Foreword

This book is a mark left in the woods. It is a sign left by two travelers who have chosen every day to share a portion of their trip, a fun one. The woods are the tall trees of the concepts, the methods, and not last, the practice of financial markets. Their journey has been long but fast-paced. At one point, they have felt a need to share and leave a mark, to tell others that they had been in these woods and they have tried to sort their way through in a manner that they have enjoyed so much to invite others along the same path. Of course, the two travelers are us, the Authors, and this corner of the woods is about portfolio management. We hope that the sense of speed through a journey and the fun we had while writing together, continuously swapping ideas and mutual cheer ups will spring to life from the pages of our book.

We are both aware that there is a chance that you, the Reader, may be using this textbook to follow one taught course, presumably at the MSc. level. This is the experience from which our joint effort stems as well. The authors crossed paths in such an environment from different sides of the desk, but their paths soon aligned to one, shared direction. We hope that you will feel what our goal has been—to tell the important apart from the unimportant, the useful from the curiosity, the feasible from the convoluted (albeit elegant).

The least youthful (we like to see the glass half-full) of the two authors carries a big debt for what he has learned from the more youthful about the real, everyday value of knowledge, its usefulness in practical situations, and a fresh taste for the simple and immediately applicable. On her turn, the most youthful of the two, has derived true inspiration from the enthusiasm, the passion, and the genuine curiosity that the least youngest still places after so many years in sharing his knowledge with students without forgetting that learning is a never ending process.

The book strives to avoid becoming one more piece in financial mathematics. Although one of the Authors stood at that gate holding an ax to prevent excesses, we cannot rule out that we may have been occasionally carried away. For the Readers who perceive being short of an adequate background, the references are classical, Simon and Blume (1994) and Wainwright and Chiang (2013) in mathematics, Mood, Graybill, Boes (1974) in statistics.[*]

Manuela and Massimo

Additional resources are available online via MyBook:
http://mybook.egeaonline.it

[*] Mood, Graybill, and Boes (1974), *Introduction to the Theory of Statistics*, McGraw Hill; Simon and Blume (1994), *Mathematics for Economists*, Norton & Co.; Wainwright and Chiang (2013), *Fundamental Methods of Mathematical Economics*, McGraw Hill.

List of Symbols and Acronyms
(in order of first appearance in the book)

HPR	Holding period return
$R_{i,t}$	Return on an asset or index i
CAGR	Compounded annual growth rate
MV	Mean-variance
$Prob$(A)	Probability of event A
E[·]	Expectation (in population)
μ_i	Expectation of asset i (in population)
Var[·]	Variance (in population)
σ_i^2	Variance of asset i
Cov[·,]	Covariance (in population)
$\sigma_{i,j}$	Covariance between assets i and j
$\rho_{i,j}$	Correlation coefficient of assets i, j
ω_i	Weight of asset i
N	Number of assets in the asset menu
R^i	Return or payoff on asset/security i
S	Number of states in the discrete case
EUT	Expected utility theorem
VNM	Von-Neumann Morgenstern (felicity function)
W	Wealth
CE	Certainty equivalent
DMU	Decreasing marginal utility
p_i	Price of a good or service
MU	Marginal utility
x_i	Quantity demanded of a good or service
H	Zero-mean bet, gamble
RRA	Relative risk aversion
ARA	Absolute risk aversion function
T	Risk tolerance function
CRA	Relative risk aversion function
Π	Risk premium
CER	Certainty equivalent rate of return
LRT	Linear risk tolerance
HARA	Hyperbolic absolute risk aversion

CARA	Constant absolute risk aversion
CRRA	Constant relative risk aversion
LHS	Left hand side
RHS	Right hand side
GMV(P)	Global minimum variance (portfolio)
FOC	First-order condition
R^f	Return on the risk free asset
CML	Capital market line
DARA	Decreasing absolute risk aversion
IARA	Increasing absolute risk aversion
FOSD	First-order stochastic dominance
$F_Y(\cdot)$	Cumulative distribution function of asset/gamble Y
CDF	Cumulative distribution function
SOSD	Second-order stochastic dominance
GARCH	Generalized autoregressive conditional heteroskedastic
PCA	Principal component analysis
CAPM	Capital asset pricing model
$R_{m,t+1}$	Rate of return on the market portfolio
SSR	Sum of squared residuals
OLS	Ordinary least squares
SD[·]	Standard Deviation
ARMA	Autoregressive Moving Average (model)
IP	Industrial production
APT	Arbitrage pricing theory
SMB	Small minus big
HML	High minus low (book to market)
WML	Winners minus losers
IRRA	Increasing Relative Risk Aversion
IID	Independent and identically distributed
TWR	Time-weighted return
NAV	Net asset value
FMV	Fund/Manager/Vehicle
SML	Security Market Line
TR	Treynor ratio
TAA	Tactical asset allocation

1 Introduction to Portfolio Analysis: Key Notions

"Uncertainty cannot be dismissed so easily in the analysis of optimizing investor behavior. An investor who knew future returns with certainty would invest in only one security, namely the one with the highest future return" (H.M. Markowitz, "Foundations of Portfolio Theory", 1991)

Summary: – 1. Financial Securities. – 2. Choice Under Risky Situations. – 3. Statistical Summaries of Portfolio Returns.

1 - Financial Securities

1.1 Definition of financial securities

Most people own a "portfolio" (i.e., a collection) of assets, such as money, houses, cars, bags, shoes, and any other durable goods that are able to retain value over time and that can be used to transfer (real) wealth and hence consumption opportunities over time. In this book, we will focus primarily on a certain type of assets, i.e., financial securities. As is generally known, financial securities (or financial assets) can be thought of as a legal contract that represents the right to receive future payoffs—usually but not exclusively in the form of monetary cash flows—under certain conditions. For instance, when you buy the stocks of a company, you are acquiring rights on a part of the future profits of that company (generally distributed in the form of *dividends*, either in cash or shares of stock). Of course, when an asset just pays out money, such monetary payoffs can be used to purchase goods and services subject to a standard budget constraint that forces a decision maker to spend only the available resources (currently or in present value terms). Our seemingly vague reference to what we have cited as "certain conditions" is due to the fact that financial

securities are generally *risky*, i.e., they pay out money, goods, or services in different, uncertain *states of the world*, as we shall see extensively throughout the rest of the book.[1]
Financial assets generally serve two main purposes:

I. To redistribute available wealth across different states of the world, to finance consumption and saving;
II. To allocate available wealth intertemporally, i.e., to allow an investor to save current income/wealth to finance future consumption or, on the opposite, to make it possible for her to borrow against her future incomes/wealth to finance current consumption.

For instance, consider an investor who uses a part of her income from wages earned in one productive sector (e.g., banking) to buy stocks issued in another sector (e.g., industrial companies such as in the automotive industry). This individual is thus making sure that her welfare is at least partially disconnected from the fortunes or mis-fortunes of the banking sector to participate in the outlook of the automotive industry. At the same time, this very investor who reduces her consumption stream to save by purchasing automotive stocks is financing her own future consumption, although the exact amount available will depend on the future, realized profitability of the sector. Nowadays there are a large number of different financial securities available, such as stocks (equity), bonds, commodities, derivatives, investment (mutual, pension, hedge) funds, etc. A broad discussion of the specific characteristics of each type of financial securities is beyond the scope of our book and can be found instead in many intermediate finance textbooks (see, e.g., Fabozzi and Markowitz, 2011). However, just to level the playing field in view of the rest of our work, in this section we offer a short review of how returns of financial assets should be computed. In the rest of the chapter we review other basic concepts that are necessary to understand the framework of portfolio analysis. First, we clarify in what sense financial securities are risky and we explain how investors can deal with choice among alternative securities in such uncertain situations. Second, we discuss how returns

[1] Throughout the rest of the book, unless stated otherwise, we shall not distinguish between the concepts of risk and uncertainty. Risk characterizes unknown events for which objective probabilities can be assigned; uncertainty applies to events for which such probabilities cannot be attributed, or for which it would not make sense to assign them because they cannot be replicated in any controlled way, thus rendering the calculation of relative frequencies difficult. One simplistic way to think about this issue is to envision all the uncertainty that we shall deal with as risk.

and risk are generally measured and how these measures can be aggregated when securities are collected to form a portfolio.

1.2 Computing the return of financial securities

Consider first an asset that does not pay any dividends or coupon interest. The simple single-period return R_t of this asset between time $t-1$ and t is defined as

$$R_t = \frac{P_t}{P_{t-1}} - 1, \tag{1.1}$$

where P_t is the price of the asset at time t and P_{t-1} is the price of the asset at time $t-1$. Therefore, an investor that has invested a monetary unit (e.g., one euro) at $t-1$ in this security will end up at time t with $1+R_t$. The return plus one is often referred to as *gross return* or, alternatively, the *holding period return* (*HPR*). If after the first period the investor reinvests her money from time t to $t+1$, at the end of her holding period she will get $(1+R_t)(1+R_{t+1})$, where R_{t+1} is the return between time t and $t+1$. This way of aggregating simple gross returns over time generalizes to any possible time interval: the gross return between time $t-h$ and t is simply given by the geometric-style (in the sense that products are considered) product:

$$(1+R_{t:h}) = (1+R_t)(1+R_{t-1}) \dots (1+R_{t-h+1}) = \prod_{i=0}^{h-1}(1+R_{t-i}). \tag{1.2}$$

Consequently, the net return over h periods, also known as *compounded* return, is simply equal to $\prod_{i=1}^{h-1}(1+R_{t-i}) - 1$. However, it generally makes little sense to discuss about returns without defining their investment horizon. Conventionally, practitioners tend to express their returns on an *annualized* basis, as this enhances comparability. For instance, consider the case in which you have invested your money for three years and you have earned a rate of return R over this period. Based on formula (1.2), the annualized return R_a (sometimes called *compounded annual growth rate* or CAGR) of your investment is simply equal to

$$R_a(3) \equiv \left[\prod_{i=1}^{3}(1+R_i)\right]^{1/3} - 1. \tag{1.3}$$

More generally, $R_a(n) \equiv [\prod_{i=1}^{n}(1+R_i)]^{1/n} - 1$. Clearly, it easier to compute arithmetic means than geometric ones. For this reason, it is also quite common to use *continuously compounded* returns, which are obtained from simple return aggregation in (1.2), when the frequency of compounding is increased towards infinity, i.e., as if we could disinvest and reinvest our accrued wealth at every moment. The continuously compounded return (also known as *log-return*) of an asset is simply defined as $R_t^c \equiv \ln(P_t/P_{t-1})$. The advantage of using continuously compounded returns is that the multi-period return is very easy to compute as it consists of the sum of the log-returns of each period:

$$R_{t:h}^c \equiv R_{t-1}^c + R_{t-2}^c + \cdots + R_{t-h+1}^c = \sum_{i=1}^{h-1} R_{t-i}^c. \tag{1.4}$$

The use of log-returns is widespread not only because they can be easily summed up to obtain multi-period returns, but also because their use simplifies the modelling of statistical properties of return time-series. Unfortunately, continuous compounding has a key drawback: while the return of a portfolio is equal to the weighted average of the simple asset returns, this statement does not hold true for log-returns, as the sum of logs is not equal to the log of the sum. However, when returns are measured on a short horizon (e.g., daily) the difference between the portfolio continuously compounded return and the weighted average of the log-returns of each asset is very small. In the rest of the book, we shall use simple returns when we are not interested in their time-series properties and log-returns in all other cases.

Finally, for assets (generally stocks) that make periodic payments (e.g., dividend) the formula in (1.1) should be slightly modified:

$$R_t \equiv \frac{P_t + D_t}{P_{t-1}} - 1, \tag{1.5}$$

where D_t is the dividend paid at time t and P_t is the *ex-dividend* price of the stock (i.e., the price of the stock immediately after the payment of the dividend). Equivalently, the continuously compounded return of a stock that pays dividends is $R_t^c \equiv \ln[(P_t + D_t)/P_{t-1}]$.

2 - Choices under Risky Situations

2.1 Choices under uncertainty: a general framework

In section 1.2, we have discussed how we can compute the realized returns of financial assets. However, we have noted that most financial securities have a fundamental characteristic: they are *risky*, meaning that their payoff depends on which of the K alternative states of the world will turn out to occur at a future point in time. The states are uncertain because they are not known in advance, when investors make their investment decisions (i.e., whether to buy, not to buy, and—when feasible—sell the securities that belong to the *asset menu* they face). However, at least under some conditions, we shall assume that investors are able to quantify such uncertainty on future states using standard *probability distributions* and the entire apparatus that classical probability theory provides them with (a brief discussion of the properties of the distribution of returns is provided in the next section). We also assume that exactly one state will occur, though investors do not know, at the outset, which one, because the states are mutually exclusive. The description of each state is complete and exhaustive, in the sense that all the relevant information is provided to an investor to tackle the decision problem being studied.
In spite of this rather rich structure imposed on the choice problem, the task that awaits us (or our investor) is a complex one and the optimal choice will result from three distinct sets of (interacting) factors:

I. how an investor's attitude toward or tolerance for risk is to be conceptualized and therefore measured;
II. how risks should be defined and measured;
III. how investors' risk attitudes interact with the subjective uncertainties associated with the available assets to determine an investor's desired portfolio holdings (demands).

First, we shall consider how the investors' beliefs about future states may be expressed. In the following example, we show how standard probability theory can be used to capture the uncertainty on the payoffs of securities through the notion that different *states* may carry different probabilities. By attaching a probability to each state, we shall be able to distinguish between a decision maker's beliefs (expressed by probabilities) about which state will occur and preferences about how she ranks the consequences of different actions.

Example 1.1. The asset menu is composed of the following three securities, A, B, and C:

State	*Security A*		*Security B*		*Security C*	
	Payoff	Prob.	Payoff	Prob.	Payoff	Prob.
i	20	3/15	18	3/15	18	3/15
ii	18	5/15	18	5/15	16	5/15
iii	14	4/15	10	4/15	12	4/15
iv	10	2/15	5	2/15	12	2/15
v	6	1/15	5	1/15	8	1/15

Security B pays 18 monetary units (say, euros) in both states *i* and *ii*. Therefore, the difference between these two states is not payoff-relevant to security B. However, it is payoff-relevant in the case of security A, in the sense that this asset pays out 20 euros in state *i* and 18 euros in state *ii*. Note that in this example, we characterize securities through their payoffs, but in future examples we shall equally use their period rate of return, computed as discussed in section 1.2.

The table above also shows the (subjectively determined) probabilities of each of the states. Because the states of the economy should be uniquely defined across the entire asset menu, the associated probabilities are simply repeated across different securities.

Of course, the table above reports redundant information because for securities B and C, one can re-define the states to consist of payoff-relevant states only. For instance, for security B there are only three payoff-relevant states, which we can call "*i+ii*", *iii*, and "*iv+v*"; in the case of security C, the payoff-relevant states are *i*, *ii*, "*iii+iv*", and *v*.

<table>
<tr><th rowspan="2">State</th><th colspan="2">Security A</th><th colspan="2">Security B</th><th colspan="2">Security C</th></tr>
<tr><th>Payoff</th><th>Prob.</th><th>Payoff</th><th>Prob.</th><th>Payoff</th><th>Prob.</th></tr>
<tr><td>i</td><td>20</td><td>3/15</td><td rowspan="2">18</td><td rowspan="2">8/15</td><td>18</td><td>3/15</td></tr>
<tr><td>ii</td><td>18</td><td>5/15</td><td>16</td><td>5/15</td></tr>
<tr><td>iii</td><td>14</td><td>4/15</td><td>10</td><td>4/15</td><td rowspan="2">12</td><td rowspan="2">6/15</td></tr>
<tr><td>iv</td><td>10</td><td>2/15</td><td rowspan="2">5</td><td rowspan="2">3/15</td></tr>
<tr><td>v</td><td>6</td><td>1/15</td><td>8</td><td>1/15</td></tr>
</table>

Example 1.1 illustrates the interplay among the three ingredients that we have listed above. First, the need to define and measure risk. For instance, if one takes notice of the potential returns, security A may be considered riskier than C because the span, the range of variation of the payoffs of security A (from a minimum of 6 to a maximum of 20), exceeds that of security C (from a minimum of 8 to a maximum of 18). Second, the usefulness of pinning down the concept of risk aversion. For instance, it is not immediately evident why a rational investor should prefer security C over security A (if any): on the one hand, security A threatens to pay out only 6 euros in state *v*; on the other hand, the same security achieves a very large payment of 20 in state *i*.[2] It is natural to ask what kind of investor would pay more for security C than for security A. Presumably such willingness would be motivated by a desire to avoid the very low payoff of 6 that the latter security may yield. Third and finally, it is unclear how such inclinations against risk—however measured—may be balanced off in the light of the probability distribution that characterizes different states.

In fact, this state-preference framework is fruitfully employed as an abstract tool for understanding the fundamentals of decision-making under uncertainty, but it is more special than it may first appear. For example, the set of states, *S*, is given exogenously and cannot be affected by the choices of the investors. In reality, many investment choices change the physical world and create chances for new outcomes and states of the world. For instance, a successful venture capital investment in cold fusion energy production will profoundly affect all other sectors and investment outlooks. Consequently, the state-preference model is not as widely applicable as it

[2] By construction, example 1.1 is perfectly symmetric: security C has a minimum payment of 8 that exceeds by 2 euros the minimum payment of security A; however, security A has a maximum payment of 20 that exceeds by 2 euros the maximum payment of security C. Hence the question in the main text stands.

might at first seem, and this should be kept in mind.

2.2 Complete and incomplete criteria of choice under uncertainty

The primary role played by the state-preference framework is to dictate how a rational investor ought to select among the different securities in her asset menu. One important distinction of criteria of choice under uncertainty, is their completeness: a *complete criterion* is always able to rank all securities or investment opportunities from top to bottom on the basis of their objective features. As such, complete criteria form a good basis for portfolio choice. For instance, an investor may (simplistically) decide to rank all available assets and to invest in some pre-determined fraction starting from the top of the resulting ranking.[3] The expected utility decision criterion to be defined in chapter 2 will satisfy this completeness property. By contrast, an *incomplete criterion* suffers from the existence of potential (usually, empirically relevant) traded combinations of the primitive assets that cannot be ranked in a precise way. As we are about to show, the celebrated mean-variance criterion is unfortunately incomplete. Paradoxically, such an incompleteness represents the reason for its success.

Although many other incomplete criteria can be defined (see Meucci, 2009, for a general framework), a first often referred to criterion is (strong) *dominance*:

Dominance: A security (strongly) dominates another security (on a state-by-state basis), if the former pays as much as the latter in all states of nature, and strictly more in at least one state.

In the absence any further indications on their behavior, we will assume that all rational individuals would prefer the dominant security to the security that it dominates. Equivalently, dominated securities will never be demanded by any rational investor. Here rational means that the investor is *non-satiated*, that is, she always prefers strictly more consumption (hence, monetary outcomes that may be used to finance such consumption) to less consumption.

However, the following example shows that the dominance criterion, although strong, is highly incomplete.

---

[3] Many "funds-of-funds" investment selection strategies are well known to be formally spelled out in this fashion, where the asset menu is composed by the (hedge or mutual) funds that can be selected.

Example 1.2. Consider the same asset menu, payoffs, and probabilities as in example 1.1:

State	*Security A*			*Security B*			*Security C*	
	Payoff	Prob.	*D*	Payoff	Prob.	*D*	Payoff	Prob.
i	20	3/15	>	18	3/15	=	18	3/15
ii	18	5/15	=	18	5/15	>	16	5/15
iii	14	4/15	>	10	4/15	<	12	4/15
iv	10	2/15	>	5	2/15	<	12	2/15
v	6	1/15	>	5	1/15	<	8	1/15

Clearly, as indicated by the signs in the "D" column (for dominance), the payoffs of security A dominate those of security B on a state-by-state basis. In this case, the exact probabilities that characterize the different states are not relevant. Even if one changes the probability distribution reported in the table, the result will stick. However, this criterion is visibly incomplete: for instance, security B does not dominate security C and, more importantly, security A does not dominate security C (and vice versa). Hence, neither security A nor C is dominated by any other security, while security B is dominated (by security A). A rational investor may then decide to select between assets A and C, ignoring B. However, she cannot find an equivalently strong and impartial rule to decide between security A and C, hence the criterion is incomplete.

The strength of dominance is that it escapes a definition of risk. Indeed, to be able to resort to such a concept may be very useful. However, in general, a security yields payoffs that in some states are larger and in some other states are smaller than under any other state. When this is the case, the best known (and yet still incomplete, as we shall see) approach at this point consists of summarizing the distributions of asset returns through their mean and variance:

$$E[R_i] = \sum_{s=1}^{S} Prob(\,state = s)R_i(s) \tag{1.6}$$

$$Var[R_i] = \sum_{s=1}^{S} Prob(\,state = s)\big[R_i(s) - E[R_i]\big]^2, \tag{1.7}$$

where *i* indicates a specific security of *N* and *S* is the number of states (e.g., 5 in examples 1.1 and 1.2). The following example shows in intuitive terms how mean and variance could be used to rank different securities, on the grounds that variance can be used to measure risk.

Example 1.3. Consider the same inputs as in examples 1.1 and 1.2:

State	*Security A*		*Security B*		*Security C*	
	Payoff	Prob.	Payoff	Prob.	Payoff	Prob.
i	20	3/15	18	3/15	18	3/15
ii	18	5/15	18	5/15	16	5/15
iii	14	4/15	10	4/15	12	4/15
iv	10	2/15	5	2/15	12	2/15
v	6	1/15	5	1/15	8	1/15
Mean	15.47		13.27		14.27	
Variance	16.78		28.46		8.46	

It is indeed the case that security C is less risky than security B.

If we decided to summarize these return distributions by their means and variances only, both securities A and C would clearly appear more attractive than asset B as they have a higher mean return and a lower variance. We therefore say that both securities A and C dominate asset B in terms of a *mean-variance dominance criterion*.

Mean-variance dominance: A security dominates another security in a mean variance (MV for short) sense, if the former is characterized by a higher expected payoff and a by a lower variance of payoffs.

However, security A fails to dominate security C (and vice versa) in a mean-variance sense. This occurs because security A has a higher mean than security C has (15.47 > 14.27), but the former also yields a higher variance (16.78 > 8.46). This shows that, as with to dominance, also the mean-variance is an incomplete criterion, that is, pairs of securities exist that cannot be simply ranked by this criterion.

Clearly, because of its incompleteness, the mean-variance criterion can at best only isolate a subset of securities that are not dominated by any other security. For instance, in example 1.3, security B, being dominated by both securities A and C, can be ruled out from the portfolio selection. However, neither security A nor C can be ruled out because they belong to the set of non-dominated assets.

Implicitly, the MV dominance criterion commits to a definition that requires an investor to dislike risk and that *identifies risk with variance*. Because the criterion implies this need to define and measure both risk aversion and risk, the mean-variance dominance is neither as strong, nor as a general concept as state-by-state dominance. In fact, we know from example 1.2 that while security A dominates state-by-state security B (and we now know that A also MV dominates B), security C does not dominate B on a state-by-state basis, while C MV dominates B.[4] Moreover, this criterion may at most identify some subset of securities (as we shall see, portfolios) that are not dominated and as such are "MV *efficient*". We shall return to these concepts in chapter 3.

3 - Statistical Summaries of Portfolio Returns

In section 2, we have introduced the idea that the returns of most financial assets (and thus of portfolios of such assets) are random variables. A random variable y is a quantity that can take a number of possible values, y_1, y_2, ..., y_n (and the case in which n diverges to infinity cannot be ruled out). The value that the random variable will assume is not known in advance, but a probability π_i is assigned to each of the possible outcomes y_i. The probability π_i can be (does not have to be) thought of as the frequency with which one would observe y_i if the experiment of observing the outcome of y could be repeated an infinite number of times.

We have already seen that the distribution of asset returns is often (and yet incompletely) characterized through their means and variances. In (1.6) and (1.7), we have shown how to compute the expected (or mean) return of an asset and its variance, respectively.[5] However, because we are also (mainly)

---

[4] Although our example does not show this feature, it is possible to build cases in which one asset dominates another on a state-by-state basis, but not in MV terms. This means that just as MV dominance does not imply state-by-state dominance, also state-by-state dominance fails to imply MV dominance. The two are merely different criteria.

[5] In the rest of the book, unless otherwise specified, we shall use the terms mean and expected return interchangeably to indicate the average value obtained by considering the probabilities as equivalent to frequencies.

interested in the risk-return profile of portfolios of assets, we now discuss how to aggregate the statistics of the individual assets to compute a portfolio mean and variance. Indeed, a portfolio is simply a linear combination of individual assets and, as a result, its return, R_P, is a random variable whose probability distribution depends on the distribution law(s) of the returns of the assets that compose the portfolio. Consequently, we can deduce some of the properties of the distribution of portfolio returns by using standard results regarding linear combinations of random variables. In particular, in what follows we focus on two-parameter distributions (sometimes called *elliptical*), of which the normal Gaussian distribution family represents the most important case, both theoretically and practically.
For instance, assume that we know that the returns R_A and R_B of two securities, A and B, are jointly normally distributed and that their means and variances are:

$$\mu_A \equiv E[R_A], \quad \sigma_A^2 \equiv Var[R_A]$$
$$\mu_B \equiv E[R_B], \quad \sigma_B^2 \equiv Var[R_B]. \tag{1.8}$$

Clearly, given these inputs, we can easily compute σ_A and σ_B, the square roots of the variances of the two assets, which are called *standard deviations* or, alternatively, *volatilities*. Yet, this information is not sufficient to compute all the required statistics characterizing the distribution of portfolio returns, and in particular portfolio variance, because asset returns are in general correlated, that is, they tend to show some form of linear dependence which goes to increase/decrease portfolio volatility above/below the variability justified by individual assets. For this reason, we need to introduce the concepts of *covariance* and of *correlation coefficient.*
The covariance σ_{AB} between two securities, A and B, is a scaled measure of the *linear association* between the two assets and it is computed as follows:

$$\sigma_{AB} \equiv Cov\left[R_{A,}R_B\right] = E\left[\left(R_{A,t} - E[R_A]\right)\left(R_{B,t} - E[R_B]\right)\right]. \tag{1.9}$$

The sign of the covariance reveals the kind of (linear) relationship that characterizes two assets. If $\sigma_{AB} > 0$, the returns of the two securities tend to move in the same direction; if $\sigma_{AB} < 0$, they tend to move in opposite directions; finally, if $\sigma_{AB} = 0$ the returns of the two securities are linearly independent (we also say they are uncorrelated). Intuitively, the covariance has to satisfy the following inequality:

$$|\sigma_{AB}| \leq \sigma_A \sigma_B. \tag{1.10}$$

Indeed, one can demonstrate that the covariance of an asset with itself is simply equal to its variance, $E\left[(R_{A,t} - E[R_A])^2\right]$; consequently, when two assets are perfectly correlated and therefore not distinguishable in a linear sense, then $\sigma_{AB} = \sigma_A \sigma_B$.

Looking at formula (1.9), it is evident that covariance is affected by the overall variability of the two assets, what statisticians call the *scales* of the two phenomena under consideration, in particular their standard deviations. As a result, if we were to rank pairs of securities based on the strength of their relationships, we would find it difficult to compare their covariances. For this reason, we usually standardize the covariance dividing it by the product of the standard deviations of the two assets:

$$\varrho_{AB} = \frac{\sigma_{AB}}{\sigma_A \sigma_B}. \tag{1.11}$$

The coefficient ϱ_{AB} is called correlation coefficient and it ranges from -1 to $+1$, as a result of the covariance bound stated in (1.9). A value of $+1$ indicates a perfect *positive* linear relationship between two assets, while -1 implies a perfect *negative* relationship between them. If two assets are completely linearly independent, they will display a correlation coefficient equal to 0. In this latter case, knowledge of the value of one variable does not give any information about the value of the other variable, at least within a linear framework.[6]

Now we have defined all the elements that allow us to compute the necessary mean-variance portfolio statistics. For the time being, we will take the values of the means, variances, and covariances of asset returns as given and show how these map in the mean and variance of portfolio returns. Later in

---

[6] Our emphasis on the fact that correlation just captures the strength of linear association may be best understood considering the following example: $R_{A,t} = R_{B,t}^2 + \eta_t$. Clearly, securities A and B are strongly associated according to a quadratic function. Yet it is easy to verify that $Cov[R_{A,}R_B] = 0$, i.e., the linear association between the two return series is zero.

the book (chapter 5) we will address how these can be empirically estimated.[7]

3.1 Portfolio mean

The computation of the mean of portfolio returns is relatively easy: the return of a portfolio simply consists of the sum of the returns of the components weighted by the fraction of wealth that is invested in each asset. Fortunately, the expected value operator enjoys a property called linearity which states that the expected (or mean) value of a sum of random variables is equal to the sum of the expected values of the random variables themselves; in addition, the expected value of a scalar multiple of a random variable is equal to the scalar coefficient applied to the expected value. Consequently, if $R_1, R_2, \ldots, R_N$ are random variables representing the returns of N securities that compose a portfolio, $\mu_1 \equiv E[R_{1,t}], \mu_2 \equiv E[R_{2,t}], \ldots, \mu_N \equiv E[R_{N,t}]$ are their expectations, and $\omega_1, \omega_2, \ldots, \omega_N$ are the weights of each security in the portfolio (expressed as a proportion of total wealth), then the portfolio mean is equal to:

$$\mu_P = E[R_P] = \sum_{i=1}^{N} \omega_i \mu_i. \tag{1.12}$$

To be more precise, consider the case of a portfolio consisting of the two securities A and B defined in (1.8), with weights ω_A and ω_B, respectively. The mean value of the portfolio can be easily computed as follows:

$$E[R_P] = E[\omega_A R_A + \omega_B R_B] = \omega_A E[R_A] + \omega_B E[R_B] = \omega_A \mu_A + \omega_B \mu_B. \tag{1.13}$$

The following example makes these simple concepts more concrete.

[7] It is important to recognize that the estimates that are obtained from actual data for means, variances, and covariances are the observable counterparts of *unobservable* theoretical concepts. Estimates of the mean and of the covariance matrix can be obtained through a variety of methods (estimation based on past data is just one common example). The way these estimates are constructed is addressed in chapter 5.

Example 1.4. Consider the two stocks A and B described below:

Market Condition	*Stock A*		*Stock B*	
	Return	Prob.	Return	Prob.
Bull	12.00%	25%	6.00%	25%
Normal	8.00%	50%	1.50%	50%
Bear	-7.00%	25%	-1.00%	25%
Mean	5.25%		2.00%	

For instance, the expected return of a portfolio composed by 30% of security A and 70% of security B is computed as follows:

$$\mu_P = \omega_A \mu_A + \omega_B \mu_B = 30\% \times 5.25\% + 70\% \times 2.00\% = 2.98\%$$

As a portfolio can be composed of a large number of assets, it may often be convenient to use a more compact matrix notation. If we indicate with $\boldsymbol{\omega}$ the $N \times 1$ vector containing the weights of the N securities that compose the portfolio and with $\boldsymbol{\mu}$ the $N \times 1$ vector of mean returns of the assets, then equation (1.12) can be rewritten as follows:

$$\mu_P = \boldsymbol{\omega}' \boldsymbol{\mu}. \tag{1.14}$$

3.2 Portfolio variance and standard deviation

As already pointed out, the computation of portfolio variance is a bit more complex than the calculation of its mean as it requires knowledge of the co-variances between each pair of asset returns. Following the standard definition of variance:

$$\begin{aligned}\sigma_P^2 = E[(R_P - \mu_P)^2] &= E\left[\left(\sum_{i=1}^{N} \omega_i R_i - \sum_{i=1}^{N} \omega_i \, \mu_i\right)^2\right] \\ &= E\left[\left(\sum_{i=1}^{N} \omega_i (R_i - \mu_i)\right)\left(\sum_{j=1}^{N} \omega_j (R_j - \mu_j)\right)\right] \\ &= E\left[\left(\sum_{i,j=1}^{N} \omega_i \omega_j (R_i - \mu_i)(R_j - \mu_j)\right)\right] = \sum_{i,j=1}^{N} \omega_i \omega_j \sigma_{i,j},\end{aligned} \tag{1.15}$$

where $\sigma_{i,j}$ is the covariance between asset *i* and asset *j*; as already discussed, the covariance of an asset with itself is simply equal to its variance. Note that also the variance formula can be rewritten using matrix notation:

$$\sigma_P^2 = \boldsymbol{\omega}'\boldsymbol{\Sigma}\boldsymbol{\omega}, \tag{1.16}$$

where $\boldsymbol{\omega}$ is again the $N \times 1$ of the weights and $\boldsymbol{\Sigma}$ is the so-called *variance-covariance matrix*, which is an $N \times N$ matrix whose main diagonal elements are the asset variances while the off-diagonal elements are the respective asset covariances. To clarify, $\boldsymbol{\Sigma}$ is a *symmetric, positive definite* matrix structured as follows:

$$\begin{bmatrix} \sigma_1^2 & \sigma_{1,2} & \cdots & \sigma_{1,N} \\ \sigma_{2,1} & \sigma_2^2 & \cdots & \vdots \\ \vdots & \vdots & \vdots & \vdots \\ \sigma_{N,1} & \cdots & \cdots & \sigma_N^2 \end{bmatrix} = \begin{bmatrix} \sigma_1^2 & \sigma_{1,2} & \cdots & \sigma_{1,N} \\ \sigma_{1,2} & \sigma_2^2 & \cdots & \vdots \\ \vdots & \vdots & \vdots & \vdots \\ \sigma_{1,N} & \cdots & \cdots & \sigma_N^2 \end{bmatrix}. \tag{1.17}$$

The second matrix clearly reflects the symmetric property of covariances and variances. Positive definiteness implies that for all *N*-component real vectors $\boldsymbol{x}$, $\boldsymbol{x}'\,\boldsymbol{\Sigma}\boldsymbol{x} > 0$. Clearly, because the vector of weights $\boldsymbol{\omega}$ is a just a special case of such a $\boldsymbol{\omega}$, it will be that $\sigma_P^2 = \boldsymbol{\omega}'\boldsymbol{\Sigma}\boldsymbol{\omega} > 0$.

To make the notion of portfolio variance more concrete, we come back to investigate in depth, a portfolio composed of only two stocks, A and B. In this application, the formula to compute portfolio variance simplifies to

$$\sigma_P^2 = \omega_A^2\sigma_A^2 + \omega_B^2\sigma_B^2 + 2\omega_A\omega_B\sigma_{AB}, \tag{1.18}$$

which can also be rewritten as

$$\sigma_P^2 = \omega_A^2\sigma_A^2 + \omega_B^2\sigma_B^2 + 2\rho_{AB}\omega_A\omega_B\sigma_A\sigma_B, \tag{1.19}$$

where $\sigma_{AB} = \rho_{AB}\sigma_A\sigma_B$.

Example 1.5. Consider again the two stocks mentioned in Example 1.4:

Market Condition	*Stock A*		*Stock B*	
	Return	Prob.	Return	Prob.
Bull	12.00%	25%	6.00%	25%
Normal	8.00%	50%	1.50%	50%
Bear	-7.00%	25%	-1.00%	25%
Mean	5.25%		2.00%	
Variance	0.0053		0.0006	
Standard Deviation	7.26%		2.52%	
Covariance	0.0015			
Correlation coefficient	0.83			

For instance, the variance of a portfolio composed of 30% of security A and 70% of security B is computed as follows

$$\begin{aligned}\sigma_P^2 &= \omega_A^2\sigma_A^2 + \omega_B^2\sigma_B^2 + 2\omega_A\omega_B\sigma_{AB} \\ &= 0.30^2\times 0.0053 + 0.70^2\times 0.0006 + 2\times 0.30\times 0.70\times 0.0015 \\ &= 0.0014.\end{aligned}$$

Clearly, the same result holds if we use matrix notation:

$$\sigma_P^2 = [0.30 \quad 0.70] \times \begin{bmatrix} 0.0053 & 0.0015 \\ 0.0015 & 0.0006 \end{bmatrix} \times \begin{bmatrix} 0.30 \\ 0.70 \end{bmatrix} = 0.0014.$$

Given such variance, it is also easy to compute also the standard deviation (or volatility) of the portfolio, which is simply equal to its square root:

$$\sigma_P = \sqrt{0.0014} = 0.0378 = 3.78\%.$$

Now consider a new asset, let's call it stock C, with the characteristics detailed below:

Market Condition	*Stock C*	
	Return	Prob.
Bull	-2.00%	25%
Normal	3.50%	50%
Bear	3.00%	25%
Mean	2.00%	
Variance	0.0005	
Standard Deviation	2.32%	

This stock has a negative covariance (hence, correlation) with both stock A and stock B. In particular, stock B and stock C have a covariance equal to −0.0005. Consequently, an equally weighted portfolio of stock B and C will have mean, variance, and standard deviation as computed below:

$$\mu_P = 0.50 \times 2.00\% + 0.50 \times 2.00\% = 2.00\%$$
$$\sigma_P^2 = 0.50^2 \times 0.0006 + 0.50^2 \times 0.0005 + 2 \times 0.50 \times 0.50 \times (-0.0005) = 0.00004$$
$$\sigma_P = \sqrt{0.00004} = 0.61\%.$$

Noticeably, this portfolio has a similar mean but a considerably lower risk (as expressed by the standard deviation) than both its component stocks. This is a consequence of the high *negative correlation* between the two assets:

$$\rho_{BC} = \frac{\sigma_{B,C}}{\sigma_B \sigma_C} = \frac{-0.0005}{0.0252 \times 0.0232} = -0.88.$$

The result is even more intuitive if we look at what happens when the market enters a bear regime. An investor holding the equally weighted portfolio defined above loses 1% of the wealth invested in stock B, but gains 3% on stock C. Overall, she gains 1% on her total wealth. Conversely, in a bull regime, she gains 6% on stock B, but loses 2% on stock C, with a total return of 2%. It is obvious that this investor would never lose money, while an investor holding only stock B or C would experience a negative return in some regimes. In practice, stock C provides a hedge to stock B in a bear regime and vice versa in a bull regime. As a result, the portfolio has a similar mean but a considerably lower risk than each of the two stocks, a point that we are about to explore in depth.

An analysis of the formula of portfolio variance leads us to a natural discovery of the concept of diversification. To illustrate this in the simplest and starkest set up, consider an equally weighted portfolio of N stocks (consequently, the weight assigned to each stock in the portfolio equals 1/N). In this case, formula (1.15) can be re-written as follows,

$$\sigma_P^2 = \sum_{i,j=1}^{N} \omega_i \omega_j \sigma_{i,j} = \sum_{i=1}^{N} \left(\frac{1}{N}\right)^2 \sigma_i^2 + \sum_{i,j=1}^{N} \left(\frac{1}{N}\right)\left(\frac{1}{N}\right) \sigma_{i,j}$$
$$= \frac{1}{N} \sum_{i=1}^{N} \frac{\sigma_i^2}{N} + \frac{N-1}{N} \sum_{i,j=1}^{N} \frac{\sigma_{i,j}}{N(N-1)}$$

$$= \frac{1}{N}\bar{\sigma}^2 + \frac{N-1}{N}\bar{\sigma}_{i,j}, \tag{1.20}$$

where $\bar{\sigma}^2$ and $\bar{\sigma}_{i,j}$ are the average portfolio variance and covariance, respectively. As N grows to infinity, the term $(1/N)\bar{\sigma}^2$ of equation (1.20) approaches zero. In other words, as N gets large the contribution of the variance of the individual stocks to the variance of the portfolio goes to zero. Therefore, the variance of a large portfolio does not depend on the individual risk of the securities, but only on their average covariance. Figures 1.1 and 1.2, illustrate this result for the US and the Italian equity markets, respectively. In the plots, the vertical axes indicate the risk of the portfolio as a percentage of the risk of an individual security. The horizontal axis represents the number of stocks included in the portfolio.[8]

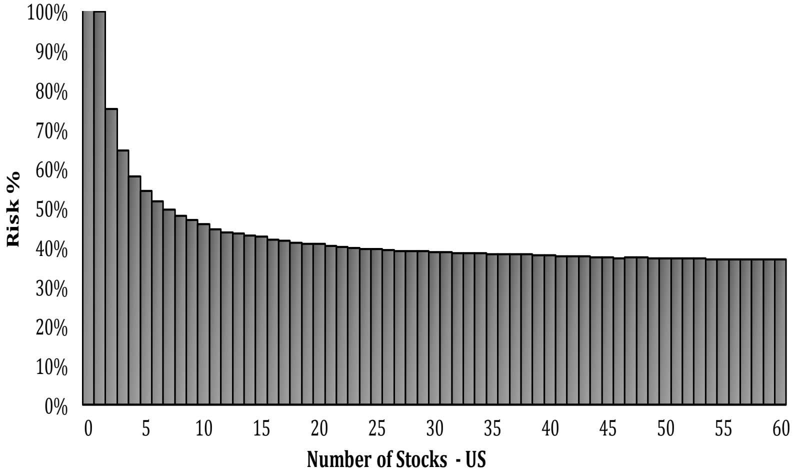

Figure 1.1

[8] The two figures were obtained as follows. For the US market, we collect monthly returns for 2,237 stocks from CRSP (Center for Research on Security Prices) over a sample spanning the period December 1994 - December 2015 and compute their variance. Then, we randomly select N stocks (with N increasing from 1 to 60) and calculate the resulting portfolio standard deviation. We repeat the exercise 1,000 times and compute the average standard deviation of all portfolios composed by N stocks. The latter is then expressed as a percentage of the average standard deviation of a single stock, randomly picked. In the case of the Italian stock market, we perform the same exercise, but with a lower number of stocks to start from (60) and a higher number of simulations (10,000) to guarantee sufficient stability in variance estimates. In this case, monthly returns are collected with for the period January 2000 - April 2016.

It is evident that in both cases the standard deviation of the portfolio sharply declines as we add the first 10 stocks, then it slowly converges towards the average covariance of the pool of stocks considered. Interestingly, the average covariance reduction is much larger for Italian stocks than for US stocks. Indeed, the total risk of a large portfolio of Italian stocks is equal to only 14% of the average risk of a single individual Italian stock, while the total risk of a US portfolio cannot be reduced below 35% of the average risk of an individual security. Clearly, the more the stocks are uncorrelated, the lower the variance of a well-diversified portfolio will be. Indeed, the second term of equation (1.20), the average covariance, depends on the average correlation coefficient among stocks. If all the stocks were uncorrelated (the average correlation coefficient would be equal to zero), a well-diversified portfolio would show zero risk, as the second term of (1.20) would be zero as well.

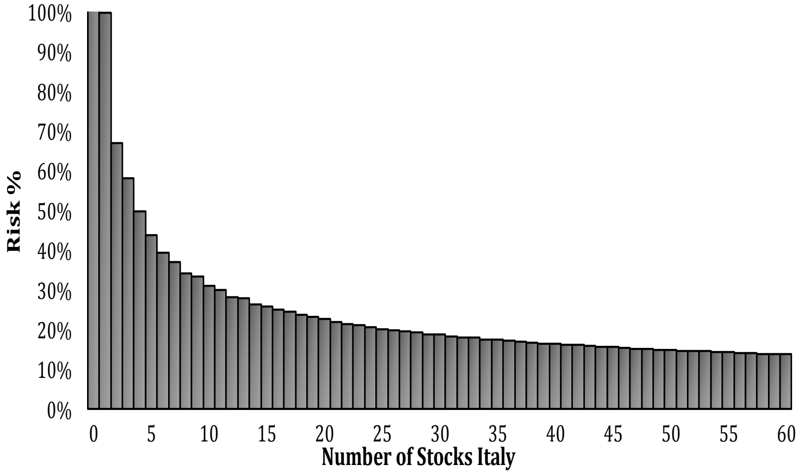

Figure 1.2

Rearranging equation (1.20) helps us understand when a portfolio has reached the minimum possible variance:

$$\sigma_P^2 = \frac{1}{N}\left(\bar{\sigma}^2 - \bar{\sigma}_{i,j}\right) + \bar{\sigma}_{i,j} \tag{1.21}$$

When the difference between the average variance and the average covariance of all stocks is equal to zero adding a new stock would not help to further decrease the portfolio variance. Since it can be eliminated by holding a large number of stocks, the risk arising from individual securities is often called *diversifiable* risk and an investor should not be rewarded for taking it. We shall examine this concept again in chapter 5.

References and Further Readings

Bailey, R., E. *The Economics of Financial Markets*. Cambridge: Cambridge University Press, 2005.

Campbell, J. Y., Lo, A. W. C., and MacKinlay, A. C. *The Econometrics of Financial Markets*. Princeton, NJ: Princeton University Press, 1997.

Cuthbertson, K., and Nitzsche, D. *Quantitative Financial Economics: Stocks, Bonds and Foreign Exchange*. John Wiley & Sons, 2005.

Danthine, J. P., and Donaldson, J. B. *Intermediate Financial Theory*. Academic Press, 2014.

Fabozzi, F., and Markowitz, H. *The Theory and Practice of Investment Management*, Second Edition, John Wiley & Sons, 2011.

Huang, C.-f., and Litzenberger, R., H., *Foundations for Financial Economics*. Amsterdam: North-Holland, 1988.

Luenberger, D. G., *Investment Science*. Oxford: Oxford University Press, 1997.

Meucci, A., *Risk and Asset Allocation*. Springer Science & Business Media, 2009.

Modigliani, F., and Pogue, G. A. An introduction to risk and return: concepts and evidence, part two. *Financial Analysts Journal*, 30, 69-86, 1974.