bookboon.com

Analysis and Linear Algebra for Finance: Part II

Patrick Roger

Patrick Roger

Analysis and Linear Algebra for Finance: Part II

Patrick ROGER
LaRGE Research Center
EM Strasbourg Business School
University of Strasbourg

Analysis and Linear Algebra for Finance: Part II
First edition

ISBN 978-87-403-0429-9

Contents

Introduction

In the second part of the book, we address several topics that come as logical extensions of the elementary tools described in part I Chapter 1 focuses on vector spaces and linear mappings. These two concepts are especially useful in economics and finance because the payoffs of financial assets are modelized as vectors, and portfolio payoffs as combinations of vectors. The usual no-arbitrage assumption leads to value securities by means of a linear mapping linking today's price to future cash-flows.

Chapter 2, devoted to functions depending on several variables, prepares the two following chapters centered on optimization. Chapter 2 is a direct extension of chapter 2 in part I. It starts by notions of topology, and then presents partial derivatives, gradients and Hessian matrices, all fundamental tools to solve optimization problems.

Hilbert spaces are also an important part of the chapter because, in general, properties demonstrated in finite-dimensional spaces are no more valid in infinite-dimensional spaces. However, some important theorems are still true in Hilbert spaces. The projection theorem and the representation Riesz theorem are such emblematic results. These theorems are used in finance when dealing with martingales[1] and valuation problems, when the number of states of nature is infinite.

Chapters 3 and 4 develop the main optimization techniques. Chapter 3 deals with the easy problems that are not constrained, and chapter 4 shows how to transform a difficult constrained problem in an equivalent easy unconstrained one. In other words, chapter shows how to solve optimization problems using Lagrangian multipliers and Kuhn-Tucker conditions. Of course, nowadays, these problems are solved by softwares or solver in spreadsheets. However, it remains important to know how to interpret the results.

[1] See P. Roger, *Stochastic Processes for Finance*, 2010, bookboon.com

Chapter 1

Vector spaces and linear mappings

Vector spaces are probably the most useful mathematical structure in economics and finance, as in many other scientific fields. Elements of vector spaces are called vectors and the reader already knows this mathematical object, at least in an intuitive way. In fact, we live in a 3-dimensional vector space and, as a good approximation, the page you are now reading is part of a 2-dimensional vector space.

In finance and economics, vectors are generally characterized by more than 2 or 3 coordinates, and in some cases they are elements of infinite-dimensional vector spaces. Whatever the case, it is fundamental to master these mathematical tools because they are important in a number of applications like arbitrage pricing, portfolio choice, and empirical studies in general.

Section I presents the definition of a vector space and its elementary properties. In the beginning of the chapter, we restrict the presentation to finite-dimensional spaces that are natural generalizations of the 2 and 3-dimensional spaces we are used to. The mathematical concept of a *vector space* is illustrated by means of the economic concept of a *complete market*.

The second section of the chapter develops the properties of linear mappings. Linear mappings are fundamental components of arbitrage pricing models. Representation of linear mappings by matrices is developed in section 3 and the special case of square matrices is addressed in more details. In particular, we present the diagonalization of square matrices and the notions of eigenvalues and eigenvectors.

Norms and inner products, arising naturally in valuation models, are developed in section 1.4 and their properties are discussed in the general framework of Hilbert spaces in section 1.5. Finally, section 1.6 presents separation theorems and Farkas lemma. In financial theory, these results allow to link the no-arbitrage assumption to the existence of a risk-neutral probability measure[1] in an economy with a finite number of states of nature.

1.1 Vector spaces : definitions and general properties

1.1.1 Definition and examples of vector spaces

Definition 1 *A* ***vector space*** *is a set E of elements, called* ***vectors****, that can be added (addition is denoted "+" as usual) and multiplied by real numbers (multiplication by a number is denoted "."). E satisfies the following properties:*

1) $(E,+)$ *is a commutative group*[2]

2) $\forall(\alpha,\beta)\in\mathbb{R}^2, \quad \forall u\in E$

$$\alpha(\beta.u) = (\alpha\beta).u \qquad \textit{(associativity)}$$

[1]See Roger, P., Probability for Finance, 2010.

[2]It means that + is associative, has an identity element denoted $\mathbf{0}$, and any element u has an inverse denoted $-u$ satisfying $u+(-u)=\mathbf{0}$

3) $\forall(\alpha,\beta)\in\mathbb{R}^2, \quad \forall u\in E,$

$(\alpha+\beta).u = \alpha.u+\beta.u$ *(distributivity with respect to the addition in* $\mathbb{R}$*)*

4) $\forall\alpha\in\mathbb{R}, \quad \forall(u,v)\in E^2,$

$\alpha.(u+v) = \alpha.u+\alpha.v$ *(distributivity with respect to the addition in* E*.)*

5) $\forall u\in E,\ 1.u = u$

Remark : The identity element for addition is the null vector denoted $\mathbf{0}$ (in bold characters for the moment to avoid possible confusion with the real number 0).

Example 2 $\mathbb{R}$ *is a vector space, endowed with the usual addition and the usual multiplication. The above remark concerning the notation of* $\mathbf{0}$ *appears to be important here because the number 0 is simultaneously the real number 0, and the identity element of the addition for the vector space* $\mathbb{R}$*. We let the reader check that* $\mathbb{R}$ *satisfies the statements of definition 1.*

Example 3 *Let $\mathbb{R}^n$ denote the set of $n-$uples*[3] $x' = (x_1, x_2, ..., x_n)$ *where $x_i \in \mathbb{R}$ for any i; $\mathbb{R}^n$ is a vector space if addition is defined by:*

$$x + y = \begin{pmatrix} x_1 \\ x_2 \\ ... \\ x_n \end{pmatrix} + \begin{pmatrix} y_1 \\ y_2 \\ ... \\ y_n \end{pmatrix} = \begin{pmatrix} x_1 + y_1 \\ x_2 + y_2 \\ ... \\ x_n + y_n \end{pmatrix}$$

and the product by a scalar is defined by:

$$\alpha.x = \alpha. \begin{pmatrix} x_1 \\ x_2 \\ ... \\ x_n \end{pmatrix} = \begin{pmatrix} \alpha x_1 \\ \alpha x_2 \\ ... \\ \alpha x_n \end{pmatrix}$$

where $\alpha \in \mathbb{R}$.

The vector space $\mathbb{R}^n$ is the natural generalization of the usual 2 and 3-dimensional spaces.

Example 4 *Let E be a vector space and $\mathcal{A}(E)$ be the set of mappings*[4] *from E to $\mathbb{R}$; $\mathcal{A}(E)$ is a vector space if addition and product by a scalar are defined as follows:*

$$\forall (f, g) \in \mathcal{A}(E), \forall u \in E, \quad \begin{cases} (f + g)(u) = f(u) + g(u) \\ (\alpha.f)(u) = \alpha f(u) \end{cases} \tag{1.1}$$

Though these definitions seem intuitive, the space $\mathcal{A}(E)$ is much more complex than the vector space $\mathbb{R}^n$; in particular, a vector $f \in \mathcal{A}(E)$ cannot

[3] Without precision, x denotes a column vector. x' (with a prime) is the corresponding row vector called the transpose of x. These notations are consistent with the notations for matrices in part I of the book.

[4] Mappings have been defined in chapter 1 of part I of the book.

be described by a finite set of real numbers because it is a mapping from E to $\mathbb{R}$.

1.1.2 Vector subspaces

Definition 5 *Let E be a vector space and F a subset of E; F is a **vector subspace** of E if, for any $\alpha \in \mathbb{R}$ and any $v \in F$, $\alpha.v \in F$, and if conditions (2) to (5) of definition 1 are satisfied for F, when addition and multiplication by a real number ("+" and ".") are restricted to F.*

Definition 5 looks complex but its meaning is simple. F is a vector subspace of E if F is itself a vector space when it is endowed with the same addition of vectors and the same product by a real number (meaning that $\alpha.v$ should stay in F if $v \in F$ and $u + v$ is in F if u and v are in F).

The following proposition provides a simple criterion to check if a subset of E is a vector subspace.

Proposition 6 *Let E be a vector space and F a subset of E; F is a vector subspace of E if and only if:*

$$\forall(\alpha, \beta) \in \mathbb{R}^2, \quad \forall(u, v) \in F^2, \alpha.u + \beta.v \in F$$

Example 7 *Let $E = \mathbb{R}^3$ and F_1 a subset of E defined by:*

$$F_1 = \{x \in E \; / \; x_1 + x_2 + x_3 = 0\}$$

It is obvious to prove that if two vectors x and y in E satisfy the condition "the sum of their coordinates is zero", any combination $\alpha.x+\beta.y$ also satisfies the condition, (α, β) being a couple of real numbers. Therefore F_1 is a vector subspace of E. In the same way, consider the subset $F_2 = \{\mathbf{0}\}$ which contains only the vector $\mathbf{0}$. It is the smallest vector subspace of E and the only one containing a single vector.

Therefore, $\mathbf{0}$ *belongs to any vector subspace of* E *and any intersection of vector subspaces contains at least the vector* $\mathbf{0}$*. This remark is generalized in the following proposition.*

Proposition 8 *Any intersection of vector subspaces of* E *is a vector subspace of* E*.*

To illustrate this proposition, let $F_3 = \{x \in E \ / \ x_1 - 2x_2 + 3x_3 = 0\}$ and show that $F_1 \bigcap F_3$ is a vector subspace de $\mathbb{R}^3$. You can use proposition 6. On the opposite, show that $F_1 \bigcup F_3$ is not a vector subspace (Hint. choose $u^1 \in F_1$ and $u^3 \in F_3$ satisfying $u^1 + u^3 \notin F_1 \bigcup F_3$).

This latter question shows that, in general, the union of two vector subspaces is not a vector subspace (denoted V.S hereafter)

On the contrary, if we define F_{13} as follows:

$$F_{13} = \left\{ x \in \mathbb{R}^3 \ / \ x = y + z \text{ with } y \in F_1 \text{ and } z \in F_3 \right\}$$

then F_{13} is a vector subspace of $E = \mathbb{R}^3$. This remark is generalized in the proposition below.

Proposition 9 *Let* $F_1, ..., F_k$ *be* k *vector subspaces of a vector space* E *and* F *be defined by:*

$$F = \left\{ \begin{array}{c} x \in E \ / \ \exists (\alpha_1, .., \alpha_k) \in \mathbb{R}^k \textit{ and } u^1 \in F_1, ..., u^k \in F_k \textit{ such that} \\ x = \sum_{i=1}^k \alpha_i u^i \end{array} \right\}$$

F *is a V.S of* E*, called the* ***sum*** *of* F_1*,* $F_2, ..., F_k$*. We write:*

$$F = F_1 + ... + F_k \tag{1.2}$$

Proposition 9 does not say that $\alpha_1, .., \alpha_k$ and the vectors $u^1 \in F_1, ..., u^k \in F_k$ are uniquely defined for a given x. In general it is not the case and an easy counter-example is given by assuming $k = 2$ and $F_1 = F_2$.

If the decomposition is unique, we use the word "direct sum" as defined below.

Definition 10 *a) The* ***direct sum*** *of k vector subspaces $F_1, ..., F_k$ of E (if it exists), is a V.S such that any x in F can be written in a unique way as $x = \sum_{i=1}^{k} \alpha_i u^i$ where $(\alpha_1, .., \alpha_k) \in \mathbb{R}^k$ et $u^1 \in F_1, ..., u^k \in F_k$. We then note:*

$$F = \bigoplus_{i=1}^{k} F_i$$

b) If E is the direct sum of two V.S F_1 and F_2, the two subspaces are said ***supplementary****.*

The direct sum does not always exist because the decomposition of vectors is not always unique. Some subspaces F_i may have common vectors different from **0**. This intuition is formalized in the following proposition.

Proposition 11 *The direct sum of vector subspaces* $F_i, i = 1, ..., k$ *is properly defined when for any pair* (i, j), $F_i \bigcap F_j = \{\mathbf{0}\}$.

Example : Completing a financial market with option contracts

Let $E = \mathbb{R}^n$ and $x \in E$ defined by:

$$\forall i = 1, ..., n, \;\; x_i = i$$

Let $\mathbf{1}$ denote the vector in $\mathbb{R}^n$ with all coordinates equal to 1 and $y^k \in E$ defined by:

$$y^k = (y_i^k, i = 1, .., n) \text{ where } y_i^k = \max(x_i - k \; ; \, 0), k = 1, ..., n-1$$

or equivalently:

$$y^k = \max(x - k\mathbf{1} \; ; \, 0)$$

F_k is the V.S containing all the vectors proportional to y_k. We then have[5]:

$$E = \bigoplus_{k=0}^{n-1} F_k$$

where, by convention $F_0 = \{\beta x, \beta \in \mathbb{R}\}$

This relation says that any vector z in E can be uniquely decomposed as follows:

$$z = \sum_{k=0}^{n-1} \alpha_k \max(x - k\mathbf{1} \; ; \; 0) \tag{1.3}$$

The financial interpretation of this example is the following. x denotes the payoffs of a financial security (a stock or an index for example) which pays $1, 2, ...$or n depending on the state of nature that occurs at the final date[6] (there is only one future date T). The vectors y^k are payoffs of **call**

[5] The proof is left as an exercise.

[6] Our reasoning is valid as soon as payoffs x_i are different in different states. Choosing

options on x with a strike price k. A call option gives the right (and not the obligation) to its holder to buy asset x at a price k at date T.

Of course, the holder of the option buys the asset x if $x > k$ and the final net cash-flow is $x - k$. But if $x < k$ the holder of the option contract does not exercise the contract and no cash-flow is exchanged at date T. The payoff is then equal to 0.

Remark that for $k = 0$, F_0 is the V.S of vectors proportional to x. In fact there are only $n-1$ option contracts with exercise prices $k = 1, ..., n-1$. The relationship 1.3 shows that any financial security can be written as a portfolio composed of x and the $n-1$ option contracts. A financial market satisfying this property is said complete. More details on this financial example can be found in our companion book *Probability for Finance*[7].

1.1.3 Basis and dimension of a vector space

In the previous example, we have shown that it is possible to construct any vector of $\mathbb{R}^n$ by combining the reference vectors x and y^k, $k = 1, .., n-1$. It is time to properly define what means "combining" and to specify the conditions under which a subset of vectors generates a given vector space.

Spanning sets of vectors

Definition 12 *Let $u^1, u^2, .., u^k$ be vectors in E and $\alpha_1, ..., \alpha_k$ be real numbers; a* ***linear combination*** *of the $u^j, j = 1, ..., k$ with coefficients α_j is the vector v defined by:*

$$v = \sum_{j=1}^{k} \alpha_j u^j$$

payoffs equal to $1, 2, ..., n$ is not crucial but simplifies the example.

[7] The idea of completing a market by traded options was initially developed by Steve Ross in a paper entitled "Options and Efficiency", published in the Quarterly Journal of Economics in 1976.

Linear combinations are really fundamental tools in financial models because, as we saw in the previous example, v is the payoff of a portfolio when the payoffs of individual securities are the vectors u^j and the α_j denote the quantities of assets.

Proposition 13 *Let $u^1, u^2, .., u^k$ be vectors in E, and F be the set of linear combinations of vectors $u^j, j = 1, ...k$, that is:*

$$F = \left\{ x \in E / \ \exists \alpha \in \mathbb{R}^k, \ x = \sum_{j=1}^{k} \alpha_j u^j \right\}$$

then F is a vector subspace of E.

In financial terms, F is the subspace of portfolios that can be built with primary securities $u^1, ..., u^k$. This result means that, using proposition 6, a linear combination of two portfolios is a portfolio.

Definition 14 *Let $u^1, u^2, .., u^k$ be vectors in E; they are* ***linearly dependent*** *if there exist coefficients $\alpha' = (\alpha_1, ..., \alpha_k)$ with $\alpha \neq \mathbf{0}$ such that:*

$$\sum_{j=1}^{k} \alpha_j u^j = 0 \tag{1.4}$$

The set $u^1, u^2, .., u^k$ is called a linearly dependent family.

This definition says that any vector u^j in the family with a weight $\alpha_j \neq 0$, can be written as a linear combination of the other $k - 1$ vectors.

Moreover, if a family of k vectors is linearly dependent, one can add any number of new vectors to the family, it stays linearly dependent. In fact, it is sufficient to give null weights to the new vectors to find the same kind of linear combination.

Remark 15 *If a vector in E represents the payoffs of a financial security in the different states of nature, a linear combination is then a vector of* *portfolio payoffs. If a family of vectors is linearly dependent it means that you can build a portfolio generating a 0 payoff in each state. In financial terms, one of the assets is a hedge for a portfolio of the other assets. The reader can easily imagine that such a situation has some consequences on the prices of these assets, the intuitive idea being: "a portfolio that pays nothing (in all states of nature) should cost nothing". We come back to this approach of arbitrage pricing at the end of the chapter.*

Linearly independent vectors and basis of a vector space

Definition 16 *Let $u^1, u^2, .., u^k$ be a set of vectors in E; they are **linearly independent** if they are not linearly dependent. The following implication is then true.*

$$\sum_{j=1}^{k} \alpha_j u^j = 0 \Rightarrow \alpha = 0 \tag{1.5}$$

In particular, two vectors u^1 and u^2 are linearly independent if there does not exist a real number β satisfying $u^2 = \beta u^1$. The two vectors cannot be colinear if they are linearly independent.

Example 17 *Let $E = \mathbb{R}^3$ and u^1, u^2, u^3 be three vectors defined by:*

$$u^1 = \begin{pmatrix} 1 \\ 1 \\ a \end{pmatrix}; u^2 = \begin{pmatrix} 2 \\ a \\ 3 \end{pmatrix}; u^3 = \begin{pmatrix} a \\ 4 \\ -1 \end{pmatrix}$$

What are the conditions on the number a under which these three vectors are linearly independent?

We need to solve the following equations:

$$\begin{aligned} \alpha_1 + 2\alpha_2 + a\alpha_3 &= 0 \\ \alpha_1 + a\alpha_2 + 4\alpha_3 &= 0 \\ a\alpha_1 + 3\alpha_2 - \alpha_3 &= 0 \end{aligned}$$

and find if there are non zero solutions for $\alpha' = (\alpha_1; \alpha_2; \alpha_3)$

The first equation leads to

$$\alpha_1 = -2\alpha_2 - a\alpha_3 \tag{1.6}$$

We replace α_1 by its expression in the two other equations. It writes:

$$\begin{aligned} (a-2)\alpha_2 + (4-a)\alpha_3 &= 0 \qquad (1.7) \\ (3-2a)\alpha_2 - (1+a^2)\alpha_3 &= 0 \end{aligned}$$

We can now write α_2 as a function of α_3 to obtain:

$$\alpha_2 = \frac{(a-4)\alpha_3}{(a-2)} \tag{1.8}$$

In equation (1.8) a must be different from 2. If $a = 2$, it is obvious that $\alpha_3 = 0$ in the first equation of system (1.7). It implies $\alpha_2 = 0$ in the second equation and finally $\alpha_1 = 0$, showing that the three vectors are linearly independent.

Assume now that $a \neq 2$; using equation (1.8) and replacing α_2 by its value in the second equation of system (1.7) gives:

$$(3-2a)\frac{(a-4)\alpha_3}{(a-2)} - (1+a^2)\alpha_3 = 0$$

For this equation to be satisfied with $\alpha_3 \neq 0$, we need:

$$(3-2a)\frac{(a-4)}{(a-2)} - (1+a^2) = 0$$

or, equivalently:

$$\begin{aligned}(3-2a)(a-4)-(a-2)(1+a^2) &= 0\\ -a^3+10a-10 &= 0\end{aligned}$$

This equation has, at least, one solution[8]*; the three vectors are then not linearly independent.*

Remark 18 *For linearly independent families, we have a property similar (or more precisely, symmetric) to the one obtained for linearly dependent families. If k vectors are linearly independent, any subset of these k vectors is also a linearly independent family.*

Definition 19 *A family $(u^1, u^2, .., u^k)$ of vectors in E is a **spanning family** if any $x \in E$ can be written as a linear combination of $(u^1, u^2, .., u^k)$.*

$$\forall x \in E, \;\; \exists \alpha \in \mathbb{R}^k \;\; \textit{tel que} \;\; x = \sum_{j=1}^{k} \alpha_j u^j$$

[8] In chapter 2 of part I of the book (devoted to limits and continuity), we saw how this result can be obtained. Intuitively, we observe that if a is positive and large, the left hand side (LHS) of the equation is negative due to the term $-a^3$. On the opposite, if a is negative and large in absolute value, this same LHS is positive. Therefore, there is at least an a for which this LHS is equal to 0 because this third-degree polynomial is a continuous function of a.

In example of subsection 1.1.2 dealing with option contracts, we showed that x and call options on x denoted y^k, $k = 1, ..., n-1$ constitute a spanning family of $\mathbb{R}^n$. When a family $\mathcal{U}$ of vectors is a spanning family of a vector space E, it is clear that any family $\mathcal{U}^*$containing $\mathcal{U}$ is also a spanning family of E. However, if $\mathcal{U}^* \supset \mathcal{U}$ and $\mathcal{U}^* \neq \mathcal{U}$ then $\mathcal{U}^*$ is a linearly dependent family.

The natural question appearing now is: what is the "smallest" spanning family of a given vector space?

Definition 20 *A family $\mathcal{U}$ of vectors in E is a **basis** of E if $\mathcal{U}$ is a spanning and linearly independent family of E.*

When $\mathcal{U}$ is linearly dependent and spans E, it is always possible to find, for a given vector x, several linear combinations in $\mathcal{U}$ that are equal to x. In fact, assume that:

$$x = \sum_{i=1}^{n} \alpha_i u^i \tag{1.9}$$

with $\mathcal{U} = \{u^1, ..., u^n\}$. If $\mathcal{U}$ is a linearly dependent family, we can write $u^1 = \sum_{i=2}^{n} \beta_i u^i$. Replacing u^1 by its value in equation 1.9 leads to:

$$x = \sum_{i=2}^{n} (\alpha_i + \alpha_1 \beta_i) u^i$$

It is a second linear combination of vectors of $\mathcal{U}$ which is equal to x.

But if $\mathcal{U}$ is linearly independent, the decomposition of x is unique. This leads to the following proposition

Proposition 21 *A family $\mathcal{U} = \{u^1, ..., u^n\}$ is a basis of a vector space E if and only if any vector $x \in E$ can be decomposed in a unique way as a linear combination of vectors of $\mathcal{U}$.*

Proof. If $\mathcal{U}$ is a basis, any vector x can be written as a linear combination of vectors of $\mathcal{U}$. Assume that there exist two decompositions as follows:

$$\begin{aligned} x &= \sum_{i=1}^{n} \alpha_i u^i \\ x &= \sum_{i=1}^{n} \gamma_i u^i \end{aligned}$$

Substracting the second equation from the first one gives:

$$0 = x = \sum_{i=1}^{n} (\alpha_i - \beta_i) u^i$$

Equation 1.5 then implies $\alpha_i = \gamma_i$ for all i.

To prove the sufficient condition, proceed as follows. If any $x \in E$ can be written as a linear combination of vectors of $\mathcal{U}$, it means that $\mathcal{U}$ spans E. But we showed that if $\mathcal{U}$ is linearly dependent, there exist several linear combinations to obtain x. Consequently, $\mathcal{U}$ is linearly independent if the combination is unique. Therefore, $\mathcal{U}$ is linearly independent and spans E, it is then a basis. ■

For any vector x and any basis $\mathcal{U}$, x is characterized by coefficients $\alpha' = (\alpha_1; ..; \alpha_n)$ satisfying $x = \sum_{i=1}^{n} \alpha_i u^i$. These coefficients do depend on the considered basis $\mathcal{U}$. The most simple basis in $E = \mathbb{R}^n$ is called the canonical basis, denoted $e^1, ..., e^n$, where the vectors e^i are defined by :

$$e^1 = \begin{pmatrix} 1 \\ 0 \\ ... \\ 0 \end{pmatrix}; e^2 = \begin{pmatrix} 0 \\ 1 \\ ... \\ 0 \end{pmatrix}; ...e^n = \begin{pmatrix} 0 \\ 0 \\ ... \\ 1 \end{pmatrix}$$

e^i has all its components equal to 0 except the i-th which is equal to 1. Therefore, for any n-tuple $x' = (x_1, ..., x_n)$, we obtain the decomposition:

$$x = \sum_{i=1}^{n} x_i e^i$$

When E is interpreted as the set of all possible portfolio payoffs, the vectors $e^1, ..., e^n$ are financial securities called **Arrow-Debreu securities**, or pure contingent securities. They pay one unit in a given state of nature and nothing in all the other states.

Definition 22 *A vector space E is finite-dimensional if there exists a spanning family composed of a finite number of vectors. In this case, the dimension of E is the number of vectors*[9] *in a basis of E.*

This definition characterizes properly the dimension of a vector space only if all the bases of a given space have the same number of vectors. The proof of this statement is left to the reader as an exercise (hint: assume it is not true and exhibit a contradiction). From this remark, we can also deduce the following proposition.

Proposition 23 *Let F denote a V.S of a finite dimensional E with $F \neq E$. Then* $\dim(F) < \dim(E)$.

Definition 22 shows that the dimension of a vector space cannot be identified by the number of vectors in a spanning family but only by the number of elements in a basis (linearly independent spanning family). Therefore, in a family of vectors included in a finite-dimensional space E, there exists a maximum number of linearly independent vectors which is equal to the dimension of the space.

[9] By convention, the vector space containing only the null vector is 0-dimensional.

Proposition 24 *Let $\mathcal{U} = (u^1, ..., u^n)$ be a basis of E and x denote a vector in E; the family $\mathcal{U}_x = (u^1, ..., u^n, x)$ is linearly dependent.*

Proof. x can be written $\sum_{i=1}^{n} x_i u^i$ since $\mathcal{U}$ is a basis. Therefore, we can find $(\alpha_1, ..., \alpha_n, \alpha_{n+1})$ with at least one of these cefficients different from 0 such that:

$$\sum_{i=1}^{n} \alpha_i u^i + \alpha_{n+1} x = 0$$

It is enough that $\alpha_i = x_i$ and $\alpha_{n+1} = -1$. Definition 14 implies that $\mathcal{U}_x$ is linearly dependent. ■

Let $\mathcal{U}$ be a set of financial securities and x a new financial contract introduced on the market. The above proposition shows that the payoffs of the new asset x can be replicated by a portfolio of securities [10] of $\mathcal{U}$. In this situation, x is called a redundant asset. Later on, we analyze the consequences of this remark on the evaluation of financial securities.

[10] Replication means that the future payoffs of x are identical to the future payoffs of the "replicating" portfolio.

Definition 25 *Let $\mathcal{U}$ be a family of vectors in E; the* ***rank*** *of $\mathcal{U}$, denoted $rk(\mathcal{U})$, is the maximum number of vectors in $\mathcal{U}$ that are linearly independent.*

In proposition 13 we showed that the set of linear combinations of a subset of vectors of E is a V.S of E. Moreover, the rank of $\mathcal{U}$ when we add to $\mathcal{U}$ a linear combination of vectors of $\mathcal{U}$ does not change. Consequently, we obtain the following proposition.

Proposition 26 *1)Let $\mathcal{U}$ be a family of vectors in E; the set of linear combinations of vectors in $\mathcal{U}$ is a V.S of dimension $p = rk(\mathcal{U})$.*

2) Let v denote a vector which is not a linear combination of vectors of $\mathcal{U}$; we then have:

$$rk(\mathcal{U} \bigcup \{v\}) = rk(\mathcal{U}) + 1$$

To illustrate the second part of the proposition, consider a vector $x \in E$ defined by :

$$x = \begin{pmatrix} 1 \\ 2 \\ ... \\ n \end{pmatrix}$$

Denote y the vector defined by $y = \max(x - k\mathbf{1};0)$ with $k \in \mathbb{N}$, and $\mathbf{1}$ the vector in E with all coordinates equal to 1. As far as $n > k > 0$, y and x are linearly independent, that is not colinear[11]. Therefore $rk(\{x, y\}) = rk(\{x\}) + 1 = 2$. One more time, if x is interpreted as the possible future prices of a stock, y is the vector of future payoffs of a call option on x with exercise price k. We observe that portfolios (linear combinations) based on x and y span a two-dimensional space when asset x only spans a 1-dimensional space. It explains why (at least in theory) options are able to improve the allocation of risk and resources in the economy. This property had been already mentioned in example of subsection 1.1.2.

[11] When two vectors are linearly dependent, they are said "colinear".

1.2 Linear mappings

In finance models, linear mappings are fundamental because of the large number of applications in which they are involved. It is especially the case in the theory of valuation based on the no-arbitrage assumption. In fact, a fundamental result of this approach is that when the market is free from arbitrage opportunities, the mapping linking future cash-flows to current prices is linear.

1.2.1 Definitions and notations

Definition 27 *Let E_1 and E_2 be two vector spaces; a mapping f from E_1 to E_2 is **linear** if:*

1) $\forall(u,v) \in E_1 \times E_1, f(u+v) = f(u) + f(v)$

2) $\forall \alpha \in \mathbb{R}, \forall u \in E, \;\; f(\alpha.u) = \alpha.f(u)$

As before, it is worth to notice that, in $f(u+v)$, the "+" sign refers the addition of vectors in the space E_1 but the same sign "+" in $f(u) + f(v)$ refers to the addition in E_2. Remember that the two additions may be quite different, depending on the characterizations of E_1 and E_2. The same remark should be done for the multiplication by a real number, even if the difference is less striking.

Linearity of a mapping f means that the image of a linear combination of vectors in E_1 is the linear combination of images in E_2 with the same coefficients. The following proposition formalizes this remark. It is sometimes used as the definition of a linear mapping.

Proposition 28 *A mapping $f : E_1 \to E_2$ is linear if and only if for any family $(u^1, ..., u^p)$ of vectors in E_1 and any p-tuple $(\alpha_1, ...\alpha_p) \in \mathbb{R}^p$, the following equality is satisfied:*

$$f\left(\sum_{i=1}^{p} \alpha_i u^i\right) = \sum_{i=1}^{p} \alpha_i f(u^i)$$

Remark 29 ***a)*** *The definition of linearity implies immediately* $f(\mathbf{0}_{E_1}) = \mathbf{0}_{E_2}$ *if* $\mathbf{0}_{E_1}$ *and* $\mathbf{0}_{E_2}$ *denote the null vectors of the two spaces. Moreover, for any* $u \in E_1$, $f(-u) = -f(u)$.

b) *If* $\mathcal{U} = \{u^1, ..., u^p\}$ *denotes a linearly dependent family in* E_1 *then* $f(\mathcal{U}) = \{f(u^1), ..., f(u^p)\}$ *is a linearly dependent family in* E_2*. On the opposite, if* $\mathcal{U} = \{u^1, ..., u^p\}$ *is a linearly independent family in* E_1, $f(\mathcal{U})$ *is not always a linearly independent family in* E_2 *(the proof is left to the reader; it is sufficient to consider the mapping* $u \to f(u) = u_1\mathbf{1}$ *where* u_1 *is the first coordinate of* u *and* $\mathbf{1}$ *is, as usual, the vector with all coordinates equal to* 1*).*

1.2.2 Kernel and image of a linear mapping

Definition 30 *1)Let* f *denote a linear mapping from* E_1 *to* E_2*; the* ***kernel*** *of* f *denoted* $Ker(f)$ *is the set of vectors* $u \in E_1$ *satisfying* $f(u) = 0$.

2) The image of f*, denoted* $\mathrm{Im}(f)$ *is the subset of* E_2 *defined by :*

$$\mathrm{Im}(f) = \{y \in E_2 \ / \ \exists \ x \in E_1 \text{ such that } y = f(x)\}$$

$Ker(f)$ is then a subset of E_1 equal to the reciprocal image of the null vector in E_2 (sometimes written $f^{-1}(\mathbf{0}_{E_2})$). On the contrary $\mathrm{Im}(f)$ is a subset of E_2, sometimes written $f(E_1)$ because it contains all vectors in E_2 that can be written $f(u)$ with $u \in E_1$.

Proposition 31 $Ker(f)$ *and* $\mathrm{Im}(f)$ *are V.S of* E_1 *and* E_2 *respectively.*

Proof. Using proposition 6, it is enough to show, for the kernel $Ker(f)$:

$$\forall (u, v) \in Ker(f) \times Ker(f), \forall (\alpha, \beta) \in \mathbb{R}^2, \alpha.u + \beta.v \in Ker(f)$$

The linearity of f implies that[12]:

$$f(\alpha.u + \beta.v) = \alpha.f(u) + \beta.f(v) = \alpha.0 + \beta.0 = 0$$

For the image, we have to prove that:

$$\forall (x, y) \in \text{Im}(f) \times \text{Im}(f), \forall (\alpha, \beta) \in \mathbb{R}^2, \alpha.x + \beta.y \in \text{Im}(f)$$

Let u and v be two vectors in E_1 such that $f(u) = x$ and $f(v) = y$. We can write:

$$\alpha.x + \beta.y = \alpha.f(u) + \beta.f(v) = f(\alpha.u + \beta.v)$$

Therefore $\alpha.x + \beta.y$ is the image of $\alpha.u + \beta.v$ through f, implying $\alpha.x + \beta.y \in \text{Im}(f)$. ■

This proposition shows in particular that the image of E by f is a V.S of E_2. With the same technique of proof as before, we can demonstrate the following proposition.

[12] We come back here to standard notations where 0 denotes the null vector in either space, assuming that the reader is now able to identify the reference space if necessary.

Proposition 32 *Let F_1 be a V.S of E_1; $f(F_1)$ is a V.S of E_2.*

Remmember that $\{0\}$ is a V.S of E_2 and the reciprocal image of $\{0\}$ by f is the V.S $Ker(f)$. This remark can be generalized as follows.

Proposition 33 *If F_2 is a V.S of E_2, $f^{-1}(F_2)$ is a V.S of E_1.*

An important property is to characterize the relationship between the kernel dimension and the properties of f. In particular, the question is to know if a vector $u \neq 0$ can satisfy $f(u) = 0$

Proposition 34 *If f is injective then $Ker(f) = \{0\}$*

Proof. f injective means that $y \neq x \Longrightarrow f(y) \neq f(x)$. As f is a linear mapping, this implication writes:

$$y - x \neq 0 \Rightarrow f(y - x) \neq 0$$

and therefore $Ker(f) = \{0\}$. The reciprocal goes as follows. If $Ker(f) = \{0\}$ and if there exist x and y, $x \neq y$ satisfying $f(x) = f(y)$, an obvious contradiction arises because $f(x - y) = 0$, meaning that $x - y \in Ker(f)$. ■

Proposition 35 *If f is surjective then* $\mathrm{Im}(f) = E_2$

Proof. This result is obvious because f surjective means that any vector in E_2 has a reciprocal image in E_1 ■

Definition 36 *A bijective linear mapping from E_1 to E_2 is called an **isomorphism**.*

This notion of isomorphism is fundamental when it comes to associate a space of linear mappings to a space of matrices, or a vector space to its dual space, as we will see in the next section.

Proposition 37 *Two vector spaces E_1 and E_2 are* ***isomorphic***[13] *if and only if their dimensions are equal.*

Proof. Denote n and p the respective dimensions of E_1 and E_2; let $\mathcal{U}$ and $\mathcal{V}$ be bases of these two spaces and f be a bijective linear mapping from E_1 to E_2.

We first show that "f injective" is equivalent to "$f(u^1), ..., f(u^n)$ are linearly independent".

If f is injective, $Ker(f) = \{0\}$. Therefore, any linear combination $\sum_{i=1}^n x_i u^i$ satisfies:

$$\sum_{i=1}^n x_i u^i = 0 \Leftrightarrow x_i = 0 \text{ for any } i \tag{1.10}$$

In this case,

$$f\left(\sum_{i=1}^n x_i u^i\right) = \sum_{i=1}^n x_i f(u^i) = 0$$

which shows that the vectors $f(u^i)$ are linearly independent. The reciprocal goes as follows: if the $f(u^i), i = 1, ..., n$ are linearly independent, we can write:

$$\sum_{i=1}^n x_i f(u^i) = 0 \Leftrightarrow x_i = 0 \text{ for any } i \tag{1.11}$$

but the linearity of f implies that $\sum_{i=1}^n x_i u^i \in Ker(f)$. Relation (1.11) then implies $Ker(f) = 0$ and f is injective. It follows directly that $n \leq p$.

As f is also surjective, $\text{Im}(f) = E_2$ and the rank of the family of vectors $f(u^1), ..., f(u^n)$ is p, meaning that $n \geq p$.

We show now that if E_1 and E_2 have equal dimensions, they are isomorphic. The basis $\mathcal{V}$ with p vectors defines a linear mapping f from E_1 to E_2 such that $\mathcal{V}$ is the image of a family $\mathcal{W}$ of E_1. We then have $rk(\mathcal{W}) = n = \dim(E_1)$ proving that f is injective.

[13] "isomorphic" means that there exists an isomorphism between E_1 and E_2.

But we also know that $\dim(E_1) = \dim(E_2)$; therefore f is also surjective. As a result, f is bijective and E_1 and E_2 are isomorphic. ■

1.2.3 The space of linear mappings

The set of linear mappings defined on a vector space E_1 and taking values in a vector space E_2 is denoted by $\mathcal{L}(E_1, E_2)$, or simply $\mathcal{L}$ when no confusion is possible.

In example 4, we showed that the set of mappings defined on a vector space E and taking values in $\mathbb{R}$ is a vector space. The property is still valid if $\mathbb{R}$ is replaced by another vector space. Therefore, $\mathcal{L}(E_1, E_2)$ is a subset of $\mathcal{A}(E_1, E_2)$. As the elements in $\mathcal{L}$ are linear mappings, we have the following property.

Proposition 38 *$\mathcal{L}$ is a V.S of $\mathcal{A}(E_1, E_2)$*

Proof. Let $(f, g) \in \mathcal{L}^2$ and $(\alpha, \beta) \in \mathbb{R}$; for any couple of vectors (x, y) of $E_1 \times E_1$, we have:

$$\begin{aligned} (\alpha f + \beta g)(x + y) &= \alpha f(x + y) + \beta g(x + y) \\ &= \alpha f(x) + \alpha f(y) + \beta g(x) + \beta g(y) \\ &= (\alpha f + \beta g)(x) + (\alpha f + \beta g)(y) \end{aligned}$$

For any $x \in E_1$ and $\gamma \in \mathbb{R}$, we also have:

$$\begin{aligned} (\alpha.f + \beta.g)(\gamma x) &= \alpha f(\gamma x) + \beta g(\gamma x) \\ &= \gamma \alpha f(x) + \gamma \beta g(x) \\ &= \gamma (\alpha f + \beta g)(x) \end{aligned}$$

■

One of the most common situations appears if E_1 is a general vector space[14] and E_2 is $\mathbb{R}$. In this situation, the elements of $\mathcal{L}(E_1, \mathbb{R})$ are called **linear functionals (or one-forms)**; we will see later on that the mapping linking a function to its integral is a linear functional. In the same way, the expectation operator is a linear functional defined on a space on random variables. In finance models, the no-arbitrage assumption implies that the valuation operator mapping future cash-flows and today prices is a linear functional[15].

Definition 39 *Let E be a vector space; the set $\mathcal{L}(E, \mathbb{R})$ of linear functionals defined on E is called the* ***dual space*** *of E.*

When E is a finite-dimensional space of dimension n, its dual space satisfies the following property.

Proposition 40 *If* $\dim(E) = n < +\infty$; $\dim(\mathcal{L}(E, \mathbb{R})) = \dim(E)$

Proof. Let $(u^1, .., u^n)$ be a basis of E and $x \in E$ written as:

$$x = \sum_{i=1}^{n} x_i u^i \tag{1.12}$$

Consider $f^1, ..., f^n$, a set of linear functionals defined by:

$$\forall i = 1, ..., n; \forall x \in E, f^i(x) = x_i$$

The family $\mathcal{F} = (f^1, ..., f^n)$ spans $\mathcal{L}(E, \mathbb{R})$. In fact, for any given linear functional g, we have:

$$g(x) = \sum_{i=1}^{n} x^i g(u^i) = \sum_{i=1}^{n} g(u^i) f^i(x) = \left(\sum_{i=1}^{n} g(u^i) f^i \right)(x) \tag{1.13}$$

[14] In particular, E_1 may be an infinite-dimensional space of random variables in probability frameworks.

[15] A financial security is defined by the future cash-flows it generates. In general, financial securities can be represented as elements of a general vector space.

It shows that g can be written as a linear combination of the f^i.

We now show that $\mathcal{F}$ is a linearly independent family. The equality $\sum_{i=1}^{n} \alpha_i f^i = 0$ means:

$$\forall x \in E, \sum_{i=1}^{n} \alpha_i f^i(x) = 0$$

but, according to the definition of f^i, this equality is equivalent to:

$$\sum_{i=1}^{n} \alpha_i x_i = 0$$

For this equality to be satisfied by any x, it is necessary that all the α_i are equal to zero[16], and this ends the proof. ■

[16] The linear mapping f^i is called the **projection** on u^i.

1.3 Finite-dimensional spaces and matrices

In this section, we assume that the two spaces E_1 and E_2 we refer to are finite-dimensional. The notations are unchanged. E_1 (E_2) denotes a vector space of dimension n (p). $\mathcal{L}(E_1, E_2)$ is the space of linear mappings from E_1 to E_2.

1.3.1 Representation of a linear mapping by a matrix

Proposition 41 *Let $\mathcal{U} = (u^1, .., u^n)$ be a basis of E_1 and $\mathcal{V} = (v^1, v^2, ..., v^p)$ be a basis of E_2; any mapping $f \in \mathcal{L}(E_1, E_2)$ is completely defined by the family of vectors $f(u^1), ..., f(u^n)$ expressed in the basis $\mathcal{V}$.*

Proof. Let $x \in E_1$ such that $x = \sum_{i=1}^n x_i u^i$; $f(x)$ can be written:

$$f(x) = f\left(\sum_{i=1}^{n} x_i u^i\right) = \sum_{i=1}^{n} x_i f(u^i)$$

Therefore, if the images $f(u^i)$ of the vectors of the family $\mathcal{U}$ are known, it is possible to characterize the image of any vector x. Each vector $f(u^i)$ belongs to E_2, it is then a p-dimensional vector. The linear mapping f is then completely specified by $n \times p$ numbers equal to the coordinates of the n vectors $f(u^i)$, $i = 1, ..., n$. ■

If we denote $f(u^i)$ as follows:

$$f(u^i) = \begin{pmatrix} a_{1i} \\ a_{2i} \\ \cdots \\ a_{pi} \end{pmatrix}$$

we can introduce the following definition.

Definition 42 *The **matrix of the linear mapping** f, denoted $M_f(\mathcal{U}, \mathcal{V})$ is a $p \times n$ matrix whose columns are the vectors $f(u^i), i = 1, ..., n$.*

According to the above notation for $f(u^i)$, we have:

$$M_f(\mathcal{U}, \mathcal{V}) = \begin{bmatrix} a_{11} & a_{12} & .... & a_{1n} \\ a_{21} & ... & ... & ... \\ ... & ... & ... & ... \\ a_{p1} & ... & ... & a_{pn} \end{bmatrix} \tag{1.14}$$

The notation $M_f(\mathcal{U}, \mathcal{V})$ is cumbersome but it is used here to emphazise that the matrix representing f depends on the two bases on E_1 and E_2. Of course, in the next sections, we will simply write M_f when no confusion can be made.

The above remarks show that being given $\mathcal{U}$ and $\mathcal{V}$, the matrix M_f is linked to f. But more generally any $p \times n$ matrix defines a linear mapping from E_1 (of dimension n) to E_2 (of dimension p). The most usual case is the one where $\mathcal{U}$ and $\mathcal{V}$ are the canonical[17] bases of E_1 and E_2.

1.3.2 Compounding linear mappings

Consider three vector spaces E_1, E_2, E_3 with dimensions n, p, m and bases $\mathcal{U}, \mathcal{V}, \mathcal{W}$; let f denote a linear mapping from E_1 to E_2 and g a linear mapping from E_2 to E_3. In general we describe this sequence as follows:

$$E_1 \xrightarrow{f} E_2 \xrightarrow{g} E_3 \tag{1.15}$$

Compounding the mappings f and g aims at defining a new mapping from E_1 to E_3.

[17] Remember that the canonical basis is the basis for which vectors have all their coordinates equal to 0, except one which is equal to 1. For example in $\mathbb{R}^3$, the canonical basis is

$$u_1 = \begin{pmatrix} 1 \\ 0 \\ 0 \end{pmatrix}; u_2 = \begin{pmatrix} 0 \\ 1 \\ 0 \end{pmatrix}; u_3 = \begin{pmatrix} 0 \\ 0 \\ 1 \end{pmatrix}$$

Definition 43 *The* ***compound mapping*** *of f and g is the mapping denoted $g \circ f$ from E_1 to E_3 defined by:*

$$\forall x \in E_1; \ g \circ f(x) = g\left[f(x)\right]$$

x writes $\sum_{i=1}^{n} x_i u^i$; but f and g are linear, so we have:

$$g \circ f(x) = g\left[f(x)\right] = g\left[\sum_{i=1}^{n} x_i f(u^i)\right] = \sum_{i=1}^{n} x_i g \circ f(u^i)$$

This equality shows that $g \circ f$ is a linear mapping.

The compounding of linear mappings is linked to the product of matrices. Denote M_f and M_g the matrices associated to the mappings f and g, defined as before. For any $x \in E_1$, $f(x) = M_f x$. The vector $f(x)$ belongs to E_2; as such it has p coordinates. Therefore, the image of $f(x)$ by g is obtained by a premultiplication of $f(x)$ by M_g. This leads to:

$$g \circ f(x) = M_g\left(M_f x\right) = M_g M_f x = M_{g \circ f}(x)$$

To calculate $M_{g \circ f}$, we apply successively f and g. Denote for example:

$$\begin{aligned} M_f &= \begin{pmatrix} a_{11} & a_{12} & \ldots & a_{1n} \\ a_{21} & \ldots & \ldots & \ldots \\ \ldots & \ldots & \ldots & \ldots \\ a_{p1} & \ldots & \ldots & a_{pn} \end{pmatrix} \\ M_g &= \begin{pmatrix} b_{11} & b_{12} & \ldots & b_{1p} \\ b_{21} & \ldots & \ldots & \ldots \\ \ldots & \ldots & \ldots & \ldots \\ b_{mp} & \ldots & \ldots & b_{mp} \end{pmatrix} \end{aligned}$$

The generic element of $M_{g\circ f}$ is c_{ik} defined as:

$$c_{ik} = \sum_{j=1}^{p} b_{kj} a_{ji}$$

We then observe that $M_{g\circ f} = M_g M_f$, the product of M_g and M_f defined in part I of the book (chapter 4).

The following proposition is a special case of this relationship.

Proposition 44 *A matrix A associated to a linear mapping f is invertible if and only if f is a bijection. The matrix representing f^{-1} is the inverse of A denoted A^{-1}.*

In fact, if B is associated to f^{-1}, the relationship $AB = I_n$ means that $B = A^{-1}$. We obtain the following corollary.

Corollary 45 *1) Let A denote a (n, n) invertible matrix. For any $u \in \mathbb{R}^n$, the system of equations $Ax = u$ has a solution $x \in \mathbb{R}^n$ defined by $x = A^{-1}u$.*

2) If a matrix A is invertible, its columns are linearly independent vectors of $\mathbb{R}^n$.

We mention these two properties as corollaries but they are equivalent to proposition 44. When the columns of A are future payoffs of financial assets in the n states of nature, the columns of A^{-1} are the quantities of securities to be held to duplicate the Arrow-Debreu securities because $AA^{-1} = I_E$.

Example 46 ***Discounting***

In chapter 4 of part I, we showed that a bank can create contracts paying a single future cash-flow by combining bonds of different maturities. We calculated the prices of the contracts using the prices of the bonds.

Suppose that the cash-flows of the three bonds are stored in a matrix M as follows:

$$M = \begin{pmatrix} 104 & 6 & 4 \\ 0 & 106 & 4 \\ 0 & 0 & 104 \end{pmatrix}$$

The price vector π *is:*

$$\pi = \begin{pmatrix} 99,5 \\ 100,4 \\ 99,6 \end{pmatrix}$$

Denote now f *the linear mapping defined from* $\mathbb{R}^3$ *to* $\mathbb{R}$ *such that* $f(M^j) = \pi_j$*, were* M^j *is the* j*-th column of* M *(the cash-flows generated by the* j*-th bond) and* π_j *is the price of bond* j[18]*. The mapping* f *is represented by a vector denoted* A *:*

$$A = \begin{pmatrix} a_1 \\ a_2 \\ a_3 \end{pmatrix}$$

So we have:

$$\begin{aligned} 100 \times a_1 + 0 \times a_2 + 0 \times a_3 &= 99,5 \\ 6 \times a_1 + 106 \times a_2 + 0 \times a_3 &= 100,4 \\ 4 \times a_1 + 4 \times a_2 + 104 \times a_3 &= 99,6 \end{aligned}$$

This system can be written as:

$$M'A = \pi$$

But M *is triangular with non-zero diagonal terms, it is then invertible. Consequently:*

$$A = (M')^{-1}\,\pi$$

The elements in A *are in fact the prices of zero-coupon bonds with respective maturities 1, 2 and 3 years.*

[18] We assume here that this mapping is linear. In fact it is true when the market is arbitrage-free.

1.3.3 The vector space of matrices

Being given two bases $\mathcal{U}$ and $\mathcal{V}$ on the vector spaces E_1 and E_2 (of dimensions n and p), the set $\mathcal{M}_{pn}$ of matrices with p rows and n columns must have the same structure as the set $\mathcal{L}(E_1, E_2)$ of linear mappings from E_1 to E_2, that is a structure of vector space. This is the only way to make coherent the structure of vector space of $\mathcal{L}(E_1, E_2)$ with the operators (addition and product by a real number) on $\mathcal{M}_{pn}$. This relationship is formalized in the following proposition.

Proposition 47 *$\mathcal{M}_{pn}$ and $\mathcal{L}(E_1, E_2)$ are isomorphic and the dimension of these two spaces is $p \times n$.*

Proof. If we prove these two spaces are isomorphic, we will be able to conclude that the dimensions are equal because of proposition 37.

The mapping which links $f \in \mathcal{L}(E_1, E_2)$ to M_f is bijective because M_f is defined by the images of the basis vectors of E_1.

Let us denote A_{ij} the matrix with all null elements, except the one on i-th row and the j-th column which is equal to 1. Any matrix $A = (a_{ij}, i = 1, ..., p; j = 1, .., n)$ can be decomposed in a unique manner as:

$$A = \sum_{i=1}^{p} \sum_{j=1}^{n} a_{ij} A_{ij}$$

Therefore the family $(A_{ij}, i = 1, ..., p; j = 1, .., n)$ is a basis of $\mathcal{M}_{pn}$ and has $p \times n$ elements. This ends the proof. ■

Example 48 *Without entering into technical details, consider a set of mutual funds, each fund being a portfolio of individual securities. In a model with one period and n states of nature, the future payoffs of the individual securities (let K be the number of traded securities) can be summarized in a matrix D with n rows and K columns. The mapping "portfolio", denoted f, from $\mathbb{R}^K$ to $\mathbb{R}^n$ associates a vector θ to a vector $f(\theta) = D\theta$ where the* *vector θ contains the quantities of securities in the portfolio and the vector $D\theta$ represents the future payoffs of the portfolio θ. Let now g be the mapping from $\mathbb{R}^n$ to $\mathbb{R}$ which links to any vector x a number $g(x) = \sum p_i x_i$ where $P = (p_i, i = 1, ..., n)$ is the vector of probabilities of occurrence of states $i = 1, ..., n$.*

If x represents the payoffs of a portfolio in the different states of nature, $g(x)$ is the expected payoff of the portfolio. Consequently, the product $PD\theta$ is the expected payoff of the mutual fund θ. It can also be written $gof(\theta)$.

1.3.4 The special case of square matrices

A square matrix has, by definition, the same number of rows and columns. In particular, the product of two (n, n) square matrices is still a (n, n) square matrix. In other words, the result of the product stays in the same space $\mathcal{M}_n$. Consider three vector spaces E_1, E_2 and E_3, all of dimension n. The product $M_g M_f$ is a (n, n) matrix if f is a mapping from E_1 to E_2 and g a mapping from E_2 to E_3.

The most important case addressed in what follows is $E_1 = E_2 = E_3 = E$.

Definition 49 *A linear mapping from E to E is called an* ***endomorphism.***

We just described the links between matrices and linear mappings. Of course, an endomorphism is represented by a matrix in $\mathcal{M}_n$. But we also know by chapter 2 of part I that if a mapping f is bijective, it has an inverse denoted f^{-1} and satisfying $f \circ f^{-1} = f^{-1} \circ f = i_E$ where i_E is the identity mapping of E defined by $i_E(x) = x$ for any $x \in E$.

i_E is obviously linear and it is easy to see that the matrix I_n represents i_E where I_n is defined by:

$$I_n = \begin{bmatrix} 1 & 0 & ... & ... \\ 0 & 1 & & \\ ... & & ... & \\ ... & & & 1 \end{bmatrix}$$

The reader can check that $I_n x = x$ for any $x \in \mathbb{R}^n$. The matrix I_n is called the **identity matrix**. Finally, we also know that the matrix associated to the compound of two mappings is the product of the matrices of the two mappings involved in the compounding. From all these remarks, we deduce that if A and B are matrices representing f and f^{-1}, we have:

$$AB = I_n$$

Proposition 50 *If f is a bijective endomorphism of E represented by a matrix M_f, f^{-1} is represented by the inverse matrix M_f^{-1}.*

Determinants

Knowing if the determinant of a square matrix is zero allows to know if this matrix is invertible. In a more geometric approach, a zero determinant means that the columns (or rows) of a given matrix are linearly dependent.

Determinant of a (2,2) matrix Denote x and y two vectors in $\mathbb{R}^2$. They are colinear if there exists $\alpha \in \mathbb{R}$ satisfying $y = \alpha x$. This equality is equivalent to:

$$\begin{aligned} y_1 &= \alpha x_1 \\ y_2 &= \alpha x_2 \end{aligned}$$

From these two equations we deduce $x_1 y_2 - x_2 y_1 = 0$; on the contrary, if $x_1 y_2 - x_2 y_1 \neq 0$, the vectors x and y are linearly independent. If x and y are the two columns of a square matrix A, A is invertible if $x_1 y_2 - x_2 y_1 \neq 0$. This remark justifies the definition of the determinant of a $(2,2)$ matrix.

Definition 51 *Let A be a $(2,2)$ matrix with generic term $a_{ij}, i,j = 1,2$. The* ***determinant*** *of A (denoted* $\det(A)$*) the number:*

$$\det(A) = a_{11}a_{22} - a_{12}a_{21}$$

Determinants of larger matrices are defined by induction. The determinant of a (n,n) matrix is a function of determinants of $(n-1,n-1)$ matrices.

The general case

Definition 52 *1) Let A be a (n,n) matrix. Let D_{ij} be the determinant of the matrix deduced from A by deleting the i-th row and the j-th column of A. The (i,j)-th* ***cofactor*** *of A is the number $C_{ij} = (-1)^{i+j} D_{ij}$.*

2) The i-th ***principal minor*** *of A is the (i,i) matrix obtained by deleting the last $n-i$ rows and columns of A.*

Part (1) helps in calculating the determinant of A as shown in the following definition. Part (2) will prove useful to characterize positive (negative) definite matrices later on in this chapter

Definition 53 *Let A be a (n,n) matrix. The* ***determinant*** *of A is defined as follows:*

$$\det(A) = \sum_{j=1}^{n} a_{ij} C_{ij}$$

for any i between 1 and n.

The determinant of a matrix A is also usually denoted as follows (the matrix is placed between two vertical bars):

$$\det(A) = \begin{vmatrix} a_{11} & a_{12} & \dots & a_{1n} \\ a_{21} & & & a_{2n} \\ \dots & & & \dots \\ a_{n1} & \dots & \dots & a_{nn} \end{vmatrix}$$

Example 54 *Let A be defined by:*

$$A = \begin{pmatrix} 3 & 1 & 2 \\ 1 & 4 & 2 \\ 2 & 6 & 4 \end{pmatrix}$$

Definition 52 for $i = 1$ gives the following development.

$$\begin{aligned} \det(A) &= 3\begin{vmatrix} 4 & 2 \\ 6 & 4 \end{vmatrix} - 1\begin{vmatrix} 1 & 2 \\ 2 & 4 \end{vmatrix} + 2\begin{vmatrix} 1 & 4 \\ 2 & 6 \end{vmatrix} \\ &= 3 \times (4 \times 4 - 6 \times 2) - 1 \times (1 \times 4 - 2 \times 2) + 2(1 \times 6 - 2 \times 4) \\ &= 8 \end{aligned}$$

Suppose now $i = 2$*. We obtain:*

$$\begin{aligned} \det(A) &= -1 \begin{vmatrix} 1 & 2 \\ 6 & 4 \end{vmatrix} + 4 \begin{vmatrix} 3 & 2 \\ 2 & 4 \end{vmatrix} - 2 \begin{vmatrix} 3 & 1 \\ 2 & 6 \end{vmatrix} \\ &= -1 \times (4 - 12) + 4 \times (12 - 4) - 2 \times (18 - 2) \\ &= 8 \end{aligned}$$

Of course, the result is the same. It is not a proof, but the proof itself is cumbersome and uninteresting on a practical point of view, so we omit it.

One of the main results concerning determinants is related to products and transposition of matrices.

Proposition 55 *1) Let A and B two (n,n) matrices. We have:*

$$\begin{aligned}\det(AB) &= \det(A)\det(B)\\ \det(A^T) &= \det(A)\end{aligned}$$

2) If a matrix B is deduced from a (n,n) matrix A by swapping two rows or two columns, the determinants of the two matrices satisfy:

$$\det(B) = -\det(A)$$

3) A square matrix A is invertible if and only if its determinant is different from zero. If $\det(A) \neq 0$, the inverse matrix A^{-1} writes:

$$A^{-1} = \frac{1}{\det(A)}C^T \tag{1.16}$$

where $C = (C_{ij}, i, j = 1, .., n)$ is the cofactor matrix. Moreover, $\det(A^{-1}) = \frac{1}{\det(A)}$.

Proposition 55 gives a way to calculate determinants but this method is not the most numerically efficient.

Equality 1.16 is illustrated in the following example 56.

Example 56 *Let A denote the $(3,3)$ matrix:*

$$A = \begin{pmatrix} 3 & 2 & 5 \\ 4 & 6 & 1 \\ 2 & 3 & 2 \end{pmatrix}$$

The determinant of A is calculated as follows:

$$\det(A) = 3\begin{vmatrix} 6 & 1 \\ 3 & 2 \end{vmatrix} - 2\begin{vmatrix} 4 & 1 \\ 2 & 2 \end{vmatrix} + 5\begin{vmatrix} 4 & 6 \\ 2 & 3 \end{vmatrix} = 15$$

Applying relation 1.16 leads to:

$$A^{-1} = \frac{1}{15}\begin{pmatrix} 9 & 11 & -28 \\ -6 & -4 & 17 \\ 0 & -5 & 10 \end{pmatrix}$$

If we apply the same technique to A^T, the cofactor matrix of A^T is the transpose of the cofactor matrix of A. Therefore we obtain:

$$\left(A^T\right)^{-1} = \frac{1}{15}\begin{pmatrix} 9 & -6 & 0 \\ 11 & -4 & -5 \\ -28 & 17 & 10 \end{pmatrix}$$

1.3.5 Changing the basis

Matrix of a linear mapping after a basis change

A linear mapping f defined on $\mathbb{R}^n$, endowed with a basis $\mathcal{U} = (u_1, ..., u_n)$, and taking values in $\mathbb{R}^m$, endowed with a basis $\mathcal{V} = (v_1, ..., v_m)$, is represented by a matrix $M_f(\mathcal{U}, \mathcal{V})$. As mentioned before in definition 42, the notation $M_f(\mathcal{U}, \mathcal{V})$ recalls that M_f depends on the two bases. In particular, the columns of M_f are the images of vectors of $\mathcal{U}$ by f, expressed in the basis $\mathcal{V}$.

It turns out that a modification of one of the two bases changes the matrix M_f. Denote $\mathcal{W}$ a second basis of E and P the matrix having in columns the vectors of $\mathcal{W}$, expressed in the initial basis $\mathcal{U}$. This matrix will be called a **change-of-basis matrix** from basis $\mathcal{U}$ to basis $\mathcal{W}$.

We can show the following proposition.

Proposition 57 *Let v be a vector of $\mathbb{R}^n$ with coordinates $x^T = (x_1, x_2, ..., x_n)$ in basis $\mathcal{U}$ and $y^T = (y_1, y_2, ..., y_n)$ in basis $\mathcal{W}$. We then have:*

$$x = Py \text{ and } y = P^{-1}x$$

$M_f(\mathcal{W},\mathcal{V})$ *is given as follows:*

$$M_f(\mathcal{W},\mathcal{V}) = P^{-1}M_f(\mathcal{U},\mathcal{V})\,P$$

Example 58 *Let* $M_f(\mathcal{U},\mathcal{V})$ *and* x *be defined by:*

$$M_f(\mathcal{U},\mathcal{V}) = \begin{pmatrix} 1 & 2 & 1 \\ 0 & 1 & 4 \\ 2 & 3 & 1 \end{pmatrix} \qquad x = \begin{pmatrix} 4 \\ 6 \\ 1 \end{pmatrix}$$

Assume that $\mathcal{U}$ *is the canonical basis of* $\mathbb{R}^3$ *and define* $\mathcal{W}$ *by:*

$$w^1 = \begin{pmatrix} 1 \\ 1 \\ 0 \end{pmatrix} \quad w^2 = \begin{pmatrix} 0 \\ 3 \\ 2 \end{pmatrix} \quad w^3 = \begin{pmatrix} 2 \\ 0 \\ 1 \end{pmatrix}$$

The image $f(x)$ *of vector* x *(in basis* $\mathcal{U}$*) is given by:*

$$f(x) = \begin{pmatrix} 1 & 2 & 1 \\ 0 & 1 & 4 \\ 2 & 3 & 1 \end{pmatrix} \begin{pmatrix} 4 \\ 6 \\ 1 \end{pmatrix} = \begin{pmatrix} 17 \\ 10 \\ 27 \end{pmatrix}$$

The matrix P *writes:*

$$P = \begin{pmatrix} 1 & 0 & 2 \\ 1 & 3 & 0 \\ 0 & 2 & 1 \end{pmatrix}$$

The inverse of P *is calculated using the cofactors (equation 1.16). We obtain:*

$$P^{-1} = \frac{1}{7}\begin{pmatrix} 3 & 4 & -6 \\ -1 & 1 & 2 \\ 2 & -2 & 3 \end{pmatrix}$$

We deduce from this formulation of P^{-1}:

$$P^{-1}M_f(\mathcal{U},\mathcal{V}) = \frac{1}{7}\begin{pmatrix} 3 & 4 & -6 \\ -1 & 1 & 2 \\ 2 & -2 & 3 \end{pmatrix}\begin{pmatrix} 1 & 2 & 1 \\ 0 & 1 & 4 \\ 2 & 3 & 1 \end{pmatrix} = \frac{1}{7}\begin{pmatrix} -9 & -8 & 13 \\ 3 & 5 & 5 \\ 8 & 11 & -3 \end{pmatrix}$$

$$P^{-1}M_f(\mathcal{U},\mathcal{V})P = \frac{1}{7}\begin{pmatrix} -9 & -8 & 13 \\ 3 & 5 & 5 \\ 8 & 11 & -3 \end{pmatrix}\begin{pmatrix} 1 & 0 & 2 \\ 1 & 3 & 0 \\ 0 & 2 & 1 \end{pmatrix} = \frac{1}{7}\begin{pmatrix} -17 & 2 & -5 \\ 8 & 25 & 11 \\ 19 & 27 & 13 \end{pmatrix}$$

Therefore, in the new basis $\mathcal{W}$, x writes :

$$x = \frac{1}{7}\begin{pmatrix} 3 & 4 & -6 \\ -1 & 1 & 2 \\ 2 & -2 & 3 \end{pmatrix}\begin{pmatrix} 4 \\ 6 \\ 1 \end{pmatrix} = \frac{1}{7}\begin{pmatrix} 30 \\ 4 \\ -1 \end{pmatrix}$$

The following section studies the case where $M_f(\mathcal{W},\mathcal{V})$ is diagonal when $\mathcal{W}$ and $\mathcal{V}$ have the same dimension.

Trace of a square matrix

Definition 59 *The **trace** of a (n,n) square matrix A is the sum of its diagonal terms and is denoted $Tr(A)$.*

$$Tr(A) = \sum_{i=1}^{n} a_{ii}$$

The elementary properties of the trace of a matrix are summarized in the following proposition.

Proposition 60 *Let A and B be two (n,n) matrices and $c \in \mathbb{R}$:*

$$\begin{aligned} Tr(cA+B) &= cTr(A)+Tr(B) \\ Tr(AB) &= Tr(BA) \end{aligned}$$

Denote A a matrix representing an endomorphism f of $\mathbb{R}^n$ when the bases are $\mathcal{U}$ and $\mathcal{V}$. The columns of A are the basis vectors of $\mathbb{R}^n$ transformed by f . If $\mathbb{R}^n$ is endowed with a new basis $\mathcal{W}$ the matrix representing f is modified (denoted B) but the trace does not change.

Proposition 61 $Tr(A) = Tr(B)$.

The reader can check that the proposition is true in example 58. The matrices with respect to the two bases were:

$$M_f(\mathcal{U},\mathcal{V}) = \begin{pmatrix} 1 & 2 & 1 \\ 0 & 1 & 4 \\ 2 & 3 & 1 \end{pmatrix} \quad \text{et} \quad M_f(\mathcal{W},\mathcal{V}) = \frac{1}{7}\begin{pmatrix} -17 & 2 & -5 \\ 8 & 25 & 11 \\ 19 & 27 & 13 \end{pmatrix}$$

The trace of the two matrices is equal to 3.

Diagonalization of square matrices

A question coming naturally to mind when dealing with changes of bases is the following: can we transform the matrix of an endomorphism in a more simple one, more precisely in a diagonal matrix, by a change of basis?

Methods of data analysis like Principal Component Analysis of Factor Analysis are based on such transformations. Even if these methods are not addressed in the present book, the reader should know that they are used in multifactor models, especially the Arbitrage Pricing Theory[19].

Eigenvalues and eigenvectors Eigenvalues and eigenvectors are the mathematical tools allowing to formalize a change of basis in such a way that the resulting matrix becomes diagonal.

Let f be an endomorphism of $\mathbb{R}^n$ and let M denote the matrix of f in basis $\mathcal{U}$.

Definition 62 *An* ***eigenvalue*** *of f (or equivalently of M) is a real number λ such that there exists a non zero vector $u \in \mathbb{R}^n$ satisfying:*

$$f(u) = Mu = \lambda u$$

u is then called an ***eigenvector*** *of f (of M) associated to the eigenvalue λ.*

Several linearly independent vectors can satisfy $Mu = \lambda u$. But if two linearly independent vectors u and v satisfy $Mu = \lambda u$ and $Mv = \lambda v$, then

[19] The seminal paper is Ross, S. (1976),"The Arbitrage Theory of Capital Asset Pricing". *Journal of Economic Theory* 13 (3): 341–360.

Two examples of papers using data analysis methods are Roll, R. and Ross, S. (1980). "An Empirical Investigation of the Arbitrage Pricing Theory". *Journal of Finance* 35 (5): 1073–1103.

Chamberlain, G. and Rothschild, M. (1983), "Arbitrage, Factor Structure, and Mean Variance Analysis on Large Asset Markets." *Econometrica* 51, 1281-1304.

the same relationship is true for any linear combination of u and v:

$$\forall (a,b) \in \mathbb{R} \times \mathbb{R}, M(au+bv) = \lambda(au+bv)$$

Of course, this relation is satisfied because f is a linear mapping:

$$\begin{aligned} f(au+bv) &= af(u)+bf(v) = M(au+bv) \\ &= a\lambda u + b\lambda v = \lambda(au+bv) \end{aligned}$$

We then obtain the following definition.

Definition 63 *The* ***eigenspace*** *of the eigenvalue λ is the vector subspace F_λ defined by:*

$$F_\lambda = \{u \in \mathbb{R}^n \text{ such that } f(u) = Mu = \lambda u\}$$

To determine the eigenvalues of a linear mapping f, we use the characteristic polynomial.

The characteristic polynomial

Definition 64 *Denote $M - \lambda I_n$ the following matrix:*

$$M - \lambda I_n = \begin{pmatrix} m_{11}-\lambda & m_{12} & \ldots & m_{1n} \\ m_{21} & m_{22}-\lambda & & m_{2n} \\ \ldots & & m_{33}-\lambda & \ldots \\ m_{n1} & m_{n2} & & m_{44}-\lambda \end{pmatrix}$$

The ***characteristic polynomial*** *of f is the polynomial $Q(\lambda)$ defined by:*

$$Q(\lambda) = \det(M - \lambda I_n)$$

$M - \lambda I_n$ is obtained by substracting λ times the identity matrix I_n to M. Solving $Q(\lambda) = 0$ provides all the eigenvalues of M.

Proposition 65 *The eigenvalues of f are the solutions of $Q(\lambda) = 0$.*

By definition, λ is an eigenvalue of M associated to the eigenvector u if:

$$Mu = \lambda u$$

This equality is equivalent to:

$$(M - \lambda I_n)\, u = 0$$

The matrix $M - \lambda I_n$ is then not invertible because $u \neq 0$. Therefore its determinant $\det(M - \lambda I_n) = Q(\lambda)$ is equal to 0.

Example 66 *Let f be a linear mapping represented by M in the canonical basis of $\mathbb{R}^3$:*

$$M = \begin{pmatrix} 1 & 0 & 2 \\ 0 & 2 & 2 \\ 0 & 1 & 1 \end{pmatrix}$$

Calculating the determinant of $M - \lambda I_n$ along the first line leads to:

$$\begin{aligned}
\det(M - \lambda I_n) &= (1-\lambda)\begin{vmatrix} 2-\lambda & 2 \\ 1 & 1-\lambda \end{vmatrix} - 0\begin{vmatrix} 0 & 2 \\ 0 & 1-\lambda \end{vmatrix} + 2\begin{vmatrix} 0 & 2-\lambda \\ 0 & 1 \end{vmatrix} \\
&= (1-\lambda)\left((2-\lambda)(1-\lambda) - 2\right) \\
&= (1-\lambda)(\lambda^2 - 3\lambda) \\
&= \lambda(1-\lambda)(\lambda - 3)
\end{aligned}$$

Equation $Q(\lambda) = 0$ has three solutions that are $\lambda_1 = 3$; $\lambda_2 = 0$ and $\lambda_3 = 1$.

What are the corresponding eigenvectors u^1, u^2, u^3? First, we need to *solve* $(M - \lambda_1 I_3)\, u^1 = 0$, *that is:*

$$\begin{pmatrix} -2 & 0 & 2 \\ 0 & -1 & 2 \\ 0 & 1 & -2 \end{pmatrix} \begin{pmatrix} u_1^1 \\ u_2^1 \\ u_3^1 \end{pmatrix} = \begin{pmatrix} 0 \\ 0 \\ 0 \end{pmatrix}$$

The solution satisfies $u_1^1 = u_3^1$ *and* $u_2^1 = 2u_3^1$. *The following vector is an example of solution:*

$$u^1 = \begin{pmatrix} 1 \\ 2 \\ 1 \end{pmatrix}$$

With the same approach for u^2 *and* u^3, *we obtain:*

$$u^2 = \begin{pmatrix} -2 \\ -1 \\ 1 \end{pmatrix} \text{ and } u^3 = \begin{pmatrix} 1 \\ 0 \\ 0 \end{pmatrix}$$

What is the matrix of f *in the basis* (u^1, u^2, u^3)?

The change-of-basis matrix P *is the matrix built with* u^1, u^2, u^3 *as columns because the initial basis was the canonical basis:*

$$P = \begin{pmatrix} 1 & -2 & 1 \\ 2 & -1 & 0 \\ 1 & 1 & 0 \end{pmatrix}$$

The inverse of P *is equal to:*

$$P^{-1} = \begin{pmatrix} 0 & 1/3 & 1/3 \\ 0 & -1/3 & 2/3 \\ 1 & -1 & 1 \end{pmatrix}$$

As a consequence, we have:

$$P^{-1}M = \begin{pmatrix} 0 & 1 & 1 \\ 0 & 0 & 0 \\ 1 & -1 & 1 \end{pmatrix}$$

$$P^{-1}MP = \begin{pmatrix} 3 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 1 \end{pmatrix}$$

Example 66 illustrates how M becomes a diagonal matrix when M is written in the basis of eigenvectors. Moreover, the elements on the diagonal are exactly the eigenvalues. We let the reader check the result of proposition 61, that is $Tr(M) = Tr(P^{-1}MP) = 4$.

When can we diagonalize a matrix?

Definition 67 *A (n, n) matrix M is* ***diagonalizable*** *if it has n eigenvalues and n linearly independent eigenvectors.*

An equivalent definition could be: M is diagonalizable if there exists a diagonal matrix D (contenant les eigenvalues) and an invertible matrix P satisfying:

$$M = P^{-1}DP$$

Of course, the diagonal elements of D are the eigenvalues of M and the columns of P are the corresponding eigenvectors of M.

It may happen that two eigenvalues are equal, for example if the characteristic polynomial is:

$$Q(\lambda) = (\lambda - 1)(\lambda - 3)^2$$

In this situation, M is diagonalizable if the dimension of the eigenspace F_3 (associated to $\lambda = 3$) is equal to 2.

On the contrary if $\dim(F_3) = 1$, M is not diagonalizable. We cannot find an invertible change-of-basis matrix P.

Symmetric matrices A non negligible part of financial theory deals with portfolio choice and portfolio management. In this framework, an important piece of information is the covariance matrix of returns which is a symmetric matrix. These matrices are special because of the following proposition.

Proposition 68 *Any square symmetric matrix M is diagonalizable and we have:*

$$\begin{aligned} P^{-1} &= P' \\ M &= PDP' \end{aligned}$$

where P denotes the matrix of eigenvectors.

The first result says that the inverse of P is its transpose P'. Such a property characterizes orhogonal[20] matrices.

1.4 Norms and inner products

1.4.1 Normed vector spaces

In the section devoted to topology in chapter 1 of Part I, we defined the concepts of distance (metric) and metric spaces. The Euclidean distance on $\mathbb{R}^n$ was defined by:

$$d(x,y) = \sqrt{\sum_{i=1}^{n} (x_i - y_i)^2} \tag{1.17}$$

If $n = 2$, $d(x,y)$ is the length of the straight line joining $x' = (x_1, x_2)$ to $y' = (y_1, y_2)$. More generally, in a finite-dimensional space, the concept of "length" of a vector is defined through a norm on the vector space under consideration.

Definition 69 *Let E be a vector space; a* ***norm*** *on E is a mapping, denoted $\|.\|$, defined on E and taking values in $\mathbb{R}^+$ satisfying:*

$$\begin{aligned} \|x\| &= 0 \Leftrightarrow x = 0 \\ \forall (x,y) &\in E, \|x+y\| \leq \|x\| + \|y\| \\ \forall x &\in E, \forall c \in \mathbb{R}^+, \|c.x\| = |c| \, \|x\| \end{aligned}$$

It appears that a norm $\|.\|$ on a vector space induces a metric d on the same space if the metric is defined by :

$$d(x,y) = \|x - y\|$$

[20] In the next section, we justify the word "orthogonal".

The Euclidean metric on $\mathbb{R}^n$ defined in relation 1.17 is induced by the following Euclidean norm on $\mathbb{R}^n$:

$$\|x\| = \sqrt{\sum_{i=1}^{n} x_i^2}$$

As for metrics, many different norms can be defined on a vector space. For example the mapping $\|x\|_{\max} = \max_i |x_i|$ can be used as a norm.

In finance, norms are associated to risk measures. For example if x_i denotes the future value of a portfolio in state i, the two abovementioned norms are interpreted differently.

Let $\mathbf{1}$ denote as usual the vector with all coordinates equal to 1 and $\overline{x} = \frac{1}{n}\sum_{i=1}^{n} x_i$ the average payoff. A usual measure of risk is the empirical variance calculated as:

$$\sigma^2(x) = \frac{1}{n}\|x - \overline{x}\mathbf{1}\|^2 = \frac{1}{n}\sum_{i=1}^{n}(x_i - \overline{x})^2$$

But in a different approach called *Value at Risk*[21], we could use $\|x - \overline{x}\mathbf{1}\|_{\max}$ as a measure of risk[22]; risk is then evaluated as the maximum difference with respect to the average payoff.

The second important tool to structure a vector space is the concept of inner product. In finite-dimensional spaces, norms and inner products are closely related. It is not always the case in infinite-dimensional spaces.

1.4.2 Inner products in vector spaces

Definition 70 *An* ***inner product*** *on a vector space E is a mapping, denoted $< .,. >$, defined on $E \times E$ and taking values in $\mathbb{R}$, symmetric, bilinear*

[21] For a detailed presentation of Value-at-Risk, see Jorion (2006), *Value at Risk: The New Benchmark for Managing Financial Risk*, McGraw-Hill Professional.

[22] This measure is not exactly what is called *Value at Risk* in the financial literature but it is in the same spirit.

and positive, that is satisfying:

1. $< x, y >=< y, x >$

2. $\forall (a, b, c, d) \in \mathbb{R}^4, \forall (x, y, z, t) \in E^4,$

$$< ax+by, cz+dt >= ac < x, z > +ad < x, t > +bc < y, z > +bd < y, t >$$

3. $< x, x >= 0 \Leftrightarrow x = 0$ *otherwise* $< x, x > > 0.$

Part 1 defines symmetry, part 2 bilinearity and part 3 positivity. The usual inner product on $\mathbb{R}^n$ is defined by:

$$< x, y >= \sum_{i=1}^{n} x_i y_i$$

Alternative notations of the inner product of two vectors x and y are (x, y) or $x'y$. The latter is consistent with the rules used to multiply matrices (see part I, chapter 4). The reason is that a column-vector is a matrix with n rows and 1 column. Consequently, x' is a matrix with 1 row and n columns. The product $x'y$ is then a matrix with 1 row and 1 column, that is a number.

Definition 70 allows for general inner products. However, we need to recall what is a positive-definite matrix to generalize inner products beyond the usual Euclidean ones.

Definition 71 *A square matrix A of dimension n is* ***positive (negative) semi-definite*** *if:*

$$\forall x \in \mathbb{R}^n, \ x'Ax \geq (\leq) \, 0$$

A square matrix A of dimension n is ***positive (negative) definite*** *if:*

$$\forall x \in \mathbb{R}^n, \ x \neq 0 \Rightarrow x'Ax > (<) 0$$

This definition allows a general characterization of inner products on $\mathbb{R}^n$.

Proposition 72 *Let A be a square symmetric positive definite matrix; the mapping $(x, y) \to x'Ay$ associating any pair of vectors of $\mathbb{R}^n$ to the product $x'Ay$ is an inner product on $\mathbb{R}^n$ denoted $< .,. >_A$. The norm associated to this inner product, denoted $\|.\|_A$ is defined by $\|x\|_A = \sqrt{x'Ax}$*

Without entering into the details of the proof, remark that the condition $< x, x >_A > 0$ for $x \neq 0$ is satisfied because A is positive definite. In the same way, A is symmetric, property ensuring that the inner product is symmetric. Moreover, if A is the identity matrix, we are back to the definition of the usual inner product.

Definition 73 *Two vectors of E are* ***orthogonal*** *if their inner product is equal to 0.*

This definition of orthogonality refers to "right angles" when the usual two-dimensional space is endowed with the standard inner product. But the definition also shows that orthogonality is a much more general concept and, mainly, that being orthogonal for a pair of vectors depends on the inner product the vector space is endowed with. For example, if A is a diagonal matrix with strictly positive numbers on the diagonal satisfying $\sum a_{ii} = 1$, A defines an inner product allowing to calculate the expectation of the product of two random variables because the diagonal terms of A define a probability measure. In this example, orthogonality is far from the usual geometric interpretation[23].

Geometric interpretation

To elaborate on geometric aspects, consider the space $\mathbb{R}^n$ endowed with the usual norm and inner product. Let x and y denote two vectors in $\mathbb{R}^n$; the norm of the normalized vector $x^* = \frac{x}{\|x\|}$ is equal to 1 by construction. Let y^o be the projection of y on the line Δ_x generated by x. y^o is proportional to x, and more precisely we have the following equality:

$$y^o =< y, x^* > x^*$$

In other words, the inner product of x and y is equal to the coordinate of the projection of y on Δ_x (apart from the standardisation factor $\|x\|$), as illustrated by figure 1.1.

Let a denote the angle between x and y, we can establish the following

[23] For example the expectation of a random variable X can be written as $\mathbf{1}'AX = \sum_{i=1}^{n} a_{ii} X_i$ where X_i is the value of X in state i and a_{ii} is the probability of occurrence of state i.

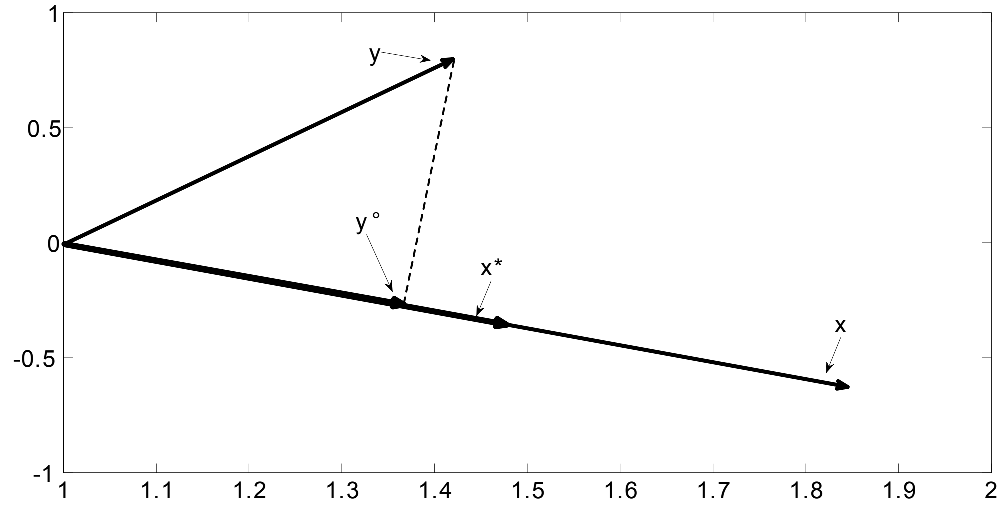

Figure 1.1: Geometric interpretation of the inner product

relationship between the inner product $< x, y >$ and the cosine of a :

$$< x, y >= \cos(a).\, \|x\|\, \|y\|$$

We are back to the well-known relationship saying that the cosine is equal to the ratio of the inner product divided by the product of norms.

1.4.3 Quadratic forms

Definition 74 *A* ***quadratic form*** f*, defined on an open subset* $D \subset \mathbb{R}^n$*, taking values in* $\mathbb{R}$*, is defined by :*

$$\forall x \in D, f(x) = x'Ax$$

where A *is a symmetric square matrix.*

Proposition 75 *A quadratic form is convex (concave) if and only if* A

is positive(negative) semi-definite. If A is positive(negative) definite, f is strictly convex (concave).

Quadratic forms are naturally present in portfolio management because the variance of the return of a portfolio x containing n stocks writes $x'\mathbf{V}x$ where $\mathbf{V}$ is the covariance matrix of returns of the n stocks.

In finance models, $\mathbf{V}$ is generally assumed positive definite, meaning that it is not possible to build a zero-variance portfolio (that is a risk-free portfolio) by combining n risky assets. It is an assumption but what is sure is that $\mathbf{V}$ is positive semi-definite because a variance of return $x'\mathbf{V}x$ cannot be negative.

The other domain where quadratic forms arise naturally is non linear optimization. We will see later on that it is easy to find the maximum (minimum) of a quadratic form when it is concave (convex).

1.5 Hilbert spaces

1.5.1 Definition

We mentioned several times that mathematical properties satisfied in finite-dimensional spaces could be false in more general spaces. However, there exists a category of infinite-dimensional vector spaces for which important properties remain valid. These spaces are called Hilbert spaces and they are well fitted to study financial problems, as it is illustrated in *Probability for Finance.*

Definition 76 *A* ***Hilbert space*** *is a vector space E whose norm is deduced from an inner product and that is complete as a metric space*[24].

[24] Remember that a metric space is complete if any Cauchy sequence converges. Therefore, speaking about a complete normed vector space is not really correct because completeness is a notion defined in metric spaces. This expression simply means that the metric d deduced from the norm on E makes (E, d) a complete metric space. This metric is defined by $d(x, y) = \|x - y\|$.

Of course, finite-dimensional spaces like $\mathbb{R}^n$ are Hilbert spaces when they are endowed with an inner product like the one defined in proposition 72.

The two essential properties for financial applications are the projection theorem and the Riesz representation theorem. Before presenting these results we first recall what is a convex set in a vector space.

Definition 77 *Let E denote a vector space and C a subset of E; C is **convex** if:*

$$\forall(x, y) \in C \times C, \forall\alpha \in [0; 1], \alpha x + (1 - \alpha)y \in C$$

First, it is important to notice that convexity can only make sense in vector spaces because, in the definition, there is a linear combination of vectors, $\alpha x + (1 - \alpha)y$. It is then necessary that this combination belongs to the vector space for the definition to make sense. As a consequence, in preceding chapters or in part 1 of the book, we could not have used convexity in a general framework. Nevertheless, the geometric interpretation of the convexity of a set is similar to what we proposed in $\mathbb{R}$ for intervals. A set is convex if, as soon as it contains two elements x and y, it also contains the segment joining these two elements.

Convexity is a standard assumption for consumption sets in microeconomics textbooks. It only means that goods are divisible. The same assumption on a set of portfolios would mean that portfolios and stocks can be combined in non integer quantities.

1.5.2 The projection theorem

Proposition 78 *Let E denote a Hilbert space and C a non empty convex set in E; any vector x in E has a unique projection on C, denoted x^* and satisfying:*

$$\forall y \in C, \langle x - x^*, y - x^* \rangle \leq 0$$

$x - x^*$ is orthogonal to the tangent to C at x^*. Consequently, the angle between $y - x^*$ and $x - x^*$ lies between $\frac{\pi}{2}$ and $\frac{3\pi}{2}$. The cosine of this angle is then negative; but it is proportional to the inner product of the two vectors, meaning that this inner product is also negative. These remarks do not prove the proposition but they provide the geometric intuition for this proposition. One of the fundamental applications of the projection theorem consists in considering the case where C is a vector subspace of E. It is exactly the proposition allowing to define the conditional expectation of a random variable as a projection on a subspace of the vector space of square integrable random variables (see *Probability for Finance*).

The representation theorem presented below can also be interpreted with the same geometric approach.

1.5.3 The Riesz representation theorem

We saw in a preceding section of this chapter that a linear mapping defined on $\mathbb{R}^n$ and taking values in $\mathbb{R}^m$ can be associated to a matrix. In particular if $m = 1$, this matrix is a vector and the mapping is a linear form. In other words, if this mapping is denoted f, we write:

$$f(x) = \sum_{i=1}^{n} a_i x_i$$

and the vector $a' = (a_1, ..., a_n)$ represents f. This result may be generalized in the framework of Hilbert spaces, provided that f is continuous.

Proposition 79 *Let E be a Hilbert space and f be a continuous linear form defined on E; there exists a unique vector $y_f \in E$ such that:*

$$\forall x \in E, \; f(x) = \langle x, y_f \rangle$$

The vector y_f represents the linear mapping f and the important result is that y_f belongs to E.

In a financial framework the vector y_f has a natural interpretation if f is a valuation operator linking the future payoffs of a financial security x to its date-0 price $f(x)$. The coordinates of y_f are linked to the prices of the Arrow-Debreu securities. We already mentioned this characteristic in finite-dimensional spaces.

1.6 Separation theorems and Farkas lemma

1.6.1 Introductive example

Let f be a linear form defined on $\mathbb{R}^2$, characterized by the relation:

$$f(x) = a_1 x_1 + a_2 x_2$$

where $x' = (x_1, x_2)$ and a_1, a_2 are real numbers. The equation $f(x) = a_1x_1 + a_2x_2 = 0$ defines a line D in $\mathbb{R}^2$. Therefore, for any linear form f, the space $\mathbb{R}^2$ is divided in three regions denoted R_1, R_2 and D. These regions are characterized by:

$$\begin{aligned}\forall x &\in R_1, f(x) > 0\\ \forall x &\in R_2, f(x) < 0\\ \forall x &\in D, f(x) = 0\end{aligned}$$

Let C denote a non empty convex set not containing 0; there exists a linear form f, that is coefficients a_1 and a_2 satisfying:

$$\forall x \in C,\ f(x) > 0$$

In other words, the convex set C is entirely in R_1. This result is intuitive because any tangent to C induces a separation such that C is on one side of the tangent. For a given tangent Δ, separing 0 and C, consider the parallel to Δ containing 0. C is entirely on one side of this parallel to Δ. This line is defined by an equation like $a_1x_1 + a_2x_2 = 0$. If the elements $x \in C$ satisfy $f(x) > 0$, the desired result is obtained. If $f(x) < 0$ for $x \in C$, it is enough to choose the linear form g defined by $g(x) = -f(x) = -a_1x_1 - a_2x_2$.

1.6.2 Separation theorems and Farkas lemma

The following proposition is a generalization of the approach illustrated in the introductive example.

Proposition 80 ***Separation theorem***

Let E be an Euclidean vector space, C be a non empty convex subset of E that does not contain the null vector. There exists a linear form f defined on E such that for any x in C, $f(x) \geq 0$.

The matrix expression of this separation theorem is called Farkas lemma. We propose hereafter two versions of this lemma, the second being usually called lemma of the alternative. This result has a beautiful financial interpretation, as we will illustrate later on.

Proposition 81 ***Farkas lemma***

Let A be a matrix with m rows and n columns; a vector $x \in \mathbb{R}^n$ satisfies $x'y \geq 0$ for any $y \in \mathbb{R}^n$ such that $Ay \geq 0$ if and only if there exists a vector $z \in \left(\mathbb{R}_+^*\right)^m$ satisfying $x' = z'A$.

Proposition 82 ***Lemma of the alternative***

Let A be a matrix with m rows and n columns; one and only one of the two following properties is true.

1. *The equation $Ax = 0$ has a solution in $\mathbb{R}^n$ with all strictly positive coordinates.*

2. *Inequality $y'A > 0$ has a solution in $\mathbb{R}^m$.*

1.6.3 Application to no-arbitrage pricing

Consider a one-period financial market on which investors trade securities at date 0, these securities providing random payoffs at date T. Recall that an arbitrage opportunity is a portfolio that costs nothing at date 0 (or maybe the cost is negative) and pays a positive amount at date T in all states of nature. Assume there are n possible states of nature and K securities traded on the market. The date-T payoffs are stored in a matrix D with n rows and K columns, each column corresponding to a security and each row to a state of nature.

Let $D = (d_{jk}, j = 1,..,n; k = 1,..,K)$, that is:

$$D = \begin{bmatrix} d_{11} & ... & d_{1K} \\ ... & d_{jk} & ... \\ ... & ... & ... \\ d_{n1} & ... & d_{nK} \end{bmatrix}$$

Denote $\Pi' = (\pi_1, ...\pi_K)$ the date-0 price vector; an arbitrage opportunity is a portfolio $\theta \in \mathbb{R}^K$ that satisfies one of the two following properties:

$$\begin{aligned} a)\ D\theta &\geq 0 \text{ and } \Pi'\theta < 0 \\ b) D\theta &> 0 \text{ and } \Pi'\theta \leq 0 \end{aligned}$$

The first case (a) means that the portfolio has a strictly negative cost ($\Pi'\theta < 0$). Moreover, $D\theta \geq 0$ means that final payoffs are never negative. Consequently, an investor characterized by a strictly increasing utility function would be ready to buy an infinite quantity of this portfolio because holding this portfolio increases date-0 utility without decreasing date-T expected utility.

Case (b) is a little bit more subtle. Remember that $D\theta > 0$.[25] Therefore, the date-T expected utility increases by holding portfolio θ. But, at the same time, $\Pi'\theta \leq 0$, meaning that date-0 utility does not decrease when buying portfolio θ. As in case (a), an investor with a strictly increasing utility function would ask an infinite quantity of θ. In a well-functioning market, arbitrage opportunities should disappear very quickly by price adjustments due to excess supply or excess demand.

At a first glance, it may be difficult to see the relationship between the definition of an arbitrage opportunity and the lemma of the alternative...except that the two use matrix notations! The difficulty comes from the fact that

[25] Being given a matrix A, writing $A \geq 0$ means that all elements of A are positive, $A > 0$ means $A \geq 0$ and at least one element is strictly positive, and finally, $A >> 0$ means that all elements of A are strictly positive.

the definition of an arbitrage opportunity takes simultaneously into account date 0 and date T.

We are going to "forget" this specificity by defining a matrix D^* which is the concatenation of D and of $-\Pi'$(minus the price vector).

$$D^* = \begin{bmatrix} d_{11} & \dots & d_{1K} \\ \dots & d_{jk} & \dots \\ \dots & \dots & \dots \\ d_{n1} & \dots & d_{nK} \\ -\pi_1 & \dots & -\pi_K \end{bmatrix}$$

This notation allows to define an arbitrage opportunity as a vector $\theta \in \mathbb{R}^K$ satisfying:

$$D^*\theta > 0 \qquad (1.18)$$

In fact, $D^*\theta > 0$ means that portfolio θ generates non negative cash-flows in all states and has a non positive cost. Moreover, at least one of the components of $D^*\theta$ is strictly positive. If the last component is strictly positive, we have a type (a) arbitrage opportunity. If this last component is zero, one of the other components is strictly positive and we face a type (b) arbitrage opportunity.

After transposing the two sides of inequality 1.18, we obtain:

$$\theta' D^{*\prime} > 0$$

This inequality corresponds to part (2) of lemma 82 when applied to D^*. We just showed that arbitrage opportunities are incompatible with equilibrium prices. We then have to assume that $D^*\theta > 0$ has no solution in θ.

As a consequence, lemma 82 implies there exists $\beta \in \mathbb{R}^{n+1}$ the components of which being all strictly positive, such that:

$$D^{*\prime}\beta = 0 \tag{1.19}$$

This equality must be true because part (b) of the lemma is false....then (a) is true!

The financial interpretation of relation 1.19 goes as follows. For the sake of clarity, focus on the first term of $D^{*\prime}\beta$; it is the inner product of β and of the first row of $D^{*\prime}$ (which corresponds to the first security). This inner product writes:

$$\sum_{j=1}^{n+1} \beta_j d^*_{j1} = 0$$

with $d^*_{j1} = d_{j1}$ if $j \leq n$ and $d^*_{j1} = -\pi_1$ if $j = n+1$. The above equality is then equivalent to:

$$\sum_{j=1}^{n} \beta_j d_{j1} = \beta_{n+1}\pi_1 \tag{1.20}$$

Denote $\gamma_j = \frac{\beta_j}{\beta_{n+1}}$ (these coefficients are well-defined because $\beta_{n+1} > 0$); relation 1.20 can be transformed in:

$$\sum_{j=1}^{n} \gamma_j d_{j1} = \pi_1 \tag{1.21}$$

In the financial approach, this equality is very important because the left-hand side contains future cash-flows and the right-hand side contains the initial price. This equality is a typical valuation model (cash-flows on one side, price on the other side). However, the economic interpretation of equation 1.21 is difficult. But if we define $\gamma_j^* = \frac{\gamma_j}{\sum_{k=1}^{n} \gamma_k}$, we obtain:

$$\left(\sum_{k=1}^{n} \gamma_k\right) \sum_{j=1}^{n} \gamma_j^* d_{j1} = \pi_1$$

In this formula, the γ_j^* are positive numbers between 0 and 1 and satisfying $\sum_{j=1}^{n} \gamma_j^* = 1$. They define a probability measure on the set of states of nature. It is also remarkable that this probability measure does not depend on the asset we considered (here we selected the first but it does not matter). It remains to give an economic interpretation to $\sum_{k=1}^{n} \gamma_k$.

To simplify this interpretation, assume that asset numbered 1 is a risk-free asset paying 1 in each state, that is a zero-coupon bond. From equation 1.21 we deduce :

$$\sum_{j=1}^{n} \gamma_j = \pi_1$$

The quantity $\sum_{j=1}^{n} \gamma_j$ is the price of a security paying 1 at date T in each state of nature. If r denotes the risk-free rate, that is the return on the risk-free asset, we can write:

$$\sum_{j=1}^{n} \gamma_j = \frac{1}{1+r}$$

It turns out that the valuation of any asset k writes:

$$\pi_k = \frac{1}{1+r} \sum_{j=1}^{n} \gamma_j^* d_{jk}$$

The price is then equal to the weighted average (an expectation using probabilistic vocabulary) of future cash-flows, discounted at the risk-free rate. Of course this interpretation can only be done when there is a risk-free asset traded on the market. But it can be generalized if there exists a portfolio generating strictly positive payoffs in any state of nature. Such a portfolio is named a **numéraire**.

Chapter 2

Functions of several variables

Most financial models include several variables. Portfolio choice depends on at least two variables, the expected return and the variance of returns of the portfolio. When n assets are traded on the market, the utility of a portfolio depends on the n weights of the assets in the portfolio. Functions of several variables also appear naturally in option pricing. The standard Black-Scholes option pricing model (1973) allows to evaluate an option contract as a function of six variables, namely the price of the underlying asset, the time to maturity, the volatility of the underlying asset, the strike price, the risk-free rate and the dividends paid on the underlying.

In this chapter, we start by concepts generalizing chapter 2 of part I which was focused on single-variable functions. First, we need to generalize some results of chapter 1 of part I. Section 2.1 defines a metric space and the notion of distance (also called metric) on a metric space. The concept of distance is very general, but in the present chapter we essentially apply it to finite-dimensional spaces like $\mathbb{R}^p$, to study functions depending on p variables. Section 2.2 presents continuity and differentiability of functions depending on several variables and some important results like multidimensional Taylor's formulas. These formulas are interesting when it comes to approximating functions by polynomials ot to stating optimality conditions (see chapters 3 and 4). Finally, section 3 deals with implicit differentiation and homogeneous functions.

2.1 Metric spaces

If you are interested in overseas races, like the Vendée Globe Challenge, you want to know the ranking of boats at regular time intervals. On the website of the race[1] you can download a map where the boats are represented on the ocean and you can also see the total remaining distance. If you think to the problem a few minutes, you see that it is not a trivial matter to calculate a distance between two points A and B on a sphere (a reasonable approximation for our planet). It becomes even more difficult when constraints are added to the problem (boats are supposed to stay on the water!). The distance is not the same as the crow flies or for people who possibly need to climb mountains or stay on oceans. Think to people walking in New York City, or in major U.S towns. As streets are orthogonal to each other it is not very useful to know that, the distance between A to B as the crow flies is 5 miles.

In mathematical terms, a distance should be a mapping linking the pair (A, B) to a positive number and satisfying some reasonable properties. But this mapping should also be sufficiently general to adapt to many different contexts.

2.1.1 Metric on a set

Definition 83 *A* ***distance (metric)*** *on a set E is a mapping d from $E \times E$ to $\mathbb{R}_+$ satisfying:*

1) $\forall (x, y) \in E \times E,\ d(x, y) = d(y, x)$ *(symmetry).*

[1] http://www.vendeeglobe.org

2) $\forall (x,y) \in E \times E,\ d(x,y) = 0$ *if* $x = y$ *and* $d(x,y) > 0$ *otherwise (positivity).*

3) $\forall (x,y,z) \in E \times E \times E,\ d(x,z) \leq d(x,y) + d(y,z)$ *(triangular inequality).*

The pair (E,d) *is called a* ***metric space.***

Part (2) says that if the distance between two elements is zero, they are identical. Though this property seems very intuitive, we are going to provide examples showing that this intuition can be misleading. Part (3) says, in everyday language, that the shortest route between two points x and z is a straight line. One more time, it seems intuitive when a distance on real numbers is defined by $d(x,y) = |x - y|$. In this case the distance between x and z is the length of the segment joining the two points. But recall the Vendée Globe Challenge! On a sphere, there are no straight lines!

Consider for example the 2-dimensional space $\mathbb{R}^2$; the usual metric on this space, called the Euclidian distance, is defined by:

$$d_0(x,y) = \sqrt{(x_1 - y_1)^2 + (x_2 - y_2)^2} \tag{2.1}$$

with $x = (x_1, x_2)$ and $y = (y_1, y_2)$.

This metric measures the "physical" distance between x and y. The reader can easily check this is the case by using the Pythagorean theorem.

Of course, the Euclidian metric is easily generalized to p-dimensional spaces as follows:

$$d_0(x,y) = \sqrt{\sum_{k=1}^{p} (x_k - y_k)^2} \tag{2.2}$$

with $x = (x_1, x_2, ..., x_p)$ and $y = (y_1, y_2, ..., y_p)$ two elements of $\mathbb{R}^p$.

However, our "sailing" example shows that there are many ways to measure distances on a sphere or in New York City. In this latter case, the real distance should be defined as d^*[2]:

$$d^*(x,y) = |x_1 - y_1| + |x_2 - y_2|$$

The reader can check that d^* satisfies the three properties of definition 83 and that any norm $\|.\|$ on a vector space E induces a metric d on E. d is defined as:

$$d(x,y) = \|x - y\|$$

where x and y are two vectors in E. Consequently, a normed vector space is also a metric space (E,d) when d is the metric induced by the norm on E.

[2] We assume that streets are either parallel or orthogonal to axes in the two-dimensional space.

Example 84 *Forecasts by financial analysts*

Financial analysts publish earnings and dividends forecasts, and target prices as well. These forecasts are important for fund managers, banks and investment advisors. Some firms also calculate a market consensus to summarize the forecasts of a set of analysts. For a given firm, the most simple summary consists in averaging the forecasts of analysts. However, such an average provides no clue about the dispersion of forecasts. Imagine that there are two firms F_1 *and* F_2*, and two analysts* A_1 *and* A_2*. The following matrix shows the earnings forecasts:*

$$M = \begin{bmatrix} & A_1 & A_2 \\ F_1 & 11 & 9 \\ F_2 & 2 & 18 \end{bmatrix} \tag{2.3}$$

The consensus (average) is 10 for the two stocks, but the forecasts are much more dispersed for the second stock. Using the consensus to take investment decisions is more risky and error-prone if individual forecasts are more dispersed. More generally, assume that the two analysts provide forecasts on N *stocks and denote* $(p_1^1, ..., p_N^1)$ *et* $(p_1^2, ..., p_N^2)$ *these forecasts. For each company,* i*, the consensus is defined as the average:*

$$\overline{p}_i = \frac{1}{2}\left(p_i^1 + p_i^2\right)$$

The risk of the consensus forecast for firm i *can then be defined as the distance* d_i *between the vector of individual forecasts* (p_i^1, p_i^2) *and the pair of consensus forecasts* $(\overline{p}_i, \overline{p}_i)$ *that would obtain if the two analysts were predicting the same earnings.*

$$d_i = \sqrt{\left(p_i^1 - \overline{p}_i\right)^2 + \left(p_i^2 - \overline{p}_i\right)^2}$$

Of course if the forecasts are actually equal, the distance is 0, meaning that there is no divergence between analysts. The geometric interpretation of this result is simply that identical forecasts lie on the first bisector of the two-

dimensional space where the first (second) axis represents the forecasts of the first (second) analyst. d_i *is in fact the distance between firm* i *and the first bisector. Of course,* d_i *is an oversimplified measure of divergence but it is important to note that all dispersion measures are built in the same spirit.*

In our example, the benchmark is the bisector where all forecasts concerning a given firm are identical. However, suppose that analysts build their forecasts using two types of factors. First, there are macroeconomic factors (or common factors) influencing all firms, and firm-specific factors. In such a framework, it could be better to neutralize the divergence on common factors to measure the forecasting risk of firm i*. For example, if the first analyst is much more optimistic than the second about common factors, her forecasts will be higher on a large part of the firms under scrutiny. The consequence is that most* p_i *will be under the first bisector and this bisector is not the good benchmark! It is necessary to find another benchmark taking into account the divergence about macroeconomic factors. One of the popular methods to do so is Principal Component Analysis which allows to find the line* D *minimizing the following quantity:*

$$\sum_{i=1}^{N} d(p_i, D)^2$$

We do not elaborate in more details this example but the reader should remember that this topic is extensively studied in finance research.

2.1.2 Open sets in metric spaces

Open and closed sets in $\mathbb{R}$ have been introduced in chapter 1 of part I. These concepts are still valid in metric spaces with minor changes in the definitions.

Definition 85 *Let* (E, d) *denote a metric space. An* ***open ball*** *centered at* x *with radius* r*, denoted* $B^{\circ}(x, r)$*, is the set of elements* $y \in E$ *satisfying*

$d(x,y) < r$. This set can be formally written as:

$$B^{\circ}(x,r) = \{y \in E \quad \text{such that } d(x,y) < r\}$$

*A **closed ball** centered at x with radius r, denoted $\overline{B}(x,r)$, is the set of elements $y \in E$ satisfying $d(x,y) \leq r$. This set can be formally written as:*

$$\overline{B}(x,r) = \{y \in E \quad \text{such that } d(x,y) \leq r\}$$

In the set $\mathbb{R}$ of real numbers, equipped with the usual metric, the open ball centered at $x \in \mathbb{R}$ with radius r is simply the interval $]x-r;x+r[$

The corresponding closed ball is the closed interval $[x-r;x+r]$. As a consequence, the concept of an open ball in a metric space is the natural generalization of open intervals in $\mathbb{R}$.

Definition 86 *Let G be a subset of a metric space E.*

*$x \in G$ is **interior** to G if there exists an open ball centered at x with radius $r > 0$, satisfying $B^{\circ}(x,r) \subset G$.*

*G is an **open** set in E if any element in G is interior to G, that is:*

$$\forall x \in G, \exists r \in \mathbb{R}_+^* \text{ such that } B^{\circ}(x,r) \subset G$$

As before, we deduce immediately the definition of a closed set.

Definition 87 *A subset F of a metric space E is **closed** if the complement $F^c = \{x \in E \quad \text{such that } x \notin F\}$ is an open set.*

The disc $G \subset \mathbb{R}^2$ that appears on figure 2.1 is open if the boundary circle is not in G; it is closed otherwise.

The following proposition is valid in any metric space.

Proposition 88 *a) Any union of open sets in E is an open set and any finite union of closed sets in E is a closed set in E.*

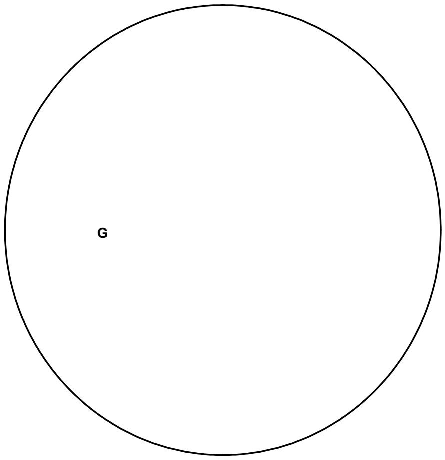

Figure 2.1: Disc in the plane

b) Any intersection of closed sets in E is a closed set and any finite intersection of open sets in E is an open set.

c) E and $\emptyset$ are simultaneously open and closed.

Proposition 88 shows that there is an asymmetry between open and closed subsets. In (a), the union is considered over any number (finite or not) of open subsets but only over a finite number of closed subsets. In (b), the intersection is over any number of closed subsets but over a finite number of open subsets. The difference may be illustrated by the following example; consider the sequence of open intervals $(\left]-\frac{1}{n};\frac{1}{n}\right[, n \in \mathbb{N}^*)$. We get:

$$\bigcap_{n\in\mathbb{N}^*}\left]-\frac{1}{n};\frac{1}{n}\right[=\{0\}$$

The set $\{0\}$ is a closed subset of $\mathbb{R}$ but is written as an infinite intersection of open subsets.

The other topological concepts are generalizations (more or less intuitive) of what we presented in chapter 1 of part I for the set $\mathbb{R}$. We briefly recall these definitions for the sake of completeness.

Definition 89 *a) The* ***closure*** *of a subset* H *of a metric space* (E, d)*, denoted* $\overline{H}$*, is the smallest closed subset such that* $H \subset \overline{H}$*. It is also the intersection of all closed subsets containing* H*.*

b) The ***interior*** *of a subset* H *of a metric space* (E, d)*, denoted* H°*, is the largest open subset such that* $H^{\circ} \subset H$*, or the union of all open subsets included in* H*.*

c) The ***exterior*** *of a subset* H *of a metric space* (E, d)*, is the interior of the complement of* H *in* E*.*

d) The ***frontier*** *of a subset* H *of a metric space* (E, d)*, is the set of elements in* E *that are neither in the interior nor in the exterior of* H*.*

All these definitions were already given in the framework of the metric space $\mathbb{R}$. We need now to add a more abstract concept which will be useful later on.

Definition 90 *a) A subset H in E is a* ***dense subset*** *of E if $\overline{H} = E$ where $\overline{H}$ is the closure of H.*

b) A metric space (E, d) is ***separable*** *if E contains a dense countable subset*[3]*.*

Another way to say the same thing is that for any x in E, there exists a sequence in H converging to x. In short:

$$\forall x \in E, \exists (y_n \in H, n \in \mathbb{N}), \lim_{n \to +\infty} d(y_n, x) = 0 \tag{2.4}$$

or

$$\forall x \in E, \forall \varepsilon > 0, \exists x^* \in H, d(x^*, x) < \varepsilon \tag{2.5}$$

Property 2.5 shows that saying a set is dense has something to do with an approximation. For any element x of E and any distance ε, it is possible to find an element of H as close as desired (at a distance lower than ε) of x.

Example 91 *The most standard (which also proves the most useful) example is the set $\mathbb{Q}$ of rational numbers which is dense in $\mathbb{R}$. For example, in any calculator, numbers like π or e are approximated by rational numbers (that is ratios of integers), with the desired level of accuracy. Doing so is relatively safe because $\mathbb{Q}$ is dense in $\mathbb{R}$.*

In $\mathbb{R}$, a set is *bounded* if it is included in an interval with finite ends. This definition can be easily adapted to metric spaces, using balls instead of intervals.

[3] Recall that a set A is countable if one can "count" its elements. In other words A is countable if there exists a bijection between the set $\mathbb{N}$ of positive integers and A.

Definition 92 *A subset G of a metric space (E, d) is said **bounded** if it is included in a ball $B(x, r)$ with $r < +\infty$.*

A counterintuitive result is that boundedness depends on the metric. A set can be bounded for a given metric and unbounded for another metric. For example, there exists on $\mathbb{R}$ a metric called the *discrete metric* , defined by :

$$d^*(x, y) = \begin{cases} 1 \text{ if } x \neq y \\ 0 \text{ otherwise} \end{cases}$$

The interval $]-\infty; 5]$ is bounded under d^* because it is included in $B^*(0, 1)$. This mapping d^* gives almost no information about the location of points in the metric space. We can only say if two elements are identical or not when we know the distance between them (0 or 1).

Finally, a compact set in $\mathbb{R}$ was a bounded and closed subset. It is still true in R^p, but it is false in general metric spaces[4]. In chapter 2 of part I we saw that a function defined on a compact set reaches its bounds and possesses a maximum and a minimum. This proposition is still valid for functions of several variables considered in this chapter.

2.1.3 Sequences in metric spaces

To define the convergence of a sequence in $\mathbb{R}$, we used absolute values $|x_n - x|$, where x denoted the limit. The same concept in metric spaces uses distances, in particular in R^p. Nothing surprising here because the mapping $(x, y) \to |x - y|$ is a metric on $\mathbb{R}$.

Definition 93 *Let (E, d) denote a metric space, $(x_n, n \in \mathbb{N})$ a sequence of elements of E and x an element of E.*

[4] In a general metric space, a set A is **compact** if, from any sequence of elements of A, it is possible to extract a convergent sub-sequence. In this book we will not need such a general definition.

$(x_n, n \in \mathbb{N})$ ***converges*** *to* x *if* $\lim_{n\to+\infty} d(x_n, x) = 0$. *We write* $\lim_{n\to+\infty} x_n = x$.

$(x_n, n \in \mathbb{N})$ *is called a* ***Cauchy sequence*** *if* $\lim_{i,j\to+\infty} d(x_i, x_j) = 0$.

Proposition 48 in chapter 1 of part 1, related to the convergence of Cauchy sequences in $\mathbb{R}$, does not remain valid in general metric spaces, but remains true in $E = \mathbb{R}^p$.

2.2 Continuity and differentiability

This section introduces continuity and differentiability of functions defined on a subset of $\mathbb{R}^p$ and taking their values in $\mathbb{R}$. The set $\mathbb{R}^p$ is endowed with the Euclidean metric, unless otherwise stated. Tools of the previous section will allow to generalize the notions of limit, continuity and differentiability presented in part I for functions of one variable. Taylor's formula is also generalized.

2.2.1 Limits and continuity

Definition 94 *Let f be a function defined on an open set $D \subset \mathbb{R}^p$. f has a* ***limit*** *$b \in \mathbb{R}$, at $a \in D$ if, for any sequence $(x_n, n \in \mathbb{N})$ in D that converges to a, the sequence $(f(x_n), n \in \mathbb{N})$ converges to b. We write:*

$$\lim_{x\to a} f(x) = b$$

This definition is almost identical to the definition of a limit in chapter 2 of part I. However, here the convergence of $(x_n, n \in \mathbb{N})$ refers to definition 93.

Definition 95 *Let f be a function defined on an open set $D \subset \mathbb{R}^p$. f is*

continuous *at* $x^* = \left(x_1^*, ..., x_p^*\right) \in D$ *if:*

$$\lim_{x \to x^*} f(x) = f(x^*)$$

Here too, the definition is very close to the definition of continuity for functions of one variable. The only difference lies in the use of a metric on R^p. Moreover, it turns out that left and right continuity are meaningless in multidimensional spaces.

Remember that $\lim_{x \to x^*} f(x) = f(x^*)$ can also be written as:

$$\lim_{x \to x^*} |f(x) - f(x^*)| = 0 \tag{2.6}$$

Therefore, continuity could be defined in a much more general way for functions f defined on an open subset D of a metric space (E, d) and taking values in another metric space (F, d_F). In this general framework, f is continuous at $x^* \in E$ if:

$$\lim_{x \to x^*} \ d_F \left(f(x), f(x^*)\right) = 0$$

We just replaced the metric on $\mathbb{R}$ (defined by the absolute value $|f(x) - f(x^*)|$) by $d_F \left(f(x), f(x^*)\right)$. In particular, we encounter this situation when considering functions defined on R^p and taking values in R^m.

2.2.2 Partial derivatives

One of the essential differences between functions of one and of several variables lies in the concept of derivative. Remember that for $f : \mathbb{R} \to \mathbb{R}$, the derivative at x_0 is defined as follows:

$$f'(x_0) = \lim_{h \to 0} \frac{f(x_0 + h) - f(x_0)}{h}$$

There is no obvious generalization for functions of several variables because $x_0 \in \mathbb{R}^p$, $p > 1$. In fact, h should also have p components, and dividing by h would be meaningless. Therefore, for a function depending on p variables, we define p partial derivatives, each one being defined as a derivative with respect to one variable, the other variables being assumed constant.

Definition 96 *The* ***partial derivative*** *of f at x^*, with respect to the i-th variable, is the limit, if it exists, defined by:*

$$\lim_{h\to 0}\frac{f(x_1^*,..,x_i^*+h,...,x_p^*)-f(x^*)}{h}$$

Alternative notations are $\frac{\partial f}{\partial x_i}(x^*)$ or sometimes $f_{x_i}(x^*)$.

We know how to derive a single-variable function. A function of p variables is a function of one variable when $p-1$ variables are held constant. So the definition works as if $p-1$ variables where not changing. In fact, let g be the function defined by:

$$g(x)=f(x_1^*,...,x_{i-1}^*,x,x_{i+1}^*,...,x_p^*)$$

the derivative of g at x_i^* writes:

$$\begin{aligned} g'(x_i^*) &= \lim_{h\to 0}\frac{g(x_i^*+h)-g(x_i^*)}{h} \\ &= \lim_{h\to 0}\frac{f(x_1^*,..,x_i^*+h,...,x_n^*)-f(x^*)}{h}=\frac{\partial f}{\partial x_i}(x^*) \end{aligned}$$

What we just wrote for the i-th variable can be written the same way for the other $p-1$ variables. Consequently, we get p partial derivatives (when the corresponding limits exist).

In financial and economic models, it is common to assume that partial derivatives are continuous functions. Such functions are said C^1-functions.

Example 97 *Let f be a function on $\mathbb{R}_+\times\mathbb{R}_+$ taking values in $\mathbb{R}$ and defined by:*

$$x=(x_1,x_2)\to f(x)=\sqrt{x_1x_2}$$

At any $x^* \in \mathbb{R}_+^* \times \mathbb{R}_+^*$ *this function has partial derivatives defined by:*

$$\frac{\partial f}{\partial x_1}(x^*) = \frac{1}{2}\sqrt{\frac{x_2^*}{x_1^*}}$$
$$\frac{\partial f}{\partial x_2}(x^*) = \frac{1}{2}\sqrt{\frac{x_1^*}{x_2^*}}$$

In fact we can write:

$$f(x) = \sqrt{x_1}\sqrt{x_2}$$

To compute $\frac{\partial f}{\partial x_1}(x^*)$, *we consider that* $\sqrt{x_2}$ *is a number equal to* $\sqrt{x_2^*}$ *(let us denote* c *this number) and we compute the derivative at* x_1^* *of the one-variable* $g(x_1) = c\sqrt{x_1}$. *This derivative is equal to:*

$$g'(x_1^*) = c \times \frac{1}{2\sqrt{x_1^*}}$$

Replace now c *by its value, that is* $\sqrt{x_2^*}$. *The result is:*

$$g'(x_1^*) = \frac{\partial f}{\partial x_1}(x^*) = \frac{1}{2}\sqrt{\frac{x_2^*}{x_1^*}}$$

The computation of $\frac{\partial f}{\partial x_2}(x^*)$ *can be done in the same way, replacing* $\sqrt{x_1}$ *by a number* b *equal to* $\sqrt{x_1^*}$. *The single-variable function is now denoted* m *and defined by* $m(x_2) = b\sqrt{x_2}$. m' *is then calculated as usual.*

$$m'(x_2) = b \times \frac{1}{2\sqrt{x_2^*}} = \frac{\partial f}{\partial x_2}(x^*) = \frac{1}{2}\sqrt{\frac{x_1^*}{x_2^*}}$$

Definition 98 *The* $p-$*dimensional vector* $\frac{\partial f}{\partial x_i}(x^*), i = 1, ..., p,$ *is called the*

gradient *of f at x^*. It is denoted $\nabla f(x^*)$ (spell nabla for ∇):*

$$\nabla f(x^*) = \begin{pmatrix} \frac{\partial f}{\partial x_1}(x^*) \\ \ldots \\ \frac{\partial f}{\partial x_p}(x^*) \end{pmatrix}$$

$\nabla f(x^*)$ is an element of the vector space R^p. Therefore, $\nabla f(x^*)$ denotes a matrix with p rows and 1 column, containing the partial derivatives of f valued at x^*.

2.2.3 Derivatives of compound functions

Calculation of partial derivatives of compound functions obeys the same rules as the ones used for functions of one variable, but the formulation is a bit more complex.

The proposition hereafter presents the case of functions depending on two variables.

Proposition 99 *Let f, g, h three continuously differentiable functions, defined on a open set $D \subset \mathbb{R}^2$. We have:*

$$\frac{\partial}{\partial x}\left[f\left(g(x,y),h(x,y)\right)\right] = \frac{\partial f}{\partial u}\frac{\partial u}{\partial x} + \frac{\partial f}{\partial v}\frac{\partial v}{\partial x}$$

where $u = g(x,y)$ and $v = h(x,y)$. The partial derivative with respect to the second variable y is defined accordingly (replacing ∂x by ∂y).

Example 100 *Let f, g and h be defined as follows:*

$$f(u,v) = \exp(uv) \quad (2.7)$$

$$g(x,y) = x + y \quad (2.8)$$

$$h(x,y) = x - y \quad (2.9)$$

First, calculation of $\frac{\partial f}{\partial u}$ and $\frac{\partial f}{\partial v}$:

$$\frac{\partial f}{\partial u} = vf(u,v) \quad (2.10)$$

$$\frac{\partial f}{\partial v} = uf(u,v) \quad (2.11)$$

Second, calculation of $\frac{\partial u}{\partial x}$ and $\frac{\partial v}{\partial x}$:

$$\frac{\partial u}{\partial x} = 1 = \frac{\partial v}{\partial x} \quad (2.12)$$

Finally

$$\frac{\partial}{\partial x}\left[f\left(g(x,y),h(x,y)\right)\right] = (x-y)\exp(x^2-y^2)+(x+y)\exp(x^2-y^2) \quad (2.13)$$
$$= 2x\exp(x^2-y^2) \quad (2.14)$$

Of course, in this example it would have been easier to directly replace u and v by their values and start with $f(x,y)=\exp(x^2-y^2)$

2.2.4 Differential of a function depending on several variables

In chapter 2 of part I, a C^1-function was approximated at $x+h$ by $f(x)+hf'(x)$ (first-order approximation). The mapping $h\to hf'(x)$ is linear. The differential of a function of several variables carries the same idea. However, starting from $x\in R^p$, we can move in different directions. In other words, writing $x+h$ refers to a vector like

$$x+h=\begin{pmatrix} x_1+h_1 \\ \ldots \\ x_p+h_p \end{pmatrix}$$

We then refer to a displacement in the direction of vector h.

The fact that partial derivatives exist is not sufficient to ensure that a function of several variables is continuous. There exist some pathological cases where the function possesses partial derivatives according to definition 95 but is not continuous. To solve this difficult question we need a little bit more, that is differentiability. We provide hereafter the definition of this word but in the sequel of the book we will in fact use a stronger (but much more intuitive) assumption.

Definition 101 *A function f defined on an open subset D of $\mathbb{R}^p$ is* ***differ-***

entiable *at $x \in D$ if there exists $\alpha \in \mathbb{R}^p$ such that:*

$$f(x+h) = f(x) + \sum_{i=1}^{p} \alpha_i \frac{\partial f}{\partial x_i}(x) + \|h\| \, \varepsilon(h)$$

with $\lim_{h \to \mathbf{0}} \varepsilon(h) = 0$.

Remember that $\|h\|$ denotes the norm of vector h. In the limit, we used bold characters for the null vector to emphasize the fact that h is a vector, not a number. In the remaining of the text, we use proposition 102 to simplify the formulation of propositions.

Proposition 102 *Any C^1 function at x is differentiable at x.*

In the following definitions and propositions we assume that functions are C^1 over the interior of their domain. Mathematicians would say that weaker assumptions are better, but assuming C^1-functions is general enough for economics and finance.

Definition 103 *Let f be a C^1-function defined on an open set $D \subset \mathbb{R}^p$. The **differential** of f at x^* is the linear form, denoted df_{x^*}, defined on $\mathbb{R}^p$ as follows:*

$$df_{x^*}(h) = \sum_{i=1}^{p} \frac{\partial f}{\partial x_i}(x^*) h_i$$

with $h^T = (h_1, ..., h_p)$.

Using notations of chapter 1, $df_{x^*}(h)$ writes as the following inner product in R^p :

$$df_{x^*}(h) =< \nabla f(x^*), h >$$

Proposition 79 of chapter 1 (Riesz representation theorem) allows to say that $\nabla f(x^*)$ represents the linear mapping df_{x^*} because, for any h:

$$df_{x^*}(h) =< \nabla f(x^*), h >$$

The cases $p = 1$ and $p = 2$ reveal the intuition behind the definition of differentials. Assume that the components of h are close to 0; $df_{x^*}(h)$ then approximates the difference $f(x^* + h) - f(x^*)$. If $p = 1, hf'(x^*)$ is a first-order approximation of $f(x^* + h) - f(x^*)$. The derivative $f'(x^*)$ also denotes the slope of the tangent to the curve representing f. If $p = 2$, the geometric interpretation of the differential is the same; the mapping $h \to df_{x^*}(h) = h_1 \frac{\partial f}{\partial x_1}(x^*) + h_2 \frac{\partial f}{\partial x_2}(x^*)$ approximates the surface f at x^* by the two-dimensional space tangent to the surface at x^*. As f is a C^1-function, it is differentiable; therefore this approximation is valid when the norm of h is small, that is when $x^* + h$ is close to x^* in the space $\mathbb{R}^p$.

Example 104 ***Interpretation of differentials***

Come back to the function $f(x) = \sqrt{x_1 x_2}$ *and define* $df_{x^*}(h)$ *for* $(x^*)^T = (1; 1)$. *Example 97 indicates that:*

$$df_{x^*}(h) = \frac{1}{2}(h_1 + h_2)$$

The set of all points such that $f(x) = 1$ *is called a* ***level curve*** *of* f[5]. *This set contains* x^* *and the equation* $\sqrt{x_1 x_2} = 1$ *implies that elements in this set satisfy:*

$$x_2 = \frac{1}{x_1}$$

On figure 2.2, x_1 (x_2) *is the coordinate on the horizontal (vertical) axis. The slope at* $(x_1, x_2) = (1, 1)$ *is* -1 *because the derivative of* $g(x_1) = 1/x_1$ *at* $x_1 = 1$ *is equal to* -1. *Moreover, the coordinates of the gradient of* f *are* $(1/2; 1/2)$. *The arrow on the figure gives the direction of the gradient; it lies on the line* $x_1 = x_2$. *This gradient is orthogonal to the tangent to the level curve. This remark is not a surprise because the level curve is the curve along which the function* $f(x_1, x_2)$ *is constant, equal to 1. Therefore, moving from*

[5] The general definition is the following: a level curve $c \in \mathbb{R}$ of a function f is the set of elements x satisfying $f(x) = c$.

x^* to $x^* + h$ along this curve keeps the differential equal to 0, that is:

$$\frac{\partial f}{\partial x_1} h_1 + \frac{\partial f}{\partial x_2} h_2 = 0$$

Using the inner product, this equation writes:

$$< \nabla f(x^*), h >= 0$$

In the neighborhood of x^*, this relationship means that the gradient and the tangent to the level curve are orthogonal.

*If f represents the utility function of an investor, the level curve $f(x_1, x_2) = 1$ defines the pairs of consumed quantities generating the same level (equal to 1) of utility. The set of all these pairs is called an **indifference curve**.*

At $(1, 1)$, the investor is indifferent if the quantity of one good marginally increases while the quantity of the other marginally decreases. At $y^* = (2, 1)$, the story is different. The differential writes:

$$df_{y^*}(h) = \sqrt{2} h_1 + \frac{1}{\sqrt{2}} h_2$$

*For $df_{y^*}(h)$ to be zero and the investor be indifferent to a substitution between the two goods, it is necessary that he obtains twice as much of good 2 than the quantity of good 1 he gives up. This ratio is the well-known **marginal rate of substitution** between the goods.*

Differentials follow the same rules as derivatives, as can be seen in the following proposition.

Proposition 105 *Let f and g be two C^1-functions, defined on an open set $D \subset \mathbb{R}^p$, and denote x^* an element of D. We have:*

1) $d(f + g)_{x^*} = df_{x^*} + dg_{x^*}$

2) $d(af)_{x^*} = adf_{x^*}$ for any $a \in \mathbb{R}$

3) $d(fg)_{x^*} = f(x^*)dg_{x^*} + g(x^*)df_{x^*}$

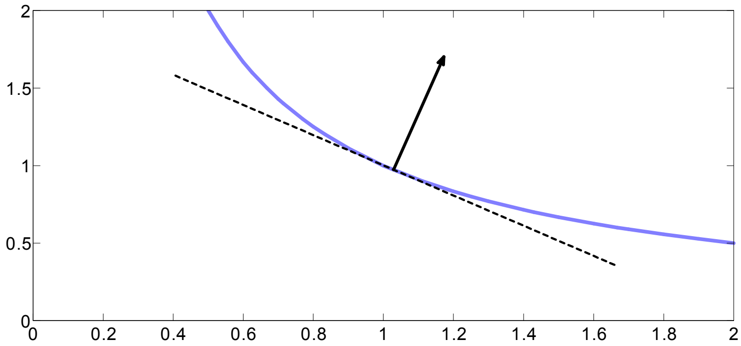

Figure 2.2: Gradient of $f(x_1, x_2) = x_1 x_2$

4) If $g(x^*) \neq 0$, *then* $d\left(\frac{f}{g}\right)_{x^*} = \frac{g(x^*)df_{x^*} - f(x^*)dg_{x^*}}{g(x^*)^2}$

(1) and (2) are obvious as consequences of the definition of partial derivatives. To prove (3), just write:

$$\begin{aligned} d(fg)_{x^*} &= \sum_{i=1}^{p} \frac{\partial(fg)}{\partial x_i}(x^*)h_i = \sum_{i=1}^{p} \left(g(x^*)\frac{\partial f}{\partial x_i}(x^*) + f(x^*)\frac{\partial g}{\partial x_i}(x^*) \right) h_i \\ &= g(x^*)\sum_{i=1}^{p} \frac{\partial f}{\partial x_i}(x^*)h_i + f(x^*)\sum_{i=1}^{p} \frac{\partial g}{\partial x_i}(x^*)h_i = f(x^*)dg_{x^*} + g(x^*)df_{x^*} \end{aligned}$$

The proof of (4) uses the same method, applying the rules of derivation for ratios of functions.

Alternate notations In most academic papers and economic textbooks, authors do not write $h \to df_x(h)$, even if it is the right way to understand that df_x is a linear mapping. In most cases, authors write:

$$df(x) = \frac{\partial f}{\partial x_1}dx_1 + ... + \frac{\partial f}{\partial x_p}dx_p$$

This simplified notation means that $dx_1, dx_2, ..., dx_p$ correspond to $h_i, i = 1, ..., p$ and $df(x)$ means $df_x(h)$, the differential of f evaluated at x.

2.2.5 The mean value theorem

In chapter 2 of part I, Rolle's theorem says that if a function g, defined on $[a; b]$, is differentiable on $]a; b[$, there exists $c \in \,]a; b[$ satisfying $g(b) - g(a) = (b - a)g'(c)$.

A similar result is valid for two-variable functions. However, one needs to be prudent in interpreting the result because of the existence of several partial derivatives.

The following example shows that the intuition is the same as in the single-variable case.

Denote f the function defined by :

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

The graph of f over the square $[1;3] \times [1;3]$ appears on figure 2.3.

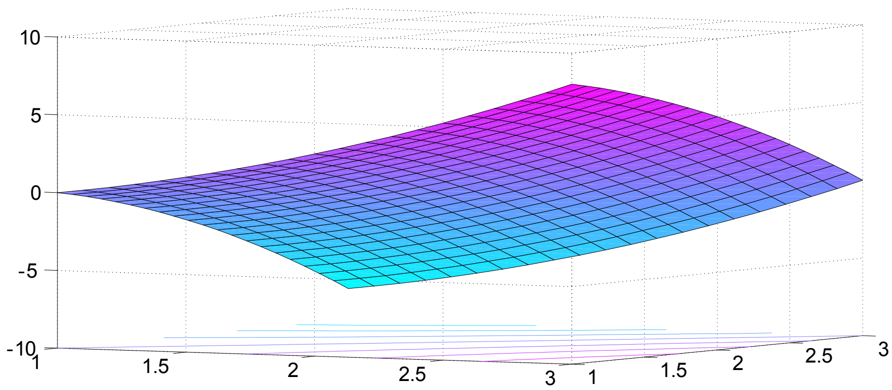

Figure 2.3: The function $f(x,y) = x^2 - y^2$

f satisfies $f(1,1) = 0$ and $f(3,2) = 5$. We can decompose $f(3,2) - f(1,1)$ as follows:

$$f(3,2) - f(1,1) = f(3,2) - f(1,2) + f(1,2) - f(1,1) \qquad (2.15)$$

Let $h(x) = f(x,2)$ and $k(y) = f(1,y)$. Equation (2.15) writes:

$$f(3,2) - f(1,1) = h(3) - h(1) + k(2) - k(1)$$

We now apply Rolle's theorem (chapter 2, part I) to the functions h and k. Therefore, there exist $c_1 \in]1;3[$ and $c_2 \in]1;2[$ such that:

$$\begin{aligned} h(3) - h(1) &= (3-1) \times h'(c_1) \\ k(2) - k(1) &= (2-1) \times k'(c_2) \end{aligned}$$

$h'(c_1)$ is the partial derivative of f with respect to x, evaluated at $(c_1; 2)$. $k'(c_2)$ is the partial derivative of f with respect to y, evaluated at $(1, c_2)$.

The mean value theorem hereafter formalizes the idea of the above example.

Proposition 106 *Let f be a C^1-function, defined on $D =]a_1; b_1[\times]a_2; b_2[\subset \mathbb{R}^2$, and $(x_1, y_1), (x_2, y_2)$ be two elements of D.*

There exists $(z_1, z_2) \in D$ such that:

$$f(x_2, y_2) - f(x_1, y_1) = (x_2 - x_1)\frac{\partial f}{\partial x}(z_1, y_2) + (y_2 - y_1)\frac{\partial f}{\partial y}(x_1, z_2)$$

In this proposition, we restrict the domain to a rectangle of $\mathbb{R}^2$. This assumption is not the most general but the key point is that (z_1, z_2) is in D.

2.2.6 Second-order partial derivatives

In the preceding section we showed that a p-variable function has p first-order partial derivatives. The calculation of second-order partial derivatives needs to derive any of the p first-order derivatives with respect to any of the p variables. More precisely, each partial derivative $\frac{\partial f}{\partial x_i}(x), i = 1, 2, ..., p$ can be derived with respect to each variable $x_j, j = 1, 2, ..., p$. As a consequence, the function possesses p^2 second-order partial derivatives. They are organized in a (p, p) matrix called the Hessian matrix or, in short, the Hessian of f.

Definition 107 *Let f a C^1-function, defined on an open set $D \subset \mathbb{R}^p$. The **Hessian matrix** (or **Hessian**) of f at x^*, denoted $H_f(x^*)$, is the (p, p) matrix defined by:*

$$H_f(x^*) = \left[\frac{\partial^2 f}{\partial x_i \partial x_j}(x^*)\right]_{\substack{i=1,\dots,p \\ j=1,\dots,p}}$$

where $\frac{\partial^2 f}{\partial x_i \partial x_j}(x^*) = \frac{\partial}{\partial x_i}\left(\frac{\partial f}{\partial x_j}(x^*)\right)$ is the partial derivative with respect to x_i (when it exists) of the partial derivative of f with respect to x_j.

Diagonal elements of $H_f(x^*)$ are denoted:

$$\frac{\partial^2 f}{\partial x_i \partial x_i}(x^*) = \frac{\partial^2 f}{\partial x_i^2}(x^*)$$

Proposition 108 *If the second-order partial derivatives of a function f are continuous at x^* (f is called a C^2-function), the Hessian matrix is symmetric, that is:*

$$\frac{\partial^2 f}{\partial x_i \partial x_j}(x^*) = \frac{\partial^2 f}{\partial x_j \partial x_i}(x^*)$$

This proposition shows that when calculating the second-order derivatives, the order you choose to derive does not matter, the final result, either $\frac{\partial^2 f}{\partial x_i \partial x_j}(x^*)$ or $\frac{\partial^2 f}{\partial x_j \partial x_i}(x^*)$, is the same. The Hessian matrix is important when solving optimization problems, especially to get sufficient conditions of optimality (chapters 3 and 4). Fortunately, in finance problems, the Hessian matrix is always symmetric.

Example 109 *Coming back to the function defined in example 97, that is*

$$f(x) = \sqrt{x_1 x_2}$$

we calculated:

$$\begin{aligned}
\frac{\partial f}{\partial x_1}(x^*) &= \frac{1}{2}\sqrt{\frac{x_2^*}{x_1^*}} = \frac{1}{2}\left(x_1^*\right)^{-\frac{1}{2}}\left(x_2^*\right)^{\frac{1}{2}} \\
\frac{\partial f}{\partial x_2}(x^*) &= \frac{1}{2}\sqrt{\frac{x_1^*}{x_2^*}} = \frac{1}{2}\left(x_1^*\right)^{\frac{1}{2}}\left(x_2^*\right)^{-\frac{1}{2}}
\end{aligned}$$

The Hessian matrix is then obtained as follows:

$$\begin{aligned}
\frac{\partial^2 f}{\partial x_1^2}(x^*) &= -\frac{1}{4}\left(x_1^*\right)^{-\frac{3}{2}}\left(x_2^*\right)^{\frac{1}{2}} = -\frac{1}{4x_1^*}\sqrt{\frac{x_2^*}{x_1^*}} \\
\frac{\partial^2 f}{\partial x_1 \partial x_2}(x^*) &= \frac{1}{4}\sqrt{\frac{1}{x_1^* x_2^*}} \\
\frac{\partial^2 f}{\partial x_2^2}(x^*) &= -\frac{1}{4}\left(x_1^*\right)^{\frac{1}{2}}\left(x_2^*\right)^{-\frac{3}{2}} = -\frac{1}{4x_2^*}\sqrt{\frac{x_1^*}{x_2^*}} \\
\frac{\partial^2 f}{\partial x_2 \partial x_1}(x^*) &= \frac{1}{4}\sqrt{\frac{1}{x_1^* x_2^*}}
\end{aligned}$$

In short we write $H_f(x^*)$*:*

$$H_f(x^*) = \begin{bmatrix} -\frac{1}{4x_1^*}\sqrt{\frac{x_2^*}{x_1^*}} & \frac{1}{4}\sqrt{\frac{1}{x_1^* x_2^*}} \\ \frac{1}{4}\sqrt{\frac{1}{x_1^* x_2^*}} & -\frac{1}{4x_2^*}\sqrt{\frac{x_1^*}{x_2^*}} \end{bmatrix} = \frac{1}{4\sqrt{x_1^* x_2^*}} \begin{bmatrix} -\frac{x_2^*}{x_1^*} & 1 \\ 1 & -\frac{x_1^*}{x_2^*} \end{bmatrix}$$

2.2.7 Taylor's formula

When a function f depends on a single variable, we know (chapter 2, part I) that the graph of f can be approximated by a straight line or by a curve representing a polynomial. Taylor's formula allows to calculate the coefficients of this polynomial.

Differentials allow to approximate C^1-functions of several variables at the first order. To approximate functions at higher orders, the right tool is a Taylor's series expansion. We restrict our presentation to second-order approximations because such a choice covers 99.9% of economic and financial models.

Definition 110 *Let β and γ two functions depending on a single variable h. We say that β has the same order of magnitude as γ in the neighborhood of 0, and we write $\beta = O(\gamma)$, if $\lim_{h\to 0}\left|\frac{\beta(h)}{\gamma(h)}\right| < +\infty$. In the same way, β is negligible with respect to γ in the neighborhood of 0 if $\lim_{h\to 0}\frac{\beta(h)}{\gamma(h)} = 0$. In this case, we note $\beta = o(\gamma)$. These two notations O and o are called* ***Landau notations.***

Using Landau notations allows to simplify formulas. $\beta = O(\gamma)$ means that $\beta(h)$ and $\gamma(h)$ are comparable in the following sense. The function β is not infinitely larger (smaller) than the function γ when h tends to 0.

$\beta = o(\gamma)$ means that β is negligible with respect to γ when h tends to 0. If such a situation occurs, that is $\beta = o(\gamma)$, the sum $\beta(h)+\gamma(h)$ is approximated by $\gamma(h)$ because $\beta(h)$ is negligible. Of course this approximation is valid only if h is close to 0.

Proposition 111 ***Taylor's formula***

Let f denote a C^2-function, defined on an open set $D \subset \mathbb{R}^p$, and $(x, x^*) \in D \times D$ such that the line joining x and x^* is in D. We have:

$$\begin{aligned} f(x) &= f(x^*) + \sum_{i=1}^{p}(x_i - x_i^*)\frac{\partial f}{\partial x_i}(x^*) \\ &+\frac{1}{2}\sum_{i=1}^{p}\sum_{j=1}^{p}(x_i - x_i^*)\left(x_j - x_j^*\right)\frac{\partial^2 f}{\partial x_i \partial x_j}(x^*) + o\left(\sum_{i=1}^{p}(x_i - x_i^*)^2\right) \end{aligned}$$

The expression $o\left(\sum_{i=1}^{n}(x_i - x_i^*)^2\right)$ means that, when the distance between x and x^* tend to zero, all terms of order greater than 2 are negligible with respect to first and second-order terms that appear in the formula as coefficients of the partial derivatives. This Taylor's formula allows to approximate a function of p variables by a second-degree polynomial.

Matrix notation Using the Hessian matrix and the gradient of f shortens the above formula as follows:

$$f(x) = f(x^*) + < \nabla f(x^*), x - x^* > + \frac{1}{2}(x - x^*)^T H_f(x^*)(x - x^*) + o(\|x - x^*\|^2)$$

This alternative formulation is based on notions presented in chapter 4 of the part I and in chapter 1 of the present book, (inner product of vectors or product of matrices).

Example 112 *Let f be a function defined on $\mathbb{R}^2$ taking values in $\mathbb{R}_+$ and defined by :*

$$f(x, y) = \exp\left(-\frac{1}{2}\left(x^2 + y^2\right)\right)$$

The partial derivatives of f are equal to:

$$\begin{aligned}
\frac{\partial f}{\partial x} &= -x \exp\left(-\frac{1}{2}\left(x^2 + y^2\right)\right); \frac{\partial f}{\partial y} = -y \exp\left(-\frac{1}{2}\left(x^2 + y^2\right)\right) \\
\frac{\partial^2 f}{\partial x^2} &= \left(x^2 - 1\right) \exp\left(-\frac{1}{2}\left(x^2 + y^2\right)\right); \frac{\partial^2 f}{\partial y^2} = \left(y^2 - 1\right) \exp\left(-\frac{1}{2}\left(x^2 + y^2\right)\right) \\
\frac{\partial^2 f}{\partial x \partial y} &= xy \exp\left(-\frac{1}{2}\left(x^2 + y^2\right)\right)
\end{aligned}$$

Applying Taylor's formula at $(0, 0)$ leads to write:

$$f(h_1, h_2) = 1 - \frac{1}{2}\left[h_1^2 + h_2^2\right] + o\left(h_1^2 + h_2^2\right)$$

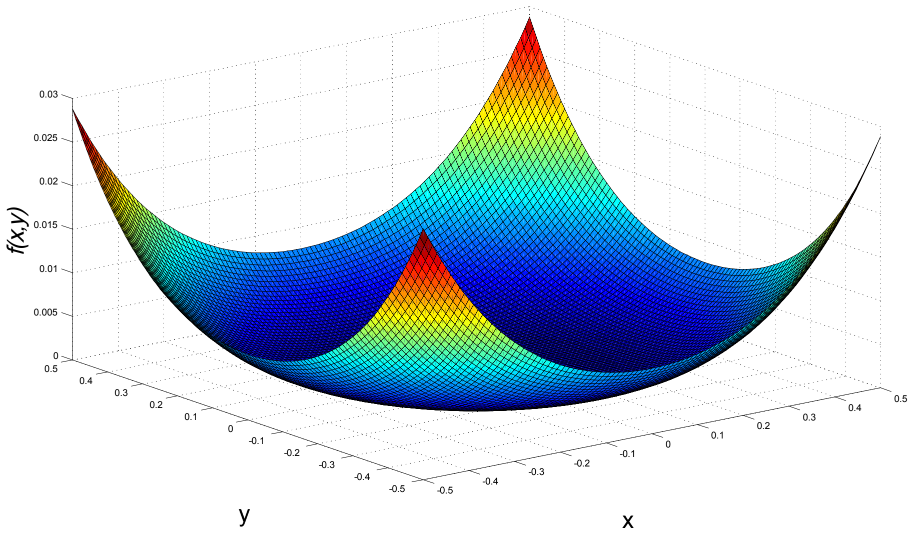

Figure 2.4: Approximation error

Figure 2.4 shows the difference between f *and the second-degree polynomial when* h_1 *and* h_2 *move between* $-0,5$ *and* $0,5$. *This difference is represented by* $o\left(h_1^2 + h_2^2\right)$. *We observe that, even "far" from (0,0), the approximation is quite good. The error is not larger than 0.03 with the function being valued 1 at 0. Of course, this case is specific; choosing a more complicated function could lead to approximations of lower quality.*

2.2.8 Convex and concave functions

We presented in the chapter 2 of part I the definition of a convex single-variable function f defined on an interval $I \subset \mathbb{R}$. You will see in the following that the definition is almost the same when f depends on p variables, except for the domain of definition D. Of course, D must be included in $\mathbb{R}^p$ but we have to be sure that the definition of a convex function is meaningful. It is the reason why we first introduce convex sets in a vector space like $\mathbb{R}^p$.

Definition 113 *Let C be a subset of $\mathbb{R}^p$. C is* ***convex*** *if:*

$$\forall (x,y) \in C \times C, \forall \alpha \in [0;1], \alpha x + (1-\alpha)y \in C$$

The geometric interpretation of this definition is simple If any two elements x and y are in the same convex set, all the segment joining x and y is also included in C. Remark in passing that if $p = 1$, C is an interval.

$\mathbb{R}^p$ being a vector space, the combination $\alpha x + (1-\alpha)y$ in definition 113 is a linear combination of the vectors x and y. This linear combination has two specific features; the coefficients α and $1-\alpha$ are positive and their sum equals 1. Such a combination of vectors is called a **convex combination**.

We can now rigorously define convex and concave functions.

Definition 114 *1) Let f be a function defined on a convex domain $D \subset \mathbb{R}^p$. f is a **convex function** on D if, for any $\alpha \in [0;1]$ and any couple $(x, y) \in D \times D$, we have:*

$$f(\alpha x + (1-\alpha) y) \leq \alpha f(x) + (1-\alpha) f(y)$$

*2) Under the same assumptions on D, f is a **concave function** on D if the inequality is reversed.*

*3) The function f is **strictly convex (concave)** if the inequality is strict in part 1 (2) of the definition.*

Assuming convex or concave functions is very common in finance or economic models. Utility functions are usually concave and cost functions are convex. These assumptions make easier solving optimization problems. These issues are addressed in chapters 3 and 4.

In chapter 2 of part I, we characterized convex (concave) functions by positive (negative) second-order derivatives. For a function f depending on p variables, the corresponding result uses second-order partial derivatives, by means of a condition on the Hessian matrix of f.

Proposition 115 *Let f be a C^2-function defined on a convex open domain $D \subset \mathbb{R}^p$.*

1) f is convex (concave) on D if and only if the Hessian matrix $H_f(x)$ is positive (negative) semi-definite at any $x \in D$.

2) If $H_f(x)$ is positive (negative) definite, f is strictly convex (concave).

Recall that positive semi-definite matrices have been presented in chapter 4 of part I (definition 52).

Example 116 *Consider a two-goods economy; an agent is characterized by the following utility function U defined on $D = \mathbb{R}_+^* \times \mathbb{R}_+^*$:*

$$U(x) = \ln(x_1 x_2)$$

where $x = (x_1,\, x_2)$ is the vector of consumed quantities and $U(x)$ measures the welfare generated by consumption. U is strictly concave on D. In fact, we have:

$$\frac{\partial U}{\partial x_1} = \frac{1}{x_1} \qquad \frac{\partial U}{\partial x_2} = \frac{1}{x_2}$$
$$\frac{\partial^2 U}{\partial x_1^2} = -\frac{1}{x_1^2} \qquad \frac{\partial^2 U}{\partial x_2^2} = -\frac{1}{x_2^2}$$
$$\frac{\partial^2 U}{\partial x_1 \partial x_2} = 0$$

It follows:

$$H_U(x) = \begin{pmatrix} -\frac{1}{x_1^2} & 0 \\ 0 & -\frac{1}{x_2^2} \end{pmatrix}$$

We can check that $H_U(x)$ is negative definite by computing $y^T H_U(x) y$, where y is a non zero vector in $\mathbb{R}^2$.

$$\begin{aligned} y^T H_U(x) &= (y_1, y_2) \begin{pmatrix} -\frac{1}{x_1^2} & 0 \\ 0 & -\frac{1}{x_2^2} \end{pmatrix} = \left(-\frac{y_1}{x_1^2}, -\frac{y_2}{x_2^2}\right) \\ y^T H_U(x) y &= \left(-\frac{y_1}{x_1^2}, -\frac{y_2}{x_2^2}\right) \begin{pmatrix} y_1 \\ y_2 \end{pmatrix} = -\frac{y_1^2}{x_1^2} - \frac{y_2^2}{x_2^2} \end{aligned}$$

We get $y^T H_U(x) y < 0$ showing that U is strictly concave. The interpretation of the concavity of U is the same as the one provided for single-variable functions. The utility obtained by consuming one more unit of a given good decreases with the quantity already consumed.

2.3 Implicit and homogeneous functions

2.3.1 The implicit function theorem

Several economic variables are often linked by complex relationships so that it is impossible to express these relationships explicitly. The most well-known example of such a relationship is the definition of the internal rate of return or, equivalently, of the yield of a coupon-bearing bond. The yield r of a bond is linked to the price p and to the future payoffs $F_1, ..., F_T$ of the bond (coupons plus reimbursment price). Though economically intuitive, it is impossible to express r as an explicit function of the variables $(p, F_1, ..., F_T)$. In the same spirit, the utility provided by the consumption of a bundle of goods $(x_1, ..., x_p)$ is measured by a utility function $U(x_1, ..., x_p)$ taking values in $\mathbb{R}$. For a given utility level u, the equation $U(x_1, ..., x_p) = u$ creates a relationship between x_1 and the $p-1$ other variables. In general, no explicit formulation exists for this relationship.

In this section, we develop some results allowing to measure the sensitivity of a given variable with respect to variations in other variables. We start by the most simple case where a function F only depends on two variables.

Definition 117 *Let F be a function defined on an open subset $D \subset \mathbb{R}^2$. The equation $F(x, y) = 0$ defines an implicit function if there exists a function $g(x) = y$, defined on an interval and taking values in an interval such that $F(x, g(x)) = 0$.*

Of course the relationship between x and y is said implicit when g cannot be defined explicitly.

Proposition 118 ***Implicit function theorem (2 variables)***

If $F : \mathbb{R}^2 \to \mathbb{R}$ is C^1 and defines an implicit function g by means of the relationship $F(x, y) = 0$, we have :

$$g'(x) = \frac{\partial y}{\partial x} = -\frac{\frac{\partial F}{\partial x}}{\frac{\partial F}{\partial y}}$$

The notation $\frac{\partial y}{\partial x}$ may seem surprising: it is not really a rigorous way to write such a derivative but this formulation is common. It is the reason why we use this expression. In the same spirit, the notations $\frac{\partial F}{\partial x}$ and $\frac{\partial F}{\partial y}$ do not look precise enough because we do not specify the values at which the partial derivatives are calculated. But these notations are commonly used as long as they do not introduce confusion or ambiguity. In our example, we know that the partial derivatives are calculated at (x, y) such that $F(x, y) = 0$.

Theorem 118 is useful to perform comparative statics. Being given the value of a function of two variables (production function, utility function, net present value, etc.), comparative statics tests the impact of the variation of one variable on the value of the other variable.

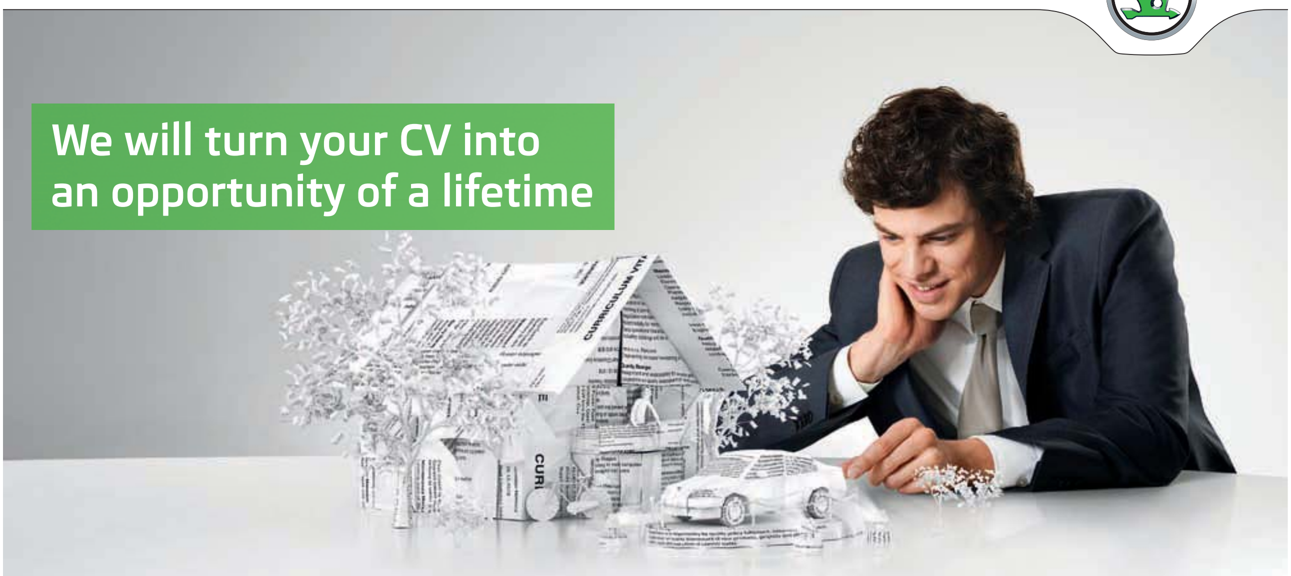

Example 119 ***The internal rate of return***

An investment project needs an initial outflow followed (in general) by inflows at future dates. Let F_0 *denotes the initial negative cashflow and* $F_t, t = 1, ..., T$ *the future positive cashflows, where* T *denotes the maturity date of the project. The* ***net present value*** *of this project, discounted at a rate* r*, is defined by:*

$$NPV(r) = \sum_{t=0}^{T} \frac{F_t}{(1+r)^t}$$

The ***internal rate of return*** *(IRR) is the rate* r^* *satisfying* $NPV(r^*) = 0$.

This equation defines an implicit relationship between the discount rate r^* *and any given cashflow of the project. This relationship is called an implicit function because you cannot write* r^* *as follows:*

$$r^* = f(T, F_t, t = 0, ..., T)$$

But the implicit function theorem allows to calculate the sensitivity of r^* *with respect to variations in any given cashflow. For example:*

$$\frac{\partial r}{\partial F_0} = -\frac{\frac{\partial NPV}{\partial F_0}}{\frac{\partial NPV}{\partial r}} = -\frac{1}{\frac{\partial NPV}{\partial r}} \qquad (2.16)$$

We can also write:

$$\frac{\partial NPV}{\partial r} = -\sum_{t=0}^{T} \frac{tF_t}{(1+r)^{t+1}} \qquad (2.17)$$

Equations (2.16) and (2.17) lead to:

$$\frac{\partial r}{\partial F_0} = \frac{1}{\sum_{t=0}^{T} tF_t(1+r)^{-t-1}}$$

A too superficial look at this formula could let the reader think that the IRR *is an increasing function of the initial cost of the project because the derivative is positive. But remember that $F_0 < 0$. As a consequence, a marginal increase in F_0 is in fact a marginal decrease of the cost of the project, everything else equal. Proposition 118 can be generalized to p-variable functions almost without modifications. Any equation $F(x) = 0$ where $x = (x_1, ..., x_p)$ defines an implicit function between components x_j and x_k for any (j, k) in $\{1, ..., p\}^2$.*

Proposition 120 ***Implicit function theorem (p variables)***

Let F be a C^1-function, defined on an open set $D \subset \mathbb{R}^p$. Assume that F defines an implicit function linking x_j and x_k by means of the equation $F(x) = 0$. It follows that:

$$\frac{\partial x_k}{\partial x_j} = -\frac{\frac{\partial F}{\partial x_j}}{\frac{\partial F}{\partial x_k}}$$

x_k is the k-th variable and x_j the j-th variable. This proposition is not very different from proposition 118 because the other variables do not play any role. It is as if we were dealing with a two-variable function, the $(p-2)$ others being kept constant.

2.3.2 Homogeneous functions and Euler theorem

Homogeneous functions are common in economics. The most well known example is the production function of a firm. When all production factors are doubled, the usual assumption is to consider that production will double. In such a situation, the function is said homogeneous of degree 1. A second example in finance is the price of an option contract when considered as a function of two variables, the strike price and the underlying price. When you double the two, the price of the option doubles. The definition below is the generalization of this intuitive example.

Definition 121 *Let f be a function defined on a set $D \subset \mathbb{R}^p$, taking values in $\mathbb{R}$, and let D^* denote a subset de D. f is* ***homogeneous*** *of degree α on*

D^* *if:*

$$\forall x \in D^*, \forall \lambda \in \mathbb{R}_+, \lambda x \in D \text{ and } f(\lambda x) = \lambda^\alpha f(x)$$

When a function is homogeneous of degree $\alpha > 1$, doubling the inputs more than doubles the output. For production functions, it is the sign of economies of scale. The unit cost of production decreases when produced quantities increase.

A homogeneous function of degree 1 satisfies $f(2x) = 2f(x)$; this equality is also true for a linear function. However, if linear forms are homogeneous of degree 1 the reciprocal is false. The function $f(x_1, x_2) = \sqrt{x_1 x_2}$ is a simple counterexample. Of course, f is not linear but is homogeneous of degree 1 because $f(2x_1, 2x_2) = \sqrt{2x_1 \times 2x_1} = 2\sqrt{x_1 x_2} = 2f(x_1, x_2)$.

The following proposition shows that a homogeneous function has homogeneous partial derivatives. Only the degree of homogeneity changes.

Proposition 122 *Let f be a C^1 function defined on an open set $D \subset \mathbb{R}^p$, homogeneous of degree α on $D^* \subset D$. The functions $\partial f / \partial x_i$ are homogeneous of degree $\alpha - 1$ on D^*.*

Proof. We can write:

$$f(\lambda x) = \lambda^\alpha f(x) \Rightarrow \frac{\partial}{\partial x_i} \left[f(\lambda x) \right] = \frac{\partial}{\partial x_i} \left[\lambda^\alpha f(x) \right] \tag{2.18}$$

Let h be defined by $x \to f(\lambda x)$. h writes $f \circ g$ where $g(x) = \lambda x$. Consequently:

$$\frac{\partial}{\partial x_i} \left[f(\lambda x) \right] = \lambda \frac{\partial f}{\partial x_i}(\lambda x)$$

The linearity of derivations leads to:

$$\frac{\partial}{\partial x_i} \left[\lambda^\alpha f(x) \right] = \lambda^\alpha \frac{\partial f}{\partial x_i}(x)$$

It implies:

$$\frac{\partial f}{\partial x_i}(\lambda x) = \lambda^{\alpha - 1} \frac{\partial f}{\partial x_i}(x)$$

This equality shows that $\frac{\partial f}{\partial x_i}$ is homogeneous of degree $\alpha - 1$. ■

When functions are homogeneous of degree 1 (for example $f(x_1, x_2) = \sqrt{x_1 x_2}$), the proposition means that partial derivatives are homogeneous of degree 0. In fact we have:

$$\frac{\partial f}{\partial x_1} = \frac{1}{2}\sqrt{\frac{x_2}{x_1}}$$

Multiplying x_1 and x_2 by a non-zero number does not change the value of $\frac{\partial f}{\partial x_1}$. If f is the utility function of an investor, the proposition shows that the marginal utility provided by the consumption of a marginal quantity of good 1 is the same when the quantity already consumed is (x_1, x_2) or when it is $(\lambda x_1, \lambda x_2)$ with $\lambda > 0$. Geometrically, this result is not surprising because, along the line $x_1 = x_2$, f is a linear function. In fact, $f(x, x) = x$ for any x.

The specific features of homogeneous functions lead to the Euler theorem that links the value of a homogeneous function at a given point to the values of its partial derivatives at this same point.

Proposition 123 ***The Euler theorem***

Let f be a C^1 function defined on $\left(\mathbb{R}_+^*\right)^p$, homogeneous of degree α. At any $x \in \left(\mathbb{R}_+^*\right)^p$, we have:

$$\sum_{i=1}^{p} x_i \frac{\partial f}{\partial x_i}(x) = \alpha f(x)$$

We only provide hereafter a sketch of the proof. By definition of homogeneity we know that:

$$f(\lambda x) = \lambda^\alpha f(x) \tag{2.19}$$

Each side of equation 2.19 is a function of λ (that is the point!). The derivatives of the two sides with respect to λ must be equal. Assume that h is a small real number such that:

$$f((\lambda + h)\, x) \simeq f(\lambda x) + \sum_{i=1}^{p} h x_i \frac{\partial f}{\partial x_i}(\lambda x)$$

We neglect the second-order terms whose order of magnitude is h^2 because they are negligible in the following limit.

$$\frac{\partial f}{\partial \lambda}(\lambda x) = \lim_{h \to 0} \frac{f((\lambda + h)\, x) - f(\lambda x)}{h} \simeq \sum_{i=1}^{p} x_i \frac{\partial f}{\partial x_i}(\lambda x) = \sum_{i=1}^{p} x_i \lambda^{\alpha-1} \frac{\partial f}{\partial x_i}(x)$$

The last equality is obtained because $\frac{\partial f}{\partial x_i}$ is homogeneous of order $\alpha - 1$ by proposition 122.

The derivative of the right-hand side of equation (2.19) writes:

$$\alpha \lambda^{\alpha-1} f(x)$$

As a consequence we obtain:

$$\sum_{i=1}^{p} x_i \lambda^{\alpha-1} \frac{\partial f}{\partial x_i}(x) = \alpha \lambda^{\alpha-1} f(x)$$

Simplifying by $\lambda^{\alpha-1}$ leads to the result:

$$\sum_{i=1}^{p} x_i \frac{\partial f}{\partial x_i}(x) = \alpha f(x)$$

Chapter 3

Optimization without constraints

When facing problems of the real life, investors, and more generally economic agents, try to do their best. This means that they try to take the best decision, taking into account the information they have.

In most cases, the real world is too complex to be entirely embedded in the formulation of the optimization problem. Models of decision making use simplified representations of the real world. In these simplified frameworks, taking a decision often means maximizing or minimizing a function depending on several variables.

In microeconomics, all the theory is based on the assumption that agents maximize their expected utility. The utility functions are assumed concave and, of course, non-linear, because of the decreasing marginal utility.

In Markowitz portfolio theory[1], investors minimize the risk (measured by the variance of returns) of their portfolio, being given a threshold of expected return they want to reach. Equivalently, the problem can be solved by maximizing the expected return, being given a level of risk the investor accepts to bear.

[1] Markowitz, H.(1952), Portfolio Selection, *Journal of Finance, 7(1), 77-91.*

In corporate finance, firms try to maximize their profits but have to take into account the inverse relationship between the prices of the products they sell and the demand for these products. The firms also try to minimize their costs which are decomposed between fixed and variable costs. In general, decisions that decrease fixed costs have a tendency to increase variable costs. Solving this kind of problem is a matter of optimization.

All these examples show that economic life is paved with the resolution of optimization problems. These problems may include constraints on the possible values of decision variables.

This chapter is devoted to the methods adapted to the resolution of non-linear optimization problems. We assume that no constraint on the decision variables makes the problem more complex to solve. The following chapter will be devoted to these constrained optimization problems.

For the sake of simplicity, we start by single-variable optimization. In principle, the reader already knows these preliminary results. They are intuitive if derivatives of functions had been well understood.

Optimizing functions of several variables is a little bit more difficult because optimality conditions are related to partial derivatives and to the Hessian matrix. Here too, these optimality conditions are natural if partial derivatives and Hessian matrix are understood.

3.1 Preliminaries

3.1.1 The domain of optimization

In general, the functions f to be optimized are defined on a domain $D \subset \mathbb{R}^n$, and take their values in $\mathbb{R}$. An optimization problem can then be written in one of the two following ways:

$$\boxed{\max_{x \in D} f(x) \text{ or } \min_{x \in D} f(x)}$$

The first formulation means that we look for $x^* \in D$ such that $f(x^*)$ is the maximum value taken by f on the domain D; in the second formulation we look for $x^* \in D$ such that $f(x^*)$ is the minimum value taken by f on D.

As already mentioned, D can be equal to $\mathbb{R}^n$, but in most cases D is a subset of $\mathbb{R}^n$, either because f is not defined on all $\mathbb{R}^n$ or because of the characteristics of the problem. For example, when minimizing the risk of a portfolio, assuming that shortsales are forbidden imply restrictions on the domain D.

Let the function f be defined by:

$$f(x) = \sqrt{x_1 x_2}$$

D is written as follows:

$$D = \left\{x \in \mathbb{R}^2 \ / \ x_1 x_2 \geq 0\right\}$$

because square roots are defined only for positive numbers.

Optimization criteria also depend on the shape of the domain D. More precisely, the fact that D is open or closed has an influence on the existence of a solution to the problem at hand.

To illustrate this remark, consider the following problem:

$$\min_{x \in D} f(x) = x^2 - 4 \tag{3.1}$$

$$D = [-3; +2] \tag{3.2}$$

This function has a minimum value equal to -4 for $x = 0$, that is $f(0) = -4 \leq f(x)$ for any $x \in D$. Figure 3.1 shows the graph of the function; you can observe that the first derivative of f at $x = 0$ is equal to 0. In fact, f' is given by:

$$f'(x) = 2x \tag{3.3}$$

Moreover, the sign of the derivative changes at $x^* = 0$ but f" is always positive (f"$(x) = 2$). The function f is then convex and "easier" to minimize.

Suppose now that you are looking for a maximum. Figure 3.1 shows that the maximum is reached for $x^{**} = -3$ with $f(x^{**}) = 5$. However, the criterion based on the value of the derivative cannot be applied because x^{**} is on the frontier of D. In this kind of situation, the solution is called a **corner solution**. The difficulty is that D is closed. On the contrary, if $D =]-3; 2[$ (D is open), f has no maximum but the minimum stays unchanged at $x^* = 0$.

These preliminary remarks show that solving an optimization problem

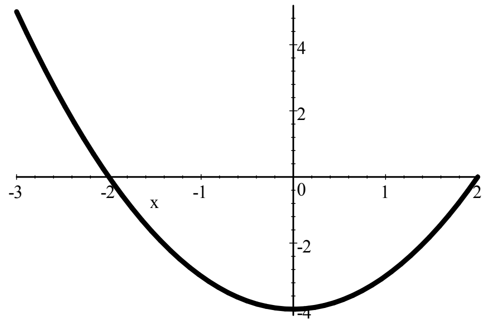

Figure 3.1: The function $f(x) = x^2 - 4$

using the successive derivatives works well if the domain is open. The solution can be more complicated to find when the domain is closed.

3.1.2 Regularity of the function to be optimized

The second feature playing a role in optimization programs is to know if the function to be optimized is sufficiently regular. Of course, if optimization criteria are based on derivatives the least we can ask is that these derivatives exist.

For example, consider the function defined on $\mathbb{R}$ by $f(x) = |x|$; f reaches its minimum for $x = 0$, but f is not differentiable at 0 (see figure 3.2). In fact, the derivative of f is nowhere equal to 0. The kind of irregularity observed at 0 is not that "wild", because f possesses at that point a right-derivative and a left-derivative. Nevertheless, no simple criterion can be found to solve the problem.

The example of $f(x) = |x|$ allows to understand why standard optimization conditions presented in the next sections require that the functions to be optimized are sufficiently regular. In financial applications, this regularity assumption is not too constraining or, more precisely, seems reasonable in a

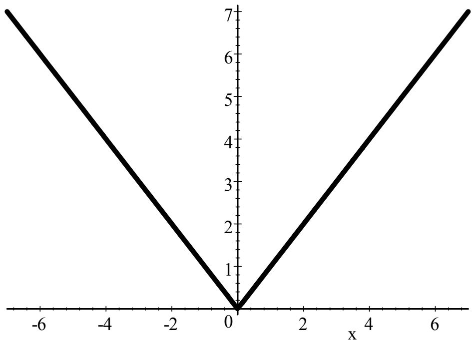

Figure 3.2: The function $f(x) = |x|$

number of circumstances.

3.1.3 Local and global optimum

Consider the function $f(x) = x\sin(x)$ depicted on figure 3.3; the graph is limited to the domain $D = [-7; +7]$. This function does not often show up in financial models but it is nevertheless interesting to understand the distinction between different types of optima.

First, f is regular and possesses all derivatives you might need. Second, it is clear that there are several points where the first derivative is equal to 0.

However, we immediately observe that the natural criterion of a null first derivative is not sufficient to distinguish maxima and minima. Looking at second-order derivatives allows to distinguish the two but only locally, that is on a short interval around the optimum under consideration.

The fact that derivatives provide local conditions to optimize functions is a serious problem in practical issues because all methods based on derivatives provide at best a local optimum (they are called *gradient methods*).

Finally, in this preliminary analysis, we have to mention that even the two

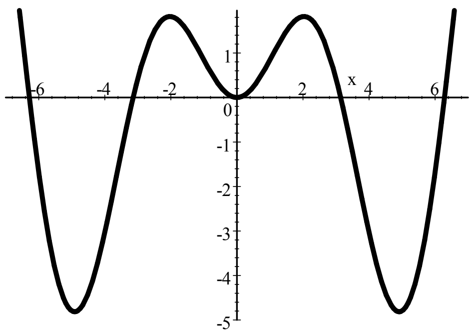

Figure 3.3: The function $f(x) = x\sin(x)$ on the domain $[-7; +7]$

first derivatives are often not sufficient to identify a maximum or a minimum.

Look at figure 3.4 that represents the function $f(x) = x^3$ over the domain $D = [-2; +2]$. The first two derivatives of f are equal to 0 at $x^* = 0$. However, the function has neither a minimum nor a maximum at x^*. One more time, a first derivative equal to 0 does not guarantee the existence of an optimum, without assuming something else.

Consider now the function $f(x) = x^4$ (figure 3.5); the function reaches its minimum at $x^* = 0$ with the two first derivatives equal to 0 at x^*.

In short, all these examples show that life may be complicated when it comes to optimizing. It is the reason why the different cases are examined in some details in the following sections.

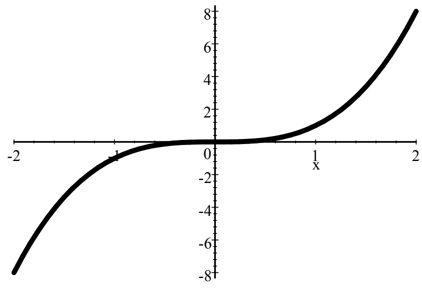

Figure 3.4: The function $f(x) = x^3$

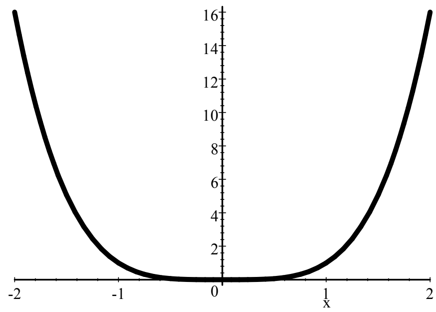

Figure 3.5: The function $f(x) = x^4$

3.2 Optimizing a single-variable function

We start by the most simple case: a function f depending on a single variable x. We are going to characterize minima and maxima of f, defined on a domain $D \subset \mathbb{R}$ and taking values in $\mathbb{R}$.

$$x \in D \to f(x) \in \mathbb{R} \tag{3.4}$$

The examples of the preceding section show that restrictions are necessary, either on D or on f, to obtain tractable optimality conditions. The first restriction, valid for the remainder of the section, is the following.

Assumption: D is an open subset of $\mathbb{R}$ and the functions considered in this section are twice continuously differentiable.

To avoid going back to part I of the book, we recall hereafter the definition of global and local optima.

Definition 124 *a) x_0 is a **local maximum** (minimum) of f if:*

$$\forall x \in]x_0 - \varepsilon; x_0 + \varepsilon, [\; f(x_0) \geq (\leq) f(x)$$

*b) x_0 is a **global maximum** (minimum) of f if there exists $\varepsilon > 0$ such that:*

$$\forall x \in D, f(x_0) \geq (\leq) f(x)$$

3.2.1 Necessary conditions of optimality

Proposition 125 *If x_0 is a local optimum of f then $f'(x_0) = 0$*

Keep in mind that this condition is necessary, not sufficient. You need to know that x_0 is an optimum to say that the first-derivative is equal to 0 at x_0. To emphazise the intuition that drives the result, consider the case of a local minimum. In a narrow interval including x_0, the function f is decreasing (increasing) on the left (right) of x_0 (otherwise x_0 would not be a minimum). Therefore, f' is negative on the left of x_0 and positive on the right. But we assumed that f is at least twice continuously differentiable; it means in particular that f' is continuous. A continuous function being negative (positive) on the left(right) of x_0 is equal to 0 at x_0.

Of course, a necessary condition is not very useful to solve empirical problems because in such problems we are looking for x_0; in most cases we do not look for properties of f at x_0 when we know that x_0 is an optimum.

It is the reason why sufficient conditions are more popular.

3.2.2 Sufficient conditions of local optimality

Proposition 126 x_0 *is a local maximum (minimum) of* f *if:*

a) $f'(x_0) = 0$

b) $f"(x_0) < (>)0$

This proposition comes from the Taylor series expansions presented in part I of the book for single-variable functions and in the preceding chapter for functions depending on several variables. In fact, we can write:

$$f(x_0 + h) = f(x_0) + hf'(x_0) + \frac{h^2}{2} f"(x_0) + \varepsilon(h^2)$$

If $f'(x_0) = 0$ and $\varepsilon(h^2)$ is negligible with respect to h^2, the difference $f(x_0+h) - f(x_0)$ has the sign of $f"(x_0)$. If $f"(x_0) < 0$, then $f(x_0) > f(x_0+h)$ meaning that x_0 is a local maximum of f. Figure 3.1 is an illustration of the proposition. The minimum is obtained at $x_0 = 2$ and the derivative is increasing over an interval including x_0. Therefore, the derivative of f' is positive but this derivative is $f"$.

Remark 127 *Proposition 126 gives a sufficient condition but this condition is not necessary. The function* $f(x) = x^4$ *represented on figure 3.5 provides a good counter-example. In fact, there is a minimum at 0 but the two first derivatives are equal to 0. In general, for polynomials like* x^n, *a minimum exists if* n *is even and an inflection point appears for* n *odd. This remark justifies the general result hereafter.*

3.2.3 Necessary and sufficient optimality conditions

Proposition 128 x_0 *is a local maximum (minimum) of* f *if and only if*[2]*:*

a) $f'(x_0) = 0$

b) The order of the first non zero derivative at x_0 *is even and the corresponding derivative is negative (positive).*

[2] "if and only if" is often shortened in "iff".

The optimality conditions presented so far are local optimality conditions. To obtain global conditions, we need to impose more assumptions on the behavior of f. The idea is that f should not be "authorized" to behave like $x\sin(x)$ with multiple changes in the sign of derivatives.

3.2.4 Global optimality conditions

Proposition 129 *If f is concave (convex) on the open convex domain D, x_0 is a global maximum (minimum) of f if $f'(x_0) = 0$.*

This result provides a very simple optimality condition only depending on the first derivative of f. Of course, the simplicity of the result comes from the concavity/convexity assumption which determines the sign of the second-order derivative. Knowing that many functions in finance or microeconomics problems satisfy this concavity/convexity assumption[3] is important. In such problems, checking if the first-order derivative is equal to 0 is enough to characterize a global optimum of f.

[3] The optimisation problems to solve are called "concave problems" in this case.

Corollary 130 *If f is strictly concave (convex), the first-order condition provides the unique optimum.*

All the propositions of this section refer to single-variable functions. However, the geometric approach underlying the results is general. If x_0 is an optimum, the first-order condition says that the tangent to the curve representing f at x_0 is horizontal (its slope is 0). In the same spirit, the second-order condition is justified by the second-order Taylor series expansion which determines the sign of the changes in f around x_0. In the next section devoted to the optimization of functions depending on several variables, the tools are different, maybe a little bit more complex, but the logic and the geometry of the problem remain the same.

There is no difficulty to address a 30-variable problem when you know how to deal with a 29-variable problem. The "difficult" step is from single-variable functions to two-variable functions. It is the reason why we introduce an intermediate section devoted to the optimization of functions of two variables.

3.3 Optimizing a function of two variables

Devoting some place to optimizing functions of two variables is justified by the fact that such functions are represented by surfaces in three-dimensional spaces. It is then possible to draw their graphs, even if we are limited to two dimensions on the paper. With more than two variables, no graph can be drawn (as far as I know!) by standard means. Figure 3.6 is an example of the graph of a two-variable function.

The function f is defined by :

$$f(x_1, x_2) = \exp\left(-x_1^2 - x_2^2\right)$$

f has a maximum at $(0, 0)$ where it is worth 1, because $\exp(0) = 1$.

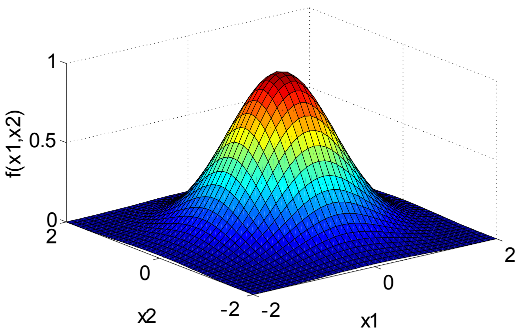

Figure 3.6: The function $f(x_1, x_2) = \exp(-x_1^2 - x_2^2)$

Suppose now that the second variable x_2 is constant, equal to 1. $f(x_1, x_2)$ becomes the single-variable function $g(x_1)$ defined by:

$$g(x_1) = f(x_1, 1) = \exp(-x_1^2 - 1)$$

g is represented on figure 3.7.

$$\exp(-x_1^2 - 1)$$

g reaches a maximum at $x_1^* = 0$ and its derivative equals 0 at x_1^*. In the same spirit we can define $h(x_2)$ by keeping x_1 constant. In such a case, h also has a maximum at $x_2 = 0$ with a null derivative at that point. But remember that keeping one variable constant is exactly what we did in the first part of the book to define partial derivatives of $f(x_1, x_2)$.

These remarks mean that partial derivatives are important in characterizing optima in multidimensional problems. The definitions of g and h and the properties of their derivatives show that two directions (along the x-axis and along the x_2-axis) should be considered when dealing with f. A maximum (x_1^*, x_2^*) of f should be a maximum for g when the value of x_2 is fixed to x_2^*

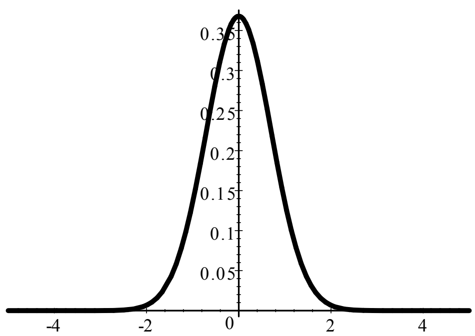

Figure 3.7: La fonction $f(x_1, 1) = \exp(-x_1^2)$

and a maximum for h when the value of x_1 is fixed to x_1^*. The geometric interpretation of this intuition is that the two-dimensional space tangent to the graph of f should be horizontal, that is parallel to the (x_1, x_2) plane. In fact, if it is not the case, we could find directions toward which f increases, a contradiction if (x_1^*, x_2^*) is a maximum.

The following subsections formalize the intuitions we just described. As in the preceding section we assume the following.

Assumption: D is an open subset of $\mathbb{R}^2$ and the functions considered in this section are twice continuously differentiable.

Continuous second-order partial derivatives ensure that the Hessian matrix is symmetric. For applications in finance and economics, it is not a restrictive assumption.

3.3.1 Local optimality conditions

Proposition 131 *If $x^* = (x_1^*, x_2^*)$ is a local optimum of f, then:*

$$\frac{\partial f}{\partial x_1}(x^*) = \frac{\partial f}{\partial x_2}(x^*) = 0$$

This proposition formalizes the intuition we just described by means of functions g and h. In the neighborhood of a maximum x^*, the values of f are lower than $f(x^*)$, especially toward the directions of x_1 and x_2 (that is if f is replaced by g or h). The conditions on partial derivatives say nothing else. These conditions can be shortened by writing $\nabla f(x^*) = 0$ where $\nabla f(x^*)$ is the gradient of f at x^*, (remember that the gradient is the vector of partial derivatives).

Of course, the gradient condition cannot be sufficient, simply because it does not allow to distinguish minima and maxima. Moreover, we already showed for single-variable functions that inflection points can exist. For functions depending on two variables, other more tricky situations can appear. Consider the function f defined by :

$$f(x) = x_1^2 - x_2^2$$

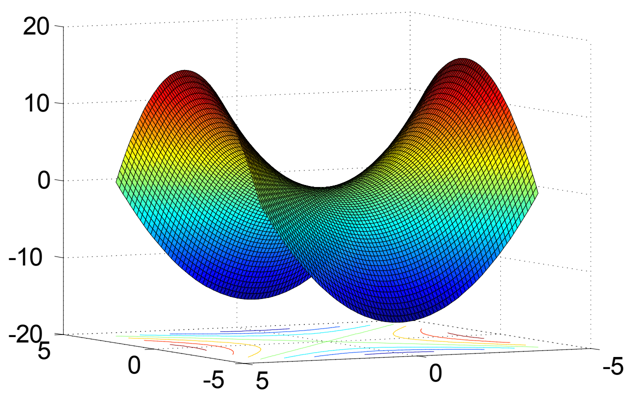

Figure 3.8: Example of a saddle point

The two first-order partial derivatives are equal to 0 at $x = (0; 0)$. In fact, these derivatives are equal to:

$$\frac{\partial f}{\partial x_1} = 2x_1 \text{ and } \frac{\partial f}{\partial x_2} = -2x_2$$

However, $x = (0; 0)$ is neither a maximum nor a minimum. The problem comes from the fact that, on one side, g (as a function of x_1 only) is convex and has a minimum at 0, but, on the other side, h (as a function of x_2 only) is concave and has a maximum at 0. This kind of situation is called a saddle point because, as you can see on figure 3.8, the graph of f in the neighborhood of $(0, 0)$ looks like a horse saddle.

This example shows that obtaining sufficient optimality conditions is going to require some precautions, even for local optima. In part I, we showed that there are n^2 second-order partial derivatives for a function depending on n variables. Therefore, we have 4 elements in the Hessian matrix for our functions depending on two variables. Eventually, the properties of the Hessian matrix are driving the sufficient conditions of optimality. They also allow to distinguish between optima and saddle points.

Proposition 132 *x^* is a local maximum (minimum) of f if the following conditions are satisfied:*
1) $\nabla f(x^*) = 0$
2) $H_f(x^*)$ is negative (positive) definite

As for single-variable functions, the proof of this proposition is based on Taylor's formula. Let us denote $h' = (h_1, h_2)$; we can approximate $f(x^* + h)$ as follows:

$$f(x^* + h) = f(x^*) + h'\nabla f(x^*) + \frac{1}{2}h'H_f(x^*)h + \varepsilon(\|h\|^2)$$

$f(x^* + h) - f(x^*)$ and $h'H_f(x^*)h$ have the same sign when condition (1) is satisfied; if $H_f(x^*)$ is negative definite, $h'H_f(x^*)h < 0$, and then $f(x^*) > f(x^* + h)$. Symetrically, if $H_f(x^*)$ is positive definite, x^* is a local minimum.

Positive and negative definite matrices have been characterized in chapter 1. Using this characterization, proposition 132 can be rewritten as follows.

Corollary 133 *x^* is a local maximum (minimum) of f if:*
1) $\nabla f(x^*) = 0$
2) $\frac{\partial^2 f}{\partial x_1^2}(x^*) < (>)0$ and $Det(H_f(x^*)) > 0$

In fact, for a matrix to be negative definite, the signs of its principal minors must alternate, the first one being negative. For a matrix to be positive definite, all principal minors must be positive.

Looking more closely to the corollary can give the false idea that variable 1 is more important than variable 2. Of course, it is not the case because the determinant of $H_f(x^*)$ writes

$$Det(H_f(x^*)) = \frac{\partial^2 f}{\partial x_1^2}(x^*)\frac{\partial^2 f}{\partial x_2^2}(x^*) - \left[\frac{\partial^2 f}{\partial x_1 \partial x_2}(x^*)\right]^2$$

If this determinant is positive, the two second-order partial derivatives $\frac{\partial^2 f}{\partial x_1^2}(x^*)$ and $\frac{\partial^2 f}{\partial x_2^2}(x^*)$ have the same sign because the product of the two is positive.

But the formulation of the corollary says nothing about what happens when the determinant is not strictly positive.

Considering the example presented at the beginning of the section ($f(x_1, x_2) = x_1^2 - x_2^2$) leads to:

$$H_f(x) = \begin{bmatrix} 2 & 0 \\ 0 & -2 \end{bmatrix}$$

and $Det\,(H_f(x)) = -4$. The signs of the principal minors alternate but the first one is positive and the determinant is negative. These features characterize a saddle point.

3.3.2 Global optimality conditions

The reasoning is exactly the same as the one we used for single-variable functions. To obtain global optimality conditions, we need to impose some restrictions (convexity or concavity) on the behavior of f.

We then obtain the following proposition.

Proposition 134 *If f, defined on a convex $D \subset \mathbb{R}^2$, is concave (convex), x^* is a global maximum(minimum) if $\nabla f(x^*) = 0$.*

This proposition is a direct generalization of proposition 129. The global optimum is obtained by means of a first-order condition because second-order conditions are automatically satisfied when f is concave (for a maximum) or convex (for a minimum). Corollary 130 can be rewritten for functions of two variables without changing a single word.

Corollary 135 *If f is strictly concave (convex), the first-order condition provides the unique maximum (minimum).*

Figure 3.9 represents the function $f(x) = x_1^2 + x_2^2$ which reaches its minimum at $x^* = (0, 0)$. The principal minors of $H_f(x)$ are positive because the Hessian matrix is diagonal, each term of the diagonal being equal to 2.

We observe on figure 3.9 that the null gradient at $(0, 0)$ is equivalent to a horizontal tangent plane at that point.

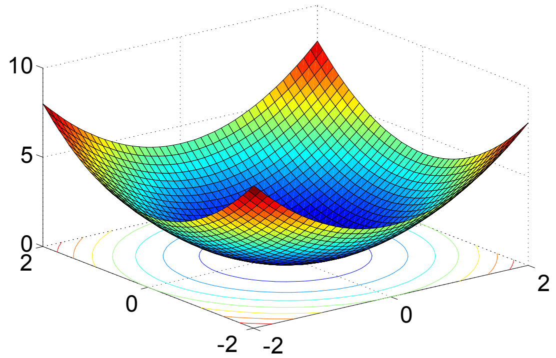

Figure 3.9: Horizontal tangent plane at the minimum of $f(x_1, x_2) = x_1^2 + x_2^2$

Many problems in finance assume the convexity (concavity) of the function to be minimized (maximized). Therefore, solutions often come by means of first-order conditions only, even if the concavity (convexity) of the function is not recalled systematically. For example, it is not always recalled that utility functions are assumed concave because it is a standard assumption in 99.9% of the models.

3.4 Functions of n variables

The general case of functions depending on n variables is not very different from the case $n = 2$ addressed in the preceding section....except that we cannot visualize the functions. The consequence of this proximity is that the statements of this section, especially the optimality conditions, are almost the same as the ones developed in the preceding section.

All the functions considered hereafter are defined on an open set $D \subset \mathbb{R}^n$. As usual, they are supposed twice continuously differentiable.

3.4.1 Local optimality conditions

Proposition 136 *If x^* is a local optimum of f, then the gradient of f at $x^*$is 0.*

Remember that, in problems with two variables, this condition means that the tangent plane at x^* is horizontal, that is parallel to the plane x_1Ox_2. In $\mathbb{R}^n$ surfaces are called hypersurfaces and planes are hyperplanes. The meaning of "horizontal" is not intuitive in higher dimensions...but the idea is still the same. If f is reduced to a single-variable function by fixing, say, the values of the last $n-1$ variables, the partial derivative of f with respect to the first variable must be 0 if x^* is an optimum. If it was not the case, we could find a direction toward which the function f increases (for a maximum) or decreases (for a minimum). It would be a clear contradiction.

Proposition 137 *x^* is a local maximum (minimum) of f if the gradient of f is zero at x^* and the Hessian matrix is negative(positive) definite at x^*.*

Corollary 138 *1) x^* is a local maximum of f if its gradient is 0 at x^* and if the signs of the principal minors of $H_f(x^*)$ alternate, the first one being negative.*

2) x^* is a local minimum of f if its gradient is 0 at x^* and if the principal minors of $H_f(x^*)$ are positive.

We let the reader check that corollary 133 is a special case of the above corollary.

Example 139 *In this example, we are going to show how to build a term structure of (continuous) interest rates in a very simple case. We assume that three bonds are traded with respective maturities 1, 2 and 4 years. Table 3.1 summarizes the data.*

We assume a simple term structure of the following form:

$$r_t = a + bt^{0.8}$$

Bond	Maturity	Coupon rate	Price(in %)
XXX	1	6%	101
YYY	2	5%	99.5
ZZZ	4	5.5%	100.5

Table 3.1: Bonds description

where t denotes the horizon under consideration, a and b are parameters to be estimated by minimizing the sum of the squares of the differences between observed prices and estimated prices. Note that if only the first two bonds are considered, a and b can be estimated without errors on prices. In fact, we should solve:

$$\begin{aligned} 101 &= 106 \exp(-a-b) \\ 99.5 &= 5\exp(-a-b) + 105\exp(-a-b2^{0.8}) \end{aligned}$$

The first equality is justified because the first bond pays a unique cash-flow of 106 in 1 year (the coupon rate is 6%). The second equality states the equality of the price and of the sum of the discounted cash-flows for the second bond. Solving these simple equations leads to:

$$\begin{aligned} a &= -2.5287 \times 10^{-2} \\ b &= .0736 \end{aligned}$$

Applying this estimation to the last bond gives a theoretical price of:

$$\sum_{t=1}^{4} 5.5 * \exp(2.5287 \times 10^{-2} - .0736 t^{0.8}) + 105.5 \exp(2.5287 \times 10^{-2} - .0736 \times 4^{0.8}) = $$

But the market price is 100.5. There is a large difference, meaning that a perfect match of the three prices is impossible.

a and b must be estimated by minimizing the following function:

$$f(a,b) = \sum_{i=1}^{3} (\pi_i - \widehat{\pi}_i)^2$$

where π_i is the market price and $\widehat{\pi}_i$ is the theoretical price.

Of course, in practice the problem is not solved manually but using a computer program or, at least, a spreadsheet. For example, the Excel Solver can easily solve this problem.

3.4.2 Global optimality conditions

We can repeat word by word what we said for functions of two variables; we just need to adapt the dimensions. To obtain global optimality conditions, we impose convexity or concavity of f.

We then obtain the following proposition.

Proposition 140 *If f, defined on a convex subset $D \subset \mathbb{R}^n$, is concave (convex), x^* is a global maximum(minimum) if* $\nabla f(x^*) = 0$.

The global optimum is obtained by means of a first-order condition because second-order conditions are automatically satisfied (f concave for a maximum and f convex for a minimum).

Corollary 141 *If f is strictly concave (convex), the first-order condition provides the unique optimum.*

This corollary is exactly the same as corollary 135. The reader understands now why it was useful to devote some place to functions depending on two variables.

Chapter 4

Constrained optimization

Most economic problems consist in finding how to optimally allocate scarce resources. This sentence describes what will be done in this chapter, that is optimizing a function when the decision variables are constrained within limits.

Chapter 3 offered only a few examples because practical problems of finance are constrained optimization problems. However, the "trick" to solve a constrained problem is to transform it in an unconstrained problem having the same solutions, then justifying chapter 3. The price to pay for this transformation is an increase in the number of decision variables.

Of course, to be interesting to study, a constrainted problem should depend at least on two decision variables[1].

Section 4.1 deals with the optimization of functions depending on two variables with one equality constraint. We introduce the Lagrangian in this simple framework. The Lagrangian is the essential tool to solve constrained problems.

Section 4.2 generalizes results of section 4.1 to problems with p variables

[1]With only one variable, two situations are possible: either the solution is in the interior of the domain limited by the constraint or it is on the frontier. In the first case, methods of chapter 3 are still valid, and in the second case, it is enough to compare the values of the function on the frontier to find the optimal one.

and m equality constraints. The final section deals with the most general problem with equality and inequality constraints.

4.1 Functions of two variables and equality constraint

In this section, we focus on the most simple constrained optimization problem (two variables, one constraint). The results are easy to interpret, and their generalization is natural afterwards. This presentation avoids losing the reader into unimportant calculation details.

4.1.1 Problem statement

The objective function f is defined on an open subset $D \subset \mathbb{R}^2$ and is twice continuously differentiable. The constraint is written by means of a function g, defined on D and also twice continuously differentiable. We develop hereafter the case of a maximization problem, but the reasoning is similar for a minimization. The two cases (maximum and minimum) will be separated when necessary.

The optimization problem, denoted $\mathcal{P}$, writes:

$$\begin{aligned} &\max_{(x_1,x_2)\in D} f(x_1,x_2) \\ &\text{u.c. } \; g(x_1,x_2) = c \end{aligned} \qquad (\mathcal{P})$$

where $c \in \mathbb{R}$ is given[2].

For example, if g is a budget constraint in a utility maximization problem, $g(x_1, x_2) = c$ means $p_1x_1 + p_2x_2 = R$ where $c = R$ is the wealth of the consumer. Such a linear constraint induces an explicit relationship between the two decision variables, that is $x_2 = (R - p_1x_1)/p_2$. In such a simple case,

[2] u.c is a shortcut for "under the constraints"

the constraint writes $x_2 = h(x_1)$ where h is a one-variable function. When this kind of transformation is possible, we are back to the single-variable (unconstrained) problem written as:

$$\max_{x_1} f(x_1, h(x_1))$$

In general, this transformation cannot be used. This is the reason why the Lagrangian has been introduced to solve optimization problems. It is another way to transform a constrained problem into an unconstrained problem (without changing the optimal values of the decision variables).

To illustrate what is going on, denote $x_2^* = \phi(x_1^*)$ where (x_1^*, x_2^*) is a local optimum of f, and x_2^* is an implicit function de x_1^* (see chapter 2), by means of the constraint $g(x_1, x_2) = c$.

The derivation of compound functions can be used to calculate the derivative of $F(x_1) = f(x_1, \phi(x_1))$ at $x_1 = x_1^*$ (chapter 2 of part I). We then write:

$$F'(x_1^*) = \frac{\partial f}{\partial x_1}(x_1^*, \phi(x_1^*)) + \phi'(x_1^*)\frac{\partial f}{\partial x_2}(x_1^*, \phi(x_1^*)) \tag{4.1}$$

The implicit function theorem allows to deduce:

$$\phi'(x_1^*) = -\frac{\frac{\partial g}{\partial x_1}(x^*)}{\frac{\partial g}{\partial x_2}(x^*)} \tag{4.2}$$

At the optimum we know that $F'(x_1^*) = 0$. Equations (4.1) and (4.2) lead to:

$$\frac{\frac{\partial g}{\partial x_1}(x^*)}{\frac{\partial g}{\partial x_2}(x^*)} = \frac{\frac{\partial f}{\partial x_1}(x^*)}{\frac{\partial f}{\partial x_2}(x^*)} \tag{4.3}$$

If λ is defined by:

$$\lambda = \frac{\frac{\partial f}{\partial x_1}(x^*)}{\frac{\partial g}{\partial x_1}(x^*)}$$

we obtain:

$$\frac{\partial f}{\partial x_1}(x^*) - \lambda \frac{\partial g}{\partial x_1}(x^*) = 0 \tag{4.4}$$

$$\frac{\partial f}{\partial x_2}(x^*) - \lambda \frac{\partial g}{\partial x_2}(x^*) = 0 \tag{4.5}$$

Equation (4.5) comes from relation (4.3). More generally, equations (4.4) and (4.5) are useful to specify the intuition behind the definition of the Lagrangian.

4.1.2 Lagrangian and optimality conditions

Definition 142 *The* ***Lagrangian*** *of problem* $\mathcal{P}$ *is the function* $\mathcal{L}(\lambda, x)$ *defined by:*

$$\mathcal{L}(\lambda, x) = f(x) + \lambda (c - g(x))$$

λ *is the* ***Lagrange multiplier*** *of the constraint* $g(x) = c$*.*

Remark 143 *In some books, the right-hand side of the constraint is* 0*. Of course, defining* $g^*(x) = g(x) - c$ *leads to write the Lagrangian as* $f(x) - \lambda g^*(x)$ *and the constraints as* $g^*(x) = 0$*.*

The following proposition shows how problem $\mathcal{P}$ is solved with the methods developed in chapter 3 by optimizing $\mathcal{L}$. It is worth noticing that if f is a function depending on two variables, $\mathcal{L}$ is a function of three variables.

Solving $\mathcal{P}$ is equivalent to optimize the Lagrangian without constraints (denote $\mathcal{P}'$ this problem):

$$\max_{(\lambda,x)} \mathcal{L}(\lambda, x) \tag{$\mathcal{P}'$}$$

If (λ^*, x^*) is a local optimum of $\mathcal{P}'$, proposition 136 of chapter 3 says that the partial derivatives of $\mathcal{L}$ are equal to 0 at (λ^*, x^*), that is:

$$\begin{aligned}
\frac{\partial \mathcal{L}}{\partial x_1}(\lambda^*, x^*) &= \frac{\partial f}{\partial x_1}(x^*) - \lambda^* \frac{\partial g}{\partial x_1}(x^*) = 0 \\
\frac{\partial \mathcal{L}}{\partial x_2}(\lambda^*, x^*) &= \frac{\partial f}{\partial x_2}(x^*) - \lambda^* \frac{\partial g}{\partial x_2}(x^*) = 0 \\
\frac{\partial \mathcal{L}}{\partial \lambda^*}(\lambda^*, x^*) &= c - g(x^*) = 0
\end{aligned}$$

The last equation $c - g(x^*) = 0$ simply means that the constraint is satisfied. A consequence is that the optimal value of $\mathcal{L}$ is also the optimal value of f. This method kills two birds with one stone. It transforms a difficult problem in an easy one by optimizing another function, but the optimal values are the same in the two problems.

Proposition 144 *If x^* is a local maximum of f under the constraint $g(x) = c$ and if the gradient of g is not zero at x^*, there exists λ^* satisfying:*

$$\frac{\partial \mathcal{L}}{\partial x_i}(\lambda^*, x^*) = \frac{\partial f}{\partial x_i}(x^*) - \lambda^* \frac{\partial g}{\partial x_i}(x^*) = 0$$

for $i = 1, 2$.

Proposition 144 is a necessary optimality condition. x^* must be an optimum for the relationship to be satisfied. As in chapter 3, second-order conditions involving the Hessian matrix of $\mathcal{L}$ are required to obtain sufficient optimality conditions.

The condition on the gradient of g (it should not be zero) is satisfied in most finance problems. In fact, the standard finance problem has linear constraints, either a budget constraint to maximize an expected utility or a portfolio constraint in portfolio choice problems. The gradient of g cannot be 0 when the constraint is linear.

Proposition 145 (λ^*, x^*) *is a local maximum (minimum) of $\mathcal{L}$ if the following conditions are satisfied:*

1) $\nabla \mathcal{L}(\lambda^*, x^*) = 0$.

2) The determinant of $H_{\mathcal{L}}(\lambda^*, x^*)$ *is positive (negative).*

Condition (2) deserves some comments, because only the determinant of $H_{\mathcal{L}}(\lambda^*, x^*)$ seems to be involved, and not all the principal minors, contrary to the sufficient optimality conditions in chapter 3. In fact, the formulation of the condition comes from the structure of $H_{\mathcal{L}}$. $\mathcal{L}$ is a linear function of λ, so the first principal minor of $H_{\mathcal{L}}$ is always 0 because it is equal to the second-order derivative of $\mathcal{L}$ with respect to λ.

The second principal minor, denoted M_2, is equal to:

$$M_2 = \begin{vmatrix} 0 & -\frac{\partial g}{\partial x_1}(x^*) \\ -\frac{\partial g}{\partial x_1}(x^*) & \frac{\partial^2 f}{\partial x_1^2}(x^*) - \lambda^* \frac{\partial^2 g}{\partial x_1^2}(x^*) \end{vmatrix} = -\left(\frac{\partial g}{\partial x_1}(x^*) \right)^2 < 0$$

M_2 is always negative, meaning that the effective optimality condition can only concern the sign of the last principal minor, that is the determinant of $H_{\mathcal{L}}(\lambda^*, x^*)$.

Notice that proposition 144 includes a condition on the gradient of g. This condition appears here in part (2) of proposition 145. In fact, if the gradient of g was 0, the first line of $H_{\mathcal{L}}(x^*)$ would be null and $\det(H_{\mathcal{L}}(x^*))$ would also be equal to 0.

Example 146 *Consider the utility maximization problem under a budget constraint (the notations are as usual):*

$$\begin{aligned} \max_{(x_1,x_2)\in\mathbb{R}^*} U(x_1,x_2) &= \sqrt{x_1x_2} \\ u.c.\ p_1x_1+p_2x_2 &= R \end{aligned}$$

The Lagrangian of the problem is:

$$\mathcal{L}(\lambda,x)=U(x_1,x_2)+\lambda(R-p_1x_1-p_2x_2) \tag{4.6}$$

The first-order conditions are the following:

$$\begin{aligned} \frac{\partial U}{\partial x_1}(x^*)-\lambda^*p_1 &= 0 \\ \frac{\partial U}{\partial x_2}(x^*)-\lambda^*p_2 &= 0 \\ R-p_1x_1^*+p_2x_2^* &= 0 \end{aligned}$$

Replacing U by its definition leads to:

$$\begin{aligned} \frac{1}{2}\sqrt{\frac{x_2^*}{x_1^*}}-\lambda^*p_1 &= 0 \\ \frac{1}{2}\sqrt{\frac{x_1^*}{x_2^*}}-\lambda^*p_2 &= 0 \\ R-p_1x_1^*+p_2x_2^* &= 0 \end{aligned}$$

We are back to the standard result of microeconomics. The ratio of marginal utilities is equal to the ratio of prices.

Consider the following parameters, $R=10$; $p_1=3$; $p_2=4$. We obtain *the following conditions:*

$$\begin{aligned} \frac{x_2^*}{x_1^*} &= \frac{3}{4} \\ 3x_1^* + 4x_2^* &= 10 \end{aligned} \tag{4.7}$$

meaning that $x_2^* = \frac{5}{4}$ *and* $x_1^* = \frac{5}{3}$.

First, we observe that R *is equally shared between the two goods because* $3 \times \frac{5}{3} = 4 \times \frac{5}{4} = 5$. *This result is in line with intuition. The utility function is symmetric, so the optimal amounts spent in each good are equal.*

Second, the Lagrange multiplier is equal to:

$$\lambda^* = \frac{1}{2p_1}\sqrt{\frac{x_2^*}{x_1^*}} = \frac{\sqrt{3}}{12} = 0.144$$

and the utility at x^* *is* $\sqrt{\frac{5}{4} \times \frac{5}{3}} = 1.4434$

Imagine now that system (4.7) is solved twice, first with $R = 9.8$ *and second with* $R = 10.2$.

If $R = 9.8$, *we obtain:*

$$R = 9.8\ ;\ x_1^* = \frac{4.9}{3}\ ;\ x_2^* = \frac{9.8}{8}\ ;\ U(x_1^*, x_2^*) = \sqrt{\frac{4.9}{3} \times \frac{9.8}{8}} = 1.4145$$

If $R = 10.2$ *the results are:*

$$R = 10.2\ ;\ x_1^* = \frac{5.1}{3}\ ;\ x_2^* = \frac{10.2}{8}\ ;\ U(x_1^*, x_2^*) = \sqrt{\frac{5.1}{3} \times \frac{10.2}{8}} = 1.4722$$

The objective function decreases by 0.0289 when R decreases by 0.2 units. A linear approximation gives a decrease in utility of 0.1445 for one unit less in the budget constraint. Symetrically, if R increases by 0.2, utility increases by 0.0288, that is an increase of 0.144 for one more unit spent. 0.144 is exactly the value of the Lagrange multiplier. It is the reason why the Lagrange multiplier measures the sensitivity of utility (objective function) with respect *to variations in available wealth (constraint). The other more direct route to come to this interpretation is to verify that the derivative of the Lagrangian with respect to wealth is exactly* λ.

It remains to check that our solution is a maximum. The Hessian matrix of the Lagrangian is:

$$H_{\mathcal{L}}\left(\lambda^*, x^*\right) = \begin{pmatrix} 0 & -1 & -1 \\ -1 & -\frac{1}{4x_1^*}\sqrt{\frac{x_2^*}{x_1^*}} & -\frac{1}{4}\sqrt{\frac{1}{x_1^* x_2^*}} \\ -1 & -\frac{1}{4}\sqrt{\frac{1}{x_1^* x_2^*}} & -\frac{1}{4x_2^*}\sqrt{\frac{x_1^*}{x_2^*}} \end{pmatrix}$$

A few calculations lead to:

$$\begin{aligned} \det(H_{\mathcal{L}}\left(\lambda^*, x^*\right)) &= \frac{1}{4}\left(\frac{\sqrt[4]{x_1^*}}{\sqrt[4]{\left(x_2^*\right)^3}} - \frac{\sqrt[4]{x_2^*}}{\sqrt[4]{\left(x_1^*\right)^3}}\right)^2 \\ &= \frac{1}{4}\frac{\left(x_1^* - x_2^*\right)^2}{\sqrt{\left(x_1^* x_2^*\right)^3}} > 0 \end{aligned}$$

This determinant is positive. x^* *then maximizes* U *under the budget constraint.*

Example 146 is a specific case of the following proposition.

Proposition 147 *If a* C^2-*function* f, *defined on an open convex subset* $D \subset \mathbb{R}^2$, *is concave (convex) and if the constraint* g *is affine on* D, *then any local maximum (minimum) is a global maximum (minimum).*

4.2 Functions of p variables with m equality constraints

We consider now twice continuously differentiable functions $f, g_1, ..., g_m$ defined on an open domain $D \subset \mathbb{R}^p$ and taking their values in $\mathbb{R}$. We also assume $m < p$. The optimization problem addressed in this section is:

$$\max_{x \in D} f(x) \qquad (\mathcal{P})$$
$$\text{u.c.} \quad g_j(x) = c_j, \; j = 1, ..., m$$

Following the approach of the preceding section, the Lagrangian of the problem is:

$$\mathcal{L}(\lambda, x) = f(x) + \sum_{j=1}^{m} \lambda_j \left(c_j - g_j(x)\right)$$

There exists one Lagrange multiplier per constraint; the initial problem with p variables and m constraints has become an unconstrained maximization problem with $p + m$ variables.

4.2.1 Local optimality conditions

Necessary conditions

We follow the same structure as before and start with a necessary optimality condition in the following proposition.

Proposition 148 *Let x^* be a local optimum of f, satisfying the constraints of problem ($\mathcal{P}$) and such that the gradients $\nabla g_j(x^*), j = 1, ..., m$ are linearly independent vectors in $\mathbb{R}^p$.*

There exists $\lambda^* \in \mathbb{R}^m$ such that the gradient of $\mathcal{L}$ is the null vector at x^*, that is:

$$\begin{aligned} \frac{\partial f}{\partial x_i}(x^*) - \sum_{j=1}^{m} \lambda_j^* \frac{\partial g_j}{\partial x_i}(x^*) &= 0 \text{ if } i = 1, ..., p \\ c_j - g_j(x^*) &= 0 \text{ if } j = 1, ..., m \end{aligned}$$

Why should the gradients be linearly independent? This condition is not intuitive at all. Consider the following example with three variables and two constraints defined as follows:

$$\begin{aligned} x_1 + 2x_2 + x_3 &= c_1 \\ 2x_1 + 4x_2 + 2x_3 &= c_2 \end{aligned}$$

The left hand side of the second equality is twice the left hand side of the first one. Therefore, we can face two situations. If $c_2 \neq 2c_1$, the problem has no solution. But if $c_1 = 2c_2$, the two constraints are redundant, one is enough. The gradients are equal to:

$$\nabla g_1(x) = \begin{pmatrix} 1 \\ 2 \\ 1 \end{pmatrix} \text{ et } \nabla g_2(x) = \begin{pmatrix} 2 \\ 4 \\ 2 \end{pmatrix}$$

These two vectors are colinear because $\nabla g_2(x) = 2 \times \nabla g_1(x)$.

In general when the gradients are not linearly independent, at least one constraint can be removed before applying proposition 148.

In problems with constraints, Lagrange multipliers are interpreted as in single-variable problems. Each multiplier measures the sensitivity of the objective function with respect to variations in the right-hand side of the constraints. These multipliers lose their significance when gradients are colinear. The problem has the same nature as the one of multicolinearity in multiple regression. When independent variables are colinear, nothing relevant can be said about the significance of the regression coefficients.

Sufficient optimality conditions

After reading chapter 3, the reader knows that sufficient optimality conditions are based on the Hessian matrix of the Lagrangian. However, this matrix is really special because $\mathcal{L}$ is a linear function of the multipliers λ_j. Therefore the second-order derivatives with respect to the multipliers λ_j are 0. In a problem with m constraints, the (m, m)-dimensional North-West corner of $H_{\mathcal{L}}(x)$ only contains zeros. For example, in a problem with 3 variables and two constraints, $H_{\mathcal{L}}(x)$ is as follows:

$$H_{\mathcal{L}}(x) = \begin{bmatrix} 0 & 0 & -\frac{\partial g_1}{\partial x_1}(x) & \frac{\partial g_1}{\partial x_2}(x) & \frac{\partial g_1}{\partial x_3}(x) \\ 0 & 0 & -\frac{\partial g_2}{\partial x_1}(x) & \frac{\partial g_2}{\partial x_2}(x) & \frac{\partial g_2}{\partial x_3}(x) \\ -\frac{\partial g_1}{\partial x_1}(x) & -\frac{\partial g_2}{\partial x_1}(x) & \frac{\partial^2 \mathcal{L}}{\partial x_1^2}(x) & \frac{\partial^2 \mathcal{L}}{\partial x_1 \partial x_2}(x) & \frac{\partial^2 \mathcal{L}}{\partial x_1 \partial x_3}(x) \\ -\frac{\partial g_1}{\partial x_2}(x) & -\frac{\partial g_2}{\partial x_2}(x) & \frac{\partial^2 \mathcal{L}}{\partial x_1 \partial x_2}(x) & \frac{\partial^2 \mathcal{L}}{\partial x_2^2}(x) & \frac{\partial^2 \mathcal{L}}{\partial x_3 \partial x_2}(x) \\ -\frac{\partial g_1}{\partial x_3}(x) & -\frac{\partial g_2}{\partial x_3}(x) & \frac{\partial^2 \mathcal{L}}{\partial x_1 \partial x_3}(x) & \frac{\partial^2 \mathcal{L}}{\partial x_3 \partial x_2}(x) & \frac{\partial^2 \mathcal{L}}{\partial x_3^2}(x) \end{bmatrix} = \begin{bmatrix} H_1 & H_2' \\ H_2 & H_3 \end{bmatrix}$$

H_1 is a $(2, 2)$ null matrix. H_2 is a $(3, 2)$ matrix containing the derivatives of the constraints with respect to the variables and H_3 is a $(3, 3)$ matrix the elements of which are the second-order derivatives of $\mathcal{L}$ with respect to the three variables.

The structure of $H_{\mathcal{L}}(x)$ implies that the first $2m$ principal minors are not significant in characterizing the optimum. In general, if the problem has p variables and m constraints, only the sign of the $p-m$ last principal minors matter. In fact, $H_{\mathcal{L}}$ is $(p+m,p+m)$-dimensional and, as just described, the $2m$ first principal minors are not significant. The number of significant minors is then $p+m-2m=p-m$.

Proposition 149 *x^* is a local maximum of f if:*

1) The constraints are satisfied at x^*.

2) There exists a vector of multipliers λ^* satisfying $\nabla\mathcal{L}(\lambda^*,x^*)=0$.

3) The signs of the last $p-m$ principal minors of $H_{\mathcal{L}}(\lambda^*,x^*)$ alternate, the first one being negative if m is even and positive if m is odd.

Part (3) means that, if the condition is satisfied, the Hessian matrix is negative semi-definite.

Proposition 150 *x^* is a local minimum of f if:*

1) The constraints are satisfied at x^*.

2) There exists a vector of multipliers λ^* satisfying $\nabla\mathcal{L}(\lambda^*,x^*)=0$.

3) The last $p-m$ principal minors of $H_{\mathcal{L}}(\lambda^*,x^*)$ have the same sign as $(-1)^m$.

Part (3) means that, if the condition is satisfied, the Hessian matrix is positive semi-definite.

4.2.2 Global optimality conditions

The global optimality conditions are quite close to the conditions proposed for functions depending on two variables. The difference comes from the existence of multiple constraints. It is the reason why we do not comment this proposition. The reasoning used for functions of two variables is still valid here.

Proposition 151 *If f, defined on an open convex set $D \subset \mathbb{R}^p$, is concave (convex), and if the constraints g_j are affine functions on D, any local maximum(minimum) is also global.*

We can now address the general case in which the two types of constraints (equalities and inequalities) coexist.

4.3 Functions of p variables with mixed constraints

4.3.1 The problem

This last section addresses the most general problem where inequality and equality constraints coexist. The functions f, g_j, h_k of problem (4.8) are defined on an open subset D in $\mathbb{R}^p$ and twice continuously differentiable. The optimization problem writes:

$$\begin{aligned} &\max f(x) \\ \text{u.c. } &g_j(x) = c_j, j = 1, ..., m \\ &h_k(x) \leq b_k, k = 1, ..., n \end{aligned} \tag{4.8}$$

Remark 152 *Choosing inequality constraints as « $\leq$ » does not matter because any inequality $h(x) \geq c$ is equivalent to $-h(x) \leq -c$.*

To gain in clarity when stating the optimality conditions we used different notations depending on the type of constraint (h for inequalities and g for equalities).

To emphasize the link with the results of the preceding section, imagine that a solution x^* to problem (4.8) has been found. The set of inequality constraints may be divided in two subsets: the first subset contains the constraints satisfying $h_k(x^*) = b_k$, and the second subset contains the constraints satisfying $h_k(x^*) < b_k$. The following definition specifies the concept of binding constraints.

Definition 153 *A constraint k is **binding** at x^* if $h_k(x^*) = b_k$.*

As mentioned in the introduction of the chapter, the constraints often refer to scarce resources. b_k then denotes the quantity of the available resource and the equality $h_k(x^*) = b_k$ means that all the resource has been consumed at x^*. It explains why the word *binding* is used.

4.3.2 The solution

Finding a solution may be difficult because, at an optimum x^* some constraints may be binding and other ones not binding. We usually interpret Lagrange multipliers as measures of the sensitivity of the objective function to variations of the right-hand side of the constraint. But in this approach, the multiplier of a constraint should be 0 when a constraint is not binding. In fact, consider the standard economic problem of utility maximization under a budget constraint, but assume that the utility function is not strictly increasing[3]. It may happen that a part of the budget is not "consumed" at the optimum x^* because the marginal utility is 0 at x^*. In this situation, one more unit of wealth would not increase utility and the multiplier would be 0.

The Lagrangian of problem (4.8) is:

$$\mathcal{L}(\lambda, \mu, x) = f(x) + \sum_{j=1}^{m} \lambda_j \left(c_j - g_j(x)\right) + \sum_{k=1}^{n} \mu_k \left(b_k - h_k(x)\right)$$

If x^* is an optimum for multipliers $\lambda^* = (\lambda_1^*, ..., \lambda_m^*)$ and $\mu^* = (\mu_1^*, ..., \mu_n^*)$,

[3] In many restaurants, the quantity of soft drinks (or sometimes appetizers) you can consume is unlimited. The reason is simply that the optimal choice of a client is not to drink an unlimited quantity of soda. The utility function for soda cannot be strictly increasing everywhere.

then:

$$\begin{aligned}\mu_k &> 0 \text{ and } b_k - h_k(x) = 0 \text{ if constraint } k \text{ is binding.}\\ \mu_k &= 0 \text{ and } b_k - h_k(x) > 0 \text{ if constraint } k \text{ is not binding.}\end{aligned}$$

In the two cases the product $\mu_k\,(b_k - h_k(x))$ is equal to zero. This remark is used to shorten the formulation of optimality conditions. The coefficients μ_k are called **Kuhn-Tucker multipliers.**

4.3.3 Necessary optimality condition

Proposition 154 *If x^* is a local maximum of f in problem 4.8 and if the gradients of all functions g_j and h_k for which $h_k(x^*) = 0$ are linearly independent, there exist $m + n$ numbers $\lambda_1^*, ..., \lambda_m^*, \mu_1^*, ..., \mu_p^*$ satisfying the three following conditions:*

$$\nabla f(x^*) - \sum_{j=1}^{m} \lambda_j^* \nabla g_j(x^*) - \sum_{k=1}^{n} \mu_k^* \nabla h_k(x^*) = 0$$

$$\forall k \in \{1, 2, ..., n\}\,, \mu_k^*(b_k - h_k(x^*)) = 0$$

$$\forall k \in \{1, 2, ..., p\}\,, \lambda_k^* \geq 0$$

In this proposition, the multipliers μ_k^* are positive or equal to 0. In fact, if a constraint is binding (think to these constraints as limitations for some resources), it means that all the resource is consumed at the optimum x^*. Obtaining one more unit of the resource would improve the optimal value of the objective function. μ_k^* then measures the variation of the objective function that would arise if one more unit of the resource k was made available.

Proposition 154 could be written for a minimization problem by simply changing the sign of coefficients μ_k^*.

4.3.4 Necessary and sufficient global optimality conditions

In the preceding section we obtained a global maximum if D is convex, f concave and the functions g_j affine. When the problem includes inequality constraints,this result is generalized as follows.

Proposition 155 *If f is concave, the functions g_j affine and the functions h_k convex, the conditions of proposition 154 mean that x^* is a global maximum of f under the constraints of problem (4.8).*

If the problem is a minimization problem, replace « f concave » by « f convex » and change the signe of the coefficients μ_k^* (they would be negative).

Index