

Financial Markets: Functions, Institutions, and Traded Assets

Providing a simple, yet exhaustive definition of finance is no quite easy task, but a possible attempt, at least from a conceptual viewpoint, is the following:¹

Finance is the study of how people and organizations allocate scarce resources over time, subject to uncertainty.

This definition might sound somewhat generic, but it does involve the two essential ingredients that we shall deal with in practically every single page of this book: *Time* and *uncertainty*. Appreciating their role is essential in understanding why finance was born in the past and is so pervasive now. The time value of money is reflected in the interest rates that define how much money we have to pay over the time span of our mortgage, or the increase in wealth that we obtain by locking up our capital in a certificate of deposit issued by a bank. It is common wisdom that the value of \$1 now is larger than the value of \$1 in one year. This is not only a consequence of the potential loss of value due to inflation.² A dollar now, rather than in the future, paves the way to earlier investment opportunities, and it may also serve as a precautionary cushion against unforeseen needs. Uncertainty is related, e.g., to the impossibility of forecasting the return that we obtain from investing in stock shares, but also to the risk of adverse movements in currency exchange rates for an import/export firm, or longevity risk for a worker approaching retirement. As we show in Chapter 2, we may model issues related to time and uncertainty within a mathematical framework, applying principles from financial economics and tools from probability, statistics, and optimization theory. Before doing that, we need a more concrete view

¹This definition is taken from [2].

²This holds under common economic conditions; the exception to the rule is deflation, which is (at the time of writing) a possibility in Euroland. In this book, we will assume that the standard economic conditions prevail.

in order to understand how financial markets work, which kinds of assets are exchanged, and which actors play a role in them and what their incentives are. We pursue this “institutional” approach to get acquainted with finance in this chapter. Some of the more mathematically inclined students tend to consider this side of the coin modestly exciting, but a firm understanding of it is necessary to put models in the right perspective and to appreciate their pitfalls and limitations.

In Section 1.1, we discuss the role of time and uncertainty in a rather abstract way that, nevertheless, lays down some essential concepts. A more concrete view is taken in Section 1.2, where we describe the fundamental classes of assets that are traded on financial markets, namely, stock shares, bonds, currencies, and the basic classes of derivatives, like forward/futures contracts and options. In order to provide a proper framework, we also hint at the essential shape of a balance sheet, in terms of assets, liabilities, and equity, and we emphasize the difference between standardized assets traded on regulated exchanges and less liquid assets, possibly engineered to meet specific client requirements, which are traded over-the-counter. In Section 1.3, we describe the classes of players involved in financial markets, such as investment/commercial banks, common/hedge/pension funds, insurance companies, brokers, and dealers. We insist on the separation between the institutional form and the role of those players: A single player may be of one given kind, in institutional terms, but it may play different roles. For instance, an investment bank can, among many other things, play the role of a prime broker for a hedge fund. Furthermore, depending on circumstances, players may act as hedgers, speculators, or arbitrageurs. The exact organization of financial markets is far from trivial, especially in the light of extensive use of information technology, and a full description is beyond the scope of this book. Nevertheless, some essential concepts are needed, such as the difference between primary and secondary markets, which is explained in Section 1.4. There, we also introduce some trading strategies, like buying on margin and short-selling, which are essential to interpret what happens on financial markets in practice, as well as to understand some mathematical arguments that we will use over and over in this book. Finally, in Section 1.5 we consider market indexes and describe some basic features explaining, for instance, the difference between an index like the Dow Jones Industrial Average and the Standard & Poor 500.

1.1 What is the purpose of finance?

If you are reading this book, chances are that it is because you would like to land a rewarding job in finance. Even if this is not the case, one of the reasons why we aim at finding a good job is because we need to earn some income in order to purchase goods and services, for ourselves and possibly other people we care about. Every month (hopefully) we receive some income, and we must plan its use. The old grasshopper and ant fable teaches that we should actually




FIGURE 1.1 Shifting consumption forward and backward in time.

plan ahead with care. Part of that income should be saved to allow consumption at some later time. Sometimes, we might need to use more income than we are earning at present, e.g., in order to finance the purchase of our home sweet home.

Now, imagine a world in which we cannot “store” money, and we have to consume whatever our income is *immediately*, no more, no less, just like we would do with perishable food, if no one had invented refrigerators and other conservation techniques. This unpleasing situation is depicted in Fig. 1.1(a). There, time is discretized in $T = 3$ time periods, indexed by $t = 1, \dots, T$.³ The income during time period t is denoted by I_t , and it is equal to the consumption C_t during the same period:

$$I_t = C_t, \quad t = 1, \dots, T.$$

³Sometimes, time discretization requires careful thinking about events. Do we earn income at the beginning or at the end of a time period? In other words, is income earned during time period t immediately available for consumption during the same time period? We may argue that income during time period t is available for consumption only during time period $t + 1$. We shall discuss more precise notation and concepts in Section 2.1.2. Here, for the sake of simplicity, we assume that every event during a time *period* is concentrated at some time *instant*. We sometimes use the rather awkward term *epoch* to refer to a specific point in time. We also often use the term time *bucket* to refer to a time period delimited by two time instants.

This state of the matter is not quite satisfactory, if we have excess income in some period and would like to delay consumption to a later time period. In Fig. 1.1(b), part of income I_2 , denoted by $L_{2,3}$, is shifted forward from time period 2 to time period 3. This results in an increase of C_3 and a decrease of C_2 . The amount of income saved can be regarded as money invested or lent to someone else. By a similar token, we might wish to anticipate consumption to an earlier time period. In Fig. 1.1(c), consumption C_2 is increased by shifting income backward in time from time period $t = 3$, which means borrowing an amount of money $B_{3,2}$, to be used in time period $t = 2$ and repaid in time period $t = 3$. Savers and borrowers may be individuals or institutions, and we may play both of these roles at different stages of our working life. Clearly, all of this may happen if there is a way to match savers and borrowers, so that all of them may improve their consumption timing. This is one of the many roles of financial markets; more specifically, we use the term **money markets** when the time span of the loan is short. In other cases, the investment may stretch over a considerable time span, especially if savers/borrowers are not just households, but corporations, innovative startups, or public administrations that have to finance the development of a new product, the building of a new hospital, or an essential infrastructure. In this case, we talk about **capital markets**.

Needless to say, if we accept to delay consumption, it is because we expect to be compensated in some way. Informally, we exchange an egg for a chicken; formally, we earn some interest rate R along the time period involved in the shift.⁴ We may interpret the shift as a flow of money over a network in time but, unlike other network flows involved in transportation over space, we do not have exact conservation of flows. With reference to Fig. 1.1(b), we have the following flow balance equations at nodes 2 and 3:

$$\begin{aligned} C_2 &= I_2 - L_{2,3}, \\ C_3 &= I_3 + L_{2,3}(1 + R), \end{aligned}$$

stating that we give up an amount $L_{2,3}$ of consumption at time 2 in exchange for an increase $(1 + R)L_{2,3}$ in later consumption. The factor $1 + R$ is a gain associated with the flow of money along the arc connecting node $t = 2$ to node $t = 3$. This is what the **time value** of money is all about. The exact value of the interest rate R , as we shall see in Chapter 3, may be related to the possibility of default (i.e., the borrower may not repay the full amount of his debt) and to inflation risk, among other things.

Clearly, there must be another side of the coin: The increase in later consumption must be paid by a counterparty in an exchange. We delay consumption while someone else anticipates it. With reference to Fig. 1.1(c), we have

⁴In financial practice, whenever an interest rate is quoted, it is always an *annual* rate. For now, let us associate the rate with an arbitrary time period.

the following flow balance equations at nodes 2 and 3:

$$\begin{aligned} C_2 &= I_2 + \frac{B_{3,2}}{1+R}, \\ C_3 &= I_3 - B_{3,2}. \end{aligned}$$

Note that we are expressing the borrowed amount $B_{3,2}$ in terms of the money at time $t = 3$, when the debt is repaid; in other words $B_{3,2}$ is a flow out of node 3. This is not essential at all: If we use money at time $t = 2$, i.e., we consider the flow $B_{3,2}^*$ into node 2, the flow balance would simply read

$$\begin{aligned} C_2 &= I_2 + B_{3,2}^*, \\ C_3 &= I_3 - B_{3,2}^*(1+R). \end{aligned}$$

The two sides of the coin must be somehow matched by a market mechanism. In practice, funds are channeled by financial intermediaries, which must be compensated for their job. In fact, there is a difference between lending and borrowing rates, called **bid-ask** (or bid-offer) **spread**, which applies to other kinds of financial assets as well. Lending and borrowing money through a bank is what we are familiar with as individuals, whereas a large corporation and a sovereign government have the alternative of raising funds by issuing securities like bonds, typically promising the payment of periodic interest, as well as the refund of the capital at some prespecified point in time, the maturity of the bond. Corporations may also raise funds by issuing stock shares. Buying a stock share does not mean that we lend money to a firm; hence, we are not entitled to the payment of any interest. Rather, we own a share of the firm and may receive a corresponding share of earnings that may be distributed in the form of dividends to stockholders. However, the amount that we will receive is random and no promise is made about dividends, as they depend on how well the business is doing, as well as the decision of reinvesting part of the earning in new business ventures, rather than distributing the whole of it.

After being first issued, securities like bonds and stock shares may be exchanged among market participants, at prices that may depend on several underlying risk factors. Since the values of these factors are not known with certainty, the future prices of bonds and stock shares are random. In fact, time is intertwined with another fundamental dimension in finance, namely, **uncertainty**. When we lend or borrow money at a given interest rate, the future cash flows are known with certainty, if we do not consider the possibility of a default on debt. However, when we buy a stock share at time $t = 0$ and plan to sell it at time $t = T$, randomness comes into play. Let us denote the initial price by $S(0)$.⁵ The future price $S(T)$ is a **random variable**, which we may denote as $S(T, \omega)$ to emphasize its dependence on the random outcome (scenario) ω . We recall that, in probability theory, a random variable is a function mapping underlying random outcomes, corresponding to future scenarios or states of nature, to numeric values. Let ω_i , $i = 1, \dots, m$, denote the i -th outcome, which

⁵Depending on notational convenience, we shall write $S(t)$ or S_t , as no ambiguity should arise.




FIGURE 1.2 Representing uncertain states of the world by a scenario fan.

occurs with probability π_i . For the sake of simplicity, we are considering a discrete and finite set of possible outcomes, whereas later we will deal extensively with continuous random variables. A simple way to depict this kind of discrete uncertainty is by a scenario fan like the one depicted in Fig. 1.2. Therefore, $S(T, \omega)$ is a random variable, and we associate a future price $S(T, \omega_i)$ with each future state of the world. The corresponding holding period return is defined as follows.

DEFINITION 1.1 (Holding period return) *Let us consider a holding period $[0, T]$, where the initial asset price is $S(0)$ and the terminal random asset price is $S(T, \omega)$. We define the **holding period return** as*

$$R(\omega) \doteq \frac{S(T, \omega) - S(0)}{S(0)} \quad (1.1)$$

and the **holding period gain** as

$$G(\omega) \doteq \frac{S(T, \omega)}{S(0)} = 1 + R(\omega). \quad (1.2)$$

The gain and the holding period return (return for short) are clearly related. A return of 10% means that the stock price was multiplied by a gain factor of 1.10.

Remark. The term *gain* is not so common in finance textbooks. Usually, terms like *total return* or *gross return* are used, rather than gain. On the contrary, terms like *rate of return* and *net return* are used to refer to (holding period) return. The problem is that these terms may ring different bells, especially to practitioners. We may use the qualifier “total” when we want to emphasize a return including dividend income, besides the capital gain related to price changes. Terms like “gross” and “net” may be related with taxation issues, which we shall always disregard. This is why we prefer using “gain,” even though this usage is less common. We shall not confuse gain, which is a multiplicative factor,

with profit/loss, which is an additive factor and is expressed in monetary terms. Furthermore, we shall reserve the term “rate of return” to the case of *annual* returns. For instance, interest rates are always quoted annually, even though they may be applied to different time periods by a proper scaling. We will use the term “return,” when the holding period may be arbitrary. The following example illustrates a further element of potential confusion when talking about return.

Example 1.1 Different shades of return

Consider a holding period consisting of two consecutive years. In year one, the return from investing in a given stock share is +10%; in year two the return is −10%. What was the “average” return?

As it turns out, the question is stated in a very imprecise way. It might be tempting to say that, trivially, the average return was 0%, the familiar arithmetic mean of +10% and −10%. However, we cannot really add returns like this. Over the two years, the gain was

$$G = (1 + 0.10) \times (1 - 0.10) = 0.99,$$

i.e., we have lost money, as the holding period return was −1% [we may recall the rule $(1 + x)(1 - x) = 1 - x^2$]. Indeed, the problem is that the very term “average” is ambiguous. If what we actually mean is the *expected value* of the annual return, which we may estimate by a *sample mean*, then we may say that the arithmetic average is, in fact

$$\bar{R}_a = \frac{0.10 - 0.10}{2} = 0.$$

But if we mean an average *over time*, we should deal with a sort of geometric average over two years:

$$(1 + 0.10) \times (1 - 0.10) = (1 + \bar{R}_g)^2 \quad \Rightarrow \quad \bar{R}_g = -0.5013\%.$$

We may also notice that, in this case, an average should refer to a standard time interval, usually one year. Indeed, we should not confuse the holding period return with an *annual* (rate of) return. We will need a way to annualize a generic holding period return.

Returns and gains are random variables. Hence, a natural question is: How should we model uncertain returns? There is a huge amount of work carried out on this subject, including plenty of empirical investigation. The next example shows that there cannot be any single convenient answer.

■ Example 1.2 One distribution does not fit all

The most familiar probability distribution is, no doubt, the normal. Can we say that the distribution of return from a stock share is normal? Empirical investigation tends to support a different view, as the normal distribution is symmetric and features thin tails, i.e., it tends to underestimate the probability of extreme events. In any case, the worst stock return we may experience is -100% , or (-1) , i.e., we lose all of our investment (this is related to the limited liability property of stock shares, discussed later). In other words, the worst gain is 0, and a stock share price can never be negative. Since the support of a normal random variable is unbounded, $(-\infty, +\infty)$, according to this uncertainty model there is always a nonzero probability of observing an impossible price.

However, let us discuss the matter from a very limited viewpoint, namely, convenience. One nice feature of a normal distribution is that if we add normal variables, we get another normal (to be precise, we should be considering *jointly normal* variables). This is nice when we add returns from different stock shares over the same time period. If we have invested 30% of our wealth in stock share a and 70% in stock share b , the holding period return for the portfolio is

$$R_p = 0.3R_a + 0.7R_b, \quad (1.3)$$

where we denote the return of stock shares a and b by R_a and R_b , respectively, and R_p is the portfolio return. To justify Eq. (1.3), let us consider:

- Initial stock prices $S_a(0)$ and $S_b(0)$
- Initial wealth $W(0)$
- Stock prices $S_a(T)$ and $S_b(T)$ at the end of the holding period
- Wealth $W(T)$ at the end of the holding period

Then, if initial wealth is split as we have assumed, we may write

$$\begin{aligned} W(0) &= \frac{0.3 \times W(0)}{S_a(0)} \cdot S_a(0) + \frac{0.7 \times W(0)}{S_b(0)} \cdot S_b(0) \\ &= N_a \cdot S_a(0) + N_b \cdot S_b(0), \end{aligned}$$

where N_a and N_b are the number of stock shares a and b , respectively, that we buy. At the end of the holding period, we have

$$\begin{aligned} W(T) &= N_a \cdot S_a(T) + N_b \cdot S_b(T) \\ &= N_a \cdot (1 + R_a) \cdot S_a(0) + N_b \cdot (1 + R_b) \cdot S_b(0) \end{aligned}$$

$$\begin{aligned}
&= W(0) \cdot \left[1 + \frac{N_a \cdot R_a \cdot S_a(0) + N_b \cdot R_b \cdot S_b(0)}{W(0)} \right] \\
&= W(0) \cdot (1 + 0.3 \cdot R_a + 0.7 \cdot R_b) \\
&= W(0) \cdot (1 + R_p),
\end{aligned}$$

which gives Eq. (1.3).

If R_a and R_b are jointly normal, then R_p is a nice normal, too, which is quite convenient. Furthermore, if the holding period return R is normal, so is the corresponding stock price $S(T) = S(0) \cdot (1 + R)$. However, imagine that we take a different perspective. Rather than considering two stock shares over one time period, let us consider one stock share over two consecutive time periods. In other words, we take a *longitudinal* view (a single variable over multiple time periods) rather than a *cross-sectional* view (multiple variables over a single time period). Let us denote by $R(1)$ and $R(2)$ the two holding period returns of that single stock share, over the two consecutive time periods. As we have mentioned, in this case we should not add returns, but rather multiply gains $G(1)$ and $G(2)$ to find the holding period gain

$$\begin{aligned}
G &= G(1) \cdot G(2) \\
&= [1 + R(1)] \cdot [1 + R(2)] \\
&= 1 + R(1) + R(2) + R(1) \cdot R(2).
\end{aligned}$$

The last expression involves a product of returns. Unfortunately, if $R(1)$ and $R(2)$ are normal, their product is not. Hence, the holding period gain G is not normal, and the same applies to the holding period return $R = G - 1$. We may only say that the holding period return is approximately normal if the single-period returns are small enough to warrant neglecting their product.

One way out is to consider the **logarithmic return**, or **log-return** for short,

$$r \doteq \log(1 + R) \equiv \log G,$$

where we use $\log$ rather than $\ln$ to denote natural logarithm. It is interesting to note that, given the well-known Taylor expansion (Maclaurin series, if you prefer)

$$\log(1 + x) \approx x - \frac{x^2}{2} + \frac{x^3}{3} - \frac{x^4}{4} + \cdots,$$

for a small x , the log-return can be approximated by the return. Since

$$\begin{aligned}\log [(1 + R(1)) \cdot (1 + R(2))] &= \log (1 + R(1)) + \log (1 + R(2)) \\ &= r(1) + r(2),\end{aligned}$$

we see that log-returns are additive, and if we assume that they are normal, we preserve normality over time.

Since

$$S(T) = S(0) \cdot G = S(0) \cdot e^r,$$

the normality of the log-return r implies that the gain and the stock prices are lognormally distributed, i.e., they may be expressed as the exponential of a normal random variable. On the one hand, this is nice, as it is consistent with the fact that we cannot observe negative stock prices. Furthermore, the product of lognormals is lognormal, which is nice in the longitudinal sense. Unfortunately, this is not nice in the cross-sectional sense, since the sum of lognormals is not lognormal, and we get in trouble when we consider the return of a portfolio of different stock shares.

To summarize, whatever modeling choice we make, some complication will arise. On the one hand, normal returns/gains (and stock prices) simplify the analysis of a portfolio over a single holding period, but they are empirically questionable and complicate the analysis over multiple time periods. On the other hand, lognormal gains (and stock prices) are fine for dynamic modeling of a single stock share, but they complicate the analysis of a portfolio. We may conclude that, whatever we choose, we have to accept some degree of approximation somewhere. The alternative is to tackle complicated distributions by numerical methods.

Beside risky assets, we shall also consider a risk-free (or riskless) asset. This is a peculiar asset for which $S(T, \omega)$ is actually a constant across states of the world. A concrete example is a safe bank account, whereby

$$B(T, \omega) = B(0) \cdot (1 + R_f),$$

for every state of the world (or scenario) $\omega \in \Omega$. The rate R_f will be referred to as **risk-free return**. If the holding period T is one year, we may refer to the annual risk-free return as the **risk-free rate**. The above framework to depict uncertainty does not only apply to stock shares, but to other financial and nonfinancial assets as well, like bonds, commodities, foreign currencies, etc. As we shall see in Chapter 2, uncertainty motivates some basic problems in finance, like portfolio optimization and risk management.

Now, if we may invest wealth at some prespecified and risk-free interest rate, why should we bother with risky investments? The answer is that risk is associated with the hope of a larger return, i.e., risky assets come with a **risk premium**. Some investors are willing to assume a limited amount of risk in exchange for the possibility of an increase in future consumption. Other investors, however, would just like to get rid of risks they bear:

- Imagine a nonfinancial firm subject to research, development, and production costs incurred in some currency, for a range of products that are exported and sold in another currency. For instance, the firm might sign a contract for the design and construction of a production plant, where the overall price offered to the client is in US dollars, but actual costs are incurred in euro. Adverse fluctuations in currency exchange rates may well wipe out profit margins. As we shall see, firms may hedge this risk away using certain derivative assets, such as forward and futures contracts.
- With reference to Fig. 1.2, let us assume that the state corresponding to outcome ω_5 is a “bad state,” i.e., a state in which we will be able only to afford a very low consumption level, possibly because an adverse event occurs (like illness, accident, or loss of job). Then, we might consider purchasing shares of an *insurance* contract, i.e., an asset whose value is strictly positive when ω_5 occurs, 0 otherwise. We assume that the insurance payoff is 1 in the bad state, but any other value will do, if assets are perfectly divisible and we may scale investments up and down at will.

More generally, an investor may shape the probability distribution of her wealth according to her taste and appetite for risk. A market participant with a given risk exposure may change it, and this is the essential function of risk management. Clearly, for any player hedging a risk exposure away, there must be another market participant willing to assume that risk or part of it. With respect to this uncertainty dimension, financial markets play the role of a **risk transfer** mechanism. For instance, insurance companies do that in exchange for a premium, and rely on *risk pooling* and *reinsurance* contracts to manage the resulting risk exposure.⁶ One of the main problems in this context is the definition of a fair insurance premium, which is a standard task in actuarial mathematics.

Note that an insurance contract is an asset from the viewpoint of the policy owner, but not a tradable one, as we cannot sell our life insurance policy. However, an insurance company, for which insurance contracts are a liability, may pool and sell them to interested investors, using a process called **securitization**, which is the creation of liquid securities from illiquid assets.⁷ By doing this, the

⁶Risk pooling may be considered as a corollary of the law of large numbers. If we aggregate a large number of small and independent risks, the overall risk should be reduced. This happens, e.g., with car insurance policies. Risk pooling may fail miserably with strongly correlated risks. Reinsurance, in a sense, is an opposite mechanism, by which a large risk is fractioned and sold to third parties.

⁷See Section 1.2.2.

risk may be fractioned and sold to investors who are willing to bear part of the risk for a given price. Securitization is a good example to illustrate the pros and cons of financial innovation. On the one hand, it allows to create securities that may offer enhanced return to holders, which is good in a regime of low interest rates. Furthermore, it may allow to insure against catastrophic risks that could not be insured otherwise. On the other one, just consider the damage inflicted to financial markets by the creation of illiquid and opaque mortgage-backed securities, bundling subprime mortgages and leading to the 2008 financial crisis. The same considerations apply to derivatives, which may be used in quite different ways by players having different views about the future or different attitudes toward risk (like hedgers and speculators, discussed later in this chapter). In all of these cases, the fundamental recurring themes are asset pricing and risk management, which we start considering in Chapter 2.

In later chapters, we will also see that making decisions under uncertainty is no trivial task, and that in real life, things are complicated by the fact that the two dimensions that we have considered, time and uncertainty, are actually intertwined. The resulting picture, illustrated in Fig. 1.3, is a **scenario tree**, where uncertainty unfolds progressively over time. The tree consists of a set of nodes n_k , $k = 0, 1, \dots, 14$. Node n_0 is the root of the tree and represents the current state of the world. Then, over three time instants, $t = 1, 2, 3$, we observe a sequence of realizations of random variables representing financial risk factors. The outcomes ω_i of the sample space are associated with **scenarios**, i.e., sequences of nodes in the tree. For instance, scenario ω_3 corresponds to the sequence of nodes

$$(n_0, n_1, n_4, n_9).$$

More formally, each scenario is a **sample path** of a stochastic process. We also see that the probability of a scenario depends on *conditional* probabilities of events. For instance, the conditional probability of node n_4 at time $t = 2$, given that we are at node n_1 at time $t = 1$, is $\pi_{4|1}$. Hence, the *unconditional* probability of scenario ω_3 is

$$P(\omega_3) = \pi_{1|0} \cdot \pi_{4|1} \cdot \pi_{9|4}.$$

Since we are at state n_0 , we may write π_1 rather than $\pi_{1|0}$, but we must be careful in distinguishing conditional and unconditional probabilities. Stochastic processes and the generation of scenario trees are discussed in Chapter 11. Dynamic policies in such a context must allow for a way to adapt a strategy to contingencies, and this leads to challenging multistage optimization models discussed in Chapter 15.

1.2 Traded assets

Finance revolves around buying and selling assets, pricing them, and assessing the involved risk. But what are assets, exactly? Open any page of a financial




FIGURE 1.3 A scenario tree generalizes the scenario fan of Fig. 1.2 by unfolding uncertainty progressively over time.

journal and you will read about assets such as stock shares, bonds, or derivatives. Indeed, these are the assets that we will mostly deal with in this book; yet, it is essential to get a broader picture. Generally speaking, an **asset** is *anything that can be transformed into money by its owner*:

- A financial institution, like a pension fund, may rely on a portfolio of bonds as an asset: The stream of coupon payments is used to pay pensions to retired workers.
- A nonfinancial firm uses machines and other equipments to produce items for sale. These items may be innovative products protected by a patent; the patent is another asset that may be sold.
- An individual may use her human capital, possibly a Ph.D. title, to land a good, rewarding, and hopefully well-paid job. Unlike other assets, a Ph.D. title is not marketable.

- An insurance policy is an asset that can be transformed into money, but only when a prespecified event occurs, since it cannot be freely traded.

We see that assets may be both tangible or intangible objects that can be transformed into a sequence of cash flows, and they come in plenty of different forms. To start putting some order, let us introduce a few basic features to help us in classifying assets:

Real vs. financial. The bonds owned by a pension fund are financial in nature, whereas manufacturing equipments are real.

Risky vs. risk-free. Stock shares are considered as risky assets, as their future price and their dividend income are not known with certainty. A certificate of deposit issued by a very solid bank is perceived as a risk-free asset, since we know exactly how much money we are going to collect at maturity, even though the risk of default cannot be ruled out with absolute certainty.

Liquid vs. illiquid. Liquidity refers to the possibility of selling an asset quickly *and* at a fair price; both sides of the coin are relevant. If we need a lot of money immediately, we may sell our home; however, if we really want to do it quickly, we may be forced to accept a price that is possibly much lower than its fair value. A similar consideration applies to manufacturing equipments, which may be very specific and difficult to sell for a fair price. On the contrary, most stock shares are very liquid and actively traded on regulated exchanges. Shares of common funds are liquid and can be redeemed on short notice, whereas shares of hedge funds may require several weeks to be liquidated.

Tradable vs. nontradable. Most financial assets are easily traded on markets, but we cannot sell our own insurance policy. The fact that an asset is nontradable does not diminish its importance. For instance, when we age, we lose a fraction of human capital, as the sheer number of future cash flows that we obtain from our job gets less and less. If our human capital is a rather safe asset, then we may initially consider tilting our strategic asset allocation toward reasonably risky stock shares. When we age, it is a common advice that we should rebalance the portfolio toward safer assets.

Exchange-traded vs. over-the-counter. Stock shares are traded on regulated exchanges, just like some simple and standardized classes of derivatives (vanilla options and futures contracts). Sometimes, we need a more specific kind of asset for risk management purposes, which may be tailored by an investment bank according to our requirements. When an investment bank engineers a very specific asset, this is sold over-the-counter (**OTC**), rather than on regulated exchanges. Plain vanilla options are examples of exchange-traded derivatives, whereas exotic options are OTC assets. Unfortunately, a tailored OTC asset will be harder to sell. Typically, we may only sell it back to the original issuer by closing the contract, and its price is less easy to quantify as it is not related to a transparent demand–offer mechanism.

In the rest of this section, we outline the most common forms of financial assets, namely, stock shares (equity), bonds, and derivatives, as well as foreign currencies and hybrid assets. Before doing so, it is useful to lay down a (very) simplified view of a balance sheet, showing the connection between assets, liabilities, and equity. We also discuss briefly the difference between assets and securities.

1.2.1 THE BALANCE SHEET

A cornerstone of corporate finance is the **balance sheet**, one of the fundamental documents periodically issued by firms, which is used by investors and stakeholders to assess the health state of the firm. The essence of a balance sheet may be schematically represented in the following tabular form:

Assets	Liabilities
	Equity

which involves three sections:

1. **Assets**, as we have seen, can be transformed to positive cash flows, i.e., future payments that the firm will receive. Hence, the asset side of the balance sheet lists what the firm “owns.”
2. **Liabilities**, on the other hand, are related to negative cash flows, i.e., future payments that will have to be covered. Hence, the liability side of the balance sheet lists what the firm “owes.”
3. **Equity** is defined as the difference between the total value of the assets and the total value of the liabilities:

$$\text{Assets} - \text{Liabilities} = \text{Equity}.$$

Equity must be positive. When equity is negative, it means that the assets will not be able to generate sufficient cash flows in order to pay the liabilities, and bankruptcy occurs.

Example 1.3 The balance sheet and financial ratios

Let us consider the extremely simplified and fictional balance sheet of a firm, reported in Table 1.1. On the asset side, we have current assets, which are liquid assets, like cash, or assets that can be converted to cash in the short term, like accounts receivable (money that will be received from customers). Fixed assets are less liquid and can be converted to cash, but not so quickly. In the case of equipment, the value may be questionable, and affected by depreciation and amortization standards, which may be chosen according to tax management

Table 1.1 Fictional balance sheet (in \$ millions) for Example 1.3.

Assets		Liabilities	
Current assets		Current liabilities	
Cash	\$80	Accounts payable	\$300
Accounts receivable	\$120	Long-term debt	\$1800
Fixed assets			
Equipment	\$2500		
Total assets	\$2700	Total liabilities	\$2100
		Total equity:	\$600

policies. The liability side can also be partitioned into short-term liabilities, like accounts payable (money that must be paid to suppliers), and long-term debt (possibly bonds). We may check that the two sides of the balance sheet, total assets and total liabilities plus equity, are matched. If ten million shares are outstanding, the **book value** of each stock share should be

$$\frac{\$600}{10} = \$60.$$

This is the book value of the firm, which need not correspond to the market value. If the market value of each share is \$40, then we say that the **book-to-market** ratio is

$$\frac{\$60}{\$40} = 1.5.$$

A ratio larger than 1 should suggest that the stock share is under-priced.

Based on the balance sheet, different ratios may be computed in order to measure the financial well-being and the solvency of a firm. A natural ratio is

$$\text{Total debt ratio} = \frac{\text{Total liabilities}}{\text{Total assets}} = \frac{\$2100}{\$2700} \approx 0.78.$$

More specific ratios consider only short-term items. In general, we aim at measuring the degree of **leverage** of a firm (or bank), i.e., the ratio of debt to equity.

Another fundamental accounting document, which we shall not discuss in detail, is the **income statement**, which links sales to net income, taking costs and taxes into account. Let us assume that net

income is \$200 (million) for our fictional firm. Net income is used to define important ratios:

- **Return on assets (ROA)**, the ratio of net income to total assets:

$$\frac{\$200}{\$2700} \approx 0.074$$

- **Return on equity (ROE)**, the ratio of net income to total equity:

$$\frac{\$200}{\$600} \approx 0.33$$

- **Earnings per share (EPS)**, the ratio of net income to shares outstanding:

$$\frac{\$200}{10} = \$20$$

- **Price-to-earnings (PE)**, the ratio of price per share to earning per share:

$$\frac{\$40}{\$20} = 2$$

These ratios are also used to classify stock shares as follows:

- **Value stocks** are stocks that look undervalued, but could deliver long-term profits to shareholders. They may feature low PE and price-to-book ratios.
- **Growth stocks**, on the contrary, look overvalued with respect to current market, but they may promise further growth opportunities due to expanding markets, new products, etc. They are generally rather volatile.

Furthermore, some of these ratios may be used in the multifactor models of Chapter 9.

We should always keep in mind that the ratios we have just defined may vary considerably across different industry sectors. Hence, rather than considering their absolute values, we should compare them against those of similar firms. By the same token, depending on the nature of the firm we are considering, the exact kind of items listed in a balance sheet may be very different, financial or nonfinancial, tangible or intangible, fairly easy or very difficult to evaluate, liquid or illiquid, as well as short or long term. It is also important to realize that the cash flows associated with assets and liabilities may be deterministic or stochastic, as the following examples illustrate.

Manufacturing firms. Specialized equipment for production is an asset, but a rather illiquid one, just as account receivables (payments to be received from clients). However, account receivables are usually short-term assets, and in this sense they contribute to firm's liquidity, even though they are not marketable. Assets may also be somewhat intangible, like customer goodwill or the portfolio of knowledge embedded in human resources. On the other hand, money that the firm owes to suppliers (accounts payable) contributes to short-term liabilities. Liabilities also include money that the firm has borrowed from a bank to finance its short-term operations or its long-term research and development programs. Alternatively, a large corporation may issue bonds to finance itself. Note that such corporate bonds are liabilities for the issuer, but they are assets from the viewpoint of bondholders, which may be financial intermediaries or individual investors.

Banks. Assets and liabilities for a bank tend to have a financial nature, but they need not be marketable. One such example is mortgages, unless they are pooled by a securitization process and sold as mortgage-backed securities. It is important to understand how the uncertain balance between assets and liabilities may be a source of risk for banks. Traditional mortgages that the bank has contracted with its clients and kept in its balance sheet are long-term assets, whereas the deposits are short-term liabilities, since the client may withdraw money whenever she feels like it. This maturity mismatch may result in considerable exposure to interest rate risk, since short- and long-term assets or liabilities react in different ways to changes in interest rates.⁸ Banks with a proprietary trading desk may hold any kind of financial asset, including bonds and stock shares. A bank may finance its operations using deposits, but since they result in short-term and uncertain liabilities, they may issue certificates of deposits or bonds, which appear in the liability side of its balance sheet.

Insurance companies. A life insurance company receives periodic payments that may be invested in financial assets, whose cash flows will be used to pay, e.g., pensions and annuities, which appear on the liability side. The financial assets may be more or less risky, just like the liabilities. A life insurer faces longevity risk and, possibly, inflation risk if pension benefits are inflation-indexed. By a similar token, a non-life insurer collects premia from policyholders and is subject to stochastic liabilities related to, e.g., loss of property and car accidents.

These examples just give a vague idea of the variety of assets and liabilities that may appear on balance sheet. The picture is complicated by the fact that the exact way in which items are listed is far from trivial, and it is affected by accounting standards and regulations, having an impact on tax payments. Moreover, there may be little agreement on how assets and liabilities are exactly valued. This results in a possibly remarkable discrepancy between the

⁸We consider interest rate risk management in Chapter 6.

book value, i.e., the value reported in the balance sheet, and the actual market value of an item. This is inevitable, when cash flows are stochastic, requiring a suitable valuation model. Whatever model we choose, model and correlation risk come into play. To understand correlation risk, let us consider an insurance company. If there is a large number of uncorrelated and relatively small risk exposures, like car accidents, the overall value of the liabilities may be fairly predictable. However, considerable risk is faced when insuring properties with large values, or when an unexpected increase in correlation among risks introduces a remarkable amount of volatility.⁹ By the same token, if a balance sheet includes derivatives, which one among the many conflicting valuation models should be used? And how should we estimate their parameters?¹⁰

Given this complexity, the accounting profession has (reasonably) given priority to standardization and consistency, rather than financial accuracy and mathematical sophistication, issuing a set of debatable guidelines and rules. As the reader can imagine, this is beyond the scope of this book, and a thorough discussion of accounting documents like balance sheet and income statement can be found in corporate finance books. Nevertheless, a bit of understanding of the balance sheet is also necessary for anyone interested in quantitative models of financial markets. The two primary assets that we describe in the following, stock shares and bonds, are clearly related to the balance sheet. Whatever assets and liabilities are listed and how exactly, a fundamental principle applies: If a firm is liquidated and closed down, assets are sold, generating funds that are used to pay the outstanding liabilities. If any equity remains, this money is distributed to stockholders (also called shareholders). This is why stock shares are referred to as equity, and stock markets as equity markets: Stock shares represent residual claims on equity.¹¹ Note that creditors, possibly bondholders, have priority over shareholders, and there is a pecking order for creditors as well. Bond indentures describe bond features like collateralization, i.e., if firm's assets are locked as a guarantee against default, and the level of seniority (priority in the pecking order) associated with the bond. Clearly, these features

⁹ An example of unexpected increase in correlation is the increase of defaults on mortgage payments, when a generalized economic crisis leads to an increase in unemployment. In this case, default is not due to strictly individual issues, like illness or delinquency. The same may apply when an increase in the interest rates makes floating-rate mortgages more expensive, making default the only possible choice for some homeowners.

¹⁰ Another issue with derivatives is their exact purpose. In fact, derivatives may be used to manage risk exposures and improve the balance sheet. However, they may also be used for quite risky speculation and, sometimes, drawing the line between the two uses is difficult. A pension fund might use derivatives in a defensive manner, but in order to prevent reckless behavior by fund managers, their use may be prohibited altogether. By the same token, at the time of writing, there is considerable controversy, here in Italy, about how public authorities have used interest rate derivatives in order to manage public debt. Many risk management strategies have backfired, which is always a possibility and, per se, is no evidence of reckless management. The problems are: (a) the appropriateness and size of the exposure that was assumed and (b) the suspiciously high prices that were paid to investment banks.

¹¹ By the way, it should be clear why sovereign governments may issue bonds, but not stock shares.

have an impact on the riskiness and value of bonds. Furthermore, the stylized structure of the balance sheet provides us with a useful representation of trading strategies¹² and an essential link between financial markets and corporate finance.

1.2.2 ASSETS VS. SECURITIES

Sometimes, it may be useful to draw a line separating assets from securities. **Securities** are assets that can be readily purchased or sold on financial markets, like stock shares and bonds; we may also say that securities are tradable or marketable financial assets. On the contrary, a health insurance policy is an asset, but it cannot be sold by its owner and cannot be considered as a security. Note that this does not mean that nontradable assets have no value. We should also note that the line between assets and securities is often not so clear. For instance a mortgage is an asset for a bank, but an illiquid one, as we said. However, pools of mortgages may be transformed into liquid securities by securitization, whereby tradable mortgage-backed securities are created. Other kinds of asset-backed securities (ABS) have been created and traded. By a similar token, a commodity like oil is not, per se, a security, even though it can be traded. The point is that an individual investor cannot really buy and store oil. However, she can take a position related to oil price by using derivatives written on oil and other nonfinancial commodities. Real estate funds have also been created to enable retail investors to take a stake in this family of alternative assets, like residential or commercial real estate. Therefore, in this book, we will not insist too much on the difference, but we will use the term “security” when the liquidity feature of an asset needs to be emphasized.

We will investigate liquid securities in some detail, but we should always keep in mind that liquidity is not only related to the specific kind of assets per se, but to market conditions as well. In conditions of stress, market liquidity may be severely reduced, putting a lot of pressure on market players in need for cash.

Example 1.4 The liquidity trap in thin markets

In a deep and liquid market, a trade has little impact on prices, but markets may get thin and, needless to say, they have a nasty habit of doing so at the least favorable moment. Consider a hedge fund financing the purchase of assets by borrowing money. We will see later that this strategy is called margin trading. In the balance sheet of the hedge fund, the borrowed money contributes to the liability side, whereas the purchased assets are on the asset side. Equity, which is

¹²See Sections 1.4.4 and 1.4.5.

the difference of the two sides of the balance sheet, will float with the value of assets, whereas the liabilities are what they are. Quite often, hedge funds purchase rather illiquid and risky assets, either to earn some additional return, or as a part of a complex trading strategy. In fact, this is why we can redeem shares of a common fund at short notice, but doing so with a hedge fund requires much more time, as complex trading strategies involving illiquid assets are not so easy to unwind.

Under market stress, a flight to quality may occur, whereby market participants sell risky assets in order to rebalance their portfolios toward safer assets, like sovereign bonds of a quite solid country. As a result, asset values may be considerably reduced, eroding equity of hedge funds. The thinner the market, the larger this effect.

Well-intended regulations specify that a minimum safety cash margin must be maintained in order to preserve equity. Hence, when equity is eroded, the fund may be forced to liquidate assets to raise additional cash. But when this happens in bad times, a vicious feedback cycle may arise. We need to sell illiquid assets to raise cash, which in turn leads to further a reduction in the market price of the assets, forcing additional sales. It may even be the case that potential buyers are aware of the state of the matter and have a strong incentive to wait for a further reduction of the price asked by a fund in desperate need of liquidity.

This liquidity trap was a key factor in the famous near-collapse of Long Term Capital Management (LTCM) in 1998. As a consequence of Russian default of bonds, market nervousness ensued, leading to a drop in the market prices of risky securities, with a huge impact on the highly leveraged portfolio of the fund. Similar issues arose in the more recent subprime mortgage crisis: Illiquid assets could not be liquidated because of a market crunch. Thus, investors in need of cash were forced to sell liquid securities, like stock shares, leading to a collapse in equity markets as well.

Example 1.5 Are you on-the-run?

Sometimes, there are slight differences in the liquidity of otherwise equivalent securities. Treasury bonds, i.e., bonds issued by sovereign governments, are issued and sold on markets at regular time intervals in order to finance public spending and debt. The most recently issued bonds are called *on-the-run*, whereas their older relatives are

called *off-the-run*. On-the-run bonds are more actively traded, and liquid, and this has an impact on their price. Some traders may try to take advantage of this price differential by suitable trading strategies, buying the cheaper bonds and short-selling the more expensive ones.

1.2.3 EQUITY

As we have pointed out, stock shares represent residual claims on the equity of a firm, i.e., what remains after liquidating assets and paying liabilities. This is why we use terms like “equity markets” and, as we shall see later, “equity derivatives.” Stock shares are risky assets, as suggested by the randomness in the holding period return of Eq. (1.1). The holding period return as defined there involves only a **capital gain**, i.e., a return related to a price change. However, there is also a possible source of income in the form of dividends distributed to shareholders. If we denote by D the dividend paid during the holding period $(0, T)$, the corresponding holding period return is

$$R(\omega) = \frac{S(T, \omega) + D - S(0)}{S(0)}. \quad (1.4)$$

Dividends may be random or not, depending on the length of the holding period. Dividends are announced with some advance with respect to the ex-dividend date,¹³ but they are uncertain for the not-so-close future. Actually, if the holding period is long enough, the exact timing with which dividends are paid is also relevant, as they may be reinvested in the stock itself or other assets. Thus, to be more precise, we should consider D in Eq. (1.4) as the value projected forward to time T . For instance, if a dividend of €0.60 will be paid in two months and the holding period is six months,

$$D = 0.60 \times e^{r \times 4/12},$$

where r is the (continuously compounded) annual interest rate, which we use to shift the cash flow four months forward.¹⁴ Care must be taken with respect to taxation, as dividend income and capital gains might be taxed in a different way. We should also mention that an important topic in corporate finance is the

¹³To be precise, the ex-dividend date does not necessarily coincide with the date on which a dividend is paid. Since stock shares change hand continuously, a rule must be established to specify who is going to receive the dividend. If we buy the stock share after the ex-dividend date, when the stock share is said to go “ex-dividend,” we are not entitled to receive the dividend, but the previous shareholder is, even if the dividend will be paid later.

¹⁴This operation is the reverse of cash flow discounting, and we will discuss such issues in Chapter 3.

dividend policy, i.e., the strategy by which a firm decides whether earning will be reinvested or distributed in the form of dividends.

When we are stockholders, we are actually the owners of shares of a firm, which is not the case for bondholders. This raises an important issue: Are we responsible for illegal behavior by the board of directors or damage caused by defective products? The answer is no, since stock shares are **limited liability** assets. This is essential, especially for large corporations, in order to enable separation between management and ownership. Apart from legal implications, this feature implies that stock prices cannot be negative and that the worst-case return from holding stock shares is -100% . From a mathematical viewpoint, as we have already mentioned, this also implies that a widely used distribution like the normal, which features an unbounded support, the whole real line $\mathbb{R} = (-\infty, +\infty)$, cannot be a model for stock returns, but an approximation at best.¹⁵

Although this will not play a major role in this book, we must keep in mind that stock shares have not only an economic nature, but a legal one as well. In practice, there may be different kinds of stock shares associated with the same firm, like **common** and **preferred** stock shares. The difference may be in voting rights, which may not be associated with preferred stock shares. On the other hand, preferred stock shares come with the “promise” of a given dividend, whereas common stocks do not have any such guarantee. The holder of a preferred share has priority over holders of common stocks in terms of dividend payments; however, if no dividend is paid, this does not involve any default on the part of the firm. On the contrary, if interest on debt is not paid, a default occurs, with the possibility of the firm being declared bankrupt. This is a relevant consideration when a firm has to decide on the best way to raise capital, by issuing either stock shares or debt. The cost of servicing debt is tax-deductible, which may yield some advantage in terms of taxation. Issuing new stock shares may dilute property, and it may not be taken well by markets, resulting in a sudden drop in the stock price. On the other hand, issuing debt increases the possibility of bankruptcy. This choice of the capital structure is a fundamental topic in corporate finance.

There are other important features of stock shares that are worth mentioning:

- Not all stock shares are publicly traded. Some may be kept under the control of original owners of a firm in order to have the final say in matters of management. Furthermore, not all firms are listed on financial markets, since this requires an expensive process, as some standard requirements must be met in order to be quoted. Private equity funds may be used to invest in privately held firms.
- Unlike other assets, like bonds or options, stock shares do not have a maturity. However, unlike energy in physics, stock shares may be created and destroyed. Sometimes, new equity is floated in order to raise

¹⁵ As we shall see later, the normal distribution may also be unsatisfactory for other reasons, as it is symmetric and thin-tailed.

additional capital. Sometimes, equity disappears when shares are repurchased by the firm itself.¹⁶ It may also be the case that a firm is delisted or acquired by another firm (not to mention the unpleasant event of bankruptcy).

- Other exceptional events that may have a relevant impact on the price of a stock share are:
 - Stock splits: Two or more stock shares of the same firm are created out of a single one. Stock splits may occur when the stock price is quite large. Increasing the number of outstanding stock shares and reducing their price may improve liquidity and lower the bid–ask spread. Sometimes, reverse splits occur, which may require some adjustments, e.g., to deal with owners of an odd number of shares when the reverse split is 1-for-2. A reverse split may occur when the stock price is very low. For instance, a low price may even preclude the listing of a share on a stock market, and a reverse split may be a corrective action to avoid delisting.
 - Spinoffs: A firm is separated in two firms, and two different stock shares are created out of each stock share of the original firm.
 - Mergers and acquisitions: Two firms are merged into a single one, with a corresponding merging of pre-existing stock shares.

Once again, all of these operations have rationales and features that are discussed in detail by books about corporate finance. We observe that their impact on stock prices must be properly accounted for. If a stock share is currently traded at a price of \$100 and a 2-for-1 split occurs, the new resulting price will be something like \$50, which clearly does not imply a return of -50% . Stock market indexes, discussed later, should take all of this into due account. By the same token, derivative contracts must clearly specify how these events are dealt with.¹⁷ A stock split has no effect on the *market capitalization* of a firm, which is given by the total number of shares outstanding, times their market price.

1.2.4 FIXED INCOME

Floating stock shares is one way a firm can raise the capital it needs. An alternative is to borrow money, which does not necessarily mean literally borrowing money from a bank. A common way to raise capital in mature financial markets is issuing a bond. Bonds are also issued by sovereign governments, as well as by local authorities: Examples are US treasury bonds and municipal bonds. A

¹⁶Stock repurchase may have different motivations, as it may be a way to compensate shareholders without issuing dividends, or a way to reduce the number of outstanding shares, when they are deemed to trade at a too low price.

¹⁷For instance, we shall see that a typical call option suffers from a drop in the underlying asset price. Usually, call options are not protected against payment of dividends, but they are against stock splits.

bond is a security that, in its simplest form, may be described by the following main features:

- The **face value** F , also called nominal or par value, which is the amount that the issuer promises to pay back to the bondholder.
- The **maturity** T , i.e., the time at which the face value will be paid back.
- The **coupon rate** c , which is the interest rate applied to the face value to define periodic interest payments that are paid to the bondholder. These payments are called “coupons” for historical reasons, as bonds were physical pieces of paper with coupons that were detached to request payment of periodic interest.

If $c = 0$, i.e., no coupon is paid along the bond life, we have a **zero-coupon** bond, often referred to as a “zero.” If $c > 0$, we have a **coupon-bearing bond**. Usually, coupons are paid twice a year, but different frequencies may be arranged.

Example 1.6 A plain coupon-bearing bond

Let us assume that $F = \$10,000$, $T = 5$, measured in years, and semiannual coupons are paid, with rate $c = 4\%$. Note that coupon rates, like all interest rates, are always quoted *annually*, but should be adjusted to the actual period they refer to. In this case, since frequency is semiannual, the actual coupon rate is 2% for six months. This means that along the bond life there will be ten cash flows to the bondholder. At times $t = k \times 0.5$, $k = 1, 2, \dots, 9$, measured in years, the cash flow will be

$$\frac{c}{2} \times F = \$200,$$

whereas the final cash flow at $T = 5$ includes both the last coupon and the face value, amounting to \$10,200.

The choice between funding alternatives depends on the circumstances. Most firms would not issue a bond for a short-term cash need,¹⁸ whereas for a long-term project, issuing bonds may be a better alternative, at least for a suitably sized firm. Debt securities are liabilities from the viewpoint of the issuing firm, which has an impact on both taxes and the risk of bankruptcy. On the one hand, the cost of servicing debt is tax-deductible; on the other one, however, this increases the risk of default. As we have hinted at before, the choice between issuing debt or equity is affected by a tradeoff related to these and other issues, such as the dilution of property, etc.

¹⁸A possible alternative is issuing commercial paper.

Usually, zeros are short-term bonds, whereas coupons are paid for longer-term maturities. For instance, US treasury bonds may be classified as:

- T-bills, zeros, with maturities up to one year
- T-notes, coupon-bearing, with maturities up to ten years
- T-bonds, coupon-bearing, with longer maturities

Some long-term zeros are in fact traded, but they are often created synthetically by stripping coupons of long-term coupon-bearing bonds. This is an example of a financial engineering practice known as **unbundling cash flows**.¹⁹

Usually, when a bond is issued, the coupon rate more or less reflects the current level of interest rates. If we compare the uncertainty in dividends of a stock share against a fixed-rate coupon bond, i.e., a bond where c is declared and fixed, we may understand why bond markets are referred to as fixed-income markets. Bonds are the basic fixed-income securities, but, as we shall see, this name also refers to quite different securities whose cash flows depend on the level of interest rates. Indeed, the term “fixed-income” is quite a bit misleading. To begin with, we may have bonds whose coupon rate is not fixed, but depends on the time-varying level of interest rates. We refer to these bonds as **floating-rate bonds**, or **floaters**. Other bonds pay coupons affected by other variables, like inflation or even a stock market index (we talk of **linkers**, in such a case). In fact, we use the term “fixed-income markets” to refer to a wide array of securities related to interest rates. They include interest rate derivatives, such as swaps and options, as well as hybrid securities, like convertible and callable bonds, discussed later.

While the cash flows of a fixed-rate bond are supposed to be known with certainty, the bond price itself is affected by the following risk factors:

- **Default risk.** The bond issuer may default on the coupon payments or even on the reimbursement of the face value, totally or partially. In fact, not all bonds are created equal: Collateral guarantees and bond indentures, which may also specify the order in which bondholders are refunded in case of bankruptcy, are relevant. In the event of default, part or all of the face value or coupons may be lost. Debt restructuring may even result in a change of maturity.
- **Inflation risk.** This is relevant for long-term bonds. Some bonds pay real-interest²⁰ coupons, i.e., the coupon rate (or the face value) is adjusted according to inflation.
- **Foreign-exchange risk.** This is obviously relevant if we invest in foreign bonds, which may be denominated in a foreign currency.
- **Interest rate risk.** We will explore the inverse relationship between bond prices and interest rates: When interest rates increase, bond prices

¹⁹See Section 1.2.6.4.

²⁰In Section 3.3, we shall see that a *nominal* interest rate may be eroded by a high inflation rate. The real interest rate is adjusted for inflation and should reflect actual purchasing power.

go down and vice versa. A rather paradoxical result is that a floating-rate bond, with uncertain cash flows, may be less risky than a fixed-rate bond.²¹

Interest rate risk is relevant when we do not plan to hold a bond until maturity. If we sell the bond, there may be considerable uncertainty about its future price. In fact, defining a holding period return for bonds is actually complicated, since we should also specify how coupons are used exactly. In an asset–liability management problem, they might just be used to pay a stream of liabilities. If they are reinvested, there is uncertainty about the future interest rates at which this will be done, resulting in **reinvestment risk**.

1.2.5 FOREX MARKETS

Another huge market is the foreign exchange market (FOREX market for short), where currencies are exchanged. The involved risk factor is the exchange rate between pairs of currencies, which is relevant also for international equity and fixed-income portfolios. Nonfinancial firms are also subject to currency exchange variability, which explains the number of FOREX derivatives available. FOREX markets are also the terrain of plenty of speculative short-term trading.

The institutional arrangements behind FOREX markets are not trivial but, given their limited role in this book, we leave the related issues to the references. There is one, somewhat annoying, detail that we have to mention. If we read a stock market quote and the price of a stock share is, say, \$12, we interpret this as the price of one share. Dimensionally, the quote is dollars per share. Hence, if we buy 3 shares at that price, from a dimensional viewpoint we spend

$$12 \frac{\text{dollar}}{\text{share}} \times 3 \text{ shares} = 36 \text{ dollars}.$$

We would probably never think of a quote in terms how many shares we can buy with 1 dollar, although sometimes, given an available budget, we must find out how many shares we may afford to buy; hence, we would not consider share per dollar as a sensible unit. When we buy commodities, the specific measurement unit plays a more explicit role. We might buy a certain kind of vegetables for, say, €3.2 per kilogram. Considering dimensions (i.e., units of measurement), if we buy 2 kg, we pay

$$3.2 \frac{\text{euro}}{\text{kg}} \times 2 \text{ kg} = 6.4 \text{ euro}.$$

In this case, too, measurement units have a straightforward interpretation when figuring out prices and cash flows.

Now, if the exchange rate between USD and EUR is quoted as

$$\text{EUR/USD} = 1.1166,$$

²¹ See Section 3.5.6.

what does it mean? At the time of writing, the value of €1 is larger than the value of \$1, and indeed the above ratio tells that we may buy 1.12 dollars with 1 euro, allowing for some rounding and neglecting transaction costs. Alternatively, we may say that the price of 1 euro is 1.12 dollars. This depends on the perspective we take, and since FOREX quotes always involve two monetary units, some clarification is in order.

Quoting exchange rates is about stating an equivalence between two amounts denominated in different countries. We might write something like

$$\text{EUR } 1 = \text{USD } 1.1166, \quad (1.5)$$

or, equivalently

$$\text{USD } 1 = \text{EUR } 1/1.1166 = \text{EUR } 0.8956. \quad (1.6)$$

Let us focus on the first case, Eq. (1.5). We say that

- EUR is the **base currency**, and we consider EUR 1 as a *fixed* number
- USD is the **quoted currency**, and we consider USD 1.1166 as a *variable* number

In a currency pair, written as EUR/USD, the currency to the left of the slash is the base currency, and the currency to the right is the quoted currency.

Depending on which currency is considered as domestic, there are two types of quotes:

- In **direct quotation** the domestic currency is the quoted currency, i.e., a variable amount of the domestic currency is quoted against a fixed amount of foreign currency. This kind of quotation is also called *normal* or *uncertain for certain*.
- In **indirect quotation** the domestic currency is the base currency, i.e., a fixed amount of domestic currency is quoted against a variable amount of foreign currency. This quotation is also called *reciprocal* or *certain for uncertain*.

For instance, a Eurozone bank quoting as in Eq. (1.5) would use an *indirect* quote. A difficulty with FOREX markets is that different quotations are used on different markets. The choice may depend on the following:

- A matter of perspective, i.e., what our domestic currency is
- A matter of convenience: for instance, a quote like EUR/JPY 115.261, stating that 1 euro corresponds to 115.261 Japanese yen is convenient, whereas the reciprocal would be less convenient
- A matter of priority, as the choice is influenced by the fact that there are some “major” currencies which are more widely traded than other ones
- A matter of local conventions, since, for instance, the conventions in the UK are different from the conventions in the USA

A further complication is introduced by bid–ask spreads. The quote in Eq. (1.5) would more likely read as

$$\text{EUR/USD } 1.1165/67,$$

stating that the bank is bidding 1.1165 dollars to buy 1 euro, and asking 1.1167 dollars to sell 1 euro. Sometimes, the “three Bs rule” is invoked:

The market maker Buys the Base currency at the Bid (low) price.

Indeed, the Eurozone bank would buy 1 euro for 1.1165 dollars. Needless to say, a Eurozone bank will also quote an exchange rate like USD/CHF, which does *not* involve any domestic currency, to add to the confusion. In this case, we have to come up with a *cross-rate*, starting with a mix of direct or indirect quotes.

An indirect quote may also be considered as a “quantity quotation,” in the sense that it gives the quantity of foreign currency needed to buy one unit of the domestic currency. The direct quote may be considered as a “price quotation,” i.e., the price of one unit of foreign currency in terms of the domestic currency. In this book, we will deal extensively with derivative pricing, including forward/futures contracts on currencies. For the sake of uniformity, we will always interpret ratios as *prices*, just as we do in commodity prices, rather than currency pairs. Hence, assuming that we are US investors, we would consider a price like

$$1.1166 \text{ dollars per euro,}$$

stating that the price of €1 is \$1.1166, so that if we want to buy €200.00, we have to pay

$$\frac{\$}{\text{€}} 1.1166 \times \text{€}200.00 = \$223.32.$$

Note that this is the contrary with respect to a base/quoted currency pair. A Eurozone investor would consider that as a price at which a euro is sold. We will neglect bid–ask spreads, and no ambiguity should arise.

1.2.6 DERIVATIVES

Stock shares and bonds are, in a sense, primary assets. They need not be primary *risk factors*, as we may build a model relating their prices (or returns) to underlying risk factors like inflation, oil price, and interest rates.²² However, the relationship between risk factors and stock share prices/returns is represented by a mathematical model, possibly estimated by statistical methods, on which there may be no general agreement.

An incredibly large class of assets has been created on top of primary assets, collectively known as **derivatives**. A derivative security is a financial asset

²²See Chapter 9 on factor models.

deriving its value from some other variable, *by an explicit formula that is written in a contract*. For instance, if S_t is the random price of an asset at time t ,²³ a typical derivative features a **payoff** $f(S_T)$ at a well-defined time T , the **maturity** of the derivative, for some well-defined function $f(\cdot)$ of the underlying asset price at maturity. We insist again on the fact that the function $f(\cdot)$ is explicitly written in the contract. More complex derivatives feature a payoff depending on the whole price path until maturity.

If the derivative is written on a stock share or a bond with price S_t , we say that the latter is the **underlying asset**. However, S_t can be something else, not necessarily a primary asset. For instance, we may consider:

- The price of a nonfinancial asset, e.g., a commodity like gold or oil, provided that a well-defined price is quoted on exchanges
- A risk factor that is not the price of a traded asset, but a financially relevant variable nevertheless, like an interest rate or a market index,²⁴ or even an elusive variable like volatility
- A risk factor that is not related to financial assets or prices, as in weather derivatives
- The price of another derivative, as in compound options or swaptions

We observe that there is room for a considerable variety of derivatives, as they may depend on a combination of underlying variables, and they also differ in terms of the function defining the payoff.

Derivatives may be used for opposite purposes, namely, risk hedging and speculation. Originally, they were meant to be risk transfer mechanisms and have quite a long history, definitely predating the development of quantitative finance. However, they have become quite controversial assets, as the volume of derivatives outstanding is so huge that it is often larger than the market of the underlying primary assets.²⁵

There are different issues related to derivatives, which may be tackled by quantitative finance models:

- Pricing²⁶: What is the fair value of a derivative, and how is it related to underlying risk factors?

²³Depending on convenience, we will write S_t or $S(t)$; we will not stick to a single notation, as no ambiguity actually arises.

²⁴See Section 1.5.

²⁵Statistics published by the Bank for International Settlements in 2014 estimated a total notional amount of OTC derivatives of about \$630 trillion. Interest rate swaps accounted for \$381 trillion. These numbers are impressive but misleading, since the notional amount of an interest rate swap, as we shall see, overestimates the value of the derivative and the actual cash flows that will occur. Nevertheless, there is no doubt that this is a huge market.

²⁶As is common in the literature, we will use the term **pricing**, even though **valuation** would be more correct. The fair value of the derivative is only a component of the actual price asked by a bank issuing derivatives, since this will include a profit margin and some buffer against residual risk that cannot be hedged away in real life.

- **Hedging:** If a market player like an investment bank writes, i.e., creates a derivative, how can it manage the ensuing risk?
- **Portfolio management:** How can a derivative be used to change the characteristics of a portfolio, increasing or reducing its exposure to selected risk factors?

A wide array of derivatives is traded, but here we just want to introduce the three basic families: Forward contracts, futures contracts, and vanilla options.

We should also notice that, beside quantitative issues, there is a host of regulatory and legal issues related to derivatives, not to mention how should they be accounted for in financial statements of banks and firms. These are, however, outside the scope of this book.

1.2.6.1 Forward contracts

A **forward contract** is an arrangement between two counterparties, which at time t_0 agree to buy and sell, respectively, an asset at a prespecified forward price $F(t_0, T)$ at a later date T , the maturity of the contract. The part agreeing to buy the asset is said to hold the **long position** in the contract, whereas the part agreeing to sell is said to hold the **short position** in the contract. Note that the contract is symmetric, in the sense that *both* parties are forced to comply with what they have agreed.

The current **spot price** of the underlying asset when the contract is written, denoted by $S(t_0)$, is a known number, whereas the spot price $S(T)$ at maturity is uncertain. At time t_0 the **forward price** $F(t_0, T)$ is established once for all. During the time interval (t_0, T) the spot price $S(t)$ will change randomly. By the same token, the forward price $F(t, T)$, observed at time t , for delivery at time T , $t \leq T$, will change as well. This is the forward price for new forward contracts written at a later time $t > t_0$, but the forward price in previously arranged contracts will *not* change. As we shall see, the value of a contract will depend on the difference between the fixed $F(t_0, T)$ and the uncertain $F(t, T)$ along the life of the contract.

Both the spot price $S(t)$ and the forward price $F(t, T)$ are stochastic processes, which are arguably correlated in some way. We should find a way to model the relationship between spot and forward price, which may be a non-trivial task. However, we can immediately see that the following spot–forward convergence condition must hold at maturity:

$$S(T) = F(T, T). \quad (1.7)$$

In fact, $F(T, T)$ is the forward price for an *immediate* delivery at time $t = T$, and it must be the same as the spot price, otherwise two prices would be quoted for the same item.²⁷

²⁷Formally, this is an example of the law of one price, which is a consequence of the no-arbitrage principle that we investigate in Section 2.3.




FIGURE 1.4 Payoff from forward contracts: (a) long position, (b) short position.

The value of the contract arises from the payoff that will result at maturity. The payoff for the long position is

$$S(T) - F(t_0, T).$$

To see why, observe that, if $S(T) > F(t_0, T)$, the long position may buy the underlying asset at the delivery price $F(t_0, T)$ and sell it at the current spot price $S(T)$ at maturity, earning a profit. Note, however, that the payoff may well be negative, since the long position has to buy at the delivery price, even when this is larger than the prevailing spot price. Going the other way around, the payoff for the short position is

$$F(t_0, T) - S(T).$$

Clearly, the sum of the two payoffs is zero: The profit for the long position is just the loss for the short position, and vice versa (this is a zero-sum bet, in some sense). The payoffs are illustrated in the two diagrams of Fig. 1.4, for the long and short positions, respectively. The long position benefits from an increase in the spot price, whereas the short position benefits from a decrease in the spot price. As we shall see in later chapters, it is common jargon to say that an investor is “long a variable” if she gains from an increase in the variable, and is “short a variable” if she gains from a decrease in the variable. The variable may be the price of an asset, an interest rate, and whatnot. When the underlying variable is not really the price of a deliverable asset, the contract is settled in cash, i.e., an amount corresponding to the payoff is exchanged (when our payoff is negative, it means that we owe money to our counterparty). In Chapter 2, we will see that, using no-arbitrage pricing principles, the forward delivery price is selected in such a way that the value of the contract is initially zero for both parties. Thus, payoff and profit coincide, as nothing is paid when entering the contract. After inception of the contract, the spot price $S(t)$ and the forward price $F(t, T)$ will change, and this will affect the value of the contract, which may drift away from zero.

The fact that the value of a forward contract is initially zero explains why it may be so attractive for a speculator, at least in principle. In the case of speculation, the profit from a successful trade in the underlying asset is limited

by the fact that we have to actually *buy* (or short-sell) it on the spot market, which may be limited by the available budget. In principle, nothing is needed to enter into a forward contract, as the forward price is determined in such a way that the value of contract at inception is zero. Thus, the return is not really defined, as the denominator is zero! In practice, it may be the case that some collateral has to be posted, and in any case there are transaction costs. Nevertheless, we will see how a considerable leverage may be obtained with derivatives in general, magnifying both profit and loss opportunities.

The other side of the coin is that a forward contract may be used to eliminate or reduce risk, i.e., for hedging purposes. Assume that we will have to buy the underlying asset at time T in the future. Since $S(T)$ is uncertain, we face some risk, but by entering into a *long* position, we will be able to buy at $F(t_0, T)$ no matter what, eliminating uncertainty altogether. If the contract is settled in cash rather than by buying the underlying asset, the net cash flow at time T will be

$$\underbrace{S(T) - F(t_0, T)}_{\text{payoff}} \quad \underbrace{-S(T)}_{\text{purchase cost}} = -F(t_0, T), \quad (1.8)$$

which is negative, since we are buying the asset, and is equivalent to a contract for physical delivery of the asset. A typical case in which derivatives are settled in cash is when the underlying is a nontradable asset like a stock market index. In other cases, physical delivery would be possible in principle, but it might be avoided because of transportation costs and the like.

Example 1.7 A long hedge

Suppose that in six months we will need 500 ounces of gold, and that the current (time $t = 0$) forward price for delivery in 0.5 years (six months) is

$$F(0, 0.5) = 1250 \text{ \$/ounce.}$$

Then, we may enter into a long position for 500 ounces to lock that price. As a practical remark, we shall see that real-life contracts may be given for standardized sizes, such as, e.g., 100 ounces. If the contract is settled by physical delivery, we shall buy gold at 1250 dollars per ounce, no matter what. The corresponding (negative) cash flow is

$$-1250 \text{ \$/ounce} \times 500 \text{ ounces} = -\$625,000.$$

If the contract is settled in cash, and the spot price at maturity turns out to be 1150 \$/ounce, our cash flow will be

$$[(1150 - 1250) - 1150] \text{ \$/ounce} \times 500 \text{ ounces} = -\$625,000,$$

the same as before. Note that, in this case, we buy at a cheaper spot price, but this is compensated by a loss on the long forward position.

If we have to sell the underlying asset, we should enter into a *short* position, which just implies a change in sign in Eq. (1.8).

A perfect hedge results if a forward contract matching both the desired maturity and the underlying asset, as well as the contract size, can be agreed. By the way, note that “perfect hedge” means that risk is completely eliminated, not that the outcome is necessarily a pleasing one. If we take the long position, as in Example 1.7, we will regret our decision if the spot price at maturity turns out to be lower than the delivery price. By the same token, a short position is not a nice place to be if the underlying spot price increases. Still, risk management should be assessed *a priori*, i.e., *ex-ante*, not *ex-post*. The main feature of forward contracts is that they are actually a private arrangement between the two counterparties, typically a firm and an investment bank. Forward contracts are not securities freely traded on regulated exchanges, but rather an OTC agreement. This implies both advantages and disadvantages. On the positive side, the details of an OTC contract may be tailored according to quite specific needs. On the negative side:

- Since there is no quoted price, which is driven by demand and offer, pricing a specific contract may be troublesome. Hence, a firm in need for a hedge might adopt a strategy of competitive pricing, which means asking around for multiple quotes to compare them. A possibly better alternative is to establish long-term relationships with a single, trustworthy bank.
- The contract is not standardized, hence it is not liquid. Unwinding the position may be difficult if the hedging needs change. This typically implies assessing the value of the contract and closing it before maturity by a cash settlement. Note that this is the result of a negotiation process, possibly implying the valuation of an illiquid contract, and not the immediate sale of a security on regulated and liquid markets.
- A further issue with forward contracts is counterparty risk. There is only one cash flow, at maturity, possibly a huge one. Imagine that we hold a short position in a forward contract written on an asset whose price is dropping dramatically. We are about to collect a remarkable payoff, but what if the long position walks away? In fact, only creditworthy firms are accepted as partners in a forward agreement, but counterparty risk is not completely eliminated.

The solution to liquidity and counterparty risk issues is represented by futures contracts, which are the exchange-traded equivalent of forward contracts.

1.2.6.2 Futures contracts

Futures contracts are quite similar to forward contracts, in the sense that the delivery of an underlying asset or commodity is arranged for a future date, at

a prespecified **futures price**²⁸ $F(t, T)$ that is continuously quoted on regulated exchanges. Futures contracts have specific features aimed at easing the difficulties with forward contracts, namely, liquidity and counterparty risk:

- Standardization, to improve liquidity
- Daily marking-to-market through a clearinghouse, to ease counterparty risk

In order to improve liquidity, futures contracts are standardized. This means that the range of available underlying assets and delivery dates is limited and cannot be arranged to perfectly suit very specific needs. For instance, certain contracts are only available for quarterly delivery, i.e., maturing at four months per year. This makes the use of futures contracts in hedging more difficult, but it results in a deeper market, where it is easy to buy and sell futures contracts. Furthermore, a liquid market is less subject to manipulation and cornering.

■ Example 1.8 Cornering in futures markets

Cornering is an illegal practice, whereby speculators accumulate a significant amount of the underlying asset. When maturity is approached, the short positions will be forced to buy the asset at large prices to honor their contracts, if the supply is limited. To circumvent this difficulty, contracts should be arranged only for underlying assets with a sufficiently deep market, or alternatively a range of underlying assets may be eligible for delivery, rather than a single one. For instance, in futures contracts on bonds, a whole range of bonds may be delivered, not only a specific one. Clear rules define the equivalence among similar, but not identical, bonds and the coefficients by which the delivery price is modified if necessary. For instance, bonds with comparable maturities, but different coupon rates may be included in the range for acceptable delivery.

The two essential features of futures contracts aimed at easing counterparty risk are:

1. The existence of a **clearinghouse**. The clearinghouse consists of a group of solid financial institutions, and it steps between the long and the short positions. The institutional arrangement is depicted in Fig. 1.5. Actually, if we hold the long position, we do not really “see” any corresponding short position in the contract. We only deal with the clearinghouse, which assumes the counterparty risk.

²⁸Please note the essential difference between the *future* spot price, which is uncertain, and the *futures price* associated with a derivative contract.




FIGURE 1.5 The institutional arrangement of a futures contract. The clearinghouse manages the margin accounts of the long and short positions.

2. The contracts are **marked to market** daily. This means that, rather than settling the contract at maturity, daily cash flows are exchanged at the end of every trading day. Indeed, both long and short positions are required to post some margin, in the form of cash or some collateral, on an account managed by the clearinghouse. If the futures price moves unfavorably, we will lose some amount of money immediately, rather than at maturity. The loss is sustained by the margin account, where daily profits are also collected in the case of a favorable movement. There is a minimum amount that must be maintained on the margin account, the **maintenance margin**. If the account level falls below the **maintenance margin**, a **margin call** is issued. Failure to comply with the margin call by posting more cash or collateral on the margin account has the consequence that our contract is immediately closed out and assumed by the clearinghouse.

We should note that the actual exposure of the clearinghouse is related to *net* position, balancing long and short positions. One proof that the mechanism does work occurred on October 19th, 1987, a day remembered as the Black Monday of 1987, when a loss in excess of 20% in the Dow Jones index occurred. The S&P500 index sustained a similar drop, with a corresponding shock on index futures. Indeed, some brokers who were members of the clearinghouse went bankrupt on that day, but the clearinghouse survived and all contracts were honored.

We will analyze later the full details of futures contracts, as well as their use for hedging and speculation. For now, we just clarify the mechanics of daily marking-to-market.²⁹ Imagine that, at time t_0 , an arbitrary moment within a trading day, we enter into a long position in a futures contract at price $F(t_0, T)$. Say that, at the end of the day, corresponding to time t_1 , when the futures prices are settled and marking-to-market takes place, the settlement price is $F(t_1, T)$. The cash flow for the long position at the end of the first day is, for each contract,

$$F(t_1, T) - F(t_0, T),$$

which is positive if there is an increase in the futures price. The corresponding cash flow for the short position is $F(t_0, T) - F(t_1, T)$. In general, if the settlement price at the end of day t_k is larger than the corresponding price of the

²⁹The picture is a bit simplified here.

previous day, i.e., if

$$F(t_k, T) > F(t_{k-1}, T),$$

money is drawn from the margin account of the short position and deposited into the margin account of the long position, and vice versa if there is a decrease in the futures price. The marking-to-market mechanism generates a series of daily cash flows for the long position:

$$\begin{aligned} &F(t_1, T) - F(t_0, T), \\ &F(t_2, T) - F(t_1, T), \\ &F(t_3, T) - F(t_2, T), \\ &\vdots, \\ &F(t_m, T) - F(t_{m-1}, T), \end{aligned}$$

where $t_m \equiv T$. The last cash flow may also be expressed as

$$S(T) - F(t_{m-1}, T),$$

since the futures price at maturity, $F(T, T)$, converges to the spot price. If we sum these cash flows, we obtain a telescoping sum:

$$\begin{aligned} \sum_{i=1}^m [F(t_i, T) - F(t_{i-1}, T)] &= F(t_m, T) - F(t_0, T) \\ &= S(T) - F(t_0, T). \end{aligned} \quad (1.9)$$

Thus, the net sum of cash flows looks just like the payoff from a forward contract. A similar expression, with a change in sign, applies to the short position. Now, in the light of this result, one could wonder whether there is a significant difference between forward and futures contracts. Indeed, there is a subtle but important difference between the two: The daily cash flows may be reinvested immediately at some interest rate, when positive. Negative cash flows, i.e., losses, may also be financed at some interest rate. We will prove in Section 12.2 that, if interest rates are deterministic, the forward and the futures price are the same. However, if the interest rate moves randomly, this will have an effect, especially if there is a definite correlation between futures prices and interest rates. This is especially the case with interest rate futures. Thus, forward and futures prices need not be identical.

Liquidity has another, possibly surprising, effect. As a general rule, futures contracts do *not* result in the actual delivery of the underlying asset, and most futures contracts are closed before maturity. To close a futures contract, all we have to do is entering into an offsetting position: A long position is closed by entering into an equivalent short position, and vice versa. This feature is essential both for hedgers and speculators, who do not really want to buy the underlying asset, especially if the price of the underlying asset is only a proxy for the actual risk factor that they are exposed to. For instance, a firm that is

Table 1.2 An illustration of the mechanics of futures markets. All data are in \$.

Day	Trade price	Settlement price	Daily gain	Cumulative gain	Account balance	Margin call
1	1350				16,000	
1		1346	−800	−800	15,200	
2		1330	−3200	−4000	12,000	
3		1334	800	−3200	12,800	
4		1315	−3800	−7000	9000	1000
5		1304	−2200	−9200	7800	2200
6		1320	3200	−6000	13,200	
7		1330	2000	−4000	15,200	
8		1328	−400	−4400	14,800	
9	1338		2000	−2400	16,800	

exposed to risk factors related to energy or transportation costs may consider using oil futures as a suitable hedging instrument, but they would certainly not be interested in the actual trade of oil.

Example 1.9 Mechanics of futures markets

Table 1.2 illustrates a possible scenario in a trade on gold futures. On day 1, when the gold futures price is \$1350 per ounce, we enter a long position for two contracts, whose unit size is 100 ounces (hence, each contract specifies the purchase of 100 ounces of gold at a total price of \$135,000). The initial margin required by the broker is \$8000 per contract, hence, we have to deposit \$16,000 on the margin account immediately. The maintenance margin is \$5000 per contract. At end of day 1, the futures is settled at \$1346. Hence, we have a cash flow

$$$(1346 - 1350) \times 200 = -\$800$$

which is actually a loss, as the futures price declined and we hold a long position. In Table 1.2, we list the settlement price for a sequence of days, resulting in daily gains, which are cumulated. The margin account falls below the maintenance margin at the end of day 4. After marking-to-market, the margin account balance is only \$9000, and \$1000 have to be posted in order to restore the maintenance margin. We get another margin call after the settlement of the next day. At some time during day 9, when the futures price is \$1338, we close the contract, with a total loss of \$2400.

1.2.6.3 Vanilla options

Options, like forward and futures contracts, concern buying or selling an asset in the future at a predetermined price. However, options are more complicated contracts, since they are asymmetric. In forward and futures contracts, the long and the short position have symmetric obligations, in the sense that both of them are forced to buy and sell, respectively, the underlying asset at the agreed price, whether they like it or not. This results in linear payoff functions, and there is a unique given price such that the value of the contract is zero at its inception. On the contrary, options feature nonlinear, possibly complicated payoffs. Furthermore, options involve two quite different roles, the option **writer** and the option **holder**, making the contract asymmetric. The option writer is the counterparty originally creating the option, which is sold to the holder. To get the point, let us focus on the simplest family of options, namely, **vanilla** options. Two kinds of vanilla options are traded, call and put options.

- In a **call option**, the option holder has the right, *but not the obligation*, to buy the underlying asset from the option writer, in the future, at a fixed price K called the **strike price**.
- In a **put option**, the option holder has the right, *but not the obligation*, to sell the underlying asset to the option writer, in the future, at a fixed strike price K .

We immediately notice the asymmetric nature of options: The holder has the right to a choice, and the option writer will have to comply, no matter what. The writer of a call option will be forced to sell the asset if the holder exercises the call option, and the writer of a put option will be forced to buy the asset if the holder exercises the put option. This immediately suggests that: (a) the payoff will be nonlinear, (b) the option writer should be compensated for this obligation, and (c) the option will have a positive value at its inception, unlike linear contracts. In the case of options, the jargon is misleadingly different from the case of futures: The option writer is said to hold the **short position** in the contract, whereas the option holder holds the **long position**. Since the option can be a call or a put, in this case the terminology does not refer to who buys or sells the underlying asset. The long position should be understood as the side of the contract that profits from an increase in the value of some asset. The long position in a futures profits from an increase in the futures price, and the long position in an option profits from an increase in the option value, for both call and put options. The short position, on the contrary, profits from a drop in the futures price or in the option value. Clearly, the option writer earns a profit if the option expire worthless, without being exercised by the holder, who paid the option premium. As with forwards/futures, the contract can be settled in cash, rather than by actual delivery of the underlying asset, if this is not tradable, or it is not convenient to do so.

If the option can be exercised only at a prespecified time T , the option **maturity**, the option is said to be **European-style**. If the option can be exercised at any time before and including a time T , which in this case is an expiration date,




FIGURE 1.6 Payoff (continuous line) and profit (dotted line) from long positions in options: (a) call option, (b) put option.

rather than a maturity, the option is said to be **American-style**. The payoff of a European-style call option, from the holder viewpoint, is

$$\max\{S(T) - K, 0\}.$$

To see why, consider that the holder will exercise only if it is convenient to do so, which is the case if the spot price $S(T)$ at maturity is larger than the strike price K . In such a case, the holder may buy the asset at K from the option writer and sell it immediately at $S(T)$ on the spot market. By the same token, the payoff of the put option, from the holder viewpoint, is

$$\max\{K - S(T), 0\}.$$

If $K > S(T)$, the option holder may buy the asset on the spot market at $S(T)$, and force the option writer to take delivery at the strike price K . The payoff and profit to the long position for a call and a put option, respectively, are depicted in the diagrams of Fig. 1.6. We immediately observe that the payoffs are nonlinear (piecewise linear, to be precise). Furthermore, payoff and profit are not the same thing, unlike the case of linear contracts with initial zero value. Since the payoff cannot be negative, it must be the case that an option has some positive value at time $t = t_0$, when the option is written, which is the fair price that the writer should ask.³⁰ Thus, the profit to holders is the payoff shifted down by the option price. While there is only one “right” forward/futures price, such that the initial value of the contract is zero, options with different strike prices are traded. We should expect that the price of a call option, all other factors being equal, is a decreasing function of the strike price, whereas the price of a put option is an increasing function of the strike price.

The diagrams for the short position are just the diagrams of Fig. 1.6 turned upside down, as shown in Fig. 1.7. The option writer is compensated by earning

³⁰We stress again that we confuse “value” and “price.” Option pricing models, as we shall see, yield a *fair value*. The actual price will account for profit and some additional fudge against the risk born by the writer.




FIGURE 1.7 Payoff (continuous line) and profit (dotted line) from short positions in options: (a) call option, (b) put option.

the option price (or premium), but it is important to realize a key difference between the two roles. If the option expires worthless, the holder will lose the whole option premium, but this is the worst that can happen. Figure 1.7(a) shows that there is no bound on the potential loss for a call writer. Thus, two essential tasks in quantitative finance are finding the fair value of options and devising ways to hedge the risk of writing options. A significant portion of this book is devoted to these two problems, and we will find out that they are tightly linked. We will also see that pricing American-style derivatives is, as a general rule, much more complicated. To see why, consider the case of an American-style put option if $S(t) < K$ at $t < T$, before maturity. The option payoff, if the option is exercised early at time t , is the same as the European-style option, with $S(T)$ replaced by $S(t)$. Hence, the option holder could earn a positive payoff, $K - S(t)$, by exercising early, but is this really an optimal choice? Should the option holder exercise immediately, or wait for a better opportunity? The answer is not really trivial, as it implies the solution of specific kind of dynamic stochastic optimization problem, an optimal stopping problem.

We observe that the option payoffs for call and put options only depend on the value of the underlying asset at maturity (or the early exercise date for American-style options). The payoff is a simple piecewise linear function that does not depend on the whole history of the underlying asset price. This is why these simple options are called *vanilla*.³¹ Vanilla options are commonly traded on regulated exchanges, but several OTC variants, involving multiple assets and more complicated payoff functions are commonly engineered. These options are often called **exotic** options.

Just like futures, options may be used for both hedging and speculation purposes. Let us illustrate these uses by two simple examples.

³¹Vanilla is the most basic ice cream flavor.

■ Example 1.10 A protective put

Let us consider a protective put strategy. We hold an asset, with value $S_0 = S(t_0)$, but we are concerned with a possible loss over the holding period $[t_0, T]$. One way to hedge risk is buying a put option with strike K . Then, the overall portfolio value at maturity is the sum of the asset value and the option value,

$$S_T + \max\{K - S_T, 0\} = \max\{K, S_T\}.$$

If we look at the total payoff, it seems that the larger the strike, the better. Clearly, this is too good to be true. Indeed, we should not forget that the protection from the put option does not come for free, and it is a safe guess that a put option with a larger strike price will be more expensive, too. On the contrary, hedging with forward or futures contracts can be achieved at no initial cost. However, we give up the whole upside potential (if S_T grows), whereas this is partially retained by hedging with options.

■ Example 1.11 A bullish speculation

The current price of an asset is $S_0 = \$100$, and we have a strong belief that it will rise in the near future. One possible strategy is simply to buy the asset. If we are right and, say, the asset price at some later time T turns out to be $S_T = \$120$, the holding period return is

$$\frac{120 - 100}{100} = 20\%.$$

Now let us assume that a call option with strike price $K = \$100$ costs \$5. If we buy the call option, the return in the above scenario is a stellar

$$\frac{\max\{120 - 100, 0\} - 5}{5} = \frac{15}{5} = 300\%.$$

Clearly, there must be some other side of the coin. To get a feeling, let us assume that we are wrong and the underlying asset price goes down by 1%. The percentage loss, if we invest in the asset itself, will be a not too painful 1%: We may be fairly disappointed, but this is a loss we may well live with. However, the call option return is

$$\frac{\max\{99 - 100, 0\} - 5}{5} = \frac{-5}{5} = -100\%,$$

since the option expires worthless and we lose the whole premium.

1.2.6.4 Hybrid securities, bundling/unbundling, and securitization

So far, we have considered simple assets like stock shares, plain bonds, and vanilla options. We have hinted at the possibility of creating more complex assets, like exotic options featuring different payoff structures. One such example is an **Asian option**, whose payoff depends on some form of average. The most natural Asian option involves an arithmetic average over time of the price of a single underlying asset, like

$$\max \left\{ \frac{1}{N} \sum_{k=1}^N S(t_k) - K, 0 \right\},$$

where $t_k, k = 1, \dots, N$, is a sequence of sampling instants. Such options are usually not traded on regulated exchanges, but sold OTC. It may sound surprising, but we may find both European- and American-style Asian options. The point is that *Asian* refers to the form of the payoff, whereas the other labels refer to the possibility of early exercise.

By assembling or disassembling assets and cash flows, a whole world of possibly quite complex assets can be created by financial engineering. The building blocks are often stock shares, bonds, and options, and the basic procedures include:

- Cash flow bundling and unbundling
- Addition of option-like features to traditional assets
- Securitization

Let us illustrate the idea with a few concrete examples.

Convertible bonds. A convertible bond is a corporate bond with an optional component: The holder has the right to exercise an option to transform it into a prespecified number of stock shares of the same firm. We may regard this kind of asset as a security bundling a bond and a sort of call option on a stock share. To be precise, the bundled derivative is often not really an option, but rather a **warrant**. The difference is that when a warrant is exercised, a brand new set of shares is created, diluting equity. Convertible bonds may be appealing to issuers as a way to raise capital when the stock share price is perceived by the management as unjustifiably low. Given the embedded option, the price of the bond will be higher than otherwise. If the stock share price rises, new stock shares will be issued and the company will stop servicing debt. Otherwise, the firm will be able to deduct the cost of debt servicing from profit, with a tax advantage. Convertible bonds may be appealing to investors as well, when assessing the company risk is difficult, and they offer upside potential if the firm grows.³²

³²Convertibles may also be interesting for sophisticated investors looking for arbitrage opportunities, which we introduce later. See, e.g., [4].

Callable bonds. A callable bond is a bond that may be bought back by the issuer, at a given price, subject to certain limitations. For instance, a bond may be declared noncallable for a given number of years after its issuance. When an investor buys a callable bond, she is essentially selling a call option back to the bond issuer. Hence, all other factors being equal, a callable bond is cheaper than a noncallable one. The bond issuer will find the call opportunity convenient if there is a drop in interest rates, since it may refinance debt by issuing brand new bonds with a reduced coupon rate. This is bad news for the bondholder, as she will be subject to reinvestment risk: She is forced to get her capital back just when the bond value is increasing because interest rates are dropping,³³ and she will have to reinvest in new bonds with lower coupons (or old bonds with higher price and corresponding lower yield).

Structured bonds. A structured bond typically offers a coupon that is not linked to interest rates, but rather to another index, like a stock market index. Even if the index return turns out to be negative, the repayment of face value of the bond is guaranteed. This shows that a structured bond includes an option element. Indeed, structured bonds were also used to circumvent regulations forbidding mutual/pension funds from investing in derivatives.

Example 1.12 A structured bond

A rather fancy, but real-life example of a structured bond is the following:

- Bond maturity is four years.
- At maturity, the payment of the face value is guaranteed, plus a single coupon; the coupon, too, will be paid at maturity, and no periodic coupon will be paid.
- The coupon is linked to the monthly average value of a basket of ten stock shares in the telecommunication industry; since maturity is 4 years, 48 monthly observations of ten stock prices are involved in the average.
- The average return of the portfolio might well be negative, but in this case the coupon will just be zero, and no loss will be sustained.
- It will be possible to ask for the anticipated payment of the coupon every six months, starting from the end of year 2.
- It is also possible to ask for the anticipated repayment of the face value, but this implies a reduction with respect to the face value.

³³We will explore the inverse relationship between bond prices and interest rates in Chapters 3 and 6.

This looks like a very complicated security, but it may be assembled by **bundling** a zero-coupon bond and an exotic option. The zero ensures the payment of the face value, which is reduced if early repayment is requested. The option is a complicated version of a call. Let $S_j(t_i)$ be the price of each underlying stock share, indexed by $j = 1, \dots, 10$, at time $t_i = i/12$, where $i = 0, 1, 2, \dots, 48$. Note that we are considering one year as the time unit, as customary in finance, and for the sake of simplicity we are assuming that one year consists of 12 *identical* months, which is not really the case. Finally, let us consider the following payoff:

$$\max \left\{ 0, \frac{1}{48} \sum_{i=1}^{48} \sum_{j=1}^{10} S_j(t_i) - K \right\},$$

where K , the strike, is just the initial value of the portfolio ,

$$K = \sum_{j=1}^{10} S_j(t_0).$$

This option has three features:

- It is a **rainbow** option, as it is written on multiple underlying assets.
- It is an **Asian** option, since its payoff is related to the average price, rather than to a single price at maturity (or early exercise).
- It is a **Bermudan-style** option, since it features early exercise opportunities, but only at a limited set of epochs, corresponding to $t = 2, 2.5, 3, 3.5$ years; thus, it is halfway between American- and European-style options.

Long-maturity zeros. As we have seen, zero-coupon bonds are typically associated with short maturities. However, we may find zeros maturing in 30 years. These zeros are often created by investment banks that hold long-term, coupon-bearing sovereign bonds, and strip the coupons creating zeros. This **coupon stripping** procedure is an example of the more general idea of **cash flow unbundling**. As we shall see, the availability of a rich array of zeros is useful in asset–liability management. Furthermore, they are quite sensitive to changes in the interest rates, and can be used for speculation and hedging purposes.

Mortgage-backed securities. When a bank issues a mortgage to a homeowner, it creates an asset in its balance sheet. This asset, however, is not liquid. In order to create a marketable security, the cash flows from a pool of mortgages can be

bundled together by a securitization procedure, in order to create a mortgage-backed security that can be traded. A mortgage-backed security can be risky, as homeowners may default on payments or, on the contrary, they might repay debt early, if interest rates move in their favor. The former issue exposes the investor to default risk, whereas the latter issue creates reinvestment risk. Despite these risks, these securities promised a larger yield than other bonds, which made them quite popular when they were introduced. Default risk should be somehow mitigated by risk pooling. It seems sensible to say that a limited amount of defaults in a pool of mortgages should be not too much of a trouble. The idea was pushed to the limit when subprime mortgages, i.e., mortgages offered to homeowners with a high chance of default, were securitized. Unfortunately, risk pooling works when risks are not quite correlated. When the subprime mortgage crisis erupted in 2007, correlations increased sharply, proving that, indeed, some of the underlying risks were not fully understood. As we discuss in Section 5.5, the matter was further complicated by tranching procedures, whereby different layers of securities with different risk levels are assembled, possibly by a second round of securitization. The ensuing crisis led to the demise of Lehman Brothers and to a revision of financial engineering practices that, at the time of writing, is not quite settled yet.

1.3 Market participants and their roles

After discussing securities that are actually traded on financial markets, let us take a more concrete look at who market participants are and their roles. In Section 1.1, we have described the role of financial markets in terms of consumption timing and risk transfer, which underlines the following functions performed by financial markets:

- To channel available funds from lenders to borrowers.
- To transfer risk, both for individuals and corporations.

Actually there are many other important functions, which includes the ones listed below:

- To provide a payment mechanism (e.g., by bank drafts and credit cards). We will not consider this side of the financial system, but we have to bear in mind that this is one of the main historical reasons behind the creation of finance during the Renaissance in Italy, when the needs of traders facing travel risk had to be met.
- To provide financial services, including the creation and sale of securities like bonds by both public and private issuers, as well as offering advice to firms in matters of financial management.
- To create market liquidity, i.e., the possibility of buying and selling assets quickly *and* at a fair price, as well as to offer portfolio adjustment facilities. The actual complexity of the information technology infrastructure,

needed to actually perform trading on markets and to take care of asset custody, should not be underestimated.

- To enable the separation between ownership and management in large corporations, which cannot always be effectively managed by family owners. Corporate growth would otherwise be impossible much beyond the size of a firm owned by the original founders. This is a controversial matter, as it may create problems related to bad incentives and agency issues. Moreover, the excess of financialization of the economy is under scrutiny, and with good reason.
- It has been claimed that the financial system also plays an information role, since the wide availability of financial data may be used to gather valuable knowledge. It may be argued that this is a bit debatable, in the light of speculation excess and some market anomalies studied by behavioral finance.

In concrete terms, all of these functions (and others) are carried out by an interconnected network of actors including

- Households and private investors
- Large corporations and smaller firms
- Governments and other public agencies, including local authorities like municipalities
- Financial intermediaries like banks, brokers, dealers, market makers, etc.
- Financial service providers like financial advisory firms, common funds, hedge funds, pension funds, insurance companies, etc.
- Regulatory and supervisory agencies, like the SEC (Security Exchange Commission) in the USA, the equivalent CONSOB³⁴ in Italy, the Basel Committee, central banks, etc.

All of these actors are connected by markets, which we may think as a platform on which transactions can be executed, either over-the-counter or on a computer network. We will discuss a bit of market structures later, but it is fundamental to immediately understand the two basic structures: **Primary markets** and **secondary markets**. When securities are created, they are first sold on primary markets. For instance, a corporation may float equity, possibly by an **initial public offering** (IPO), which needs some support from investment banks, under the scrutiny of regulatory bodies. By a similar token, a government may sell bonds to institutional investors using an auction mechanism. Households do not have direct access to primary markets, but they operate on secondary markets, where securities may be freely traded after they are issued.

³⁴Commissione Nazionale per le Società e la Borsa.

Example 1.13 Selling vs. writing options

A common source of confusion is the nature of the obligations related to selling an option, like a vanilla call or put. If we sell a call option, are we obligated to sell the underlying asset to the holder if the option is exercised? The source of confusion is the kind of market on which the option is sold. The *writer* is the one selling the option first on *primary* markets. Then the option, assuming it is an exchange-traded one, may change hands on secondary markets, but the obligation is only assumed by the original writer. If we buy and then sell an option, we are just selling the rights to a new holder, incurring a profit or a loss. To avoid any ambiguity, we shall always use the term “option writing” when we mean “selling on primary markets,” collecting the option premium and assuming the obligations stated in the contract. When we talk about “selling” an option, it will always refer to secondary markets, as part of a trading strategy.

Some specific market players, like brokers and dealers, make sure that there is sufficient liquidity on markets. Before discussing some of these actors in more detail, let us underline that they play different, but not mutually exclusive, roles. For instance, governments may be net savers or net borrowers, just like households. However, we know well that our role can change over time, since we may borrow money at the beginning of our working life (e.g., under the form of a mortgage) and, hopefully, we become savers as our career progresses. Of particular interest are some key roles that may be played by investors, non-financial firms, and financial intermediaries: Hedgers, speculators, and arbitrageurs. These will be discussed later.

1.3.1 COMMERCIAL VS. INVESTMENT BANKS

Banks come in many forms, including retail banks mostly dealing with households, commercial banks offering services to small-medium firms, and large investment banks. Investment banks are often involved in mergers and acquisitions, and they also act as underwriters to help corporations in raising capital by floating equity or issuing bonds. Usually these securities are bought by investment banks on primary markets, and then sold on secondary markets. Furthermore, there are different legal entities, like banks floating their own equity and credit unions. Here, we just want to draw the line between **deposit-** and **non-deposit taking** banks.

A deposit-taking institution, like a retail or commercial bank, may also collect funds from households, who deposit money on accounts that may be more or less protected against bankruptcy. Bankruptcy may result from careless credit distribution decisions by the bank, from the difficulty to collect loans back due to economic stagnation, or, in extreme cases, from risky proprietary

trading, i.e., trading that the bank carries out itself, rather than on behalf of its clients. The recent trend has been to reduce this kind of government-backed protection.

Large investment banks are non-deposit taking, and must raise capital by other means, like floating equity, issuing bonds, or borrowing money from other banks. As we shall see in Section 1.4.4, using debt rather than equity has the results of boosting ROE by a leverage mechanism. The leverage ratio measures the amount of debt used with respect to own equity. Roughly speaking, if ROA is 5%, a leverage ratio of 2 will double ROE. Unfortunately, the same happens in the case of loss, and when excessive debt is used, bankruptcy risk is considerably increased.³⁵

Since risky proprietary trading activity may hurt clients of deposit-taking banks, a line between the two types of banks was drawn in the USA after the 1929 crash (Glass–Steagall Act), which is why investment banks could not take deposits. This line has been blurred in the last decades. Furthermore, the increasing interconnection among market players has increased systemic risk, i.e., the possibility that the collapse of a large institution affects many others by a domino effect.³⁶

1.3.2 INVESTMENT FUNDS AND INSURANCE COMPANIES

Individual investors may feel that they lack the information required to make sound investment decisions. Furthermore, it may be difficult to properly diversify the risk exposure with a limited budget, as transaction costs preclude the possibility of many small investments in a broad set of securities. Hence, they may purchase shares of **mutual funds**, that are supposedly managed by skilled professionals who, in exchange for a fee, should provide good return opportunities to their clients. Shares are continuously created in the case of an open-end fund. Shares are destroyed when a client redeems her shares of a mutual fund. On the contrary, closed-end funds have a given number of shares that may be traded.

There are two basic kinds of fund manager. **Active** managers try to earn extra return by skill and by pursuing, for instance, stock-picking and market-timing strategies. The actual performance of active managers is the subject of a good amount of controversy. As an alternative, we may consider a **passive** manager, who will not try to do any better than the market as a whole, but will just provide a diversified portfolio tracking a broad index. We shall see that the passive view has some theoretical support by equilibrium models like the capital asset pricing model. Clearly, the fee required by a passive manager should be definitely small with respect to the cost of an active fund.

³⁵ Apparently, the leverage ratio of Lehman Brothers before their collapse was something like 20. LTCM, too, had reduced equity by forcing investors out before their near collapse.

³⁶ The collapse of Lehman Brothers affected hedge funds, among other things, as they acted as prime brokers for these funds.

The ultimate active fund is a **hedge fund**. Despite the misleading name,³⁷ hedge fund managers pursue possibly very risky and nonstandard investment strategies in order to earn extra return. When we buy a share of a mutual fund, we are client of the fund. On the contrary, when we buy a share of a hedge fund, we are *partners* of the hedge fund, which has a different legal nature. This is related to the high risk involved, and in fact only wealthy individuals are allowed to expose themselves to this level of risk. Furthermore, due to the complex trading strategies and the use of possibly illiquid assets, it may take a considerable amount of time to redeem shares of a hedge fund. These barriers are somewhat circumvented by funds of hedge funds.

The ultimate passive fund is an **exchange-traded fund**, **ETF** for short, which is a just passive fund tracking an index. In order to reduce costs, ETF shares are not distributed through a commercial network, unlike passive mutual funds, but they traded on exchanges, just like stock shares. This opens a thorny issue, since the ETF is supposed to track an index, but an uncontrolled demand–offer mechanism might cause its value to drift away from the fair one.³⁸ Market-makers guarantee the necessary liquidity and make sure that the ETF value is kept in line. Furthermore, a deviation from the fair value would create an arbitrage opportunity, which will be exploited by skilled investors, assuming liquid and well-functioning markets.

There are other non-deposit taking financial intermediaries that are engaged in fund management, namely, pension funds and insurance companies. These intermediaries face difficult asset–liability management problems, as they collect pension contributions and insurance premia that must be properly invested in assets, in order to generate cash flows and meet an *uncertain* stream of liabilities. A **non-life insurance** company may deal with, e.g., a stream of car accidents or other property damage. A **life insurance** company faces a similar task as a pension fund. Liabilities are uncertain because of longevity risk and, possibly, inflation-indexation. A **defined-benefit pension fund** must guarantee to retired workers an income that depends on the received wages, according to prespecified rules. The contribution level may be increased over time, depending on contingencies. Recently, because of decreasing interest rates and increasing life expectancy, there has been a shift toward **defined-contribution pension funds**, in which there is no defined income, and considerable risk is borne by the retired worker. It may seem that a defined-benefit fund is much preferable from the worker’s viewpoint. However, we should also consider that with a company defined-benefit fund it is more difficult to transfer vested benefits, if a worker changes the employer. Furthermore, a firm with a well-funded pension fund might become the target of hostile takeovers.

³⁷Risk hedging means *reducing* risk.

³⁸This cannot happen with a mutual fund, whose **net asset value (NAV)** is evaluated and reported daily by the fund management.

1.3.3 DEALERS AND BROKERS

A fundamental requirement of financial markets is liquidity. Hence, there is a need for institutional players that are continuously available to buy and sell an asset. These market-makers are also referred to as **specialists**. This role may be played by **dealers** and **brokers**. It may be the case that the same institution is both a dealer and a broker, but the two functions are different. To understand the difference, think of a real estate agent. Her role is to connect two counterparties, but she does not really own an inventory of houses and apartments. This is the role of a broker. The broker has no inventory of assets, and as such she does not suffer from inventory risk. A commission on the trade is paid to the broker to compensate her. There are also primary brokers associated with hedge funds, which may need large trades. On the contrary, when we travel around the world and exchange currencies at an airport, we do business with a dealer. The dealer *does* keep an inventory of the assets she trades. Clearly, this inventory entails some risk. In fact, the dealer is compensated by enforcing a bid–ask spread:³⁹

- The **bid price** is the price at which the dealer is willing to buy the asset from us.
- The **ask price** is the price at which the dealer is willing to sell the asset to us.

Needless to say, the ask price is larger than the bid price, and their difference is a measure of market liquidity. We have seen an example of bid–ask spread in Section 1.2.5 on foreign exchange, and we will see similar examples in the case of stock shares when we discuss market mechanisms. The same applies to interest rates, as the rates at which we may lend or borrow money, when dealing with a bank, are quite different.

Bid–ask spreads are a form of **market friction**. Other market frictions are represented by taxes and by transaction costs associated with trades. These fees may have a fixed and/or a variable component. In general, thanks to the use of information technology, transaction costs have been reduced over the years.⁴⁰ For the sake of simplicity, we will usually ignore such frictions, which may be a sensible approximation for large institutional investors.

1.3.4 HEDGERS, SPECULATORS, AND ARBITRAGEURS

Market participants are often engaged in risk transfer, which is the traditional purpose of insurance contracts. More recently, a huge market of derivative assets has been developed, connecting **hedgers** and **speculators**. Hedgers are exposed to risk factors, like interest rates and currency exchange rates, and would

³⁹You may also hear the term bid–offer, but I personally prefer bid–ask, since the difference between bid and ask sounds much clearer to me than the difference between bid and offer.

⁴⁰Some argue that the reduction of market frictions and the related increase of transaction frequency is far from being a blessing, as it may lead to market instability. High-frequency algorithmic trading strategies are often blamed for this.

like to reduce or eliminate that exposure. One possibility is to lock a rate for the future by forward contracts, for instance. Speculators, on the contrary, have a definite view about the direction that markets will take, and they are willing to take a bet on it. Thus, speculators may be willing to “buy volatility” from hedgers. We should realize that these two roles are not mutually exclusive. Indeed, there are multiple sources of risk that affect the value of a portfolio of assets, and a market participant might want to shape her portfolio in such a way that it is made less sensitive to risk factors on which she does not feel like betting, and more sensitive to other factors on whose direction she is more confident. Hence, she will increase the exposure to some risk factors, behaving as a speculator, and at the same time she will reduce the exposure to other risk factors, behaving as a hedger. As a concrete example, an investor may feel that she is good at picking stock shares that will perform better than the market as a whole, but she is unsure about the market direction. In the case of a market crash, being good at stock-picking and lose less than the market may only be a partial consolation. As we shall see, she will be interested in taking risks that are specific to some firms, while getting rid of systematic market risk. Hedgers and speculators need models to quantify uncertainty in risk factors and to understand how different sources of risk affect asset prices. On the one hand, we need tools to measure risk. On the other hand, we also need risk management approaches and decision models to find the best hedging strategy.⁴¹

Pricing models are also needed to check the consistency of the prices of assets that depend on common risk factors. For instance, derivatives written on the same underlying asset should be somehow related. If prices are inconsistent, trading strategies may be devised in order to take advantage of price misalignment. In technical terms, we talk of arbitrage opportunities, which are exploited by **arbitrageurs**. In liquid and well-functioning markets, it may be argued that arbitrage opportunities should not last long, as arbitrageurs will be quick in detecting and exploiting them, bringing prices back in line. We will investigate the mathematics of arbitrage in Section 2.3. There, we shall take a simplistic view of markets, ignoring market frictions, modeling errors, and liquidity issues. Nevertheless, we will be able to develop powerful pricing models based on the idea that there should be no arbitrage opportunity in market equilibrium. As usual, market reality is definitely more complex, and the actual arbitrage strategies may be not so sharp and may fail to work for an array of reasons. It may also be argued that arbitrageurs are sort of parasites taking advantage of what other market participants do, without really contributing to any growth in the real economy. Nevertheless, arbitrageurs play a vital role to the correct market functioning by keeping prices in line. One example that we have already hinted at is the need to ensure consistency between prices of an ETF share and the index that the fund is supposed to track.

⁴¹See Section 2.2.

1.4 Market structure and trading strategies

Quantitative finance relies on mathematical models that are, by necessity, an abstraction of market reality. Finding the right level of abstraction and detail simplification is definitely an art rather than a science, and we are engaged in a quest for those models that are inevitably wrong, but hopefully useful. While we may not be interested in an overly detailed view of market structures and the institutional mechanisms by which a trade is executed, we must be aware of some fundamental features that we outline here.

1.4.1 PRIMARY AND SECONDARY MARKETS

We have already hinted at the difference between primary and secondary markets. A primary market is where a security is first traded. The exact mechanism depends on the kind of security. For instance, an auction mechanism is used to introduce new government bonds on the market, but the auction is restricted to institutional investors. In the case of a stock, we should distinguish an **IPO**, i.e., the initial public offering of shares of a firm that is first quoted on the market, from a **seasoned offering**, where further equity is floated by an already quoted firm. An IPO may be a costly business, as several requirements are typically set by regulators and must be met by a firm floating equity on exchanges. A pool of investment banks is involved in the process, which may also include so-called “road shows” to present the offering to investors. Anyway, we must keep in mind that shares need not be traded on an exchange. Some firms are kept private and possibly owned by private equity funds. A significant part of equity can also be kept by the original owners to maintain control over management, and only the rest are floated and are outstanding on secondary markets.

1.4.2 OVER-THE-COUNTER VS. EXCHANGE-TRADED DERIVATIVES

Not all assets are traded on regulated exchanges, as some are traded OTC. For instance, forward agreements are negotiated directly between the two counterparts, unlike futures. Another example of OTC derivatives are exotic options with possibly quite complicated payoffs. The advantage of an OTC agreement is that it may be tailored to meet specific risk hedging requirements. The disadvantage is that the lack of a quoted price may put a firm or a public administration at disadvantage. Furthermore, nonstandardized assets are rather illiquid, which means that unwinding the position may be expensive, if not impossible.

1.4.3 AUCTION MECHANISMS AND THE LIMIT ORDER BOOK

As we have mentioned, auctions are used, for instance, when selling sovereign bonds on primary markets. Here, we consider secondary markets for stock shares and describe an auction mechanism based on the **limit order book**.

Table 1.3 A five-level limit order book for a liquid stock share.

Bid		Ask	
Quantity	Price	Price	Quantity
130	77.26	77.28	8881
137	77.25	77.31	273
5855	77.23	77.33	115
300	77.22	77.34	272
13,080	77.16	77.35	738

When a market participant issues an order, she may specify a limit price, i.e., the maximum (minimum) price at which she is willing to buy (sell) an asset. The limit order book is structured on two columns:

- On the left, we observe the buy orders, associated with limit prices sorted in decreasing order. The top level line reports the highest bid price, as well as the related quantity (possibly related to different orders).
- On the right, we observe the sell orders, associated with limit prices sorted in increasing order. The top level orders are associated with the smallest ask price.

A five-level limit order book is reported in Table 1.3. The two top quotes are called the **inside quotes**. When limit prices cross each other, a trade takes place. Otherwise, no trade is executed. Orders need not specify a limit price, as an investor may just issue an order to be executed at the best available price. It may happen that a large order is executed at different prices, when its size exceeds the quantity available in an inside quote. The spread between the inside quotes reflects liquidity. We may notice that the bid–ask spread in Table 1.3 is quite small. Table 1.4 tells a rather different story, as there is a much larger spread, especially in percentage terms, between the inside quotes. A large spread typically comes with less trades during a day and lower volumes.⁴²

Quite often, price-contingent orders are used. There are two features: (1) the kind of order, which may be buy or sell, and (2) the activation condition, which is related to the price going above or below a threshold level. Therefore, we have four basic types of price-contingent orders.

- The **stop-loss** order is a selling order to be activated when the price goes below a limit. The rationale behind the order is clear: We hold an asset, and in case of a drop in price we want to cut losses and get rid of it.

⁴²The data reported here are not quite recent but real. They refer to the Paris stock exchange in 2010, and the first share is a large and well-known French cosmetics producer, whereas the second one is a less traded producer of containers for overseas shipping.

Table 1.4 A five-level limit order book for a rather illiquid stock share.

Bid		Ask	
Quantity	Price	Price	Quantity
108	21.91	22.20	100
55	21.90	22.30	206
110	21.85	22.35	232
260	21.88	22.40	100
54	21.77	22.50	100

- The **limit-sell** order is a sell order activated by a price going above the threshold. In this case, the idea is that we hold an asset, and we sell it when a target profit has been achieved.
- The **limit-buy** order is a buy order activated when the price goes below a limit. This means that the asset is cheap enough to be bought. This is related to a contrarian strategy, a strategy trying to buy undervalued stocks.
- The **stop-buy** order is a buy order activated when the price goes above the threshold. The rationale is that the price is high enough to signal an increasing trend. This is related to a momentum strategy, i.e., a strategy trying to chase increasing trends.

High-frequency analysts build models at the limit order level, considering both prices and volume, with the aim of developing algorithmic trading strategies. Other models at this microstructure level concern the optimal execution of a large trade in order to minimize market impact. We shall not consider this operational level in this book.

1.4.4 BUYING ON MARGIN AND LEVERAGE

Buying on margin is a leveraged strategy aimed at boosting returns using debt. In corporate finance, **leverage** refers to the ratio of debt over equity. Here we have a similar use, as leverage means buying an asset by only partially using our own capital, and borrowing the rest from a broker or a bank.

Imagine that we have a strong view about a specific stock share, a bullish one in particular. As we have already mentioned, we may use derivatives, rather than just going long the asset to take advantage of our view, but it may very well be the case that derivatives written on that specific stock share are not available. Hence, we may resort to leverage, more specifically, to **buying on margin**. We have already met the term *margin* when dealing with futures contracts and margin accounts. Here we refer to posting the asset itself as a collateral of our debt, plus some cash acting as a buffer and guaranteeing the broker that we will repay the debt even if the asset price drops.

The mechanism revolves around the concept of **margin ratio**. Margin requirements specify an initial margin ratio, as well as a **maintenance margin**. To grasp the idea, it is useful to refer to the asset–liability–equity triad. In this case:

- The asset is the amount of stock shares that we have purchased.
- The liability is the money we owe to the broker.
- Equity, as usual, is their difference.

In this case, the margin ratio is defined as the ratio of the value of equity to the value of assets. If the margin ratio falls below the maintenance margin, we get a margin call, which means that we have to post additional cash (or other collateral). Failure to do so will result in our position being liquidated by the broker. This is best illustrated by an example.

Example 1.14 Margin trading

Say that the current price of a stock share of Boom Corp is \$100, and we buy 100 shares, for a total amount of \$10,000. To finance the trade, we borrow \$4,000 from the broker. The initial situation is as follows:

Assets	Liabilities
Stock \$10,000	Loan from broker \$4000
	Equity
	\$6000

The initial margin ratio is

$$\frac{\text{Equity}}{\text{Assets}} = \frac{\$6000}{\$10,000} = 60\%,$$

and let us assume that the maintenance margin ratio is 30%. Note that, for the sake of simplicity, we are not considering the interest payment to the broker. If things turn sour and the stock price falls to \$70 per share, the new balance sheet will be

Assets	Liabilities
Stock \$7000	Loan from broker \$4000
	Equity
	\$3000

and the margin ratio now is just

$$\frac{\$3000}{\$7000} = 43\%.$$

A natural question is: How far can a stock price fall before getting a margin call? If we let P be price of the stock, the margin ratio is

$$\frac{100P - \$4000}{100P}.$$

The limit price is obtained by setting this ratio to 30% and solving for P , which yields $P_{\text{lim}} = \$57.14$.

The effect of leverage is to boost both profit and loss. To see this, imagine that our view in the previous example is correct and that the Boom Corp price rises by 30%. The relevant return is return on equity (ROE), rather than return on assets (ROA). ROA is the usual rate of return that we consider when dealing with portfolio management (30% in this case), but ROE is boosted by the fact that we need to invest only a fraction of the value of the assets. To be a bit more realistic, let us assume that the broker requires an interest rate of 3% (over the holding period of the trade). ROE is

$$\frac{\$10,000 \times 0.30 - \$4000 \times 0.03}{\$6000} = 48\%.$$

The more leverage we apply, the better, in the rosy scenario. For instance, if we increase initial leverage to 50%, ROE is

$$\frac{\$10,000 \times 0.30 - \$5000 \times 0.03}{\$5000} = 57\%,$$

to be compared with the 30% of a normal trade. In practice, with a 50% leverage we double return, which is eroded by the 3% interest. Clearly, there must be another side of the coin. Imagine that we are wrong and price plummets to \$70. With a 50% leverage, ROE is

$$\frac{-\$10,000 \times 0.30 - \$5000 \times 0.03}{\$5000} = -63\%,$$

i.e., we double loss and on top of it we have to pay interest on debt.

The example we have considered is rather stylized. The understanding of margin trading arrangements is essential to interpret what happens in real life, since it is one of the factors of relevant events like the LTCM demise. The required cash can also be obtained by posting securities, typically through a repo agreement,⁴³ i.e., a repurchase agreement. This is a sort of collateralized loan, as relatively safe securities are sold to a counterparty, with the agreement to repurchase them later at a given price. It is easy to see that this boils down

⁴³See Section 5.4.

to borrowing money for a given interest rate. Things may really go badly as feedback effects may arise for large trades on illiquid assets. Imagine that the asset involved in a margin trade loses value because of a market crash, and we start getting margin calls. This may be associated with a reduction in market liquidity, which implies that we have to sell illiquid securities in order to raise cash. Unfortunately, selling assets in an illiquid market may have a large impact, leading to a further reduction in asset values, so that we are caught in a feedback cycle, potentially leading to bankruptcy. The larger the leverage, the more difficult it is to get out of such a situation.

1.4.5 SHORT-SELLING

Short-selling, like buying on margin, is a strategy that can be used for speculative purposes. In this case, however, the bet is a bearish one, as short-selling profits from a drop in the asset value. The mechanics of a short sale is as follows:

- At time $t = 0$ we borrow the asset through a dealer/broker, then we sell it and deposit the proceeds plus required margin into an account.
- At time $t = T$ we close out the position by buying the stock and returning it to the party from which it was borrowed.

If the asset is a stock share that pays a dividend in the interval $(0, T)$, a corresponding cash amount must be paid as well. A similar consideration applies to bonds and coupons. Therefore, when considering profit from a short sale, there is a change in sign with respect to the usual case:

$$\text{Profit} = \text{initial price} - (\text{ending price} + \text{dividends}).$$

Of course, the trade will result in a loss if the asset price increases. Furthermore, borrowing the asset may be expensive, as we must compensate the broker and/or the asset holder, and possibly limited to a short time interval. Sometimes, a **short-squeeze** occurs, i.e., the short position is forced to close the trade at a very unfavorable time, just when the asset price is rocketing.

In margin trading, we borrow cash to buy an asset, whereas in short-selling we borrow the asset itself and raise cash, which must be kept into an account with the broker until the trade is closed. As usual, the lender of the asset protects herself by requiring the deposit of a margin, in addition to the proceeds of the short sale. In this case, the margin ratio is defined as equity divided by the value of the assets owed, which is a liability in this case. This definition of margin ratio differs from the one we used in the case of buying on margin. As a mnemonic help, the margin ratio is always defined by dividing equity by the side that is sensitive to the current value of the traded stock shares, i.e., the asset side when buying on margin and the liability side when short-selling.

■ Example 1.15 A short trade

We are strongly bearish about stock shares of DotBomb. Hence, we sell 1000 shares at the current price of \$100. The proceeds, \$100,000, are deposited in our margin account, together with some collateral. If the initial margin required is 50%, we must deposit a corresponding amount in cash or rather safe securities, e.g., T-bills. The initial situation is as follows:

Assets	Liabilities
Cash + T-bills \$150,000	Short position in stock \$100,000
	Equity
	\$50,000

If we are right, and DotBomb falls to \$70, we can close our position out and earn \$30,000 (neglecting commissions and interest). If, however, DotBomb rises to \$110, the new situation is

Assets	Liabilities
Cash + T-bills \$150,000	Short position in stock \$110,000
	Equity
	\$40,000

and the margin ratio drops to

$$\frac{\$40,000}{\$110,000} = 36\%.$$

How much can the stock price increase, before we get a margin call? If the maintenance ratio is 30%, we must find a stock price P such that

$$\frac{\$150,000 - 1000P}{1000P} = 30\%.$$

By solving for P , we obtain $P_{\text{lim}} = \$115.38$.

This example, too, is somewhat stylized. We are not considering the cost of the trade, and the fact that it may be expensive to keep the short position open for a long time. Just like with buying on margin, a realistic assessment of the return of a trade should be based on ROE. In this case, the actual investment is the additional margin that has to be posted, which may be rewarded at a given interest rate. When short-selling an asset is difficult or expensive, a short position may also be created by taking a position in a derivative. For instance, we may sell futures contracts written on the asset, as we shall see in more detail in Chapter 12.

Short-selling is a highly controversial strategy, as it is believed by some to be a way to manipulate and depress the markets. A famous case in point is George Soros' bet against the GB pound in 1992, when the UK was forced to leave the European Exchange Rate Mechanism. When markets crash, short sales are often prohibited, as they are considered by some as a nasty way to generate a vicious feedback cycle.⁴⁴ A different view maintains that short-selling is essential to preserve liquidity, as well as to keep prices in check when markets roar. We should also consider that short-selling may play a role in hedging and is not necessarily a speculative strategy. As it happens with many matters in finance, the jury is out. We have to stress the fact that in later chapters, especially when discussing pricing by no-arbitrage, we will assume that unlimited short sales are possible. Clearly, this is a rather idealized view of real markets.

1.5 Market indexes

We are all familiar with stock market indexes like Dow Jones Industrial Average (DJIA), NASDAQ, Dax, and Nikkei, which are often mentioned on newspapers and TV news. Some of them have a quite long history (the DJIA has been computed since 1896), and they are also the underlying factor of several traded derivatives, like index futures and options. Indexes are tracked and replicated by index funds and ETFs.⁴⁵ Most widely known indexes are related to a geographic area, such as a national stock market, but some indexes, like MSCI (Morgan Stanley Capital International), refer to world markets. On the other hand, we may also use more specific indexes, related to a given industry sector, or even nonfinancial markets, as is the case of indexes for the real estate market.

Indeed, there is a wide variety of indexes, beyond the familiar ones for stock markets. For instance, the EURIBOR and LIBOR rates are actually indexes, as they are an average of a set of interbank offered rates. We may also use bond market indexes, which are a bit more problematic, since bonds, unlike stock shares, have a maturity. Thus, the pool of bonds in an index must be continuously updated. Furthermore, we shall see that the volatility of a bond price gets smaller and smaller as the maturity is approached. An increasingly important index, VIX, tracks stock market volatility. Intuition would suggest that such an index should be built by estimating the standard deviation of stock market return by familiar descriptive statistics. Unfortunately, the usefulness of such a backward-looking index would be questionable. A forward-looking index, may rely on the implied volatility of a set of traded options.⁴⁶

⁴⁴Short-selling strategies have always been controversial, as illustrated by significant historical cases reported in [12]. In the USA, in the midst of political discussions about possible prohibition of short-selling, the practice was even deemed "unAmerican."

⁴⁵ETFs may also be short or leveraged. A short ETF allows to take a short position in the index, profiting from a market drop. A leveraged ETF multiplies profits and losses by a given factor.

⁴⁶We will discuss implied volatility in Section 13.6. Here, it suffices to say that it is a volatility such that option prices predicted by a mathematical model match the actual market prices.

As we can imagine, the definition of a suitable index is far from trivial and involves managing extraordinary events as well, such as stock splits, mergers/acquisitions, delistings, etc. In order to get a feeling for the involved issues, it is interesting to compare two well-known indexes:

- The Dow Jones Industrial Average, which includes 30 blue chips, and is a **price-based** index.
- The Standard & Poor's S&P500 index, which is a broadly based index involving 500 stock shares, and is a **market-value-weighted** index.

Generally speaking, we may consider a set of m stock share prices S_k , $k = 1, \dots, m$, and define an index

$$I = \frac{1}{D} \sum_{k=1}^m w_k S_k,$$

for a given set of **weights** w_k , and a **divisor** D . Usually, when defining an average, we assume that weights add up to 1, which in this case would be obtained by choosing D as the sum of weights. Actually, the divisor is initially chosen in such a way that the resulting index assumes a “nice” value, say, 100 or 1000. More importantly, the divisor is changed when the index composition is changed to reflect new market conditions,⁴⁷ or when events such as spinoffs or mergers/acquisitions take place. The defining features of the aforementioned indexes are:

- Weights w_k are all set to 1 for the DJIA, i.e., the index essentially tracks a portfolio consisting of one stock share for each name in the index.
- Weights in the S&P500 index correspond to the number of outstanding stock shares (free-float only); hence, the portfolio reflects the actual market capitalization of each firm.

Example 1.16 Price-based vs. market-value-weighted indexes

Consider the following scenario:

- Stock share A has an initial price $S_A(0) = \$25$, at time $t = 0$, which is increased by 20% to $S_A(T) = \$30$ at time $t = T$. The total market capitalization is \$500 million (hence, 20 million shares are outstanding).

⁴⁷ An interesting market anomaly is the plunge in the price of stock shares that are dropped from an index. Rationally, this should not imply anything in terms of intrinsic firm value, but the consequent reduction in trading activity on that stock share may have a significant effect.

- Stock share B has an initial price $S_B(0) = \$100$, which drops by 10% to $S_B(T) = \$90$ at time $t = T$. The total market capitalization is \$100 million (hence, one million shares are outstanding).

A price-based index would initially be

$$\frac{25 + 100}{2} = 62.5,$$

where we assume a divisor $D = 2$, which is really inconsequential when considering percentage changes in the index. At the end of the time horizon, the new index value would be

$$\frac{30 + 90}{2} = 60,$$

with a drop of 4%. Note that the price drop of the more expensive stock share dominates here, but this does not reflect the true market weights. Let us consider a market-value-weighted index, with initial value

$$\frac{25 \times 20 \cdot 10^6 + 100 \times 1 \cdot 10^6}{10^6} = 600,$$

where we set $D = 10^6$. The new index value would be

$$\frac{30 \times 20 \cdot 10^6 + 90 \times 1 \cdot 10^6}{10^6} = 690,$$

with an increase of 15%.

We notice a relevant difference in the behavior of the two indexes. The difference may also be reflected in the way the index is adjusted when something new happens. Consider, for instance, a 2-for-1 stock split. Clearly, a market-value-weighted index would not be affected, but an adjustment would be needed for the price-based index in order to preserve the continuity in its value.

■ Example 1.17 Index adjustments

Let us consider how to manage an index for a stock market on which two stocks are traded. Company A has 50 shares outstanding, with current price \$2, and company B has 10 shares outstanding, with current price \$10. The current value of a price-based index is 6, whereas the value of a market-value-weighted index is 100. Let us consider the following scenario: The price of Company A 's stock increases to

\$4 per share, and Company *B*'s stock splits 2 for 1 and is priced at \$5. How will the values of the price-based and market-value-weighted indexes change?

To begin with, we have to find the divisors. The current divisor for the price-based index is clearly $D = 2$, since

$$\frac{2 + 10}{2} = 6.$$

Then, it is important to notice that, actually, the second stock price did not change. The drop from \$10 to \$5 merely reflects the split. After the change in price of the first share, without considering the stock split, the new index would be

$$\frac{4 + 10}{2} = 7.$$

The new divisor is changed in order to reflect the split without introducing a discontinuity in the index:

$$\frac{4 + 5}{D'} = 7 \quad \Rightarrow \quad D' = \frac{9}{7}.$$

The divisor for the market-value-weighted index is found as follows:

$$\frac{2 \times 50 + 10 \times 10}{D} = 100 \quad \Rightarrow \quad D = 2.$$

However, the stock split is inconsequential for this index and does not require any adjustment in the divisor. Hence, the new index value is

$$\frac{50 \times 4 + 20 \times 5}{2} = 150.$$

An important observation is that indexes are not adjusted when dividends are paid. This is relevant, as a stock share price experiences a corresponding drop when a dividend is paid, and this will also affect the index, as well as derivatives written on the index. In option pricing models, the collective dividend behavior of the stock shares in the index may be approximated by a continuous-time dividend yield. Also note that the index is nondimensional and should be regarded as a number, rather than as a price. When defining derivative payoffs, the index must be multiplied by a given number in order to define a monetary payoff. For instance, the S&P500 index is multiplied by 250 to be converted into a monetary value.

Problems

1.1 Consider assets A_1 and A_2 , whose holding period returns R_1 and R_2 , in five possible scenarios, are given in the following table:

Scenario	Probability	R_1	R_2
ω_1	0.2	0.03	0.09
ω_2	0.2	0.17	0.16
ω_3	0.3	0.28	0.10
ω_4	0.2	0.05	0.02
ω_5	0.1	-0.04	0.16

Note that the probabilities are not equal, and that returns are not given as a percentage (if you prefer, you might also write, e.g., $R_1(\omega_1) = 3\%$). Find the expected value and the standard deviation of the returns of the two assets, as well as their (Pearson) coefficient of correlation.

1.2 We are pursuing a short-selling strategy, where we have shorted 300 shares of XYZ, at price €40. The initial margin required by the broker is 50% of the overall value, and the maintenance margin is 25%. What is the limit price of the stock before we are slapped with a margin call?

1.3 Consider a European-style call option maturing in five months, with strike price $K = €40$, written on a stock share with current price $S(0) = €35$. We (very unrealistically) assume that the uncertainty about the stock price at maturity $T = 5/12$ may be represented by eight equally likely scenarios: $S(T) \in \{20, 25, 30, 35, 40, 45, 50, 55\}$. Find the expected value of the option payoff.

1.4 Let us consider a market index for a tiny market, on which just 3 stocks are traded. In this market, 50,000 shares of the first firm are outstanding, 100,000 of the second one, and 80,000 of the third one. The index is a weighted-average of the three stock prices, reflecting the capitalization of the three firms. The current stock prices are €50, €30, and €45, respectively. To make the index easy to read and nondimensional, it is divided by a divisor (established once for all and kept constant in time; we rule out exceptional events like those described in Example 1.17); assume that with that choice of divisor, the index now is 118. We also assume that the stock shares do not pay any dividend.

The following table lists the stock prices (in EUR) for a three-day scenario (a single sample path):

Day	1	2	3
Price of stock 1	52	48	45
Price of stock 2	28	25	30
Price of stock 3	43	40	39

Find the corresponding scenario for the index value.

Further reading

- Most general books on financial asset management, like [1] and [5], have one or more sections on financial institutions and market mechanisms. More detail is provided in specific texts like [8]; see also [14].
- Many useful pieces of information on financial institutions are also given in [11], with a nice twist toward risk management.
- Market microstructure is dealt with in [7].
- The first chapters of [13] specifically cover the market structure for bond and debt markets. Bond markets, including bonds with embedded options, are also treated in [6].
- An adequate discussion of FOREX markets is provided by [15].
- A full understanding of how financial markets and institutions work cannot be achieved without some knowledge of real stories. The case of Long Term Capital Management is described, among others, in [9]. Another very useful reading is [3].
- While this book is concerned with financial markets, it is also essential to acquire some background knowledge on corporate finance, which is provided, among many others, by [10].

Bibliography

- 1 Z. Bodie, A. Kane, and A. Marcus. *Investments* (9th ed.). McGraw-Hill, New York, 2010.
- 2 Z. Bodie and R. Merton. *Finance*. Prentice Hall, Upper Saddle River, NJ, 1999.
- 3 R. Bookstaber. *A Demon of Our Own Design: Markets, Hedge Funds, and the Perils of Financial Innovation*. Wiley, New York, 2008.
- 4 N.P. Calamos. *Convertible Arbitrage: Insights and Techniques for Successful Hedging*. Wiley, Hoboken, NJ, 2003.
- 5 K. Cuthbertson and D. Nitzsche. *Investments*. Wiley, Chichester, 2008.
- 6 F.J. Fabozzi. *Bond Markets: Analysis and Strategies* (9th ed.). Pearson, Upper Saddle River, NJ, 2015.
- 7 L. Harris. *Trading and Exchanges: Market Microstructure for Practitioners*. Oxford University Press, New York, 2003.
- 8 P. Howells and K. Bain. *Financial Markets and Institutions* (5th ed.). Pearson Education, Harlow, 2007.
- 9 R. Lowenstein. *When Genius Failed: The Rise and Fall of Long-Term Capital Management*. Random House Trade Paperbacks, New York, 2001.
- 10 S.A. Ross, R.W. Westerfield, J.F. Jaffe, and B.D. Jordan. *Corporate Finance: Core Principles and Applications* (4th ed.). McGraw-Hill, New York, 2014.

- 11 A. Saunders and M.M. Cornett. *Financial Markets and Institutions* (6th ed.). McGraw-Hill, New York, 2014.
- 12 R. Sloan. *Don't Blame the Shorts: Why Short Sellers Are Always Blamed for Market Crashes and How History Is Repeating Itself*. McGraw-Hill, New York, 2009.
- 13 S.M. Sundaresan. *Fixed Income Markets and Their Derivatives* (2nd ed.). South Western College Publishing, Cincinnati, OH, 2002.
- 14 S.R. Veale, editor. *Stocks, Bonds, Options, Futures: Investments and their Markets* (2nd ed.). Prentice Hall, Paramus, NJ, 2001.
- 15 J. Walmsey. *The Foreign Exchange and Money Markets Guide* (2nd ed.). Wiley, New York, 2000.