbol{h}.$$

Thus, $\| f(\boldsymbol{a}+\boldsymbol{h}) - f(\boldsymbol{a}) \| \to 0$ as $\boldsymbol{h} \to 0$. That is, f is continuous at $\boldsymbol{a}$. □

However, there is a nice connection between directional derivative and derivative.

THEOREM 1.28. *Let $A \subset \mathbb{R}^m$; let $f\colon A \to \mathbb{R}^n$. If f is differentiable at $\boldsymbol{a}$, then all the directional derivatives of f at $\boldsymbol{a}$ exist, and*

$$f'(\boldsymbol{a}; \boldsymbol{u}) = \mathbb{D}f(\boldsymbol{a}) \cdot \boldsymbol{u}.$$

PROOF. Let $\mathbf{B}_{\boldsymbol{a}} = \mathbb{D}f(\boldsymbol{a})$. Set $\boldsymbol{h} = t\boldsymbol{u}$ in the definition of differentiability, where $t \neq 0$. Then by hypothesis,

$$\begin{aligned} 0 &= \lim_{t \to 0} \frac{\| f(\boldsymbol{a}+t\boldsymbol{u}) - f(\boldsymbol{a}) - \mathbf{B}_{\boldsymbol{a}} \cdot t\boldsymbol{u} \|}{\|t\boldsymbol{u}\|} \\ &= \lim_{t \to 0} \frac{\| f(\boldsymbol{a}+t\boldsymbol{u}) - f(\boldsymbol{a}) - t \cdot (\mathbf{B}_{\boldsymbol{a}} \cdot \boldsymbol{u}) \|}{|t| \cdot \|\boldsymbol{u}\|}. \end{aligned} \tag{1.2}$$

If $t \downarrow 0$, we multiply (1.2) by $\|\boldsymbol{u}\|$ to conclude that

$$\lim_{t \downarrow 0} \frac{f(\boldsymbol{a} + t\boldsymbol{u}) - f(\boldsymbol{a})}{t} - \mathbf{B}_{\boldsymbol{a}} \cdot \boldsymbol{u} = \mathbf{0}.$$

If $t \uparrow 0$, we multiply (1.2) by $-\|\boldsymbol{u}\|$ to reach the same conclusion. Thus, $f'(\boldsymbol{a}; \boldsymbol{u}) = \mathbf{B}_{\boldsymbol{a}} \cdot \boldsymbol{u}$. □

► EXERCISE 1.29. *Define* $f : \mathbb{R}^2 \to \mathbb{R}$ *by setting*

$$f(x, y) = \begin{cases} 0 & \textit{if } (x, y) = (0, 0) \\ \dfrac{x^2 y}{x^4 + y^2} & \textit{if } (x, y) \neq (0, 0). \end{cases}$$

Show that all directional derivatives of f *exist at* $(0, 0)$*, but that* f *is not differentiable at* $(0, 0)$*.*

► EXERCISE 1.30. *Let* $f : \mathbb{R}^2 \to \mathbb{R}$ *be defined by* $f(x, y) = \sqrt{|xy|}$*. Show that* f *is not differentiable at* $(0, 0)$*.*

1.3 Partial Derivatives and the Jacobian

We now introduce the notion of the "partial derivatives" of a real-valued function. Let $(\boldsymbol{e}_1, \ldots, \boldsymbol{e}_m)$ be the stand basis of $\mathbb{R}^m$, i.e.,

$$\begin{aligned} \boldsymbol{e}_1 &= (1, 0, 0, \ldots, 0), \\ \boldsymbol{e}_2 &= (0, 1, 0, \ldots, 0), \\ &\cdots \\ \boldsymbol{e}_m &= (0, 0, \ldots, 0, 1). \end{aligned}$$

DEFINITION 1.31 (Partial Derivatives). Let $A \subset \mathbb{R}^m$; let $f : A \to \mathbb{R}$. We define the j^{th} *partial derivative of* f *at* $\boldsymbol{a}$ to be the directional derivative of f at $\boldsymbol{a}$ with respect to the vector $\boldsymbol{e}_j$, provided this derivative exists; and we denote it by $\mathbb{D}_j f(\boldsymbol{a})$. That is,

$$\mathbb{D}_j f(\boldsymbol{a}) = \lim_{t \to 0} \frac{f(\boldsymbol{a} + t\boldsymbol{e}_j) - f(\boldsymbol{a})}{t}.$$

REMARK 1.32. It is important to note that $\mathbb{D}_j f(\boldsymbol{a})$ is the ordinary derivative of a certain function; in fact, if $g(x) = f(a_1, \ldots, a_{j-1}, x, a_{j+1}, \ldots, a_m)$, then $\mathbb{D}_j f(\boldsymbol{a}) = g'(a_j)$. This means that $\mathbb{D}_j f(\boldsymbol{a})$ is the slope of the tangent line at $(\boldsymbol{a}, f(\boldsymbol{a}))$ to the curve obtained by intersecting the graph of f with the plane $x_i = a_i$ with $i \neq j$. See Figure 1.6.

We now relate partial derivatives to the derivative in the case where f is a real-valued function.

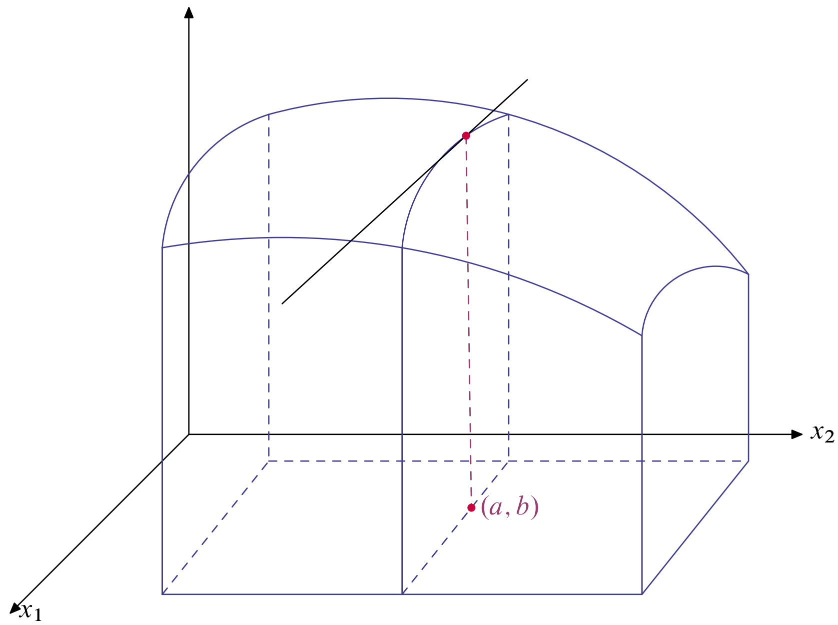

FIGURE 1.6. $\mathbb{D}_1 f(a, b)$.

THEOREM 1.33. *Let $A \subset \mathbb{R}^m$; let $f : A \to \mathbb{R}$. If f is differentiable at $\boldsymbol{a}$, then*

$$\mathbb{D} f(\boldsymbol{a}) = \begin{bmatrix} \mathbb{D}_1 f(\boldsymbol{a}) & \mathbb{D}_2 f(\boldsymbol{a}) & \cdots & \mathbb{D}_m f(\boldsymbol{a}) \end{bmatrix}.$$

PROOF. If f is differentiable at $\boldsymbol{a}$, then $\mathbb{D} f(\boldsymbol{a})$ is a $(1 \times m)$-matrix. Let

$$\mathbb{D} f(\boldsymbol{a}) = \begin{bmatrix} \lambda_1 & \lambda_2 & \cdots & \lambda_m \end{bmatrix}.$$

It follows from Theorem 1.28 that

$$\mathbb{D}_j f(\boldsymbol{a}) = f'(\boldsymbol{a}; \boldsymbol{e}_j) = \mathbb{D} f(\boldsymbol{a}) \cdot \boldsymbol{e}_j = \lambda_j. \qquad \square$$

Theorem 1.33 can be generalized as follows:

THEOREM 1.34. *Let $A \subset \mathbb{R}^m$; let $f : A \to \mathbb{R}^n$. Suppose A contains a neighborhood of $\boldsymbol{a}$. Let $f_i : A \to \mathbb{R}$ be the i^{th} component function of f, so that*

$$f(\boldsymbol{x}) = \begin{bmatrix} f_1(\boldsymbol{x}) \\ \vdots \\ f_n(\boldsymbol{x}) \end{bmatrix}.$$

a. *The function f is differentiable at $\boldsymbol{a}$ iff each component function f_i is differentiable at $\boldsymbol{a}$.*

b. *If f is differentiable at $\boldsymbol{a}$, then its derivative is the $(n \times m)$-matrix whose i^{th} row is the derivative of the function f_i. That is,*

$$\mathbb{D} f(\boldsymbol{a}) = \begin{bmatrix} \mathbb{D} f_1(\boldsymbol{a}) \\ \vdots \\ \mathbb{D} f_n(\boldsymbol{a}) \end{bmatrix} = \begin{bmatrix} \mathbb{D}_1 f_1(\boldsymbol{a}) & \cdots & \mathbb{D}_m f_1(\boldsymbol{a}) \\ \vdots & \ddots & \vdots \\ \mathbb{D}_1 f_n(\boldsymbol{a}) & \cdots & \mathbb{D}_m f_n(\boldsymbol{a}) \end{bmatrix}.$$

► EXERCISE 1.35. *Prove Theorem 1.34.*

DEFINITION 1.36 (Jocobian Matrix). Let $A \subset \mathbb{R}^m$; let $f : A \to \mathbb{R}^n$. If the partial derivatives of the component functions f_i of f exist at $\boldsymbol{a}$, then one can form the matrix that has $\mathbb{D}_j f_i(\boldsymbol{a})$ as its entry in row i and column j. This matrix, denoted by $\mathbf{J} f(\boldsymbol{a})$, is called the *Jacobian matrix* of f. That is,

$$\mathbf{J} f(\boldsymbol{a}) = \begin{bmatrix} \mathbb{D}_1 f_1(\boldsymbol{a}) & \cdots & \mathbb{D}_m f_1(\boldsymbol{a}) \\ \vdots & \ddots & \vdots \\ \mathbb{D}_1 f_n(\boldsymbol{a}) & \cdots & \mathbb{D}_m f_n(\boldsymbol{a}) \end{bmatrix}.$$

REMARK 1.37. The Jacobian encapsulates all the essential information regarding the linear function that best approximates a differentiable function at a particular point. For this reason it is the Jacobian which is usually used in practical calculations with the derivative

REMARK 1.38. If f is differentiable at $\boldsymbol{a}$, then $\mathbf{J} f(\boldsymbol{a}) = \mathbb{D} f(\boldsymbol{a})$. However, it is possible for the partial derivatives, and hence the Jacobian matrix, to exist, without it following that f is differentiable at $\boldsymbol{a}$ (see Exercise 1.29).

1.4 GRADIENT AND ITS GEOMETRIC INTERPRETATION

DEFINITION 1.39 (Gradient). Let $A \subset \mathbb{R}^m$; let $f : A \to \mathbb{R}$. Suppose A contains a neighborhood of $\boldsymbol{a}$. The *gradient* of f, denoted by $\nabla f(\boldsymbol{a})$, is defined by

$$\nabla f(\boldsymbol{a}) := \sum_{i=1}^{m} \mathbb{D}_i f(\boldsymbol{a}) \cdot \boldsymbol{e}_i = \begin{bmatrix} \mathbb{D}_1 f(\boldsymbol{a}) & \mathbb{D}_2 f(\boldsymbol{a}) & \cdots & \mathbb{D}_m f(\boldsymbol{a}) \end{bmatrix}.$$

REMARK 1.40. It follows from Theorem 1.33 that if f is differentiable at $\boldsymbol{a}$, then $\nabla f(\boldsymbol{a}) = \mathbb{D} f(\boldsymbol{a})$. The inverse does not hold; see Remark 1.38.

Let us now present a very important fact about gradient: the gradient is orthogonal to the *level set*.[6] You will see this fact again and again. For simplicity, we shall restrict ourselves on the case that $f : \mathbb{R}^2 \to \mathbb{R}$.

Consider the gradient of f at $(a, b) \in \mathbb{R}^2$:

$$\nabla f(a, b) = (\mathbb{D}_1 f(a, b), \mathbb{D}_2 f(a, b)).$$

We show that $\nabla f(a, b)$ is orthogonal to the level set $L_f(f(a, b)) = \{(x, y) \in \mathbb{R}^2 : f(x, y) = f(a, b)\}$ at (a, b), which means that $\nabla f(a, b)$ is orthogonal to the tangent line at (a, b). Let us begin with an example.

EXAMPLE 1.41. Let $f(x, y) = x^2 + y^2$. Then $\nabla f(1, 3) = (2x, 2y)|_{(x,y)=(1,3)} = (2, 6)$. The level set of $f(1, 3) = 10$ is given by $x^2 + y^2 = 10$. Calculus yields

$$\left.\frac{\mathrm{d}y}{\mathrm{d}x}\right|_{(1,3)} = -\frac{1}{3}.$$

Hence, the tangent line at $(1, 3)$ is given by

$$y = 3 - \frac{x - 1}{3}.$$

Then the result follows immediately; see Figure 1.7.

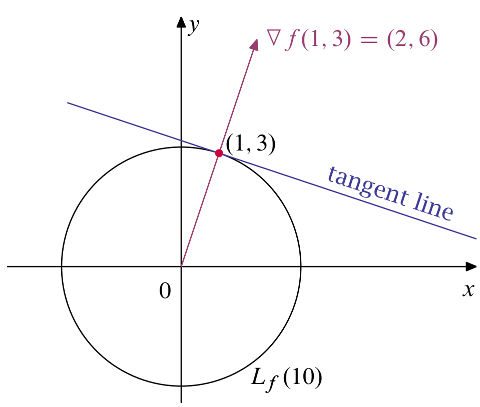

FIGURE 1.7. The geometric interpretation of gradient.

---

[6] Given a real-valued function f on $A \subset \mathbb{R}^m$, the *level set* of f through c, where c is in the range, is

$$L_f(c) := \{\boldsymbol{x} \in A : f(\boldsymbol{x}) = c\}.$$

We use level sets to help analyze functions in higher-dimensional spaces.

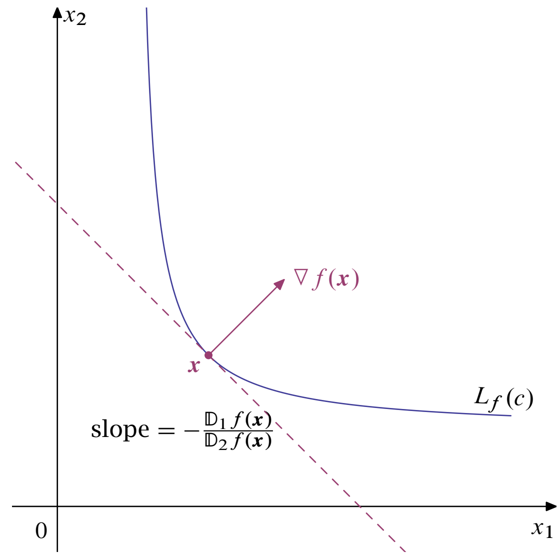

FIGURE 1.8. The geometric interpretation of $\nabla f(\boldsymbol{x})$.

We next turn to the more general analysis. Fix c in the range of f and take an arbitrary point $\boldsymbol{x} = (x_1, x_2)$ on the level set $L_f(c)$. If we change x_1 and x_2, and are to remain on the level set, $\mathrm{d}x_1$ and $\mathrm{d}x_2$ must be such as to leave the value of f unchanged at c. They must therefore satisfy

$$f'(\boldsymbol{x}; (\mathrm{d}x_1, \mathrm{d}x_2)) = \mathbb{D}_1 f(\boldsymbol{x})\,\mathrm{d}x_1 + \mathbb{D}_2 f(\boldsymbol{x})\,\mathrm{d}x_2 = 0. \tag{1.3}$$

By solving (1.3) for $\mathrm{d}x_2/\mathrm{d}x_1$, the slope of the level set through $\boldsymbol{x}$ will be (see Figure 1.8)

$$\frac{\mathrm{d}x_2}{\mathrm{d}x_1} = -\frac{\mathbb{D}_1 f(\boldsymbol{x})}{\mathbb{D}_2 f(\boldsymbol{x})}.$$

Since the slope of the vector $\nabla f(\boldsymbol{x}) = (\mathbb{D}_1 f(\boldsymbol{x}), \mathbb{D}_2 f(\boldsymbol{x}))$ is $\mathbb{D}_2 f(\boldsymbol{x})/\mathbb{D}_1 f(\boldsymbol{x})$, we obtain the desired result.

REMARK 1.42. We will provide an physical interpretation of gradient in page 30.

1.5 Continuously Differentiable Functions

We know that mere existence of the partial derivatives does not imply differentiability (see Exercise 1.29). If, however, we impose the additional condition that these partial derivatives are continuous, then differentiability is assured.

THEOREM 1.43. *Let A be open in $\mathbb{R}^m$. Suppose that the partial derivatives $\mathbb{D}_j f_i(\boldsymbol{x})$ of the component functions of f exist at each point $\boldsymbol{x} \in A$ and are continuous on A. Then f is differentiable at each point of A.*

A function satisfying the hypotheses of this theorem is often said to be *continuously differentiable*, or *of class* $\mathscr{C}^1$, on A.

PROOF OF THEOREM 1.43. It suffices to show that each component function of f is differentiable. Therefore we may restrict ourselves to the case of a real-valued function $f : A \to \mathbb{R}$. Let $\boldsymbol{a} \in A$. Then

$$\begin{aligned} f(\boldsymbol{a}+\boldsymbol{h}) - f(\boldsymbol{a}) = & f(a_1+h_1, a_2, \dots, a_m) - f(a_1, \dots, a_m) \\ & + f(a_1+h_1, a_2+h_2, a_3, \dots, a_m) - f(a_1+h_1, a_2, \dots, a_m) \\ & + \cdots \\ & + f(a_1+h_1, \dots, a_m+h_m) - f(a_1+h_1, \dots, a_{m-1}+h_{m-1}, a_m). \end{aligned}$$

Note that $\mathbb{D}_1 f$ is the derivative of the function g defined by

$$g(x) = f(x, a_2, \dots, a_m).$$

Applying the mean-value theorem (Rudin, 1976, Theorem 5.10) to g we obtain

$$f(a_1+h_1, a_2, \dots, a_m) - f(a_1, \dots, a_m) = h_1 \cdot \mathbb{D}_1 f(b_1, a_2, \dots, a_m)$$

for some $b_1 \in (a_1, a_1+h_1)$. Similarly the i^{th} term in the sum equals

$$h_i \cdot \mathbb{D}_i f(a_1+h_1, \dots, a_{i-1}+h_{i-1}, b_i, a_{i+1}, \dots, a_m) = h_i \cdot \mathbb{D}_i f(\boldsymbol{c}_i)$$

for some $\boldsymbol{c}_i$. Then

$$\begin{aligned} \lim_{\boldsymbol{h}\to\boldsymbol{0}} \frac{\left| f(\boldsymbol{a}+\boldsymbol{h}) - f(\boldsymbol{a}) - \sum_{i=1}^m \mathbb{D}_i f(\boldsymbol{a}) \cdot h_i \right|}{\|\boldsymbol{h}\|} &= \lim_{\boldsymbol{h}\to\boldsymbol{0}} \frac{\left| \sum_{i=1}^m \left[\mathbb{D}_i f(\boldsymbol{c}_i) - \mathbb{D}_i f(\boldsymbol{a}) \right] \cdot h_i \right|}{\|\boldsymbol{h}\|} \\ &= \lim_{\boldsymbol{h}\to\boldsymbol{0}} \sum_{i=1}^m |\mathbb{D}_i f(\boldsymbol{c}_i) - \mathbb{D}_i f(\boldsymbol{a})| \cdot \frac{|h_i|}{\|\boldsymbol{h}\|} \\ &\leqslant \lim_{\boldsymbol{h}\to\boldsymbol{0}} \sum_{i=1}^m |\mathbb{D}_i f(\boldsymbol{c}_i) - \mathbb{D}_i f(\boldsymbol{a})| \\ &= 0, \end{aligned}$$

since $\mathbb{D}_i f$ is continuous at $\boldsymbol{a}$. □

REMARK 1.44. It follows from Theorem 1.43 that $\sin(xy)$ and $xy^2 + ze^{xy}$ are both differentiable since they are of class $\mathscr{C}^1$.

Let $A \subset \mathbb{R}^m$ and $f : A \to \mathbb{R}^n$. Suppose that the partial derivative $\mathbb{D}_j f_i$ of the component functions of f exist on A. These then are functions from A to $\mathbb{R}$, and we may consider their partial derivatives, which have the form

$$\mathbb{D}_k(\mathbb{D}_j f_i) =: \mathbb{D}_{jk} f_i$$

and are called the *second-order partial derivatives of* f. Similarly, one defines the third-order partial derivatives of the functions f_i, or more generally the *partial derivatives of order* r for arbitrary r.

DEFINITION 1.45. If the partial derivatives of the function f_i of order less than or equal to r are continuous on A, we say f is *of class* $\mathcal{C}^r$ on A. We say f is *of class* $\mathcal{C}^\infty$ on A if the partials of the functions f_i of all orders are continuous on A.

DEFINITION 1.46 (Hessian). Let $\boldsymbol{a} \in A \subset \mathbb{R}^m$; let $f : A \to \mathbb{R}$ be twice-differentiable at $\boldsymbol{a}$. The $m \times m$ matrix representing the second derivative of f is called the *Hessian* of f, denoted $\mathbf{H} f(\boldsymbol{a})$:

$$\mathbf{H} f(\boldsymbol{a}) = \begin{bmatrix} \mathbb{D}_{11} f(\boldsymbol{a}) & \mathbb{D}_{12} f(\boldsymbol{a}) & \cdots & \mathbb{D}_{1m} f(\boldsymbol{a}) \\ \mathbb{D}_{21} f(\boldsymbol{a}) & \mathbb{D}_{22} f(\boldsymbol{a}) & \cdots & \mathbb{D}_{2m} f(\boldsymbol{a}) \\ \vdots & \vdots & \ddots & \vdots \\ \mathbb{D}_{m1} f(\boldsymbol{a}) & \mathbb{D}_{m2} f(\boldsymbol{a}) & \cdots & \mathbb{D}_{mm} f(\boldsymbol{a}) \end{bmatrix} = \mathbb{D}(\nabla f).$$

REMARK 1.47. If $f : A \to \mathbb{R}$ is of class $\mathcal{C}^2$, then the Hessian of f is a symmetric matrix, i.e., $\mathbb{D}_{ij} f(\boldsymbol{a}) = \mathbb{D}_{ji} f(\boldsymbol{a})$ for all $i, j = 1, \dots, m$ and for all $\boldsymbol{a} \in A$. See Rudin (1976, Corollary to Theorem 9.41, p. 236).

► EXERCISE 1.48. *Find the Hessian of the Cobb-Douglas function*

$$f(x, y) = x^\alpha y^\beta.$$

1.6 The Chain Rule

We now extend the familiar chain rule to the current setting.

THEOREM 1.49 (Chain Rule). *Let* $A \subset \mathbb{R}^m$*; let* $B \subset \mathbb{R}^n$*. Let* $f : A \to \mathbb{R}^n$ *and* $g : B \to \mathbb{R}^p$*, with* $f(A) \subset B$*. Suppose* $f(\boldsymbol{a}) = \boldsymbol{b}$*. If* f *is differentiable at* $\boldsymbol{a}$*, and if* g *is differentiable at* $\boldsymbol{b}$*, then the composite function* $g \circ f : A \to \mathbb{R}^p$ *is differentiable at* $\boldsymbol{a}$*. Furthermore,*

$$\mathbb{D}(g \circ f)(\boldsymbol{a}) = \mathbb{D}g(\boldsymbol{b}) \cdot \mathbb{D}f(\boldsymbol{a}).$$

PROOF. Omitted. See Spivak (1965, Theorem 2-2), Rudin (1976, Theorem 9.15), or Munkres (1991, Theorem 7.1). □

1.7 Quadratic Forms: Definite and Semidefinite Matrices

DEFINITION 1.50 (Quadratic Form). Let $\mathbf{A}$ be a symmetric $n \times n$ matrix. A *quadratic form* on $\mathbb{R}^n$ is a function $Q_{\mathbf{A}}: \mathbb{R}^n \to \mathbb{R}$ of the form

$$Q_{\mathbf{A}}(\boldsymbol{x}) = \boldsymbol{x} \cdot \mathbf{A}\boldsymbol{x} = \sum_{i=1}^{n} \sum_{j=1}^{n} a_{ij} x_i x_j .$$

Since the quadratic form $Q_{\mathbf{A}}$ is completely specified by the matrix $\mathbf{A}$, we henceforth refer to $\mathbf{A}$ itself as the quadratic form. Observe that if f is of class $\mathcal{C}^2$, then the Hessian $\mathbf{H}f$ of f defines a quadratic form; see Remark 1.47.

DEFINITION 1.51. A quadratic form $\mathbf{A}$ is said to be

- *positive definite* if we have $\boldsymbol{x} \cdot \mathbf{A}\boldsymbol{x} > 0$ for all $\boldsymbol{x} \in \mathbb{R}^n \smallsetminus \{\mathbf{0}\}$;
- *positive semidefinite* if we have $\boldsymbol{x} \cdot \mathbf{A}\boldsymbol{x} \geqslant 0$ for all $\boldsymbol{x} \in \mathbb{R}^n$;
- *negative definite* if we have $\boldsymbol{x} \cdot \mathbf{A}\boldsymbol{x} < 0$ for all $\boldsymbol{x} \in \mathbb{R}^n \smallsetminus \{\mathbf{0}\}$;
- *negative semidefinite* if we have $\boldsymbol{x} \cdot \mathbf{A}\boldsymbol{x} \leqslant 0$ for all $\boldsymbol{x} \in \mathbb{R}^n$.

1.8 The Implicit Function Theorem

Here is a typical problem:

> "Assume that the equation $x^3 y + 2e^{xy} = 0$ determines y as a differentiable function of x. Find $\mathrm{d}y/\mathrm{d}x$."

One solves this calculus problem by "looking at y as a function of x," and differentiating with respect to x. One obtains the equation

$$3x^2 y + x^3 \frac{\mathrm{d}y}{\mathrm{d}x} + 2e^{xy}\left(y + x\frac{\partial y}{\partial x}\right) = 0,$$

which one solves for $\mathrm{d}y/\mathrm{d}x$. The derivative $\mathrm{d}y/\mathrm{d}x$ is of course expressed in terms of x and the unknown function y.

The case of an arbitrary function f is handled similarly. Supposing that the equation $f(x, y) = 0$ determines y as a differentiable function of x, say $y = g(x)$, the equation $f(x, g(x)) = 0$ is an identity. One applies the chain rule to calculate

$$\frac{\partial f}{\partial x} + \frac{\partial f}{\partial y} g'(x) = 0,$$

so that

$$g'(x) = -\frac{\partial f/\partial x}{\partial f/\partial y},$$

where the partial derivatives are evaluated at the point $(x, g(x))$. Note that the solution involves a hypothesis not given in the statement of the problem. In order to find $g'(x)$, it is *necessary* to assume that $\partial f/\partial y \neq 0$ at the point in question.

It in fact turns out that $\partial f/\partial y \neq 0$ is also *sufficient* to justify the assumptions we made in solving the problem. That is, if the function $f(x, y)$ has the property that $\partial f/\partial y \neq 0$ at a point (a, b) that is a solution of the equation $f(x, y) = 0$, then this equation *does* determine y as a function of x, for x near a, and this function of x is differentiable.

This result is a special case of a theorem called the *implicit function theorem*, which we consider in this section.

Example 1.52. Consider the function $f : \mathbb{R}^2 \to \mathbb{R}$ defined by

$$f(x, y) = x^2 + y^2 - 1.$$

If we choose (a, b) with $f(a, b) = 0$ and $a \neq \pm 1$, there are (Figure 1.9) open intervals A containing a and B containing b with the following property: if $x \in A$, there is a unique $y \in B$ with $f(x, y) = 0$. We can therefore define a function $g: A \to \mathbb{R}$ by the condition $g(x) \in B$ and $f(x, g(x)) = 0$ (if $b > 0$, as indicated in Figure 1.9, then $g(x) = \sqrt{1 - x^2}$). For the function f we are considering there is another number b_1 such that $f(a, b_1) = 0$. There will also be an interval B_1 containing b_1 such that, when $x \in A$, we have $f(x, g_1(x)) = 0$ for a unique $g_1(x) \in B_1$ (here $g_1(x) = -\sqrt{1 - x^2}$). Both g and g_1 are differentiable. These functions are said to be defined *implicitly* by the equation $f(x, y) = 0$.

If we choose $a = 1$ or -1 it is impossible to find any such function g defined in an open interval containing a.

We now introduce the *Implicit Function Theorem*. Let E be open in $\mathbb{R}^{k+n}$; let $f: E \to \mathbb{R}^n$. Write f in the form $f(\boldsymbol{x}, \boldsymbol{y})$ for $\boldsymbol{x} \in \mathbb{R}^k$ and $\boldsymbol{y} \in \mathbb{R}^n$. Think of the equation $f(\boldsymbol{x}, \boldsymbol{y}) = \mathbf{0}$ as a system of n equations in $k + n$ variables. With $(\boldsymbol{a}, \boldsymbol{b}) \in \mathbb{R}^k \times \mathbb{R}^n$ as a given solution, i.e. $f(\boldsymbol{a}, \boldsymbol{b}) = \mathbf{0}$, the theorem tells us under what condition, we can solve for the variables $\boldsymbol{y}$ near $\boldsymbol{b}$ in terms of the variables $\boldsymbol{x}$, to obtain a unique continuously differentiable solution. The new function $g(\boldsymbol{x})$ so obtained is said to be given by the equation $f(\boldsymbol{x}, \boldsymbol{y}) = \mathbf{0}$ implicitly. That is why the theorem is so named.

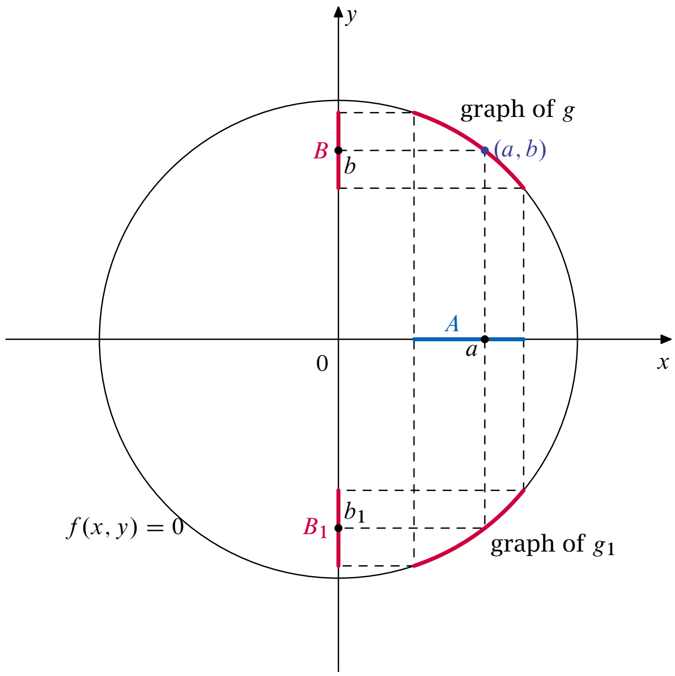

FIGURE 1.9. Implicit function theorem.

THEOREM 1.53 (Implicit Function Theorem). *Let $E \subset \mathbb{R}^{k+n}$ be open; let $f\colon E \to \mathbb{R}^n$ be of class $\mathscr{C}^r$. Write f in the form $f(\boldsymbol{x}, \boldsymbol{y})$ for $\boldsymbol{x} \in \mathbb{R}^k$ and $\boldsymbol{y} \in \mathbb{R}^n$. Suppose that $(\boldsymbol{a}, \boldsymbol{b})$ is a point of E such that $f(\boldsymbol{a}, \boldsymbol{b}) = \mathbf{0}$. Let $\mathbf{M}$ be the $n \times n$ matrix*

$$\mathbf{M} = \begin{bmatrix} \mathbb{D}_{k+1} f_1(\boldsymbol{a}, \boldsymbol{b}) & \mathbb{D}_{k+2} f_1(\boldsymbol{a}, \boldsymbol{b}) & \cdots & \mathbb{D}_{k+n} f_1(\boldsymbol{a}, \boldsymbol{b}) \\ \mathbb{D}_{k+1} f_2(\boldsymbol{a}, \boldsymbol{b}) & \mathbb{D}_{k+2} f_1(\boldsymbol{a}, \boldsymbol{b}) & \cdots & \mathbb{D}_{k+n} f_2(\boldsymbol{a}, \boldsymbol{b}) \\ \vdots & \vdots & \ddots & \vdots \\ \mathbb{D}_{k+1} f_n(\boldsymbol{a}, \boldsymbol{b}) & \mathbb{D}_{k+2} f_n(\boldsymbol{a}, \boldsymbol{b}) & \cdots & \mathbb{D}_{k+n} f_n(\boldsymbol{a}, \boldsymbol{b}) \end{bmatrix}.$$

If $\det(\mathbf{M}) \neq 0$, then there is a neighborhood A of $\boldsymbol{a} \in \mathbb{R}^k$ and a unique continuous function $g\colon A \to \mathbb{R}^n$ such that $g(\boldsymbol{a}) = \boldsymbol{b}$ and

$$f(\boldsymbol{x}, g(\boldsymbol{x})) = \mathbf{0}$$

for all $\boldsymbol{x} \in A$. The function g is in fact of class $\mathscr{C}^r$.

PROOF. The proof is too long to give here. You can find it from, e.g., Spivak (1965, Theorem 2-12), Rudin (1976, Theorem 9.28), or Munkres (1991, Theorem 2.9.2). □

EXAMPLE 1.54. Let $f: \mathbb{R}^2 \to \mathbb{R}$ be given by the equation

$$f(x,y) = x^2 - y^3.$$

Then $(0,0)$ is a solution of the equation $f(x,y) = 0$. Because $\partial f(0,0)/\partial y = 0$, we do not expect to be able to solve this equation for y in terms of x near $(0,0)$. But in fact, we can; and furthermore, the solution is unique! However, the function we obtain is not differentiable at $x = 0$. See Figure 1.10.

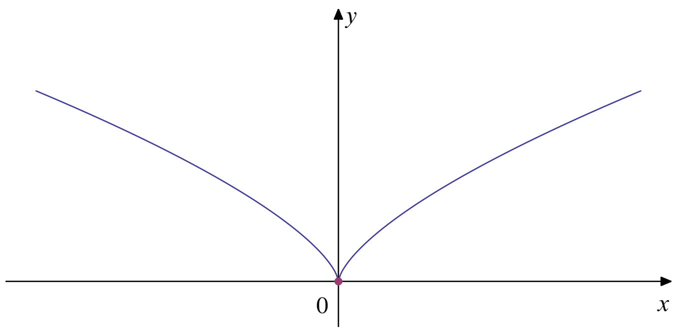

FIGURE 1.10. y is not differentiable at $x = 0$.

EXAMPLE 1.55. Let $f : \mathbb{R}^2 \to \mathbb{R}$ be given by the equation

$$f(x,y) = -x^4 + y^2.$$

Then $(0,0)$ is a solution of the equation $f(x,y) = 0$. Because $\partial f(0,0)/\partial y = 0$, we do not expect to be able to solve for y in terms of x near $(0,0)$. In fact, however, we can do so, and we can do so in such a way that the resulting function is differentiable. However, the solution is not unique. See Figure 1.11.

Now the point $(1,1)$ is also a solution to $f(x,y) = 0$. Because $\partial f(1,1)/\partial y = 2$, one can solve this equation for y as a continuous function of x in a neighborhood of $x = 1$. See Figure 1.11.

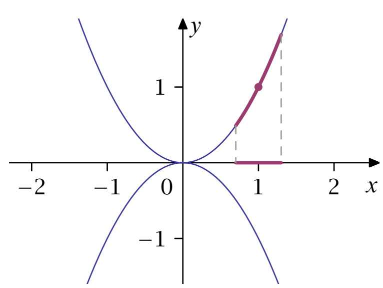

FIGURE 1.11. Example 1.55.

REMARK 1.56. We will use the Implicit Function Theorem in Theorem 2.10. The theorem will also be used to derive comparative statics for economic models, which we perhaps do not have time to discuss.

1.9 HOMOGENEOUS FUNCTIONS AND EULER'S FORMULA

DEFINITION 1.57 (Homogeneous Function). A function $f : \mathbb{R}^n \to \mathbb{R}$ is *homogeneous of degree r* (for $r = \ldots, -1, 0, 1, \ldots$) if for every $t > 0$ we have

$$f(tx_1, \ldots, tx_n) = t^r f(x_1, \ldots, x_n).$$

► EXERCISE 1.58. *The function*

$$f(x, y) = Ax^\alpha y^\beta, \quad A, \alpha, \beta > 0,$$

is known as the Cobb-Douglas *function. Check whether this function is homogeneous.*

THEOREM 1.59 (Euler's Formula). *Suppose that $f : \mathbb{R}^n \to \mathbb{R}$ is homogeneous of degree r (for some $r = \ldots, -1, 0, 1, \ldots$) and differentiable. Then at any $\boldsymbol{x}^* \in \mathbb{R}^n$ we have*

$$\nabla f(\boldsymbol{x}^*) \cdot \boldsymbol{x}^* = r f(\boldsymbol{x}^*).$$

PROOF. By definition we have

$$f(t\boldsymbol{x}^*) - t^r f(\boldsymbol{x}^*) = 0.$$

Differentiating with respect to t using the chain rule, we have

$$\nabla f(t\boldsymbol{x}^*) \cdot \boldsymbol{x}^* = r t^{r-1} f(\boldsymbol{x}^*).$$

Evaluating at $t = 1$ gives the desired result. □

LEMMA 1.60. *If f is homogeneous of degree r, its partial derivatives are homogeneous of degree $r - 1$.*

► EXERCISE 1.61. *Prove Lemma 1.60.*

► EXERCISE 1.62. *Let $f(x, y) = Ax^\alpha y^\beta$ with $\alpha + \beta = 1$ and $A > 0$. Show that Theorem 1.59 and Lemma 1.60 hold for this function.*

2

OPTIMIZATION IN $\mathbb{R}^N$

THIS CHAPTER IS based on Luenberger (1969, Chapters 8 & 9), Mas-Colell, Whinston and Green (1995, Sections M.J & M.K), Sundaram (1996), Duggan (2010), and Jehle and Reny (2011, Chapter A2).

2.1 INTRODUCTION

An *optimization problem in* $\mathbb{R}^n$, or simply an *optimization problem*, is one where the values of a given function $f : \mathbb{R}^n \to \mathbb{R}$ are to be maximized or minimized over a given set $X \subset \mathbb{R}^n$. The function f is called the *objective function* and the set X the *constraint set*. Notationally, we will represent these problems by

$$\text{Maximize } f(\boldsymbol{x}) \quad \text{subject to } \boldsymbol{x} \in X,$$

and

$$\text{Minimize } f(\boldsymbol{x}) \quad \text{subject to } \boldsymbol{x} \in X,$$

respectively. More compactly, we shall also write

$$\max\{f(\boldsymbol{x}) : \boldsymbol{x} \in X\} \quad \text{and} \quad \min\{f(\boldsymbol{x}) : \boldsymbol{x} \in X\}.$$

EXAMPLE 2.1. **(a)** Let $X = [0, \infty)$ and $f(x) = x$. Then the problem $\max\{f(x) : x \in X\}$ has no solution; see Figure 2.1(a).

(b) Let $X = [0, 1]$ and $f(x) = x(1-x)$. Then the problem $\max\{f(x) : x \in X\}$ has exactly one solution, namely $x = 1/2$; see Figure 2.1(b).

(c) Let $X = [-1, 1]$ and $f(x) = x^2$. Then the problem $\max\{f(x) : x \in X\}$ has two solutions, namely $x = -1$ and $x = 1$; see Figure 2.1(c).

◇

Example 2.1 suggests that we shall talk of the *set of solutions* of the optimization problem, which is denoted

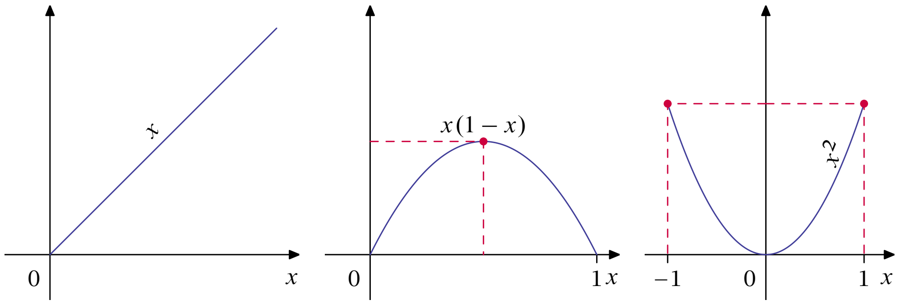

(a) No solution. (b) Exactly one solution. (c) Two solutions.

FIGURE 2.1. Example 2.1.

$$\operatorname{argmax}\{f(\boldsymbol{x}) : \boldsymbol{x} \in X\} = \{\boldsymbol{x} \in X : f(\boldsymbol{x}) \geqslant f(\boldsymbol{y}) \text{ for all } \boldsymbol{y} \in X\}.$$

We close this section by considering an optimization problem in economics.

EXAMPLE 2.2. There are n commodities in an economy. There is a consumer whose *utility* from consuming $x_i \geqslant 0$ units of commodity i ($i = 1, \dots, n$) is given by $u(x_1, \dots, x_n)$, where $u\colon \mathbb{R}^n_+ \to \mathbb{R}$ is the consumer's *utility function.* The consumer's income is $I \geqslant 0$, and faces the price vector $\boldsymbol{p} = (p_1, \dots, p_n)$. His budget set is given by (see Figure 2.2)

$$B(\boldsymbol{p}, I) := \{\boldsymbol{x} \in \mathbb{R}^n_+ : \boldsymbol{p} \cdot \boldsymbol{x} \leqslant I\}.$$

The consumer's objective is to maximize his utility over the budget set, i.e.,

$$\text{Maximize } u(\boldsymbol{x}) \quad \text{subject to } \boldsymbol{x} \in B(\boldsymbol{p}, I).$$

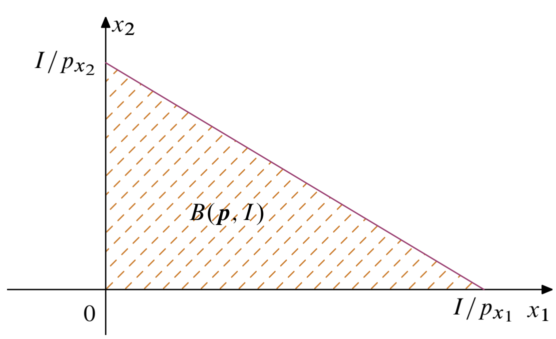

FIGURE 2.2. The budget set $B(p_{x_1}, p_{x_2}, I)$.

2.2 Unconstrained Optimization

Definition 2.3 (Maximizer). Given $X \subset \mathbb{R}^n$, $f : X \to \mathbb{R}$, $\boldsymbol{x} \in X$, we say $\boldsymbol{x}$ is a *maximizer* of f if

$$f(\boldsymbol{x}) = \max\{f(\boldsymbol{y}) : \boldsymbol{y} \in X\}.$$

We say $\boldsymbol{x}$ is a *local maximizer* of f if there is some $\varepsilon > 0$ such that for all $\boldsymbol{y} \in X \cap \mathcal{B}(\boldsymbol{x};\varepsilon)$ we have $f(\boldsymbol{x}) \geqslant f(\boldsymbol{y})$. And $\boldsymbol{x}$ is a *strict local maximizer* of f if the latter inequality holds strictly.

First-Order Analysis

Recall that X° is the *interior* of $X \subset \mathbb{R}^n$ (Definition 1.13), and $f'(\boldsymbol{x};\boldsymbol{u})$ is the *directional derivative* of f at $\boldsymbol{x}$ with respect to $\boldsymbol{u}$ (Definition 1.21).

Theorem 2.4. *Let $X \subset \mathbb{R}^n$; let $\boldsymbol{x} \in X^\circ$; let $f : X \to \mathbb{R}$ be differentiable at $\boldsymbol{x}$. If $\boldsymbol{x}$ is a local maximizer of f, then for every direction $\boldsymbol{u} \in X$ we have $f'(\boldsymbol{x};\boldsymbol{u}) = 0$.*

Proof. Suppose that $\boldsymbol{x}$ is an interior local maximizer and let $\boldsymbol{u} \in X$. Take $\varepsilon > 0$ such that $\mathcal{B}(\boldsymbol{x};\varepsilon) \subset X$ and $f(\boldsymbol{x}) \geqslant f(\boldsymbol{y})$ for all $\boldsymbol{y} \in \mathcal{B}(\boldsymbol{x};\varepsilon)$. In particular, $f(\boldsymbol{x}) \geqslant f(\boldsymbol{x} + \alpha\boldsymbol{u})$ for $\alpha \in \mathbb{R}$ small. Since f is differentiable at $\boldsymbol{x}$, we have

$$f'(\boldsymbol{x};\boldsymbol{u}) = \lim_{\alpha \downarrow 0} \frac{f(\boldsymbol{x} + \alpha\boldsymbol{u}) - f(\boldsymbol{x})}{\alpha} \leqslant 0,$$

and

$$f'(\boldsymbol{x};\boldsymbol{u}) = \lim_{\alpha \uparrow 0} \frac{f(\boldsymbol{x} + \alpha\boldsymbol{u}) - f(\boldsymbol{x})}{\alpha} \geqslant 0.$$

Therefore, $f'(\boldsymbol{x};\boldsymbol{u}) = 0$, as claimed. □

Remark 2.5. If f is differentiable at $\boldsymbol{x}$ and $\boldsymbol{x}$ is an interior local maximizer of f, then since $f'(\boldsymbol{x};\boldsymbol{u}) = \nabla f(\boldsymbol{x}) \cdot \boldsymbol{u}$ (Theorem 1.28, Theorem 1.33 and Definition 1.39), we know that for all $\boldsymbol{u}$

$$\nabla f(\boldsymbol{x}) \cdot \boldsymbol{u} = 0,$$

which implies that $\nabla f(\boldsymbol{x}) = \boldsymbol{0}$.

Definition 2.6 (Critical Point). A vector $\boldsymbol{x} \in \mathbb{R}^n$ such that $\nabla f(\boldsymbol{x}) = 0$ is called a *critical point*.

Example 2.7. Let $X = \mathbb{R}^2$ and $f(x,y) = xy - 2x^4 - y^2$. The first order condition is

$$\nabla f(x,y) = (y - 8x^3, x - 2y) = (0,0).$$

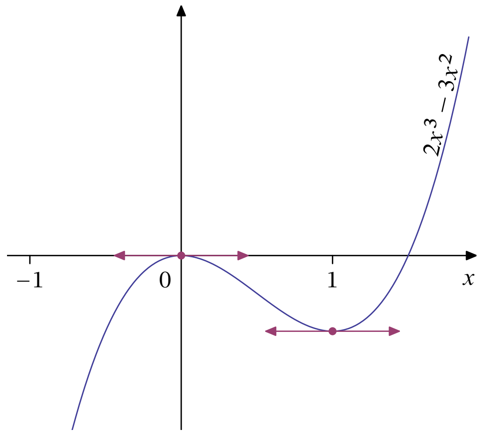

FIGURE 2.3. $x = 0$ and $x = 1$ are local optima but not global optima.

Thus, the critical points are $(x, y) = (0, 0), (1/4, 1/8), (-1/4, -1/8)$.

Second-Order Analysis

The first-order conditions for unconstrained local optima do not distinguish between *maxima* and *minima* (see the following Example 2.9). To obtain such a distinction in the behavior of f at an optimum, we need to examine the behavior of the Hessian $\mathbf{H}f$ of f (see Definition 1.46).

THEOREM 2.8. *Suppose f is of class $\mathscr{C}^2$ on $X \subset \mathbb{R}^n$, and $\boldsymbol{x} \in X^\circ$.*

a. *If f has a local maximum at $\boldsymbol{x}$, then $\mathbf{H}f(\boldsymbol{x})$ is negative semidefinite.*

b. *If f has a local minimum at $\boldsymbol{x}$, then $\mathbf{H}f(\boldsymbol{x})$ is positive semidefinite.*

c. *If $\nabla f(\boldsymbol{x}) = 0$ and $\mathbf{H}f(\boldsymbol{x})$ is negative definite at some $\boldsymbol{x}$, then $\boldsymbol{x}$ is a strict local maximum of f on X.*

d. *If $\nabla f(\boldsymbol{x}) = 0$ and $\mathbf{H}f(\boldsymbol{x})$ is positive definite at some $\boldsymbol{x}$, then $\boldsymbol{x}$ is a strict local minimum of f on X.*

PROOF. See Sundaram (1996, Section 4.6). □

EXAMPLE 2.9. Let $f : \mathbb{R} \to \mathbb{R}$ be defined by $f(x) = 2x^3 - 3x^2$. It is easy to check that $f \in \mathscr{C}^2$ on $\mathbb{R}$ and there are two critical points: $x = 0$ and $x = 1$. Invoking the second-order conditions, we get $f''(0) = -6$ and $f''(1) = 6$. Thus, the point $x = 0$ is a strict local maximum of f on $\mathbb{R}$, and the point $x = 1$ is a strict local minimum of f on $\mathbb{R}$; see Figure 2.3.

However, there is nothing in the first- or second-order conditions that will help determine whether these points are *global* optima. In fact, they are not: global optima do not exist in this example, since $\lim_{x\to+\infty} f(x) = +\infty$ and $\lim_{x\to-\infty} f(x) = -\infty$.

2.3 Equality Constrained Optimization: Lagrange's Method

First-Order Analysis

Lagrange's Theorem

Theorem 2.10 (Lagrange's Theorem). *Let $f : \mathbb{R}^n \to \mathbb{R}$, and $g_i : \mathbb{R}^n \to \mathbb{R}$ be $\mathscr{C}^1$ functions, where $i = 1, \ldots, k$. Suppose that $\boldsymbol{x}^*$ is a local maximizer or minimizer of f on the set*

$$X := U \cap \left\{\boldsymbol{x} \in \mathbb{R}^n : g_i(\boldsymbol{x}) = 0, i = 1, \ldots, k\right\},$$

where $U \subset \mathbb{R}^n$ is open. Suppose also that the list of vectors $(\nabla g_1(\boldsymbol{x}^*), \ldots, \nabla g_k(\boldsymbol{x}^*))$ is linearly independent (this is called the constraint qualification*). Then, there exists a vector $\boldsymbol{\lambda}^* = (\lambda_1^*, \ldots, \lambda_k^*) \in \mathbb{R}^k$ such that*

$$\nabla f(\boldsymbol{x}^*) = \sum_{i=1}^{k} \lambda_i^* \cdot \nabla g_i(\boldsymbol{x}^*).$$

Here we give a proof for the case of *two variables* and *one constraint* (This proof is from Duggan 2010, Theorem 5.1). We are interested in this proof partly because we will use the Implicit Function Theorem (Theorem 1.53). For a general proof, see Sundaram (1996, Section 5.6).

Proof (Two variables, one constraint). We show that if

$$\boldsymbol{x}^* \in \operatorname{argmax}\{f(\boldsymbol{x}) : g(\boldsymbol{x}) = 0\},$$

and $\nabla g(\boldsymbol{x}^*) \neq \boldsymbol{0}$, then there exists $\lambda \in \mathbb{R}$ such that

$$\nabla f(\boldsymbol{x}^*) = \lambda \nabla g(\boldsymbol{x}^*). \tag{2.1}$$

Without loss of generality, we assume that $\boldsymbol{x}^* = \boldsymbol{0}$ and $\mathbb{D}_2 g(\boldsymbol{x}^*) \neq 0$. The Implicit Function Theorem (Theorem 1.53) implies that in an open interval I around $x_1^* = 0$, we may then view the level set $L_g(0)$ as the graph of a function $\varphi : I \to \mathbb{R}$ such that for all $z \in I$ we have

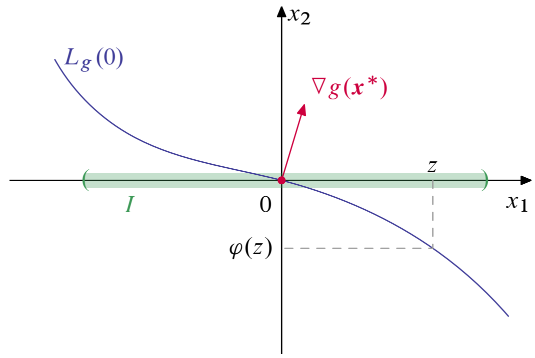

FIGURE 2.4. Proof of Lagrange's Theorem.

$$g(z, \varphi(z)) = 0. \tag{2.2}$$

See Figure 2.4. Notice that

$$\mathbf{0} = \boldsymbol{x}^* = (0, \varphi(0)).$$

Furthermore, φ is continuously differentiable with derivative (by (2.2))

$$\varphi'(z) = -\frac{\mathbb{D}_1 g(z, \varphi(z))}{\mathbb{D}_2 g(z, \varphi(z))}. \tag{2.3}$$

Because $\boldsymbol{x}^* \in U$ and U is open, we can choose the interval I small enough that each $(z, \varphi(z)) \in U$. Then $z = 0$ is a local maximizer of the unconstrained problem

$$\max_{z \in I} f(z, \varphi(z)).$$

Then, by the first-order condition, we have

$$\mathbb{D}_1 f(\mathbf{0}) + \mathbb{D}_2 f(\mathbf{0}) \cdot \varphi'(0) = 0,$$

which implies (by (2.3))

$$\mathbb{D}_1 f(\mathbf{0}) - \mathbb{D}_2 f(\mathbf{0}) \cdot \frac{\mathbb{D}_1 g(\mathbf{0})}{\mathbb{D}_2 g(\mathbf{0})} = 0. \tag{2.4}$$

Defining

$$\lambda = \frac{\mathbb{D}_2 f(\mathbf{0})}{\mathbb{D}_2 g(\mathbf{0})},$$

we have

$$\begin{aligned}
\lambda \nabla g(\mathbf{0}) &= \frac{\mathbb{D}_2 f(\mathbf{0})}{\mathbb{D}_2 g(\mathbf{0})} \cdot \begin{bmatrix} \mathbb{D}_1 g(\mathbf{0}) & \mathbb{D}_2 g(\mathbf{0}) \end{bmatrix} \\
&= \begin{bmatrix} \frac{\mathbb{D}_2 f(\mathbf{0}) \mathbb{D}_1 g(\mathbf{0})}{\mathbb{D}_2 g(\mathbf{0})} & \mathbb{D}_2 f(\mathbf{0}) \end{bmatrix} \\
&= \begin{bmatrix} \mathbb{D}_1 f(\mathbf{0}) & \mathbb{D}_2 f(\mathbf{0}) \end{bmatrix} \\
&= \nabla f(\mathbf{0}),
\end{aligned}$$

where the third equality follows from (2.4). □

Geometric Interpretation

Let us consider an optimization problem on $\mathbb{R}^2$ and with one constraint:

$$\max_{x_1, x_2} f(x_1, x_2) \quad \text{subject to } g(x_1, x_2) = 0. \tag{2.5}$$

Consider the level sets $L_f(c)$ and $L_g(0)$. Recall that (see Section 1.4) for every point $\boldsymbol{x} \in L_f(c)$, we have

$$\left.\frac{\mathrm{d}x_2}{\mathrm{d}x_1}\right|_{\text{along } L_f(c)} = -\frac{\mathbb{D}_1 f(\boldsymbol{x})}{\mathbb{D}_2 f(\boldsymbol{x})}, \tag{2.6}$$

and for every point $\boldsymbol{y} \in L_g(0)$ we have

$$\left.\frac{\mathrm{d}y_2}{\mathrm{d}y_1}\right|_{\text{along } L_g(0)} = -\frac{\mathbb{D}_1 g(\boldsymbol{y})}{\mathbb{D}_2 g(\boldsymbol{y})}. \tag{2.7}$$

It follows from Theorem 2.10 that if $\boldsymbol{x}^*$ is a solution to (2.5), then

$$\begin{aligned}
\mathbb{D}_1 f(\boldsymbol{x}^*) &= \lambda^* \mathbb{D}_1 g(\boldsymbol{x}^*), \\
\mathbb{D}_2 f(\boldsymbol{x}^*) &= \lambda^* \mathbb{D}_2 g(\boldsymbol{x}^*), \\
g(\boldsymbol{x}^*) &= 0.
\end{aligned}$$

Suppose that $\lambda^* \neq 0$. Then we can rewrite the above conditions as follows:

$$\frac{\mathbb{D}_1 f(\boldsymbol{x}^*)}{\mathbb{D}_2 f(\boldsymbol{x}^*)} = \frac{\mathbb{D}_1 g(\boldsymbol{x}^*)}{\mathbb{D}_2 g(\boldsymbol{x}^*)} \tag{2.8}$$

$$g(\boldsymbol{x}^*) = 0. \tag{2.9}$$

The first equation (2.8) says that solution values of x_1 and x_2 will be at a point where the slope of the level set for the objective function and the slope of the level set for the constraint are equal. The second equation (2.9) tells us we must also be on the level set of the constraint equation. See Figure 2.5.

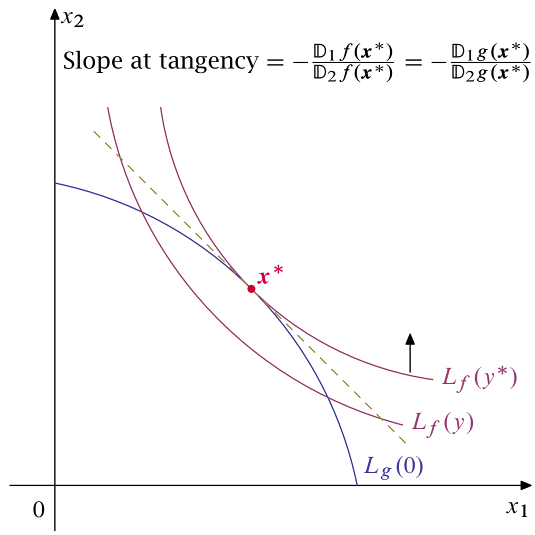

FIGURE 2.5. The first-order conditions for a solution to Lagrange's problem identify a point of tangency between a level set of the objective function and the constraint.

The Constraint Qualification

We show that Theorem 2.10 fails without the constraint qualification.

EXAMPLE 2.11. Let $X = \mathbb{R}$, $f(x) = (x+1)^2$ and $g(x) = x^2$. Consider the problem of maximizing f subject to $g(x) = 0$. The maximizer is clearly $x = 0$. But $\mathbb{D}g(0) = 0$ and $\mathbb{D}f(0) = 2$, so there is no λ such that $\mathbb{D}f(0) = \lambda\mathbb{D}g(0)$.

The Lagrange Multipliers

The vector $\boldsymbol{\lambda}^* = (\lambda_1^*, \dots, \lambda_k^*)$ in Theorem 2.10 is called the vector of *Lagrange multipliers* corresponding to the local optimum $\boldsymbol{x}^*$. The i^{th} multiplier λ_i^* measures the sensitivity of the value of the objective function at $\boldsymbol{x}^*$ to a small relaxation of the i^{th} constraint g_i.

Physical Interpretation of Gradient ⋆

We now can provide a physical interpretation of gradient (Shastri, 2011). Consider any linear functional $\varphi\colon \mathbb{R}^n \to \mathbb{R}$ defined by $\varphi(\boldsymbol{x}) = \sum_{i=1}^n \alpha_i x_i$, and the problem of finding its maxima on the unit sphere $\mathbb{S}^{n-1} := \{\boldsymbol{x} \in \mathbb{R}^n : \|\boldsymbol{x}\| = 1\}$. The Lagrange multiplier function in this case is

$$\mathcal{L} = \sum_{i=1}^n \alpha_i x_i - \lambda\left(\sum_{i=1}^n x_i^2 - 1\right).$$

Thus,

$$\alpha_i = 2\lambda x_i \quad \text{for all } i = 1, \dots, n.$$

Now fix $i \in \{1, \dots, n\}$, and we have

$$x_j = \alpha_j \frac{x_i}{\alpha_i} \quad \text{for all } j = 1, \dots, n.$$

Therefore, for each $i \in \{1, \dots, n\}$

$$x_i = \pm \frac{\alpha_i}{\left(\sum_{i=1}^n \alpha_i^2\right)^{1/2}}. \tag{2.10}$$

Now let $f : U \to \mathbb{R}$ be a $\mathscr{C}^1$ function in a neighborhood U of $\mathbf{0} \in \mathbb{R}^n$. Let $\varphi(\boldsymbol{x}) = \mathbb{D} f(\mathbf{0}) \cdot \boldsymbol{x}$. To each $\boldsymbol{v} \in \mathbb{S}^{n-1}$ we can consider the path $(-\varepsilon, \varepsilon) \to U$ given by $t \mapsto t\boldsymbol{v}$ and look at the function $t \mapsto f(t\boldsymbol{v})$. The derivative of this map at 0 is nothing by $\mathbb{D} f(\mathbf{0}) \cdot \boldsymbol{v}$. Therefore, form (2.10), it follows that the extrema of the function

$$\boldsymbol{v} \mapsto \mathbb{D} f(\mathbf{0}) \cdot \boldsymbol{v} = \sum_{i=1}^n \mathbb{D}_i f(\mathbf{0}) \cdot v_i$$

occur at $\pm \nabla f(\mathbf{0}) / \|\nabla f(\mathbf{0})\|$. Thus, ∇f is the direction in which the increment in f is the maximum.

Lagrange's Method

Let an equality-constrained optimization problem of the form

$$\text{Maximize } f(\boldsymbol{x}) \quad \text{subject to } \boldsymbol{x} \in X = U \cap \left\{\boldsymbol{x} \in \mathbb{R}^n : g(\boldsymbol{x}) = \mathbf{0}\right\}, \tag{2.11}$$

be give, where $f : \mathbb{R}^n \to \mathbb{R}$ and $g : \mathbb{R}^n \to \mathbb{R}^k$ are of $\mathscr{C}^1$ functions, and $U \subset \mathbb{R}^n$ is open. We describe a procedure for using Theorem 2.10 to solve (2.11).

Step 1. Set up a function $\mathscr{L} : X \times \mathbb{R}^k \to \mathbb{R}$, called the *Lagrangian*, defined by

$$\mathscr{L}(\boldsymbol{x}, \boldsymbol{\lambda}) = f(\boldsymbol{x}) - \sum_{i=1}^k \lambda_i g_i(\boldsymbol{x}).$$

The vector $\boldsymbol{\lambda} = (\lambda_1, \dots, \lambda_k) \in \mathbb{R}^k$ is called the vector of *Lagrange multipliers*.

Step 2. Find all critical points of $\mathscr{L}(\boldsymbol{x}, \boldsymbol{\lambda})$:

$$\frac{\partial \mathscr{L}}{\partial x_i}(\boldsymbol{x}, \boldsymbol{\lambda}) = \mathbb{D}_i f(\boldsymbol{x}) - \sum_{\ell=1}^k \lambda_\ell \mathbb{D}_i g_\ell(\boldsymbol{x}) = 0, \quad i = 1, \dots, n \tag{2.12}$$

$$\frac{\partial \mathscr{L}}{\partial \lambda_j}(\boldsymbol{x}, \boldsymbol{\lambda}) = g_j(\boldsymbol{x}) = 0, \quad j = 1, \dots, k. \tag{2.13}$$

Define

$$M := \left\{ (\boldsymbol{x}, \boldsymbol{\lambda}) : \boldsymbol{x} \in U, \text{ and } (\boldsymbol{x}, \boldsymbol{\lambda}) \text{ satisfies (2.12) and (2.13)} \right\}.$$

Step 3. Evaluate f at each point $\boldsymbol{x}$ in the set

$$\left\{ \boldsymbol{x} \in \mathbb{R}^n : \exists\, \boldsymbol{\lambda} \in \mathbb{R}^k \text{ such that } (\boldsymbol{x}, \boldsymbol{\lambda}) \in M \right\}.$$

Thus, we see that Lagrange's method is a clever way of converting a maximization problem with constraints, to another maximization problem without constraint, by increasing the number of variables.

Why the Lagrange's Method typically succeeds in identifying the desired optima? This is because the set of all critical points of $\mathcal{L}$ contains the set of all local maximizers and minimizers of the objective function f on X at which the constraint qualification is met. That is, if $\boldsymbol{x}^*$ is a maximizer or minimizer of f on X, and if the constraint qualification holds at $\boldsymbol{x}^*$, then there exists $\boldsymbol{\lambda}^*$ such that $(\boldsymbol{x}^*, \boldsymbol{\lambda}^*)$ is a critical point of $\mathcal{L}$.

We are not going to explain why the Lagrange's method could fail (but see Sundaram 1996, Section 5.4 for details).

EXAMPLE 2.12. Consider the problem

$$\max_{(x,y)\in\mathbb{R}^2} \left\{ f(x,y) = -x^2 - y^2 \right\} \quad \text{subject to } g(x,y) = x + y - 1 = 0.$$

First, form the Lagrangian,

$$\mathcal{L}(x, y, \lambda) = -x^2 - y^2 - \lambda(x + y - 1).$$

Then set all of its first-order partials equal to zero:

$$\frac{\partial \mathcal{L}}{\partial x} = -2x - \lambda = 0$$
$$\frac{\partial \mathcal{L}}{\partial y} = -2y - \lambda = 0$$
$$\frac{\partial \mathcal{L}}{\partial \lambda} = x + y - 1 = 0.$$

So the critical points of $\mathcal{L}$ is

$$(x^*, y^*, \lambda^*) = (1/2, 1/2, -1).$$

Hence, $f(x^*, y^*) = -1/2$; see Figure 2.6.

► EXERCISE 2.13. *A consumer purchases a bundle (x, y) to maximize utility. His income is $I > 0$ and prices are $p_x > 0$ and $p_y > 0$. His utility function is*

$$u(x, y) = x^a y^b,$$

where $a, b > 0$. Find his optimal choice (x^*, y^*).

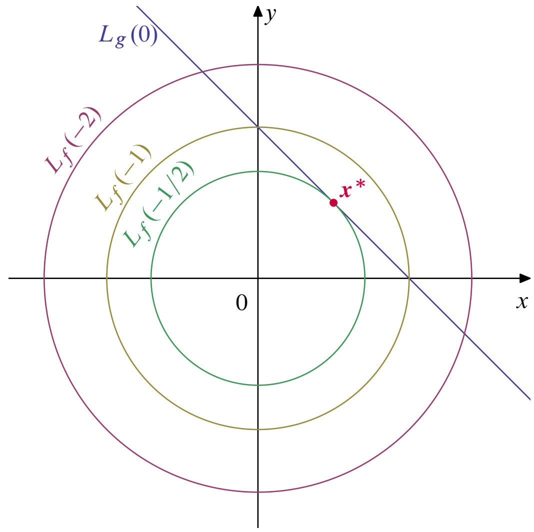

FIGURE 2.6. Lagrange's method.

Lagrange's Theorem Is Not Sufficient

Lagrange's Theorem (Theorem 2.10) only gives us a *necessary*—not a *sufficient*—condition for a constrained local maximizer. To see why the first order condition is not generally sufficient, consider the following example.

EXAMPLE 2.14. Let $U = \mathbb{R}^2$, $f(x, y) = x + y^2$, and $g(x, y) = x - 1$. Consider the problem

$$\max_{(x,y)\in\mathbb{R}^2} f(x, y)$$
$$\text{s.t.} \quad g(x, y) = 0.$$

Observe that $(x^*, y^*) = (1, 0)$ satisfies the constraint $g(x^*, y^*) = 0$, and the constraint qualification is also satisfied. Furthermore, the first-order condition from Lagrange's Theorem is satisfied at $(x^*, y^*) = (1, 0)$. This is because

$$\nabla f(1, 0) = (1, 0) \quad \text{and} \quad \nabla g(1, 0) = (1, 0).$$

Hence, by letting $\lambda = 1$ we have $\nabla f(1, 0) = \lambda \nabla g(1, 0)$.

However, $(1, 0)$ is NOT a constrained local maximizer: for arbitrarily small $\varepsilon > 0$, we have $g(1, \varepsilon) = 0$ and $f(1, \varepsilon) = 1 + \varepsilon^2 > 1 = f(1, 0)$. See Figure 2.7.

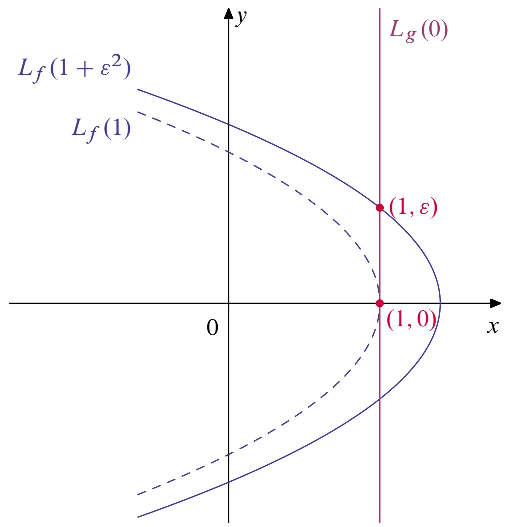

FIGURE 2.7. The Lagrange's Theorem is not sufficient.

Second-Order Analysis

We probably do not have time to discus the second-order conditions. See Jehle and Reny (2011, Section A2.3.4) and Sundaram (1996, Section 5.3).

2.4 INEQUALITY CONSTRAINED OPTIMIZATION: KUHN-TUCKER THEOREM

We now consider maximization defined by *inequality constraints*. The constraint set will now be

$$X = U \cap \left\{ \boldsymbol{x} \in \mathbb{R}^n : h_i(\boldsymbol{x}) \leqslant 0, i = 1, \ldots, \ell \right\},$$

where $U \subset \mathbb{R}^n$ is open, and $h_i : \mathbb{R}^n \to \mathbb{R}$ for every $i = 1, \ldots, \ell$. Given $\boldsymbol{x} \in \mathbb{R}^n$, we say the i^{th} constraint is *binding* if $h_i(\boldsymbol{x}) = 0$, and *slack* if $h_i(\boldsymbol{x}) < 0$.

First-Order Analysis

EXAMPLE 2.15. Figure 2.8 illustrates a problem with two inequality constraints and depicts three possibilities, depending on whether none, one, or two constraints are binding.

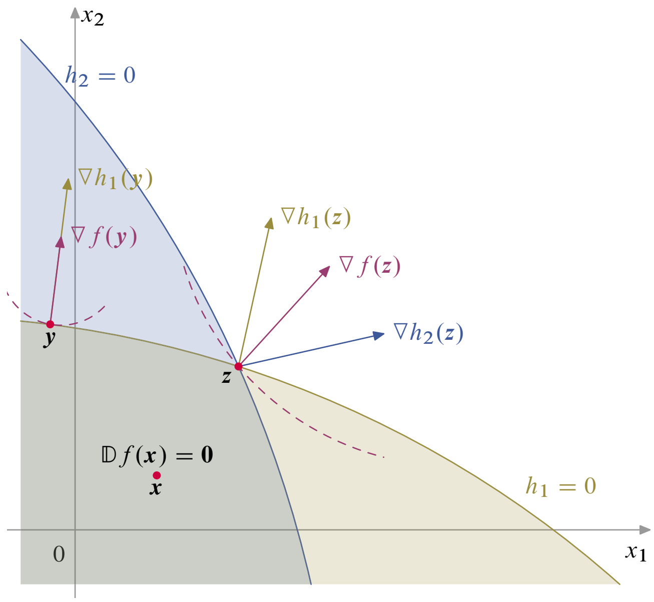

FIGURE 2.8. Inequality constrained optimization.

- In the first case, we could have a constrained local maximizer such as $\boldsymbol{x}$, for which no constraints bind. Such a vector must be a critical point of the objective function.
- In the second case, we could have a single constraint binding at a constrained local maximizer such as $\boldsymbol{y}$, and here the gradients of the objective and constraint are collinear. As we will see, these gradients actually point in the same direction.
- Lastly, we could have a constrained local maximizer such as $\boldsymbol{z}$, where both constraints bind. Here, the gradient of the objective is not collinear with the gradient of either constraint, and it may appear that no gradient restriction is possible. But in fact, $\nabla f(\boldsymbol{z})$ can be written as a linear combination of $\nabla h_1(\boldsymbol{z})$ and $\nabla h_2(\boldsymbol{z})$ with non-negative weights.

The restrictions evident in Figure 2.8 are formalized in the next theorem.

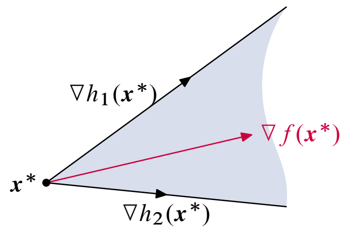

FIGURE 2.9. Kuhn-Tucker Theorem.

THEOREM 2.16 (Kuhn-Tucker Theorem). *Let $f : \mathbb{R}^n \to \mathbb{R}$ and $h_i : \mathbb{R}^n \to \mathbb{R}$ be of $\mathscr{C}^1$ class functions, $i = 1, \ldots, \ell$. Suppose $\boldsymbol{x}^*$ is a local maximizer of f on*

$$X = U \cap \left\{\boldsymbol{x} \in \mathbb{R}^n : h_i(\boldsymbol{x}) \leqslant 0, i = 1, \ldots, \ell\right\},$$

where U is an open subset in $\mathbb{R}^n$. Suppose that the first k, where $k \leqslant \ell$, constraints are the binding ones at $\boldsymbol{x}^*$, and assume the gradients of the binding constraints, $\{\nabla h_1(\boldsymbol{x}^*), \ldots, \nabla h_k(\boldsymbol{x}^*)\}$, are linear independent (this is called the constraint qualification*). Then there exists a vector $\boldsymbol{\lambda}^* = (\lambda_1^*, \ldots, \lambda_\ell^*) \in \mathbb{R}^\ell$ such that*

$$\lambda_i^* \geqslant 0 \quad \textit{and} \quad \lambda_i^* h_i(\boldsymbol{x}^*) = 0 \quad \textit{for all } i = 1, \ldots, \ell, \tag{KT-1}$$

$$\nabla f(\boldsymbol{x}^*) = \sum_{i=1}^{\ell} \lambda_i \nabla h_i(\boldsymbol{x}^*). \tag{KT-2}$$

PROOF. See Sundaram (1996, Section 6.5). □

REMARK 2.17. Geometrically, the first-order condition from the Kuhn-Tucker Theorem means that the gradient of the objective function, $\nabla f(\boldsymbol{x}^*)$, is contained in the "semi-positive cone" generated by the gradients of binding constraints, i.e., it is contained in the set

$$\left\{\sum_{i=1}^{\ell} \alpha_i \nabla h_i(\boldsymbol{x}^*) : \alpha_1, \ldots, \alpha_\ell \geqslant 0\right\},$$

depicted in Figure 2.9.

REMARK 2.18. Condition (KT-1) in Theorem 2.16 is called the condition of *complementary slackness*: if $h_i(\boldsymbol{x}^*) < 0$ then $\lambda_i^* = 0$; if $\lambda_i^* > 0$ then $h_i(\boldsymbol{x}^*) = 0$.

The Kuhn-Tucker Multipliers

The vector $\boldsymbol{\lambda}^*$ in Theorem 2.16 is called the vector of *Kuhn-Tucker multipliers* corresponding to the local maximizer $\boldsymbol{x}^*$. The Kuhn-Tucker multipliers measure the sensitivity of the objective function at $\boldsymbol{x}^*$ to relaxations of the various constraints:

- If $h_i(\boldsymbol{x}^*) < 0$, then the i^{th} constraint is already slack, so relaxing it further will not help raise the value of the objective function, and λ_i^* must be zero.
- If $h_i(\boldsymbol{x}^*) = 0$, then relaxing the i^{th} constraint may help increase the value of the maximization exercise, so we have $\lambda_i^* \geqslant 0$.

Two Differences

There are two important differences from the case of equality constraints (see Theorem 2.10 and Theorem 2.16):

- The constraint qualification now holds only for the gradients of binding constraints. (With equality constraints, every constraint is binding, but now some may not be.)
- The multipliers are non-negative. This difference comes from the fact that now only the inequality $h_i(\boldsymbol{x}) \leqslant 0$ needs to be maintained, so relaxing the constraint never hurts.

The Constraint Qualification

As with the analogous condition in Theorem 2.10 (see Example 2.11), here we show that the constraint qualification in Theorem 2.16 is essential.

EXAMPLE 2.19. Consider the following maximization problem

$$\begin{aligned}
&\max\{f(x,y) = x\}\\
\text{s.t. } &h_1(x,y) = -(1-x)^3 + y \leqslant 0\\
&h_2(x,y) = -x \leqslant 0\\
&h_3(x,y) = -y \leqslant 0.
\end{aligned}$$

See Figure 2.10. Clearly the solution is $(x^*, y^*) = (1,0)$. At this point we have

$$\nabla h_1(1,0) = (0,1), \quad \nabla h_3(1,0) = (0,-1) \quad \text{and} \quad \nabla f(1,0) = (1,0).$$

Since $x^* > 0$, it follows from the complementary slackness condition (KT-1) that $\lambda_2^* = 0$. But now (KT-2) fails: for any $\lambda_1 \geqslant 0$ and $\lambda_3 \geqslant 0$, we have

$$\lambda_1 \nabla h_1(1,0) + \lambda_3 \nabla h_3(1,0) = (0, \lambda_1 - \lambda_3) \neq \nabla f(1,0).$$

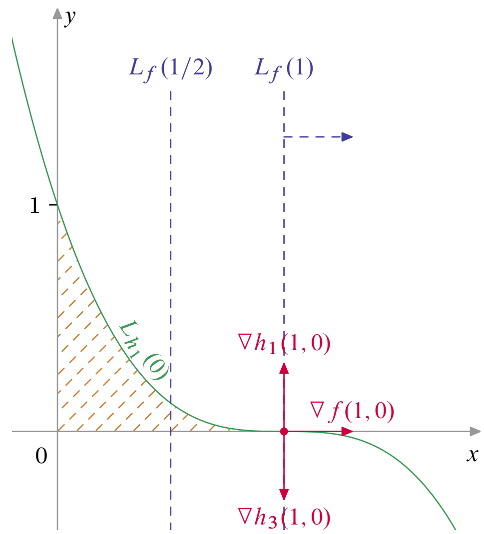

FIGURE 2.10. The constraint qualification fails at $(1,0)$.

This is because the constraint qualification fails at $(1,0)$: There are two binding constraints at $(x^*, y^*) = (1,0)$, namely h_1 and h_3, and the gradients $\nabla h_1(1,0)$ and $\nabla h_3(1,0)$ are colinear. Certainly $\nabla f(1,0)$ cannot be contained in the cone generated by $\nabla h_1(1,0)$ and $\nabla h_3(1,0)$; see Remark 2.17.

The Lagrangian

As with equality constraints, we can define the Lagrangian $\mathcal{L}\colon \mathbb{R}^n \times \mathbb{R}^\ell \to \mathbb{R}$ by

$$\mathcal{L}(\boldsymbol{x}, \boldsymbol{\lambda}) = f(\boldsymbol{x}) - \sum_{i=1}^{\ell} \lambda_i h_i(\boldsymbol{x}),$$

and then condition (KT-2) from Theorem 2.16 is the requirement that $\boldsymbol{x}$ is a critical point of the Lagrangian given multipliers $\lambda_1, \dots, \lambda_\ell$.

Let us consider a numerical example.

EXAMPLE 2.20. Consider the problem

$$\max \left\{ f(x,y) = x^2 - y \right\}$$
$$\text{s.t. } h(x,y) = x^2 + y^2 - 1 \leqslant 0.$$

Set the Lagrangian:

$$\mathcal{L}(x,y,\lambda) = x^2 - y - \lambda(x^2 + y^2 - 1).$$

The critical points of $\mathcal{L}$ are the solutions (x, y, λ) to

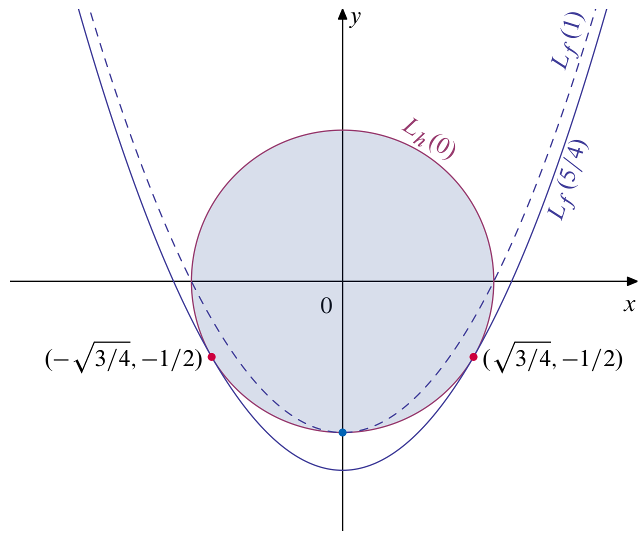

FIGURE 2.11. A numerical example.

$$2x - 2\lambda x = 0 \tag{2.14}$$

$$-1 - 2\lambda y = 0 \tag{2.15}$$

$$\lambda \geqslant 0 \tag{2.16}$$

$$x^2 + y^2 - 1 \leqslant 0 \tag{2.17}$$

$$\lambda(x^2 + y^2 - 1) = 0. \tag{2.18}$$

For (2.14) to hold, we must have $x = 0$ or $\lambda = 1$.

- If $\lambda = 1$, then from (2.15) we have $y = -1/2$, and from (2.18) we have $x = \pm\sqrt{3}/2$. That is,
$$(x, y, \lambda) = \left(\pm\frac{\sqrt{3}}{2}, -\frac{1}{2}, 1\right).$$
We thus have $f(x, y) = 5/4$; see Figure 2.11.

- Now suppose that $x = 0$. Then $\lambda > 0$ by (2.15). Hence h is binding: $x^2 + y^2 - 1 = 0$, and so $y = \pm 1$. Since (2.15) implies that $y = 1$ is impossible, we have
$$(x, y, \lambda) = \left(0, -1, \frac{1}{2}\right).$$
At this critical point, we have $f(0, -1) = 1 < 5/4$, which means that $(0, -1, 1/2)$ cannot be a solution. Since there are no other critical points, it follows that there are exactly two solutions to the maximization problem, namely $(x^*, y^*) = (\pm\sqrt{3}/2, -1/2)$.

► EXERCISE 2.21. *Let* $U = \mathbb{R}^2$*; let*

$$f(x, y) = (x-1)^2 + y^2,$$

and let

$$h(x, y) = 2kx - y^2 \leqslant 0, \quad k > 0.$$

Solve the maximization problem

$$\max\{f(x, y) : h(x, y) \leqslant 0\}.$$

EXAMPLE 2.22. Consider a consumer's problem:

$$\begin{aligned} \max_{(x,y)\in\mathbb{R}^2} \ & \{u(x, y) = x + y\} \\ \text{s.t.} \quad & h_1(x, y) = -x \leqslant 0 \\ & h_2(x, y) = -y \leqslant 0 \\ & h_3(x, y) = p_x x + p_y y - I \leqslant 0, \end{aligned} \tag{2.19}$$

where $p_x, p_y, I > 0$.

We first identify all possible combinations of constraints that can, in principle, be binding at the optimum. There are eight combinations to be check:

$$\varnothing, h_1, h_2, h_3, (h_1, h_2), (h_1, h_3), (h_2, h_3), \text{ and } (h_1, h_2, h_3).$$

Of these, the last one can be ruled out, since $h_1 = h_2 = 0$ implies that $h_3 < 0$. Moreover, since u is strictly increasing in both arguments, it is obvious that $h_3 = 0$. So we only need to check three combinations: (h_1, h_3), (h_2, h_3), and h_3.

- If the optimum occurs at a point where only h_1 and h_3 are binding, then

  $$\left[\nabla h_1(x, y), \nabla h_3(x, y)\right] = \left[(-1, 0), (p_x, p_y)\right]$$

  is linear independent. So the constraint qualification holds at such a point.

- Similarly, the constraint qualification holds if only (h_2, h_3) bind.

- If h_3 is the only binding constraint, then $\nabla h_3(x, y) = (p_x, p_y) \neq \mathbf{0}$; that is, the constraint qualification holds.

► EXERCISE 2.23. *Solve the problem of* (2.19).

Second-Order Analysis

We probably do not have time to discus the second-order conditions. See Duggan (2010, Section 6.3).

2.5 Envelop Theorem

Let $A \subset \mathbb{R}$. The graph of a real-valued function f on A is a curve in the $\mathbb{R}^2$ plane, and we shall also refer to the curve itself as f. Given a one-dimensional parametrized family of curves $f_\alpha \colon A \to \mathbb{R}$, where α runs over some interval, the curve $h \colon A \to \mathbb{R}$ is the *envelope of the family* if

- each point on the curve h is tangent to the graph of one of the curves f_α and
- each curve f_α is tangent to h.

(See, e.g., Apostol 1967, p. 342 or Zorich 2004b, p. 252 for this definition.) That is, for each α, there is some q and also for each q, there is some α, satisfying

$$\begin{cases} f_\alpha(q) = h(q) \\ f'_\alpha(q) = h'(q). \end{cases}$$

We may regard h as a function of α if the correspondence between curves and points on the envelope is one-to-one.

An Envelopment Theorem for Unconstrained Maximization

Consider now an unconstrained parametrized maximization problem. Let $x^*(q)$ be the value of the control variable x that maximizes $f(x, q)$, where q is our parameter of interest. For some fixed x, the function

$$\varphi_x(q) := f(x, q)$$

defines a curve. We also define the *value function*

$$V(q) := f\left(x^*(q), q\right) = \max_x \varphi_x(q).$$

Under appropriate conditions, the graph of the value function V will be the envelope of the curves φ_x. “Envelope theorems” in maximization theory are concerned with the tangency conditions this entails.

Example 2.24. Let

$$f(x, q) = q - (x - q)^2 + 1, \quad x, q \in [0, 2].$$

Then given q, the maximizing x is given by $x^*(q) = q$, and $V(q) = q + 1$.

For each x, the function φ_x is given by

$$\varphi_x(q) = q - (x - q)^2 + 1.$$

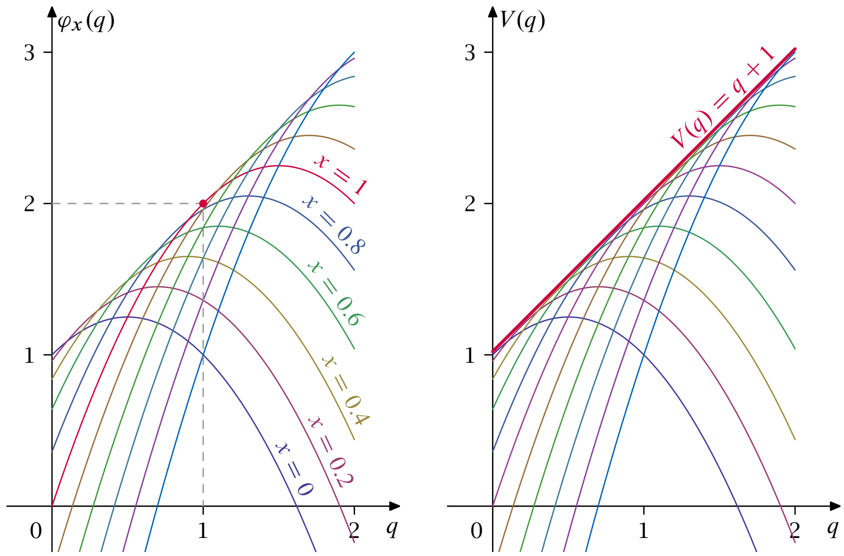

(a) The family of curves $\{\varphi_x(q) : x = 0, 0.2, \ldots, 2\}$.

(b) The envelop.

FIGURE 2.12. The graph of V is the envelope of the family of graphs of the functions φ_x.

The graphs of these functions and of V are shown for selected values of x in Figure 2.12. Observe that the graph of V is the envelope of the family of graphs of the functions φ_x. Consequently the slope of V is the slope of the φ_x to which it is tangent, that is,

$$V'(q) = \frac{\partial \varphi_x}{\partial q}\bigg|_{x=x^*(q)=q} = \frac{\partial f}{\partial q}\bigg|_{x=x^*(q)=q} = 1 + 2(x-q)|_{x=x^*(q)=q} = 1.$$

This last observation is one version of the Envelope Theorem.

An Envelope Theorem for Constrained Maximization

Consider the maximization problem,

$$\max_{\boldsymbol{x} \in \mathbb{R}^n} f(\boldsymbol{x}, \boldsymbol{q}) \quad \text{s.t. } g_i(\boldsymbol{x}, \boldsymbol{q}) = 0, \quad i = 1, \ldots, m, \tag{2.20}$$

where $\boldsymbol{x}$ is a vector of choice variables, and $\boldsymbol{q} = (q_1, \ldots, q_\ell) \in \mathbb{R}^\ell$ is a vector of parameters that may enter the objective function, the constraints, or both.

Suppose that for each $\boldsymbol{q}$ there exists a unique solution $\boldsymbol{x}(\boldsymbol{q})$. Furthermore, we assume that the objective function $f : \mathbb{R}^n \to \mathbb{R}$, constraints $g_i : \mathbb{R}^n \times \mathbb{R}^\ell \to \mathbb{R}$ $(i = 1, \ldots, m)$, and solutions $\boldsymbol{x} : \mathbb{R}^\ell \to \mathbb{R}^n$ are differentiable in the parameter $\boldsymbol{q}$.

Then, for every parameter $\boldsymbol{q}$, the maximized value of the objective function is $f(\boldsymbol{x}(\boldsymbol{q}), \boldsymbol{q})$. This defines a new function, $V : \mathbb{R}^{\ell} \to \mathbb{R}$, called the *value function*. Formally,

$$V(\boldsymbol{q}) := \max_{\boldsymbol{x} \in \mathbb{R}^n} \{ f(\boldsymbol{x}, \boldsymbol{q}) : g_i(\boldsymbol{x}, \boldsymbol{q}) = 0, i = 1, \ldots, m \} . \tag{2.21}$$

THEOREM 2.25 (Envelope Theorem). *Consider the value function $V(\boldsymbol{q})$ for the problem (2.20). Let $(\lambda_1, \ldots, \lambda_m)$ be values of the Lagrange multipliers associated with the maximizer solution $\boldsymbol{x}(\bar{\boldsymbol{q}})$ at $\bar{\boldsymbol{q}}$. Then for each $k = 1, \ldots, \ell$,*

$$\frac{\partial V(\bar{\boldsymbol{q}})}{\partial q_k} = \frac{\partial f(\boldsymbol{x}(\bar{\boldsymbol{q}}), \bar{\boldsymbol{q}})}{\partial q_k} - \sum_{j=1}^{m} \lambda_j \frac{\partial g_j(\boldsymbol{x}(\bar{\boldsymbol{q}}), \bar{\boldsymbol{q}})}{\partial q_k} \tag{2.22}$$

PROOF. By definition, $V(\boldsymbol{q}) = f(\boldsymbol{x}(\boldsymbol{q}), \boldsymbol{q})$ for all $\boldsymbol{q}$. Using the chain rule, we have

$$\frac{\partial V(\bar{\boldsymbol{q}})}{\partial q_k} = \sum_{i=1}^{n} \left[\frac{\partial f(\boldsymbol{x}(\bar{\boldsymbol{q}}), \bar{\boldsymbol{q}})}{\partial x_i} \frac{\partial x_i(\bar{\boldsymbol{q}})}{\partial q_k} \right] + \frac{\partial f(\boldsymbol{x}(\bar{\boldsymbol{q}}), \bar{\boldsymbol{q}})}{\partial q_k} .$$

It follows from Theorem 2.10 that

$$\frac{\partial f(\boldsymbol{x}(\bar{\boldsymbol{q}}), \bar{\boldsymbol{q}})}{\partial x_i} = \sum_{j=1}^{m} \lambda_j \frac{\partial g_j(\boldsymbol{x}(\bar{\boldsymbol{q}}), \bar{\boldsymbol{q}})}{\partial x_i} .$$

Hence,

$$\begin{aligned}
\frac{\partial V(\bar{\boldsymbol{q}})}{\partial q_k} &= \sum_{i=1}^{n} \left[\frac{\partial f(\boldsymbol{x}(\bar{\boldsymbol{q}}), \bar{\boldsymbol{q}})}{\partial x_i} \frac{\partial x_i(\bar{\boldsymbol{q}})}{\partial q_k} \right] + \frac{\partial f(\boldsymbol{x}(\bar{\boldsymbol{q}}), \bar{\boldsymbol{q}})}{\partial q_k} \\
&= \sum_{i=1}^{n} \left[\left(\sum_{j=1}^{m} \lambda_j \frac{\partial g_j(\boldsymbol{x}(\bar{\boldsymbol{q}}), \bar{\boldsymbol{q}})}{\partial x_i} \right) \frac{\partial x_i(\bar{\boldsymbol{q}})}{\partial q_k} \right] + \frac{\partial f(\boldsymbol{x}(\bar{\boldsymbol{q}}), \bar{\boldsymbol{q}})}{\partial q_k} \\
&= \sum_{j=1}^{n} \lambda_j \sum_{i=1}^{n} \left[\frac{\partial g_j(\boldsymbol{x}(\bar{\boldsymbol{q}}), \bar{\boldsymbol{q}})}{\partial x_i} \frac{\partial x_i(\bar{\boldsymbol{q}})}{\partial q_k} \right] + \frac{\partial f(\boldsymbol{x}(\bar{\boldsymbol{q}}), \bar{\boldsymbol{q}})}{\partial q_k} .
\end{aligned}$$

Finally, since $g_j(\boldsymbol{x}(\boldsymbol{q}), \boldsymbol{q}) = 0$ for all $\boldsymbol{q}$, we have

$$\sum_{i=1}^{n} \left[\frac{\partial g_j(\boldsymbol{x}(\bar{\boldsymbol{q}}), \bar{\boldsymbol{q}})}{\partial x_i} \frac{\partial x_i(\bar{\boldsymbol{q}})}{\partial q_k} \right] + \frac{\partial g_j(\boldsymbol{x}(\bar{\boldsymbol{q}}), \bar{\boldsymbol{q}})}{\partial q_k} = 0.$$

Combining, we get (2.22). □

Let us consider an example.

EXAMPLE 2.26. We are given the problem

$$\max_{(x,y) \in \mathbb{R}^2} \{ f((x,y), q) = xy \} \quad \text{s.t.} \quad g((x,y), q) = 2x + 4y - q = 0.$$

Forming the Lagrangian, we get

$$\mathcal{L} = xy - \lambda(2x + 4y - q),$$

with first-order conditions:

$$\begin{aligned} y - 2\lambda &= 0 \\ x - 4\lambda &= 0 \\ q - 2x - 4y &= 0. \end{aligned} \tag{2.23}$$

These solve for $x(q) = q/4$, $y(q) = q/8$ and $\lambda(q) = q/16$. Thus,

$$V(q) = x(q)y(q) = \frac{q^2}{32}.$$

Differentiating $V(q)$ with respect to q we get

$$V'(q) = \frac{q}{16}.$$

Now let us verify this using the Envelope Theorem. The theorem tells us that

$$V'(q) = \frac{\partial f((x(q), y(q)), q)}{\partial q} - \lambda(q)\frac{\partial g((x(q), y(q)), q)}{\partial q} = \lambda(q) = \frac{q}{16}.$$

EXAMPLE 2.27. Consider a consumer whose utility function $u\colon \mathbb{R}^n_+ \to \mathbb{R}$ is strictly increasing in every commodity $i = 1, \ldots, n$. Then this consumer's problem is

$$\max u(\boldsymbol{x}) \quad \text{s.t.} \sum_{i=1}^{n} x_i p_i = I.$$

The Lagrangian is

$$\mathcal{L}(\boldsymbol{x}, \lambda) = u(\boldsymbol{x}) - \lambda\left(I - \sum_{i=1}^{n} x_i p_i\right).$$

It follows from Theorem 2.25 that

$$V'(I) = \frac{\partial \mathcal{L}(\boldsymbol{x}, \lambda)}{\partial I} = \lambda.$$

That is, λ measures the *marginal utility of income*.

Integral Form Envelope Theorem

The Envelope theorems we introduced so far rely on assumptions that are not satisfactory for applications, e.g., mechanism design. Unfortunately, it is too technique to develop the more advanced treatment of the Envelope Theorem. We refer the reader to Milgrom and Segal (2002) and Milgrom (2004, Chapter 3) for the integral form Envelope Theorem.

3

CONVEX ANALYSIS IN $\mathbb{R}^N$

ROCKAFELLAR (1970) IS THE classical reference for finite-dimensional convex analysis. As for infinite-dimensional convex analysis, Luenberger (1969) is an excellent text.

To understand the material what follows, it is necessary that the reader have a good background in Multivariable Calculus (Chapter 1) and Linear Algebra.

In this chapter, we will exclusively consider convexity in $\mathbb{R}^n$ for concreteness, but much of the discussion here generalizes to infinite dimensional vector spaces. You can consult Berkovitz (2002), Ok (2007, Chapter G) and Royden and Fitzpatrick (2010, Chapter 6).

3.1 Convex Sets

DEFINITION 3.1 (Convex Set). A subset $C \subset \mathbb{R}^n$ is *convex* if for every pair of points $\boldsymbol{x}_1, \boldsymbol{x}_2 \in C$, the line segment

$$[\boldsymbol{x}_1, \boldsymbol{x}_2] := \{\boldsymbol{x} : \boldsymbol{x} = \lambda \boldsymbol{x}_1 + (1-\lambda)\boldsymbol{x}_2, 0 \leqslant \lambda \leqslant 1\}$$

belongs to C.

► EXERCISE 3.2. *Sketch the following sets in $\mathbb{R}^2$ and determine from figure which sets are convex and which are not:*

a. $\left\{(x, y) : x^2 + y^2 \leqslant 1\right\}$,

b. $\left\{(x, y) : 0 < x^2 + y^2 \leqslant 1\right\}$,

c. $\left\{(x, y) : y \geqslant x^2\right\}$,

d. $\left\{(x, y) : |x| + |y| \leqslant 1\right\}$, *and*

e. $\left\{(x, y) : y \geqslant 1/1 + x^2\right\}$.

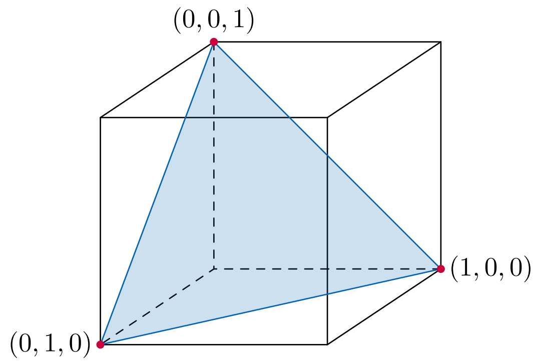

FIGURE 3.1. Δ^2 in $\mathbb{R}^3$.

LEMMA 3.3. *Let $\{C_\alpha\}$ be a collection of convex sets such that $C := \bigcap_\alpha C_\alpha \neq \varnothing$. Then C is convex.*

PROOF. Let $x_1, x_2 \in C$. Then $x_1, x_2 \in C_\alpha$ for all α. Since C_α is convex, we have $[x_1, x_2] \subset C_\alpha$ for all α. Hence, $[x_1, x_2] \subset C$, so C is convex. □

NOTATION. For each positive integer n, define

$$\Delta^{n-1} := \left\{ (\lambda_1, \ldots, \lambda_n) \in [0,1]^n : \sum_{i=1}^{n} \lambda_i = 1 \right\}. \tag{3.1}$$

For $n = 1$, the set Δ^0 is the singleton $\{1\}$. For $n = 2$, the set Δ^1 is the closed line segment joining $(0, 1)$ and $(1, 0)$. For $n = 3$, the set Δ^2 is the closed triangle with vertices $(1, 00)$, $(0, 1, 0)$ and $(0, 0, 1)$ (see Figure 3.1).

DEFINITION 3.4. A point $x \in \mathbb{R}^n$ is a *convex combination* of points $x_1, \ldots, x_k$ if there exists $\lambda \in \Delta^{k-1}$ such that

$$x = \sum_{i=1}^{k} \lambda_i x_i.$$

LEMMA 3.5. *A set $C \subset \mathbb{R}^n$ is convex iff every convex combination of points in C is also in C.*

PROOF. The "if" part is evident. So we shall prove the "only if" statement by induction on k. It holds for $k = 2$ by definition. Suppose the statement is true for $k = n$. Now consider $k = n + 1$. Let $x_1, \ldots, x_{k+1} \in C$ and $\lambda \in \Delta^n$ with $\lambda_{k+1} \in (0, 1)$. Then

$$
\begin{aligned}
\boldsymbol{x} &= \sum_{i=1}^{n+1} \lambda_i \boldsymbol{x}_i \\
&= \left(\sum_{j=1}^{n} \lambda_j \right) \left[\sum_{i=1}^{n} \left(\frac{\lambda_i}{\sum_{j=1}^{n} \lambda_j} \right) \boldsymbol{x}_i \right] + \lambda_{k+1} \boldsymbol{x}_{k+1} \\
&\in C.
\end{aligned}
$$
□

Let $A \subset \mathbb{R}^n$, and let $\mathscr{A}$ be the class of all convex subsets of $\mathbb{R}^n$ that contain A. We have $\mathscr{A} \neq \varnothing$—after all, $\mathbb{R}^n \in \mathscr{A}$. Then, by Lemma 3.3, $\bigcap \mathscr{A}$ is a convex set in $\mathbb{R}^n$ that contains A. Clearly, this set is the smallest (that is, $\subset$-minimum) convex subset of $\mathbb{R}^n$ that contains A.

Definition 3.6. The *convex hull* of A, denoted by $\operatorname{cov}(A)$, is the intersection of all convex sets containing A.

► Exercise 3.7. *For a given set A, let $K(A)$ denote the set of all convex combinations of points in A. Show that $K(A)$ is convex and $A \subset K(A)$.*

Theorem 3.8. *Let $A \subset \mathbb{R}^n$. Then* $\operatorname{cov}(A) = K(A)$.

Proof. Let $\mathscr{A}$ be the family of convex sets containing A. Since $\operatorname{cov}(A) = \bigcap \mathscr{A}$ and $K(A) \in \mathscr{A}$ (Exercise 3.7), we have $\operatorname{cov}(A) \subset K(A)$.

To prove the reverse inclusion relation, take an arbitrary $C \in \mathscr{A}$. Then $A \subset C$. It follows from Lemma 3.5 that $K(A) \subset C$. Hence $K(A) \subset \bigcap \mathscr{A} = \operatorname{cov}(A)$. □

3.2 Separation Theorem

This section is devoted to the establishment of *separation theorems*. In some sense, these theorems are the fundamental theorems of optimization theory. For simplicity, we restrict our analysis on $\mathbb{R}^n$.

Definition 3.9. A *hyperplane* $\mathscr{H}_{\boldsymbol{a}}^{\beta}$ in $\mathbb{R}^n$ is defined to be the set of points that satisfy the equation $\langle \boldsymbol{a}, \boldsymbol{x} \rangle = \alpha$. Thus,

$$
\mathscr{H}_{\boldsymbol{a}}^{\beta} := \left\{ \boldsymbol{x} \in \mathbb{R}^n : \langle \boldsymbol{a}, \boldsymbol{x} \rangle = \beta \right\}.
$$

The vector $\boldsymbol{a}$ is said to be a *normal* to the hyperplane.

Remark 3.10. Geometrically, a hyperplane $\mathscr{H}_{\boldsymbol{a}}^{\beta}$ in $\mathbb{R}^n$ is a translation of an $(n-1)$-dimensional subspace (an *affine manifold*). Algebraically, it is a level set of a linear functional. For an excellent explanation about hyperplanes, see Luenberger (1969, Section 5.12).

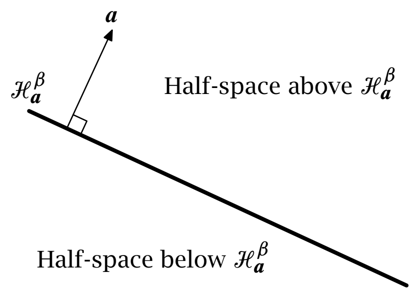

FIGURE 3.2. Hyperplane and half-spaces.

A hyperplane $\mathcal{H}_{\boldsymbol{a}}^\beta$ divides $\mathbb{R}^n$ into two *half spaces*, one on each side of $\mathcal{H}_{\boldsymbol{a}}^\beta$. The set $\{\boldsymbol{x} \in \mathbb{R}^n : \langle \boldsymbol{a}, \boldsymbol{x} \rangle \geqslant \beta\}$ is called the *half-space above* the hyperplane $\mathcal{H}_{\boldsymbol{a}}^\beta$, and the set $\{\boldsymbol{x} \in \mathbb{R}^n : \langle \boldsymbol{a}, \boldsymbol{x} \rangle \leqslant \beta\}$ is called the *half-space below* the hyperplane $\mathcal{H}_{\boldsymbol{a}}^\beta$; see Figure 3.2.

It therefore seems natural to say that two sets X and Y are separated by a hyperplane $\mathcal{H}_{\boldsymbol{a}}^\beta$ if they are contained in different half spaces determined by $\mathcal{H}_{\boldsymbol{a}}^\beta$. We will introduce two separation theorems.

THEOREM 3.11 (A First Separation Theorem). *Let C be a closed and convex subset of $\mathbb{R}^n$; let $\boldsymbol{y} \in \mathbb{R}^n \smallsetminus C$. Then there exists a vector $\boldsymbol{a} \in \mathbb{R}^n$ with $\boldsymbol{a} \neq \boldsymbol{0}$, and a scalar $\beta \in \mathbb{R}$ such that $\langle \boldsymbol{a}, \boldsymbol{y} \rangle > \beta$ and $\langle \boldsymbol{a}, \boldsymbol{x} \rangle < \beta$ for all $\boldsymbol{x} \in C$.*

To motivate the proof, we argue heuristically from Figure 3.3, where C is assumed to have a tangent at each boundary point. Draw a line from $\boldsymbol{y}$ to $\boldsymbol{x}^*$, the point of C that is closest to $\boldsymbol{y}$. The vector $\boldsymbol{y} - \boldsymbol{x}^*$ is orthogonal to C in the sense that $\boldsymbol{y} - \boldsymbol{x}^*$ is orthogonal to the tangent line at $\boldsymbol{x}^*$. The tangent line, which is exactly the set $\{\boldsymbol{z} \in \mathbb{R}^n : \langle \boldsymbol{y} - \boldsymbol{x}^*, \boldsymbol{z} - \boldsymbol{x}^* \rangle = 0\}$, separates $\boldsymbol{y}$ and C. The point $\boldsymbol{x}^*$ is characterized by the fact that $\langle \boldsymbol{y} - \boldsymbol{x}^*, \boldsymbol{x} - \boldsymbol{x}^* \rangle \leqslant 0$ for all $\boldsymbol{x} \in C$. If we move the tangent line parallel to itself so as to pass through a point $\boldsymbol{x}_0 \in (\boldsymbol{x}^*, \boldsymbol{y})$, we are done. We now justify these steps in a series of claims.

CLAIM 1. Let C be a convex subset of $\mathbb{R}^n$ and let $\boldsymbol{y} \in \mathbb{R}^n \smallsetminus C$. If there exists a point in C that is closest to $\boldsymbol{y}$, then it is unique.

PROOF. Suppose that there were two points $\boldsymbol{x}_1$ and $\boldsymbol{x}_2$ of C that were closest to $\boldsymbol{y}$. Then $(\boldsymbol{x}_1 + \boldsymbol{x}_2)/2 \in C$ since C is convex, and so[1]

[1] For a set $X \subset \mathbb{R}^n$, the function $d(\boldsymbol{y}, X)$ is the distance from $\boldsymbol{y}$ to X defined by

$$d(\boldsymbol{y}, X) = \inf_{\boldsymbol{x} \in X} \|\boldsymbol{y} - \boldsymbol{x}\|.$$

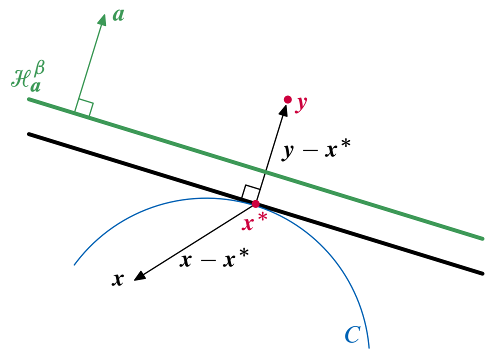

FIGURE 3.3. The separating hyperplane theorem.

$$\begin{aligned} d(\boldsymbol{y},C) &\leqslant \left\|\frac{\boldsymbol{x}_1+\boldsymbol{x}_2}{2}-\boldsymbol{y}\right\| \\ &= \left\|\frac{1}{2}\left[(\boldsymbol{x}_1-\boldsymbol{y})+(\boldsymbol{x}_2-\boldsymbol{y})\right]\right\| \\ &\leqslant \frac{1}{2}\|\boldsymbol{x}_1-\boldsymbol{y}\|+\frac{1}{2}\|\boldsymbol{x}_2-\boldsymbol{y}\| \\ &= d(\boldsymbol{y},C). \end{aligned}$$

Hence the triangle inequality holds with equality. It follows from Exercise 1.8 that there exists $\kappa \geqslant 0$ such that $\boldsymbol{x}_1-\boldsymbol{y}=\kappa(\boldsymbol{x}_2-\boldsymbol{y})$. Clearly, $\kappa \neq 0$; for otherwise $\boldsymbol{x}_1-\boldsymbol{y}=\boldsymbol{0}$ implies that $\boldsymbol{y}=\boldsymbol{x}_1 \in C$. Then $\kappa=1$ since $\|\boldsymbol{x}_1-\boldsymbol{y}\|=\|\boldsymbol{x}_2-\boldsymbol{y}\|=d(\boldsymbol{y},C)$. But then $\boldsymbol{x}_1=\boldsymbol{x}_2$.

CLAIM 2. Let C be a closed subset of $\mathbb{R}^n$ and let $\boldsymbol{y} \in \mathbb{R}^n \setminus C$. Then there exists a point $\boldsymbol{x}^* \in C$ that is closest to $\boldsymbol{y}$.

PROOF. Take an arbitrary point $\boldsymbol{x}_0 \in C$. Let $r > \|\boldsymbol{x}_0-\boldsymbol{y}\|$. Then $C_1 := \overline{\mathcal{B}(\boldsymbol{y};r)} \cap C$ is nonempty (at least $\boldsymbol{x}_0$ is in the intersection), closed, and bounded and hence is compact. The function $\boldsymbol{x} \mapsto \|\boldsymbol{x}-\boldsymbol{y}\|$ is continuous on C_1 and so attains its minimum at some point $\boldsymbol{x}^* \in C_1$, i.e.,

$$\|\boldsymbol{x}^*-\boldsymbol{y}\| \leqslant \|\boldsymbol{x}-\boldsymbol{y}\| \quad \text{for all } \boldsymbol{x} \in C_1.$$

For every $\boldsymbol{x} \in C \setminus C_1$, we have

$$\|\boldsymbol{x}-\boldsymbol{y}\| > r > \|\boldsymbol{x}_0-\boldsymbol{y}\| \geqslant \|\boldsymbol{x}^*-\boldsymbol{y}\|,$$

since $\boldsymbol{x}_0 \in C_1$.

CLAIM 3. Let C be a convex subset of $\mathbb{R}^n$ and let $\boldsymbol{y} \in \mathbb{R}^n \setminus C$. Then $\boldsymbol{x}^* \in C$ is a closest point in C to $\boldsymbol{y}$ iff

$$\left\langle \boldsymbol{y}-\boldsymbol{x}^*, \boldsymbol{x}-\boldsymbol{x}^* \right\rangle \leqslant 0 \quad \text{for all } \boldsymbol{x} \in C. \tag{3.2}$$

PROOF. Let $\boldsymbol{x}^* \in C$ be a closest point to $\boldsymbol{y}$ and let $\boldsymbol{x} \in C$. Since C is convex, we have

$$[\boldsymbol{x}^*, \boldsymbol{x}] := \left\{ \boldsymbol{z}(t) \in \mathbb{R}^n : \boldsymbol{z}(t) = \boldsymbol{x}^* + t(\boldsymbol{x} - \boldsymbol{x}^*), t \in [0,1] \right\} \subset C.$$

Let

$$\begin{aligned} g(t) := \|\boldsymbol{z}(t) - \boldsymbol{y}\|^2 &= \left\langle \boldsymbol{x}^* + t(\boldsymbol{x} - \boldsymbol{x}^*) - \boldsymbol{y}, \boldsymbol{x}^* + t(\boldsymbol{x} - \boldsymbol{x}^*) - \boldsymbol{y} \right\rangle \\ &= \sum_{i=1}^{n} \left[x_i^* - y_i + t(x_i - x_i^*) \right]^2 . \end{aligned}$$

Observe that $g(0) = \boldsymbol{x}^*$. Since g is continuously differentiable on $(0,1]$ and $\boldsymbol{x}^* \in \operatorname{argmin}_{\boldsymbol{x} \in C} \|\boldsymbol{x} - \boldsymbol{y}\|$, we have $g'_+(0) \geqslant 0$. Since

$$\begin{aligned} g'(t) &= 2 \sum_{i=1}^{n} \left[x_i^* - y_i + t(x_i - x_i^*) \right] (x_i - x_i^*) \\ &= 2 \left[- \sum_{i=1}^{n} (y_i - x_i^*)(x_i - x_i^*) + t \sum_{i=1}^{n} (x_i - x_i^*)^2 \right] \qquad (3.3) \\ &= 2 \left[- \left\langle \boldsymbol{y} - \boldsymbol{x}^*, \boldsymbol{x} - \boldsymbol{x}^* \right\rangle + t \|\boldsymbol{x} - \boldsymbol{x}^*\|^2 \right]. \end{aligned}$$

Letting $t \downarrow 0$ we get (3.2).

Conversely, suppose that (3.2) holds. Take an arbitrary $\boldsymbol{x} \in C \setminus \{\boldsymbol{x}^*\}$. It follows from (3.3) that if $t \in (0,1]$ then

$$g'(t) = 2 \left[- \left\langle \boldsymbol{y} - \boldsymbol{x}^*, \boldsymbol{x} - \boldsymbol{x}^* \right\rangle + t \|\boldsymbol{x} - \boldsymbol{x}^*\|^2 \right] \geqslant 2t \|\boldsymbol{x} - \boldsymbol{x}^*\|^2 > 0.$$

That is, g is strictly increasing on $[0,1]$. Thus, $g(1) = \|\boldsymbol{x} - \boldsymbol{y}\| > \|\boldsymbol{x}^* - \boldsymbol{y}\| = g(0)$.

PROOF OF THEOREM 3.11. We now can complete the proof of Theorem 3.11. Let $\boldsymbol{x}^* \in C$ be the closest point to $\boldsymbol{y}$ (by Claim 1 and Claim 2). Let $\boldsymbol{a} = \boldsymbol{y} - \boldsymbol{x}^*$. Then for all $\boldsymbol{x} \in C$, we have $\langle \boldsymbol{a}, \boldsymbol{x} - \boldsymbol{x}^* \rangle \leqslant 0$ (by Claim 3), i.e., $\langle \boldsymbol{a}, \boldsymbol{x} \rangle \leqslant \langle \boldsymbol{a}, \boldsymbol{x}^* \rangle$, with equality occurring when $\boldsymbol{x} = \boldsymbol{x}^*$. Hence,

$$\max_{\boldsymbol{x} \in C} \langle \boldsymbol{a}, \boldsymbol{x} \rangle = \left\langle \boldsymbol{a}, \boldsymbol{x}^* \right\rangle.$$

On the other hand, $\langle \boldsymbol{a}, \boldsymbol{y} - \boldsymbol{x}^* \rangle = \|\boldsymbol{a}\|^2 > 0$, so

$$\langle \boldsymbol{a}, \boldsymbol{y} \rangle = \left\langle \boldsymbol{a}, \boldsymbol{x}^* \right\rangle + \|\boldsymbol{a}\|^2 > \left\langle \boldsymbol{a}, \boldsymbol{x}^* \right\rangle.$$

Finally, take an arbitrary $\beta \in (\langle \boldsymbol{a}, \boldsymbol{x}^* \rangle, \langle \boldsymbol{a}, \boldsymbol{y} \rangle)$. We thus have $\langle \boldsymbol{a}, \boldsymbol{y} \rangle > \beta$ and $\langle \boldsymbol{a}, \boldsymbol{x} \rangle \leqslant \langle \boldsymbol{a}, \boldsymbol{x}^* \rangle < \beta$ for all $\boldsymbol{x} \in C$. $\square$

THEOREM 3.12 (A Second Separation Theorem). *Let X and Y be two disjoint convex subsets of $\mathbb{R}^n$. Then there exists $\boldsymbol{a} \in \mathbb{R}^n$ with $\boldsymbol{a} \neq \boldsymbol{0}$ and a scalar $\beta \in \mathbb{R}$ such that $\langle \boldsymbol{a}, \boldsymbol{x} \rangle \geqslant \beta$ for all $\boldsymbol{x} \in X$ and $\langle \boldsymbol{a}, \boldsymbol{y} \rangle \leqslant \beta$ for all $\boldsymbol{y} \in Y$. That is, there is a hyperplane $\mathcal{H}_{\boldsymbol{a}}^{\beta}$ that separates X and Y.*

PROOF. We leave the proof to the reader. See Berkovitz (2002, Theorem 3.3) and Jehle and Reny (2011, Theorem A2.24). □

3.3 Convex Functions

Throughout this section we will assume the subset $C \subset \mathbb{R}^n$ is convex and f is a real-valued function defined on C, that is, $f : C \to \mathbb{R}$. When we take $\boldsymbol{x}_1, \boldsymbol{x}_2 \in C$, we will let $\boldsymbol{x}_t := t\boldsymbol{x}_1 + (1-t)\boldsymbol{x}_2$, for $t \in [0,1]$, denote the convex combination of $\boldsymbol{x}_1$ and $\boldsymbol{x}_2$.

Definitions

DEFINITION 3.13. A function $f : C \to \mathbb{R}$ is *convex* if for all $\boldsymbol{x}_1, \boldsymbol{x}_2 \in C$ and $t \in [0,1]$,

$$f[t\boldsymbol{x}_1 + (1-t)\boldsymbol{x}_2] \leqslant tf(\boldsymbol{x}_1) + (1-t)f(\boldsymbol{x}_2).$$

The function f is *strictly convex* if the above inequality holds strictly.

A function $f : C \to \mathbb{R}$ is *concave* if for all $\boldsymbol{x}_1, \boldsymbol{x}_2 \in C$ and $t \in [0,1]$,

$$f[t\boldsymbol{x}_1 + (1-t)\boldsymbol{x}_2] \geqslant tf(\boldsymbol{x}_1) + (1-t)f(\boldsymbol{x}_2).$$

The function f is *strictly concave* if the above inequality holds strictly.

DEFINITION 3.14. A function $f : C \to \mathbb{R}$ is *quasi-convex* if, for all $\boldsymbol{x}_1, \boldsymbol{x}_2 \in C$ and $t \in [0,1]$,

$$f[t\boldsymbol{x}_1 + (1-t)\boldsymbol{x}_2] \leqslant \max\{f(\boldsymbol{x}_1), f(\boldsymbol{x}_2)\}.$$

The function f is *strictly quasi-convex function* if the above inequality holds strictly.

A function $f : C \to \mathbb{R}$ is *quasi-concave* if, for all $\boldsymbol{x}_1, \boldsymbol{x}_2 \in C$ and $t \in [0,1]$,

$$f[t\boldsymbol{x}_1 + (1-t)\boldsymbol{x}_2] \geqslant \min\{f(\boldsymbol{x}_1), f(\boldsymbol{x}_2)\}.$$

The function f is *strictly quasi-concave* if the above inequality holds strictly.

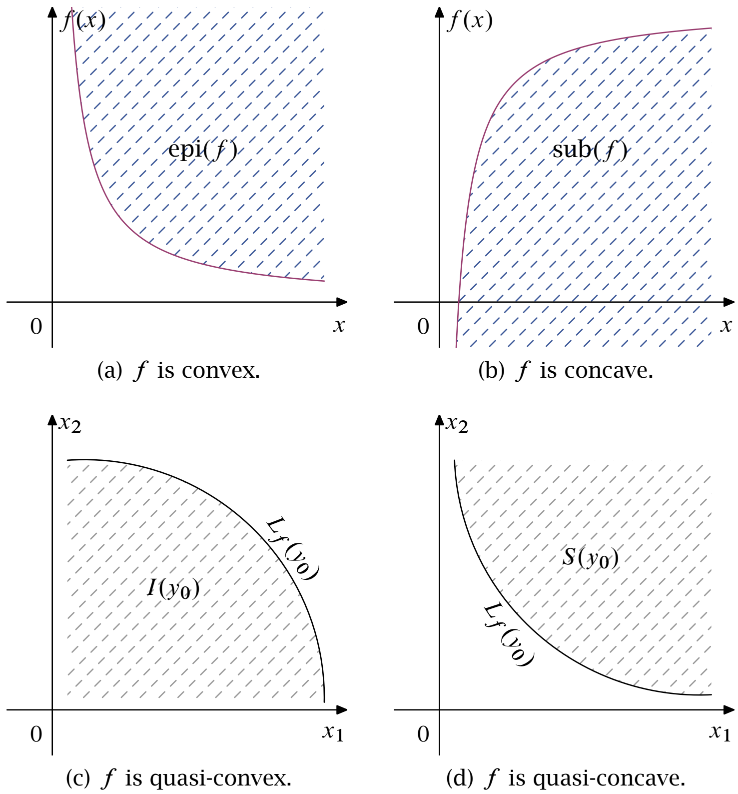

FIGURE 3.4.

Geometric Interpretation

Given a function $f : C \to \mathbb{R}$ and $y_0 \in \mathbb{R}$, let us define

The *epigraph* of f: $\operatorname{epi}(f) := \{(\boldsymbol{x}, y) \in C \times \mathbb{R} : f(\boldsymbol{x}) \leqslant y\}$;

The *subgraph* of f: $\operatorname{sub}(f) := \{(\boldsymbol{x}, y) \in C \times \mathbb{R} : f(\boldsymbol{x}) \geqslant y\}$;

The *superior set* for level y_0: $S(y_0) := \{\boldsymbol{x} \in C : f(\boldsymbol{x}) \geqslant y_0\}$;

The *inferior set* for level y_0: $I(y_0) := \{\boldsymbol{x} \in C : f(\boldsymbol{x}) \leqslant y_0\}$.

We then have (see Figure 3.4)

- f is convex $\iff$ epi(f) is convex.
- f is concave $\iff$ sub(f) is convex.
- f is quasi-convex $\iff$ $I(y_0)$ is convex.
- f is quasi-concave $\iff$ $S(y_0)$ is convex.

Convexity and Quasi-convexity

It is a simple matter to show that concavity (convexity) implies quasi-concavity (quasi-convexity).

THEOREM 3.15. *Let $C \subset \mathbb{R}^n$ and $f : C \to \mathbb{R}$. If f is concave on C, it is also quasi-concave on C. If f is convex on C, it is also quasi-convex on C.*

PROOF. We only prove the first claim, and leave the second as an exercise. Suppose f is concave on C. Take any $\boldsymbol{x}, \boldsymbol{y} \in C$ and $t \in [0, 1]$. Without loss of generality, we let

$$f(\boldsymbol{x}) \geqslant f(\boldsymbol{y}).$$

By the definition of concavity, we have

$$\begin{aligned} f\left[t\boldsymbol{x} + (1-t)\boldsymbol{y}\right] &\geqslant tf(\boldsymbol{x}) + (1-t)f(\boldsymbol{y}) \\ &= t\left[f(\boldsymbol{x}) - f(\boldsymbol{y})\right] + f(\boldsymbol{y}) \\ &\geqslant f(\boldsymbol{y}) \\ &= \min\{f(\boldsymbol{x}), f(\boldsymbol{y})\}. \end{aligned}$$

Hence, f is quasi-concave. □

► EXERCISE 3.16. *Prove the second claim in Theorem 3.15: if f is convex, then it is also quasi-convex.*

Concavity and Hessian

We now characterize concavity of a function using the Hessian matrix.

THEOREM 3.17. *Let $A \subset \mathbb{R}^n$. The (twice continuously differentiable) function $f : A \to \mathbb{R}$ is concave if and only if $\mathbf{H}f(\boldsymbol{x})$ is negative semidefinite for every $\boldsymbol{x} \in A$.*

PROOF. See Mas-Colell et al. (1995, Theorem M.C.2). □

EXAMPLE 3.18. Let $A := (0, \infty) \times (-5, \infty)$. Let $f(x, y) = \ln x + \ln(y + 5)$. For each point $(x, y) \in A$, the Hessian of f is

$$\mathbf{H}f(x, y) = \begin{bmatrix} -1/x^2 & 0 \\ 0 & -1/(y+5)^2 \end{bmatrix}.$$

Then for each $(u, v) \in \mathbb{R}^2$, we have

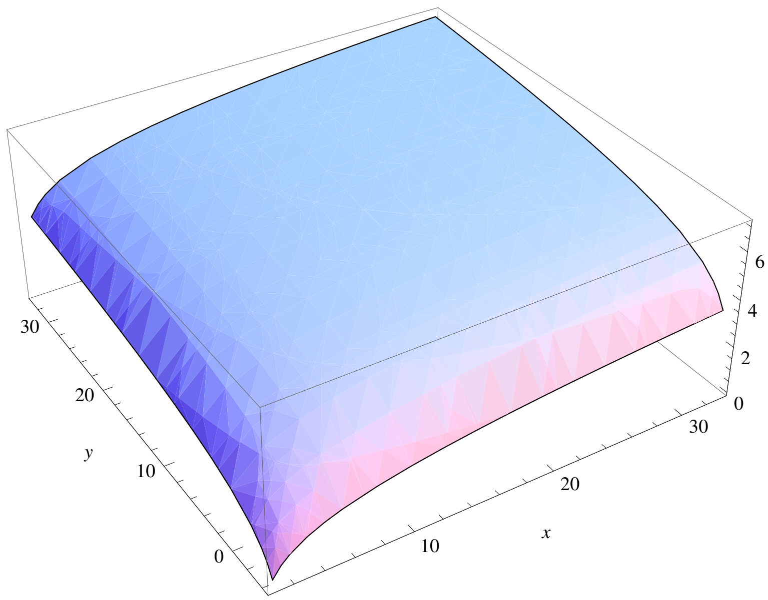

FIGURE 3.5. The function $\ln x + \ln(y+5)$ is concave.

$$\begin{bmatrix} u & v \end{bmatrix} \cdot \begin{bmatrix} -1/x^2 & 0 \\ 0 & -1/(y+5)^2 \end{bmatrix} \begin{bmatrix} u \\ v \end{bmatrix} = -\frac{u^2}{x^2} - \frac{v^2}{(y+5)^2} \leqslant 0.$$

That is, $\mathbf{H} f(x,y)$ is negative semidefinite. Hence, $f(x,y)$ is concave. See Figure 3.5.

Jensen's Inequality

THEOREM 3.19 (Jensen's Inequality). *Let $f : C \to \mathbb{R}$, where $C \subset \mathbb{R}^n$ is convex. Then f is convex iff for every finite set of points $\boldsymbol{x}_1, \dots, \boldsymbol{x}_k \in C$ and every $\boldsymbol{t} = (t_1, \dots, t_k) \in \Delta^{k-1}$ (see (3.1) for the definition of Δ^{k-1}),*

$$f\left(\sum_{i=1}^{k} t_i \boldsymbol{x}_i\right) \leqslant \sum_{i=1}^{k} t_i f(\boldsymbol{x}_i). \tag{3.4}$$

PROOF. The ”If” part is evident. So we only prove the “only if” part. Suppose that f is convex. We shall prove (3.4) by induction on k. For $k = 1$ the relation is trivial (remember that $\Delta^0 = \{1\}$ when $k = 1$). For $k = 2$ the relation follows from the definition of a convex function. Suppose that $k > 2$ and that (3.4) has been established for $k - 1$. We show that (3.4) holds for k.

If $t_k = 1$, then there is nothing to prove. If $t_k < 1$, set $T = \sum_{i=1}^{k-1} t_i$. Then $T + t_k = 1$, $T = 1 - t_k > 0$ and

$$\sum_{i=1}^{k-1} \frac{t_i}{T} = \frac{T}{T} = 1.$$

Hence,

$$\begin{aligned} f\left(\sum_{i=1}^{k} t_i \boldsymbol{x}_i\right) &= f\left(\sum_{i=1}^{k-1} t_i \boldsymbol{x}_i + t_k \boldsymbol{x}_k\right) \\ &= f\left(T\left[\sum_{i=1}^{k-1}\left(\frac{t_i}{T}\right)\boldsymbol{x}_i\right] + t_k \boldsymbol{x}_k\right) \\ &\leqslant Tf\left(\sum_{i=1}^{k-1}\left(\frac{t_i}{T}\right)\boldsymbol{x}_i\right) + t_k f(\boldsymbol{x}_k) \\ &\leqslant T\left[\sum_{i=1}^{k-1}\left(\frac{t_i}{T}\right) f(\boldsymbol{x}_i)\right] + t_k f(\boldsymbol{x}_k) \\ &= \sum_{i=1}^{k} t_i f(\boldsymbol{x}_i), \end{aligned}$$

where the first inequality follows from the convexity of f and the second inequality follows from the inductive hypothesis. □

3.4 Convexity and Optimization

Concave Programming

We first present two results which indicates the importance of convexity for optimization theory.

- In convex optimization problems, all *local optima* must also be *global optima*.
- If a *strictly* convex optimization problem admits a solution, the solution must be *unique*.

THEOREM 3.20. *Let $X \subset \mathbb{R}^n$ be convex and $f: X \to \mathbb{R}$ be concave. Then*

a. *Any local maximizer of f is a global maximizer of f.*

b. *The set* $\operatorname{argmax}\{f(\boldsymbol{x}) : \boldsymbol{x} \in X\}$ *of maximizers of f on X is either empty or convex.*

PROOF. **(a)** Suppose $\boldsymbol{x}$ is a local maximizer but not a global maximizer of f. Then there exists $\varepsilon > 0$ such that

$$f(\boldsymbol{x}) \geqslant f(\boldsymbol{y}), \quad \text{for all } \boldsymbol{y} \in X \cap \mathcal{B}(\boldsymbol{x};\varepsilon),$$

and there exists $\boldsymbol{z} \in X$ such that

$$f(\boldsymbol{z}) > f(\boldsymbol{x}). \tag{3.5}$$

Since X is convex, $[t\boldsymbol{x} + (1-t)\boldsymbol{z}] \in X$ for all $t \in (0,1)$. Take $\bar{t}$ sufficiently close to 1 so that $\bar{t}\boldsymbol{x} + (1-\bar{t})\boldsymbol{z} \in \mathcal{B}(\boldsymbol{x};\varepsilon)$. By the concavity of f and (C1), we have

$$f\left[\bar{t}\boldsymbol{x} + (1-\bar{t})\boldsymbol{z}\right] \geqslant \bar{t} f(\boldsymbol{x}) + (1-\bar{t}) f(\boldsymbol{z}) > f(\boldsymbol{x}).$$

A contradiction.

(b) Suppose that $\boldsymbol{x}_1$ and $\boldsymbol{x}_2$ are both maximizers of f on X. Then $f(\boldsymbol{x}_1) = f(\boldsymbol{x}_2)$. For any $t \in (0,1)$, we have

$$f(\boldsymbol{x}_1) \geqslant f\left[t\boldsymbol{x}_1 + (1-t)\boldsymbol{x}_2\right] \geqslant t f(\boldsymbol{x}_1) + (1-t) f(\boldsymbol{x}_2) = f(\boldsymbol{x}_1).$$

That is, $f[t\boldsymbol{x} + (1-t)\boldsymbol{x}_2] = f(\boldsymbol{x}_1)$. Thus, $t\boldsymbol{x}_1 + (1-t)\boldsymbol{x}_2$ is a maximizer of f on X. ◻

THEOREM 3.21. *Let $X \subset \mathbb{R}^n$ be convex and $f: X \to \mathbb{R}$ is strictly concave. Then* $\operatorname{argmax}\{f(\boldsymbol{x}) : \boldsymbol{x} \in X\}$ *either is empty or contains a single point.*

► EXERCISE 3.22. *Prove Theorem 3.21.*

We now present a extremely important theorem, which says that the first-order conditions of the Kuhn-Tucker Theorem (Theorem 2.16) are both *necessary* and *sufficient* to identify optima of convex inequality-constrained optimization problem, provided a mild regularity condition is met.

THEOREM 3.23 (Kuhn-Tucker Theorem under Convexity). *Let $U \subset \mathbb{R}^n$ be open and convex. Let $f : U \to \mathbb{R}$ be a concave $\mathcal{C}^1$ function. For $i = 1, \dots, \ell$, let $h_i : U \to \mathbb{R}$ be convex $\mathcal{C}^1$ functions. Suppose there is some $\bar{\boldsymbol{x}} \in U$ such that*

$$h_i(\bar{\boldsymbol{x}}) < 0, \quad i = 1, \dots, \ell.$$

(This is called the Slater's condition.*) Then $\boldsymbol{x}^*$ maximizes f over*

$$X := \{\boldsymbol{x} \in U : h_i(\boldsymbol{x}) \leqslant 0, i = 1, \dots, \ell\}$$

if and only if there is $\boldsymbol{\lambda}^* \in \mathbb{R}^\ell$ such that the Kuhn-Tucker first-order conditions hold:

$$\boldsymbol{\lambda}^* \geqslant \mathbf{0}, \quad \sum_{i=1}^{\ell} \lambda_i^* h_i(\boldsymbol{x}^*) = 0. \tag{KTC-1}$$

$$\nabla f(\boldsymbol{x}^*) = \sum_{i=1}^{\ell} \lambda_i^* \nabla h_i(\boldsymbol{x}^*). \tag{KTC-2}$$

PROOF. See Sundaram (1996, Section 7.7). □

EXAMPLE 3.24. Let $U = (0, \infty) \times (-5, \infty)$. Let $f(x, y) = \ln x + \ln(y + 5)$, $h_1(x, y) = x + y - 4$ and $h_2(x, y) = -y$. Consider the problem

$$\begin{aligned} &\max_{(x,y) \in U} f(x, y) \\ &\text{s.t.} \quad h_1(x, y) \leqslant 0, \quad h_2(x, y) \leqslant 0. \end{aligned} \tag{3.6}$$

► EXERCISE 3.25. *Show that $(x^*, y^*) = (4, 0)$ is the unique point satisfying the first-order condition for a local maximizer of the problem (3.6).*

Clearly, the Slater's condition holds. Then, combining Theorem 3.23, Exercise 3.25 and Example 3.18, we conclude that $(4, 0)$ is a global maximizer. See Figure 3.6.

Slater's Condition

For a formal demonstration of the need for Slater's condition, let us consider the following example.

EXAMPLE 3.26. Let $U = \mathbb{R}$; let $f(x) = x$ and $g(x) = x^2$. The only point in $\mathbb{R}$ satisfying $g(x) \leqslant 0$ is $x = 0$, so this is trivially the constrained maximizer of f. But

$$f'(0) = 1 \quad \text{and} \quad g'(0) = 0,$$

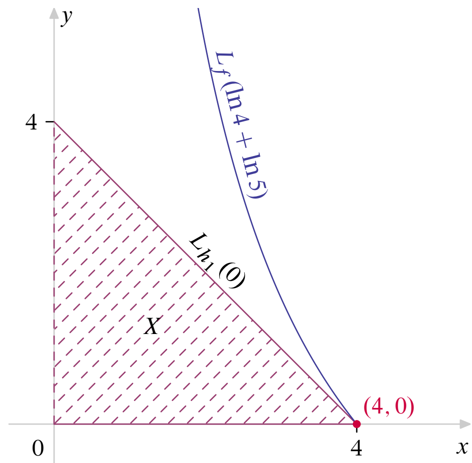

FIGURE 3.6. **(4, 0)** is the global maximizer.

so there is no $\lambda \geqslant 0$ such that $f'(0) = \lambda g'(0)$.

Quasi-concavity and Optimization

Quasi-concave and quasi-convex functions fail to exhibit many of the sharp properties that distinguish concave and convex functions. As an example, we show that Theorem 3.23 fails for quasi-concave objective functions.

EXAMPLE 3.27. Let $f : \mathbb{R} \to \mathbb{R}$ and $h : \mathbb{R} \to \mathbb{R}$ be quasi-concave continuously differentiable functions, where

$$f(x) = \begin{cases} x^3 & \text{if } x \in (-\infty, 0) \\ 0 & \text{if } x \in [0, 1] \\ (x-1)^2 & \text{if } x \in (1, \infty), \end{cases}$$

and $h(x) = -x$; see Figure 3.7.

► EXERCISE 3.28. *Show that for every point $x \in [0, 1]$, there exists $\lambda \geqslant 0$ such that the pair (x, λ) satisfies (KT-1) and (KT-2) (see p. 36).*

Furthermore, the Slater's condition holds. However, it is clear that no point $x \in [0, 1]$ can be a solution to the problem (Why?).

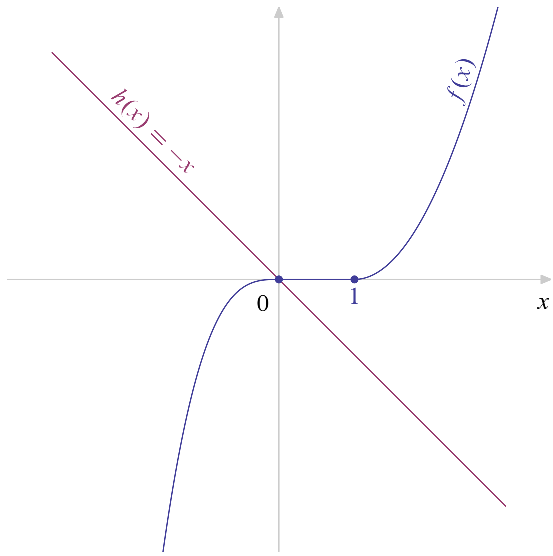

FIGURE 3.7. Quasi-concavity and optimization.

4

DYNAMIC OPTIMIZATION THEORY

DYNAMIC OPTIMIZATION THEORY is widely used in repeated games (Mailath and Samuelson, 2006), macroeconomics (Ljungqvist and Sargent, 2004) and finance (Duffie, 2001).

The classical reference for this topic is Stokey and Lucas (1989). We will follow Ok (2007, Section C.6, C.7.1 and E).

4.1 CORRESPONDENCES

Correspondences arise quite frequently in optimization theory and theoretical economics. This section introduces basic notion and properties of correspondences. For an advanced discussion, see Aubin and Frankowska (1990) and Aliprantis and Border (2006, Chapter 17).

Basic Definition

DEFINITION 4.1 (Correspondence). Let X and Y be subsets of $\mathbb{R}^m$ and $\mathbb{R}^n$, respectively. A *correspondence* Γ from X to Y is a map that associates with each element $\boldsymbol{x} \in X$ a nonempty subset $\Gamma(\boldsymbol{x}) \subset Y$. We write $\Gamma\colon X \twoheadrightarrow Y$ to denote that Γ is a correspondence from X into Y.

EXAMPLE 4.2. For any $n \in \mathbb{N}$, $\boldsymbol{p} \in \mathbb{R}^n_{++}$ and $I > 0$, defnie

$$B(\boldsymbol{p}, I) := \left\{ \boldsymbol{x} \in \mathbb{R}^n_+ : \sum_{i=1}^{n} x_i p_i \leqslant I \right\},$$

which is called the *budget set* of a consumer with income I at prices $\boldsymbol{p}$. If we treat $\boldsymbol{p}$ and I as variables, then it would be necessary to view B as a correspondence. We have $B := \mathbb{R}^{n+1}_{++} \twoheadrightarrow \mathbb{R}^n_+$.

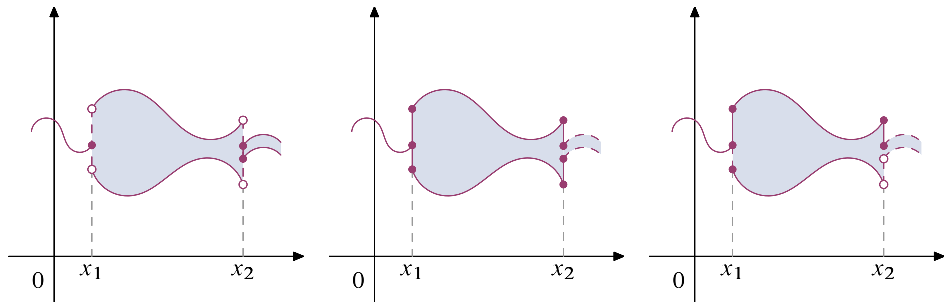

(a) Not upper hemicontinuous at x_1. Not upper hemicontinuous at x_2. Lower hemicontinuous.

(b) Not lower hemicontinuous at x_1. Not lower hemicontinuous at x_2. Upper hemicontinuous.

(c) Now upper hemicontinuous at x_1 and x_2. Not lower hemicontinuous at x_1 and x_2.

FIGURE 4.1. Upper hemicontinuity and lower hemicontinuity.

Continuity of Correspondence

DEFINITION 4.3. For any two subsets $X \subset \mathbb{R}^m$ and $Y \subset \mathbb{R}^n$, a correspondence $\Gamma\colon X \twoheadrightarrow Y$ is said to be

- *upper hemicontinuous at* $\boldsymbol{x} \in X$ if, for every open subset $\mathcal{O}$ of Y with $\Gamma(\boldsymbol{x}) \subset \mathcal{O}$, there exists a $\delta > 0$ such that
$$\Gamma(\mathcal{B}(\boldsymbol{x};\delta)) \subset \mathcal{O}.$$
Γ is called *upper hemicontinuous* if it is upper hemicontinuous on the entire X.

- *lower hemicontinuous at* $\boldsymbol{x} \in X$ if, for every open set $\mathcal{O}$ in Y with $\Gamma(\boldsymbol{x}) \cap \mathcal{O} \neq \varnothing$, there exists a $\delta > 0$ such that
$$\Gamma(\boldsymbol{x}') \cap \mathcal{O} \neq \varnothing \quad \text{for all } \boldsymbol{x}' \in \mathcal{B}(\boldsymbol{x};\delta).$$
Γ is called *lower hemicontinuous* if it is lower hemicontinuous on the entire X.

- *continuous at* $\boldsymbol{x} \in X$ if it is both upper and lower hemicontinuous at $\boldsymbol{x}$. It is called *continuous* if it is continuous on the entire X.

Intuitively speaking, upper hemicontinuity of a correspondence $\Gamma\colon X \twoheadrightarrow Y$ guarantees that the image set $\Gamma(\boldsymbol{x})$ of a point $\boldsymbol{x} \in X$ does not *explode* consequent on a small perturbation of $\boldsymbol{x}$, and lower hemicontinuity of Γ guarantees that the image set $\Gamma(\boldsymbol{x})$ of a point $\boldsymbol{x} \in X$ does not *implode* consequent on a small perturbation of $\boldsymbol{x}$; see Figure 4.1.

► EXERCISE 4.4. *Define* $\Phi, \Psi, \Gamma\colon [0,1] \twoheadrightarrow [0,1]$ *by*

$$\Phi(x) = \begin{cases} \{0\} & \textit{if } x \in [0,1) \\ [0,1] & \textit{if } x = 1, \end{cases}$$

$$\Psi(x) = \begin{cases} [0,1] & \textit{if } x \in [0,1) \\ \{0\} & \textit{if } x = 1, \end{cases}$$

and

$$\Gamma(x) = [0, x].$$

Show that (i) Φ is upper hemicontinuous everywhere, but it is not lower hemicontinuous at the point 1*. (ii) Ψ is lower hemicontinuous everywhere, but it is not upper hemicontinuous at* 1*. (iii) Γ is continuous.*

Definition 4.5. For any two subsets $X \subset \mathbb{R}^m$ and $Y \subset \mathbb{R}^n$, a correspondence $\Gamma\colon X \twoheadrightarrow Y$ is said to be *compact-valued* if $\Gamma(\boldsymbol{x})$ is a compact subset of Y for each $\boldsymbol{x} \in X$.

4.2 The Standard Dynamic Programming Problem

Chapter 2 deals essentially with *static optimization*, that is optimal choice at a single point of time. Many optimization problems that arise in economic models often possess a structure that is inherently *dynamic*—then involve optimization over time. Such problems are significantly more complex than the static ones in general.

The basic dynamic programming problem is to find a sequence $\{x_t\}$ that would

$$\begin{aligned} &\text{Maximize} \quad \varphi(x_0, x_1) + \sum_{t=1}^{\infty} \delta^t \varphi(x_t, x_{t+1}) \\ &\text{s.t.} \quad \begin{cases} x_{t+1} \in \Gamma(x_t), \quad t = 0, 1, 2, \dots \\ x_0 \in X \text{ given.} \end{cases} \end{aligned} \tag{DP}$$

Here,

- $X \subset \mathbb{R}^n$ is called the *state space* of the problem;
- x_0 is called the *initial state*;
- $\Gamma\colon X \twoheadrightarrow X$ is called the *transition correspondence*, which tells which states are possible “tomorrow” given the state of the system “today”;
- $\varphi\colon \mathrm{Gr}(\Gamma) \to \overline{\mathbb{R}}$ is referred to as the *return function*;[1]

[1] Here $\mathrm{Gr}(\Gamma)$ is the *graph* of Γ, i.e.,

$$\mathrm{Gr}(\Gamma) := \{(\boldsymbol{x}, \boldsymbol{y}) \in X \times Y : y \in \Gamma(\boldsymbol{x})\}.$$

- $\delta \in (0,1)$ is called the *discount factor*.

DEFINITION 4.6. A sequence $\{x_t\} \subset X$ such that $x_1 \in \Gamma(x_0)$ and $x_{t+1} \in \Gamma(x_t)$ for each t is called a *feasible plan*.

We call the series

$$\sum_{t=0}^{\infty} \delta^t \varphi(x_t, x_{t+1})$$

the *present value* of the stream of returns that obtain every period along the feasible plan $\{x_t\}$.

ASSUMPTION 1. For every feasible plan $\{x_t\} \subset X$,

$$\lim_{k\to\infty} \sum_{t=0}^{k} \delta^t \varphi(x_t, x_{t+1}) \in \overline{\mathbb{R}}.$$

Assumption 1 says that we can compute the present value of the intertemporal stream of returns that is induced by any feasible plane; that is, the series $\sum_{t=0}^{\infty} \delta^t \varphi(x_t, x_{t+1})$ converges. Throughout this chapter we suppose that Assumption 1 holds, so we consider a *standard dynamic programming problem* $(X, x_0, \Gamma, \varphi, \delta)$.

We need two further assumptions.

ASSUMPTION 2. The return function $\varphi\colon \mathrm{Gr}(\Gamma) \to \bar{\mathbb{R}}$ is continuous and bounded.

ASSUMPTION 3. The transition correspondence $\Gamma\colon X \twoheadrightarrow X$ is compact-valued and continuous.

DEFINITION 4.7. Let $(X, x_0, \Gamma, \varphi, \delta)$ be a standard dynamic programming problem. We define the class

$$\mathcal{D}(X, \Gamma, \varphi, \delta) := \{(X, x, \Gamma, \varphi, \delta) : x \in X\}, \tag{4.1}$$

which is the collection of all dynamic programming problems that differ from our original problem only in their initial states. The class of all such collections of standard dynamic programming problems is denoted by $\mathcal{DP}$.

We can rewrite our optimization problem (DP) in more familiar terms. Let $\Omega_\Gamma(x)$ stand for the set of all feasible plans for the problem $(X, x, \Gamma, \varphi, \delta)$. That

Further, $\overline{\mathbb{R}} := \mathbb{R} \cup \{-\infty, +\infty\}$.

is, define the correspondence $\Omega_\Gamma : X \twoheadrightarrow X^\infty$ by

$$\Omega_\Gamma(x) := \left\{\{x_t\} \in X^\infty : x_1 \in \Gamma(x) \text{ and } x_{t+1} \in \Gamma(x_t), t \in \mathbb{N}\right\}.$$

Define next the map $F_{\Gamma,\varphi} := \{\Omega_\Gamma(x) \times \{x\} : x \in X\} \to \bar{\mathbb{R}}$ by

$$F_{\Gamma,\varphi}\left(\{x_t\}, x\right) = \varphi(x, x_1) + \sum_{t=1}^{\infty} \delta^t \varphi(x_t, x_{t+1}).$$

Because φ is bounded (Assumption 2), the function $F_{\Gamma,\varphi}$ is bounded: there exists a $K > 0$ with $|\varphi| \leqslant K$, so

$$\left|F_{\Gamma,\varphi}\left(\{x_t\}, x\right)\right| \leqslant |\varphi(x, x_1)| + \sum_{t=1}^{\infty} \delta^t |\varphi\left(x_t, x_{t+1}\right)| \leqslant K + \sum_{t=1}^{\infty} \delta^t K = \frac{K}{1-\delta}$$

for all $x \in X$ and $\{x_t\} \in \Omega_\Gamma(x)$.

We may now rewrite our optimization problem (DP) as

$$\begin{aligned} \text{Maximize} \quad & F_{\Gamma,\varphi}\left(\{x_t\}, x\right) \\ \text{subject to} \quad & \{x_t\} \in \Omega_\Gamma(x). \end{aligned} \tag{DP$'$}$$

In the next section, we shall develop a *recursive method* to solve the problem (DP′). For simplicity, we suppose throughout this section that a solution to (DP′) exists:

ASSUMPTION 4. A solution to the problem (DP′) exists. That is,

$$\operatorname{argmax}\left\{F_{\Gamma,\varphi}\left(\{x_t\}, x\right) : \{x_t\} \in \Omega_\Gamma(x)\right\} \neq \varnothing.$$

4.3 The Principle of Optimality

The Value Function

Fix any $\mathcal{D} := \mathcal{D}(X, \Gamma, \varphi, \delta) \in \mathcal{DP}$. Define a function $V : X \to \bar{\mathbb{R}}$ by letting

$$V(x) := \max\left\{F_{\Gamma,\varphi}\left(\{x_t\}, x\right) : \{x_t\} \in \Omega_\Gamma(x)\right\}. \tag{4.2}$$

The function V is called the *value function* for the collection $\mathcal{D}$.

We will fix X, Γ, φ and δ from now on. To save notation, let us denote Ω_Γ by Ω and $F_{\Omega,\delta}$ by F.

Theorem 4.8 (Bellman). *Let $\mathcal{D} \in \mathcal{DP}$, take any $x_0 \in X$ and $\{x_t^*\} \in \Omega(x_0)$ and define $V\colon X \to \overline{\mathbb{R}}$ by (4.2). If $V(x_0) = F(\{x_t^*\}, x_0)$, then*

$$V(x_0) = \varphi(x_0, x_1^*) + \delta V(x_1^*), \tag{4.3}$$
$$V(x_t^*) = \varphi(x_t^*, x_{t+1}^*) + \delta V(x_{t+1}^*), \tag{4.4}$$

for each $t = 1, 2, \ldots$. If Assumption 2 holds, the converse is also true.

The second part of Theorem 4.8 says that if a solution to (DP′) exists, we can deduce the optimal plan from the value function V under quite general circumstances.

Proof of Theorem 4.8. **(i)** We first have

$$\begin{aligned} V(x_0) = F(\{x_t^*\}, x_0) &= \varphi(x_0, x_1^*) + \sum_{t=1}^{\infty} \delta^t \varphi(x_t^*, x_{t+1}^*) \\ &\geqslant \varphi(x_0, x_1) + \sum_{t=1}^{\infty} \delta^t \varphi(x_t, x_{t+1}) \end{aligned}$$

for all $\{x_t\} \in \Omega(x_0)$. Notice that if $\{x_2, x_3, \ldots\} \in \Omega(x_1^*)$, then $\{x_1^*, x_2, x_3, \ldots\} \in \Omega(x_0)$. We thus have

$$\varphi(x_0, x_1^*) + \sum_{t=1}^{\infty} \delta^t \varphi(x_t^*, x_{t+1}^*) \geqslant \varphi(x_0, x_1^*) + \delta\varphi(x_1^*, x_2) + \sum_{t=2}^{\infty} \delta^t \varphi(x_t, x_{t+1}),$$

that is,

$$F(\{x_2^*, x_3^*, \ldots\}, x_1^*) \geqslant F(\{x_2, x_3, \ldots\}, x_1^*),$$

for all $\{x_2, x_3, \ldots\} \in \Omega(x_1^*)$. But then $V(x_1^*) = F(\{x_2^*, x_3^*, \ldots\}, x_1^*)$, and hence

$$\begin{aligned} V(x_0) = F(\{x_t^*\}, x_0) &= \varphi(x_0, x_1^*) + \sum_{t=1}^{\infty} \delta^t \varphi(x_t^*, x_{t+1}^*) \\ &= \varphi(x_0, x_1^*) + \delta \left[\varphi(x_1^*, x_2^*) + \sum_{t=2}^{\infty} \delta^{t-1} \varphi(x_t^*, x_{t+1}^*) \right] \\ &= \varphi(x_0, x_1^*) + \delta \left[\varphi(x_1^*, x_2^*) + \sum_{t=1}^{\infty} \delta^{t} \varphi(x_{t+1}^*, x_{t+2}^*) \right] \\ &= \varphi(x_0, x_1^*) + \delta F(\{x_2^*, x_3^*, \ldots\}, x_1^*) \\ &= \varphi(x_0, x_1^*) + \delta V(x_1^*). \end{aligned}$$

This gives (4.3). By induction we obtain (4.4).

(ii) Conversely, assume that Assumption 2 holds. Let $x_0 \in X$ and $\{x_t^*\} \in \Omega(x_0)$ satisfy (4.3) and (4.4). Then

$$\begin{aligned} V(x_0) &= \varphi(x_0, x_1^*) + \delta V(x_1^*) \\ &= \varphi(x_0, x_1^*) + \delta \left[\varphi(x_1^*, x_2^*) + \delta V(x_2^*)\right] \\ &= \varphi(x_0, x_1^*) + \delta\varphi(x_1^*, x_2^*) + \delta^2 V(x_2^*) \\ &= \cdots \\ &= \varphi(x_0, x_1^*) + \sum_{t=1}^{k} \delta^t \varphi(x_t^*, x_{t+1}^*) + \delta^{k+1} V(x_{k+1}^*) \end{aligned}$$

for all $k \in \mathbb{N}$. It follows from Assumption 2 that V is bounded, so there exists a $K > 0$ such that $|V| \leqslant K$, and this clearly entails that $\delta^k V(x_k) \to 0$. Thus, letting $k \to \infty$, we obtain $V(x_0) = F(\{x_t^*\}, x_0)$. □

The Optimal Policy Correspondence

The second part of Theorem 4.8 tells us how to go from the value function V of a standard dynamic programming problem with Assumptions 2 and 4 to its *optimal path*.

Definition 4.9 (Optimal Policy Correspondence). We define the *optimal policy correspondence* for $\mathcal{D}(X, \Gamma, \varphi, \delta) \in \mathcal{DP}$ as the correspondence $P : X \twoheadrightarrow X$ with

$$P(x) := \operatorname{argmax}\{\varphi(x, y) + \delta V(y) : y \in \Gamma(x)\}.$$

It follows from Theorem 4.8 that if Assumptions 2 and 4 hold, then a sequence $\{x_t\} \subset X$ is a solution to (DP′) iff

$$\begin{aligned} x_1 &\in P(x_0), \\ x_2 &\in P(x_1), \\ &\cdots, \\ x_{t+1} &\in P(x_t), \\ &\cdots \end{aligned}$$

The Principle of Optimality

We then turn to the existence problem by the following important result:

THEOREM 4.10 (Principle of Optimality, Bellman). *For every* $\mathcal{D}(X, \Gamma, \varphi, \delta) \in \mathcal{DP}$ *and every bounded real-valued function* $W : X \to \mathbb{R}$*, if*

$$W(x) = \max\{\varphi(x, y) + \delta W(y) : y \in \Gamma(x)\} \quad \text{for all } x \in X, \tag{4.5}$$

then

$$W(x) = \max\left\{F_{\Gamma,\varphi}\left(\{x_t\}, x\right) : \{x_t\} \in \Omega_\Gamma(x)\right\} \quad \text{for all } x \in X.$$

Theorem 4.10 says that if there exists a bounded real-valued function $W : X \to \mathbb{R}$ satisfying (4.5), then W is the value function for our dynamic programming problem. Thus, by Theorem 4.8, we can deduce the optimal plan from W.

PROOF OF THEOREM 4.10. Suppose that W satisfies (4.5), and fix an arbitrary $x \in X$. By (4.5), for every $\{x_t\} \in \Omega(x)$,

$$\begin{aligned} W(x) &\geqslant \varphi(x, x_1) + \delta W(x_1) \\ &\geqslant \varphi(x, x_1) + \delta\varphi(x_1, x_2) + \delta^2 W(x_2) \\ &\cdots \\ &\geqslant \varphi(x, x_1) + \sum_{t=1}^{k} \delta^k \varphi(x_t, x_{t+1}) + \delta^{k+1} W(x_{k+1}). \end{aligned}$$

Letting $k \to \infty$ and noticing that $\lim_k \delta^{k+1} W(x_{k+1}) = 0$, we have

$$W(x) \geqslant F\left(\{x_t\}, x\right) \quad \text{for all } \{x_t\} \in \Omega(x).$$

We then show that there exists a feasible plan $\{x_t^*\} \in \Omega(x)$ such that

$$W(x) = F\left(\{x_t^*\}, x\right). \tag{4.6}$$

It follows from (4.5) that there is a sequence $\{x_t^*\} \in \Omega(x)$ such that $W(x) = \varphi(x, x_1^*) + \delta W(x_1^*)$ and

$$W(x_t^*) = \varphi(x_t^*, x_{t+1}^*) + \delta W(x_{t+1}^*) \quad m = 1, 2, \ldots$$

Hence, for each $k \in \mathbb{N}$ we have

$$\begin{aligned} W(x) &= \varphi(x, x_1^*) + \delta W(x_1^*) \\ &= \varphi(x, x_1^*) + \left[\delta\varphi(x_1^*, x_2^*) + \delta^2 W(x_2^*)\right] \\ &\cdots \\ &= \varphi(x, x_1^*) + \sum_{t=1}^{k} \delta^t \varphi(x_t^*, x_{t+1}^*) + \delta^{k+1} W(x_{k+1}^*). \end{aligned}$$

Taking limit yields (4.6). □

We close this chapter with an example.

EXAMPLE 4.11. Let $x_0 \in [0,1]$. Consider the problem of choosing a sequence $\{x_t\} \subset \mathbb{R}$ in order to

$$\begin{aligned} &\max \sum_{t=0}^{\infty} \frac{1}{2^t} \ln\left(\sqrt{x_t} - x_{t+1}\right) \\ &\text{s.t.} \quad x_{t+1} \in \left[0, \sqrt{x_t}\right], \quad t = 0, 1, \ldots \end{aligned}$$

We may view this problem as a standard dynamic programming problem $(X, x_0, \Gamma, \varphi, \delta)$, where

- $X = [0,1]$
- $\Gamma(x) = [0, \sqrt{x}]$
- $\varphi(x_t, x_{t+1}) = \ln(\sqrt{x_t} - x_{t+1})$
- $\delta = 1/2$.

We first compute the value function. Theorem 4.10 says that it is sufficient to solve the following functional equation:

$$W(x) = \max\left\{\ln\left(\sqrt{x} - y\right) + \frac{1}{2}W(y) : y \in \left[0, \sqrt{x}\right]\right\}, \quad x \in [0,1]. \tag{4.7}$$

We *guess* that

$$W(x) = \alpha \ln x + \beta, \quad \alpha \in \mathbb{R}_+ \text{ and } \beta \in \mathbb{R}.$$

With this guess, the problem is to find $(\alpha, \beta) \in \mathbb{R}_+ \times \mathbb{R}$ such that

$$\alpha \ln x + \beta = \max\left\{\ln\left(\sqrt{x} - y\right) + \frac{\alpha}{2}\ln y + \frac{\beta}{2} : y \in \left[0, \sqrt{x}\right]\right\}, \quad x \in [0,1].$$

Calculus gives

$$y = \frac{\alpha}{2+\alpha}\sqrt{x}. \tag{4.8}$$

Hence,

$$\begin{aligned} \alpha \ln x + \beta &= \ln\left(\sqrt{x} - \frac{\alpha}{2+\alpha}\sqrt{x}\right) + \frac{\alpha}{2}\ln\left(\frac{\alpha}{2+\alpha}\sqrt{x}\right) + \frac{\beta}{2} \\ &= \underbrace{\left(\frac{1}{2} + \frac{\alpha}{4}\right)}_{\alpha} \ln x + \underbrace{\left[\ln\left(1 - \frac{\alpha}{2+\alpha}\right) + \frac{\alpha}{2}\ln\left(\frac{\alpha}{2+\alpha}\right) + \frac{\beta}{2}\right]}_{\beta}. \end{aligned}$$

This gives us

$$\alpha = \frac{2}{3} \quad \text{and} \quad \beta = \ln 9 - \frac{8}{3}\ln 4.$$

Therefore, the value function is

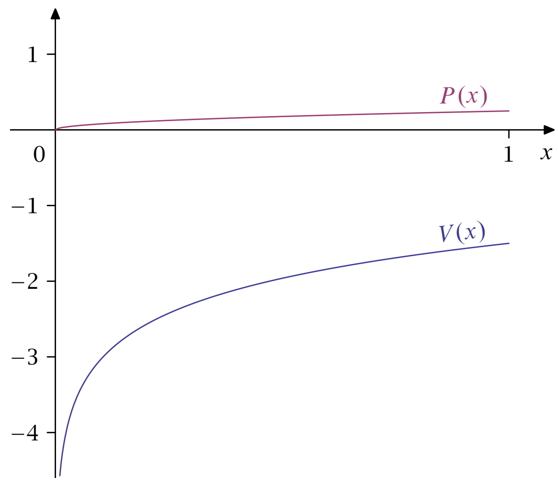

FIGURE 4.2. The value function $V(x)$ and the optimal policy correspondence $P(x)$.

$$V(x) = \frac{2}{3}\ln x + \ln 9 - \frac{8}{3}\ln 4, \quad x \in [0,1].$$

Then let us consider the optimal policy correspondence, $P(x)$. We have the following problem

$$\max\{\varphi(x,y) + \delta V(y) : y \in \Gamma(x)\},$$

which gives

$$P(x) = \left\{\frac{1}{4}\sqrt{x}\right\}, \quad x \in [0,1].$$

See Figure 4.2.

5

METRIC SPACES

THE IDEAS OF "metric" and "metric space" are abstractions of the concept of distance in Euclidean space. Most of the ideas about metric spaces in general are motivated by geometric ideas about sets in $\mathbb{R}$ or $\mathbb{R}^n$, where $n > 1$.

There are a lot of excellent introductory references for the analysis of metric spaces, e.g., Kolmogorov and Fomin (1970), Apostol (1974), Rudin (1976), Gamelin and Greene (1999, Chapter 1), Carothers (2000), Zorich (2004a), Shirali and Vasudeva (2006), Ok (2007), etc. For more advanced references, Folland (1999), Dudley (2002) and Royden and Fitzpatrick (2010) are classic texts.

This chapter is based on Carothers (2000, Part One) and Ok (2007, Chapter C). Our discussion is very elementary and incomplete. Anyone wants to learn metric spaces should refer a text from the above list.

5.1 Basic Notions

Given a set $X \neq \varnothing$, our first order of business is to define a *distance function* on X. Suppose $X = \mathbb{R}^2$. What would we want a "reasonable" distance to do?

DEFINITION 5.1 (Metric). Let X be a nonempty set. A function $d\colon X \times X \to [0, \infty)$ is called a *metric* provided for all $x, y, z \in X$,

a. $d(x, y) = 0$ iff $x = y$;

b. (Symmetry) $d(x, y) = d(y, x)$;

c. (Triangle Inequality) $d(x, y) \leqslant d(x, z) + d(z, y)$.

If d is a metric on X, we say that (X, d) is a *metric space*. If (X, d) is a metric space and Y is a subset of X, then the restriction d' of d to $Y \times Y$ is clearly a metric on Y. The metric space (Y, d') is called a *subspace* of (X, d).[1]

► EXERCISE 5.2. *Prove that $d(x, y) \geqslant 0$ for all $x, y \in X$ by (a)—(c) in Definition 5.1.*

[1] Precisely, $d' = d{\restriction}(Y \times Y)$.

REMARK 5.3. Sometimes we suppress mention of the metric d and refer to X itself as being a metric space.

EXAMPLE 5.4. **(a)** The real line $\mathbb{R}$ with the metric $d(x, y) = |x - y|$ is a metric space. More generally, $\mathbb{R}^n$ is a metric space when provided with the metric

$$d\left((x_1, \ldots, x_n), (y_1, \ldots, y_n)\right) = \sqrt{\sum_{k=1}^{n} (x_k - y_k)^2},$$

called the *usual metric* on $\mathbb{R}^n$. The space $(\mathbb{R}^n, d)$ is called the *n-dimensional Euclidean space.*

(b) Let $X \neq \varnothing$. Then X admits at least one metric: the *discrete metric* d defined by setting

$$d(x, y) = \begin{cases} 1 & \text{if } x \neq y \\ 0 & \text{if } x = y. \end{cases}$$

▶ EXERCISE 5.5. *Prove that the functions defined in Example 5.4 are metrics.*

5.2 Convergent Sequences

Open Balls

Our aim is to define and discuss the notions of convergence and continuity in metric space. Before this, we need to introduce some notation for "small" sets. Throughout this section, we will assume that we are always dealing with a generic metric space (X, d), unless otherwise specified.

DEFINITION 5.6 (Open Ball). For a point $x \in X$ and $r > 0$, the set

$$\mathcal{B}(x; r) := \left\{x' \in X : d(x', x) < r\right\}$$

is called the *open ball* centered at x of radius r.

EXAMPLE 5.7. In $\mathbb{R}$ we have $\mathcal{B}(x; r) = (x - r, x + r)$; in $\mathbb{R}^2$ the set $\mathcal{B}(x; r)$ is the open disk of radius r centered at x.

▶ EXERCISE 5.8. *Consider the discrete metric space defined in Example 5.4(b). Find $\mathcal{B}(x; 1)$ and $\mathcal{B}(x; 2)$, where $x \in X$.*

A subset $A \subset X$ is *bounded* if there is $x \in X$ and $r > 0$ such that $A \subset \mathcal{B}(x; r)$; that is, A is bounded if it is contained in some ball of large enough radius. But exactly which x and r does not much matter.

Convergence

In order to prove the existence of a vector satisfying a desired property, it is common to establish an appropriate sequence of vectors converging to a limit. In many cases the limit of this sequence can be shown to satisfy the required property. It is for this reason that the concept of convergence plays an important role in analysis. Thus, we have the convergence of sequences of real numbers, the convergence of sequences of complex numbers, the convergence of sequence of functions, etc.

We will use the following terminology to define convergence:

DEFINITION 5.9 (Neighborhood). A *neighborhood* (nhood, for short) of x is any set containing an open ball about x. Intuitively, you should think of a nhood of x as a "thick" set of points near x.

We now discuss the *convergence* of sequences that live in metric space. Its value will be made clear through seeing how often it will be called in subsequent work.

DEFINITION 5.10 (Convergence). A sequence $\{x_n\}$ is said to *converge* to the point $x \in X$ if

$$\lim_{n\to\infty} d(x_n, x) = 0.$$

In this case, x is the *limit* of $\{x_n\}$ and we write $x_n \to x$, or $\lim_{n\to\infty} x_n = x$.

REMARK 5.11. Since the definition of convergence is stated in terms of the sequence of real numbers $\{d(x_n, x)\}$, we can easily derive the following equivalent reformulations:

$$\begin{aligned}&x_n \to x \text{ iff, given any } \varepsilon > 0, \text{ there is an } N \in \mathbb{N}\\ &\text{such that } d(x_n, x) < \varepsilon \text{ whenever } n \geqslant N,\end{aligned} \tag{C1}$$

or

$$\begin{aligned}&x_n \to x \text{ iff, given any } \varepsilon > 0, \text{ there is an } N \in \mathbb{N}\\ &\text{such that } \{x_n : n \geqslant N\} \subset \mathcal{B}(x; \varepsilon).\end{aligned} \tag{C2}$$

If it should happen that $\{x_n : n \geqslant N\} \subset A$ for some $N \in \mathbb{N}$, we say that the sequence $\{x_n\}$ is *eventually* in A. Thus, our last formulation (C2) can be written

$$x_n \to x \text{ iff, given any } \varepsilon > 0, \text{ the sequence } \{x_n\} \text{ is eventually in } \mathcal{B}(x; \varepsilon), \tag{C3}$$

or, in yet another incarnation,

$$x_n \to x \text{ iff } \{x_n\} \text{ is eventually in every nhood of } x. \tag{C4}$$

This final version (C4) is blessed by a total lack of Ns and εs!

LEMMA 5.12. *The limit of a convergent sequence in a metric space is unique.*

PROOF. Suppose that $x, y \in X$ are both limits of a sequence $\{x_n\}$ in X. Then for all n,

$$d(x, y) \leqslant d(x, x_n) + d(x_n, y).$$

Let $n \to \infty$. Then the RHS of the above display tends to 0, so that $d(x, y) = 0$. Consequently, $x = y$. □

EXAMPLE 5.13. **(a)** A sequence $\{x_n\}$ is convergent in a discrete space iff it is eventually constant.

(b) Take any $n \in \mathbb{N}$, and let $\{x_m\} = \{(x_{1,m}, \dots, x_{n,m})\}$ be a sequence in $\mathbb{R}^n$. It is easy to show that $x_m \to (x_1, \dots, x_n)$ iff $x_{i,m} \to x_i$ for each $i = 1, \dots, n$.

5.3 Open Sets and Closed Sets

Open Sets

We now make precise the vague notion of a "thick" set in a metric space: a "thick" set is one that contains an entire nhood of each of its points. Let us give it a better name: open set.

DEFINITION 5.14 (Open Set). A subset U of X is said to be *open* provided for every point $x \in U$, there is an open ball centered at x that is contained in U. In other words, U is open if, given $x \in U$, there is some $r > 0$ such that $\mathcal{B}(x; r) \subset U$.

EXAMPLE 5.15. We show that an open ball is open. Take an arbitrary $x' \in \mathcal{B}(x; r)$ and define $r' = r - d(x', x)$. Let $y \in \mathcal{B}(x'; r')$. Then $d(y, x') < r'$, so that, by the triangle inequality,

$$d(y, x) \leqslant d(y, x') + d(x', x) < r' + d(x', x) = r.$$

Therefore, $\mathcal{B}(x'; r') \subset \mathcal{B}(x; r)$. See Figure 5.1.

THEOREM 5.16. *The open sets in a metric space (X, d) have the following properties:*

a. *Any union of open sets is open.*

b. *Any finite intersection of open sets is open.*

c. *$\varnothing$ and X are both open.*

PROOF. **(a)** Let $\{O_\lambda\}_{\lambda \in \Lambda}$ be a family of open subsets of X and let $O = \bigcup_{\lambda \in \Lambda} O_\lambda$. Suppose $x \in O$. Then there exists an index λ such that $x \in O_\lambda$. Since O_λ is open, there exists some $r > 0$ such that $\mathcal{B}(x; r) \subset O_\lambda$. Then $\mathcal{B}(x; r) \subset O$.

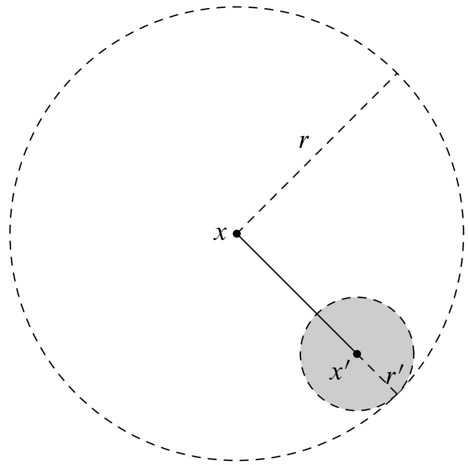

FIGURE 5.1. An open ball is open.

(b) Let $O_1, \dots, O_m$ be open subsets of X and let $O = O_1 \cap \cdots O_m$. Let $y \in O$. Since each O_i is open, there exists $r_i > 0$ such that

$$\mathcal{B}(y; r_i) \subset O_i, \quad \forall\, i = 1, \dots, m.$$

Set $r = \min\{r_1, \dots, r_m\}$.

(c) $\varnothing$ is open since $\varnothing$ has no elements and, obviously, X is open. □

REMARK 5.17. The finiteness assumption in Theorem 5.16(2) is essential. In $\mathbb{R}$, we have

$$\bigcap_{n=1}^{\infty} \mathcal{B}(0; 1/n) = \{0\},$$

which is not open in $\mathbb{R}$.

Closed Sets

DEFINITION 5.18 (Closed Set). A set F in a metric space (X, d) is said to be a *closed set* if its complement $X \smallsetminus F$ is open (see Definition 5.14).

All statements concerning open sets may be translated into statements concerning closed sets.

► EXERCISE 5.19. *Prove:*

a. $\varnothing$ *and* X *are closed.*

b. *An arbitrary intersection of closed sets is closed. A finite union of closed sets is closed.*

c. *Any finite set is closed.*

Observe that the Definition 5.18 depends on a knowledge of open sets. We would like to define a closed set in terms of, say, sequences. First observe that F is closed iff $X \smallsetminus F$ is open, and so F is closed iff

$$x \in X \smallsetminus F \Longrightarrow \mathcal{B}(x;r) \subset X \smallsetminus F \quad \text{for some } r > 0.$$

But this is the same as saying: F is closed iff

$$\mathcal{B}(x;r) \cap F \neq \varnothing \quad \text{for every } r > 0 \Longrightarrow x \in F. \tag{5.1}$$

Remark 5.20. Intuitively, a point x that satisfies $\mathcal{B}(x;r) \cap F \neq \varnothing$ for every $r > 0$ is evidently "very close" to F in the sense that x cannot be separated from F by any positive distance. At worst, x might be on the "boundary" of F. Thus, condition (5.1) is telling us that a set is closed iff it contains all such "boundary" points.

Let us now translate condition (5.1) into a sequential characterization of closed sets.

Theorem 5.21. *Given a set F in (X,d), the following are equivalent:*

a. *F is closed; that is, $X \smallsetminus F$ is open.*

b. *If $\mathcal{B}(x;r) \cap F \neq \varnothing$ for every $r > 0$, then $x \in F$.*

c. *If a sequence $\{x_n\} \subset F$ converges to some point $x \in X$, then $x \in F$.*

Proof. **(a $\Longleftrightarrow$ b)** This is clear from our observations above and the definition of an open set.

(b $\Longrightarrow$ c) Suppose that $\{x_n\} \subset F$ and $x_n \to x \in X$. Then $\mathcal{B}(x;r)$ contains infinitely many x_n for any $r > 0$, and hence $\mathcal{B}(x;r) \cap F \neq \varnothing$ for any $r > 0$. Thus $x \in F$, by (b).

(c $\Longrightarrow$ b) If $\mathcal{B}(x;r) \cap F \neq \varnothing$ for all $r > 0$, then for each n there is an $x_n \in \mathcal{B}(x;1/n)$. The sequence $\{x_n\}$ satisfies $\{x_n\} \subset F$ and $x_n \to x$. Hence, by (c), $x \in F$. □

Interior and Closure

Sets are not "doors"! For example, X is open and closed, but $(0,1]$ is neither open nor closed in $\mathbb{R}$. However, it is possible to describe the "open part" of a set and the "closure" of a set.

DEFINITION 5.22 (Interior). Given a set E in (X, d), we define the *interior* of E, written E°, to be the largest open set contained in E. That is,

$$\begin{aligned} E^\circ &= \bigcup \{U : U \text{ is open and } U \subset E\} \\ &= \bigcup \{\mathcal{B}(x;r) : \mathcal{B}(x;r) \subset E \text{ and for some } x \in E, r > 0\} \\ &= \{x \in E : \mathcal{B}(x;r) \subset E \text{ for some } r > 0\}. \end{aligned}$$

DEFINITION 5.23 (Closure). For a subset E of a metric space X, the *closure* of E, written $\bar{E}$, is the smallest closed set containing E. That is,

$$\bar{E} = \bigcap \{F : F \text{ is closed and } E \subset F\}.$$

Clearly, E° is an open subset of E and $\bar{E}$ is a closed set containing E.

You will need the following result in Mas-Colell, Whinston and Green (1995).

LEMMA 5.24. *If A is convex, then so is $\bar{A}$.*

PROOF. Take arbitrary points $\boldsymbol{x}, \boldsymbol{y} \in \bar{A}$. We need to show that $[\boldsymbol{x}, \boldsymbol{y}] \subset \bar{A}$. Let $\boldsymbol{x} \in [\boldsymbol{x}, \boldsymbol{y}]$, so $\boldsymbol{z} = \lambda\boldsymbol{x} + (1-\lambda)\boldsymbol{y}$ for some $\lambda \in [0,1]$. There exist sequences $\{\boldsymbol{x}_n\} \subset A$ and $\{\boldsymbol{y}_n\} \subset A$ such that $\boldsymbol{x}_n \to \boldsymbol{x}$ and $\boldsymbol{y}_n \to \boldsymbol{y}$. Since A is convex, we have $\boldsymbol{z}_n := \lambda\boldsymbol{x}_n + (1-\lambda)\boldsymbol{y}_n \in A$ for all $n \in \mathbb{N}$. Therefore, $\boldsymbol{z}_n \to \lambda\boldsymbol{x} + (1-\lambda)\boldsymbol{y} = \boldsymbol{z}$, so $\boldsymbol{z} \in \bar{A}$. □

► EXERCISE 5.25. *Show that if $A \subset B$, then $\bar{A} \subset \bar{B}$.*

5.4 CONTINUOUS FUNCTIONS

Throughout this section, unless otherwise specified, (X, d_X) and (Y, d_Y) are arbitrary metric spaces.

DEFINITION 5.26 (Continuous Mapping). We say that the map $f : X \to Y$ is *continuous at a point* $x \in X$ if, for any $\varepsilon > 0$, there exists a $\delta > 0$ (which may depend on f, x and ε) such that for each $y \in Y$ (see Figure 5.2)

$$d_X(x,y) < \delta \quad \text{implies} \quad d_Y\big(f(x), f(y)\big) < \varepsilon.$$

Put differently, f is continuous at x if, for any $\varepsilon > 0$, there exists a $\delta > 0$ such that

$$f\big(\mathbb{B}_X(x;\delta)\big) \subset \mathbb{B}_Y(f(x);\varepsilon).$$

If f is continuous at every point of X, we simply say that f is *continuous on* X, or often just that f is continuous.

We now turn to characterize continuity in terms of open sets, closed sets and sequences.

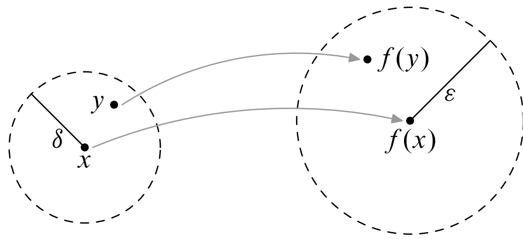

FIGURE 5.2. Inequalities in Definition 5.26 illustrated in the case of Euclidean planes $X = \mathbb{R}^2$ and $Y = \mathbb{R}^2$.

THEOREM 5.27. *Given $f : X \to Y$, the following are equivalent:*

a. *f is continuous on X (by the ε-δ definition).*

b. *For every $x \in X$, if $x_n \to x$ in X, then $f(x_n) \to f(x)$ in Y.*

c. *If E is closed in Y, then $f^{-1}(E)$ is closed in X.*

d. *If V is open in Y, then $f^{-1}(V)$ is open in X.*

PROOF. **(a $\Longrightarrow$ b)** Suppose that $x_n \to x \in X$. Given $\varepsilon > 0$, let $\delta > 0$ be such that $f(\mathbb{B}_X(x;\delta)) \subset \mathbb{B}_Y(f(x);\varepsilon)$. Since $\{x_n\}$ is eventually in $\mathbb{B}_X(x;\delta)$, the sequence $\{f(x_n)\}$ is eventually in $\mathbb{B}_Y(f(x);\varepsilon)$. Since ε is arbitrary, this means that $f(x_n) \to f(x) \in Y$.

(b $\Longrightarrow$ c) Let $E \subset Y$ be closed. To show $f^{-1}(E) \subset X$ is closed, we only need to show that for every convergent sequence $\{x_n\} \subset f^{-1}(E)$ with $x_n \to x$, we have $x \in f^{-1}(E)$. Observe that $\{x_n\} \subset f^{-1}(E)$ implies that $\{f(x_n)\} \subset E$. It follows from (b) that $f(x_n) \to f(x)$. Since E is closed, we have $f(x) \in E$ and so $x \in f^{-1}(E)$.

(c $\Longleftrightarrow$ d) It is evident since for every set $A \subset Y$ we have $f^{-1}(Y \setminus A) = X \setminus f^{-1}(A)$.

(d $\Longrightarrow$ a) Given $x \in X$ and $\varepsilon > 0$, the set $\mathbb{B}_Y(f(x);\varepsilon)$ is open in Y and so, by (d), the set $f^{-1}(\mathbb{B}_Y(f(x);\varepsilon))$ is open in X. But then there exists some $\delta > 0$ such that $\mathbb{B}_X(x;\delta) \subset f^{-1}(\mathbb{B}_Y(f(x);\varepsilon))$ since $x \in f^{-1}(\mathbb{B}_Y(f(x);\varepsilon))$. □

EXAMPLE 5.28. Let $f : X \to \mathbb{R}$ is a continuous real-valued function. Then the set $E := \{x \in X : f(x) = 0\}$ is closed.

5.5 Complete Metric Spaces

Definition 5.29 (Cauchy Sequence). A sequence $\{x_n\}$ in a metric space (X, d) is said to be a *Cauchy sequence* if for every $\varepsilon > 0$, there exists an $M \in \mathbb{N}$ (which may depend on ε) such that $d(x_k, x_\ell) < \varepsilon$ for all $k, \ell \geqslant M$.

Metric spaces in which every Cauchy sequence is convergent are of particular interest in analysis; in such spaces it is possible to identify convergent sequences without explicitly identifying their limits.

Example 5.30. **(a)** The sequence $\{1/n\}$ is a Cauchy sequence in $\mathbb{R}$ because

$$\left|\frac{1}{k} - \frac{1}{\ell}\right| \leqslant \left|\frac{1}{k}\right| + \left|\frac{1}{\ell}\right| \to 0 \quad \text{as } k, \ell \to \infty.$$

(b) Every convergent sequence $\{x_n\} \subset X$ is Cauchy (Prove it). A Cauchy sequence need not converge. For example, $\{1/n\}$ is Cauchy in $(0, 1]$, but it does not converge in $(0, 1]$ since $1/n \to 0 \notin (0, 1]$.

Remark 5.31. The difference between the definition of convergence (Definition 5.10) and the definition of a Cauchy sequence is that the limit is explicitly involved in the former, but not in the latter.

Definition 5.32 (Completeness). A metric space X is said to be *complete* if every Cauchy sequence in X converges to a point in X.

Example 5.33. All Euclidean n-spaces are complete. One should remember this important fact though we are not going to prove it.[2]

Recall that $(0, 1]$ is not a complete metric subspace of $\mathbb{R}$, while $[0, 1]$ is. This suggests a tight connection between the closedness of a set and its completeness as a metric subspace.

Proposition 5.34. *Let X be a metric space, and Y a metric subspace of X. If Y is complete, then it is closed in X. Conversely, if Y is closed in X and X is complete, then Y is complete.*

Proof. **(i)** Let Y be complete, and take an arbitrary $\{x_n\} \subset Y$ that converges in X. Then $\{x_n\}$ is Cauchy, and thus $\lim x_n \in Y$. It follows from Theorem 5.21 that Y is closed.

(ii) Suppose that X is complete and Y is closed in X. If $\{x_n\}$ is a Cauchy sequence in Y, then it is Cauchy in X. So $x_n \to x \in X$. But since Y is closed, it must be the case that $x \in Y$ (Theorem 5.21). It follows that Y is complete. □

[2] The proof can be found in Rudin (1976, Theorem 3.11).

5.6 Compact Metric Spaces

A compact space is the best of all possible worlds. Compactness is one of the most fundamental concepts of real analysis and one that plays an important role in optimization theory.

Definition 5.35 (Open Cover). Let X be a metric space and $S \subset X$. A class $\mathcal{O}$ of subsets of X is said to *cover* S if $S \subset \bigcup \mathcal{O}$. If all members of such a class $\mathcal{O}$ are open in X, then we say that $\mathcal{O}$ is an *open cover* of S.

Definition 5.36 (Compactness). A metric space X is said to be *compact* if every open cover of X has a finite subset that also covers X. A subset S of X is said to be *compact in* X (or a *compact subset of* X) if every open cover of S has a finite subset that also covers S.

Example 5.37. We claim that $(0, 1)$ is not compact in $\mathbb{R}$. Consider the collection $\mathcal{O} := \{(1/n, 1) : n = 1, 2, \ldots\}$ and observe that $(0, 1) = (1/2, 1) \cup (1/3, 1) \cup \cdots$, that is, $\mathcal{O}$ is an open cover of $(0, 1)$. However, $\mathcal{O}$ does not have a finite subset covering $(0, 1)$.

Theorem 5.38 (Characterization of Compactness for a Metric Space). *For a metric space X, the following three assertions are equivalent:*

a. *X is complete and totally bounded;*

b. *X is compact;*

c. *X is sequentially compact.*

Proof. See Royden and Fitzpatrick (2010, Section 9.5). □

References

[1] Aliprantis, Charalambos D. and Kim C. Border (2006) *Infinite Dimensional Analysis: A Hitchhiker's Guide*, Berlin: Springer-Verlag, 3rd edition. [61]

[2] Apostol, Tom M. (1967) *Calculus, Volume I: One-Variable Calculus, with an Introduction to Linear Algebra*, New York: John Wiley & Sons, Inc. [41]

[3] ——— (1974) *Mathematical Analysis*: Pearson Education, 2nd edition. [71]

[4] Aubin, Jean-Pierre and Hélène Frankowska (1990) *Set-Valued Analysis*, Boston: Birkhäuser. [61]

[5] Berkovitz, Leonard D. (2002) *Convexity and Optimization in* $\mathbb{R}^n$, Pure and Applied Mathematics: A Wiley-Interscience Series of Texts, Monographs and Tracts, New York: Wiley-Interscience. [45, 51]

[6] Carothers, N. L. (2000) *Real Analysis*, Cambridge, Massachusetts: Cambridge University Press. [71]

[7] Crossley, Martin D. (2005) *Essential Topology*, Springer Undergraduate Mathematics Series, London: Springer-Verlag. [5]

[8] Dudley, Richard Mansfield (2002) *Real Analysis and Probability*, **74** of Cambridge Studies in Advanced Mathematics, New York: Cambridge University Press, 2nd edition. [71]

[9] Duffie, Darrell (2001) *Dynamic Asset Pricing Theory*, New Jersey: Princeton University Press, 3rd edition. [61]

[10] Duggan, John (2010) "Notes on Optimization and Pareto Optimality," Rochester University, http://www.johnduggan.net/lecture[23, 27, 40]

[11] Folland, Gerald B (1999) *Real Analysis: Modern Techniques and Their Applications*, Pure and Applied Mathematics: A Wiley-Interscience Series of Texts, Monographs and Tracts, New York: John Wiley & Sons, Inc. 2nd edition. [71]

[12] Gamelin, Theodore W. and Robert Everist Greene (1999) *Introduction to Topology*, New York: Dover Publications, Inc. 2nd edition. [71]

[13] Jehle, Geoffrey A. and Philip J. Reny (2011) *Advanced Microeconomic Theory*, New York: Prentice Hall, 3rd edition. [23, 34, 51]

[14] Kolmogorov, A. N. and S. V. Fomin (1970) *Introductory Real Analysis*, New York: Dover Publications, Inc. [71]

[15] Ljungqvist, Lars and Thomas J. Sargent (2004) *Recursive Macroeconomic Theory*, Cambridge, Massachusetts: MIT Press, 2nd edition. [61]

[16] Luenberger, David G. (1969) *Optimization by Vector Space Methods*, New York: John Wiley & Sons, Inc. [23, 45, 47]

[17] Mailath, George J. and Larry Samuelson (2006) *Repeated Games and Reputations: Long-Run Relationships*, New York: Oxford University Press. [61]

[18] Mas-Colell, Andreu, Michael D. Whinston, and Jerry R. Green (1995) *Microeconomic Theory*, New York: Oxford University Press. [23, 53, 77]

[19] Milgrom, Paul (2004) *Putting Auction Theory to Work*, Cambridge, U.K.: Cambridge University Press. [44]

[20] Milgrom, Paul and Ilya Segal (2002) "Envelope Theorems for Arbitrary Choice Sets," *Econometrica*, **70** (2), pp. 583–601. [44]

[21] Munkres, James R. (1991) *Analysis on Manifolds*, Boulder, Colorado: Westview Press. [1, 18, 20]

[22] Ok, Efe A. (2007) *Real Analysis with Economic Applications*, New Jersey: Princeton University Press. [45, 61, 71]

[23] Rockafellar, R. Tyrrell (1970) *Convex Analysis*, Princeton Mathematical Series, New Jersey: Princeton University Press, 2nd edition. [45]

[24] Royden, Halsey and Patrick Fitzpatrick (2010) *Real Analysis*, New Jersey: Prentice Hall, 4th edition. [45, 71, 80]

[25] Rudin, Walter (1976) *Principles of Mathematical Analysis*, New York: McGraw-Hill Companies, Inc. 3rd edition. [16, 17, 18, 20, 71, 79]

[26] Shastri, Anant R. (2011) *Elements of Differential Topology*, Boca Raton, FL: CRC Press. [30]

[27] Shirali, Satish and Harkrishan L. Vasudeva (2006) *Metric Spaces*, London: Springer-Verlag. [71]

[28] Spivak, Michael (1965) *Calculus on Manifolds: A Modern Approach to Classical Theorems of Advanced Calculus*, Boulder, Colorado: Westview Press. [1, 18, 20]

[29] Stokey, Nancy L. and Robert E. Lucas (1989) *Recursive Methods in Economic Dynamics*, Cambridge, Massachusetts: Harvard University Press. [61]

[30] Sundaram, Rangarajan K. (1996) *A First Course in Optimization Theory*, Cambridge, Massachusetts: Cambridge University Press. [23, 26, 27, 32, 34, 36, 57]

[31] Willard, Stephen (2004) *General Topology*, New York: Dover Publications, Inc. [4]

[32] Zorich, Vladimir A. (2004a) *Mathematical Analysis I*, Universitext, Berlin: Springer-Verlag. [71]

[33] ——— (2004b) *Mathematical Analysis II*, Universitext, Berlin: Springer-Verlag. [41]

Index