Chapter 3

Classical Financial Mathematics

Contents

In the fifteenth century mathematics was mainly concerned
with questions of commercial arithmetic and
the problems of the architect.
—Joseph Alois Schumpeter (1883–1950),
Capitalism, Socialism and Democracy [687]

To valuate any security, one looks at its cash flows. As we are most interested in the present value of expected cash flows, three features stand out: magnitude and direction of the cash flows, times when each cash flow occurs (**dating dollar flows**, so to speak), and an appropriate factor to discount future cash flows. The security is more valuable when the payment is higher, the payment is received earlier, and/or the discounting factor is smaller, other things being equal. One can sort through many hurdles once this idea is grasped firmly.

The current chapter deals with elementary financial mathematics essential to any financial calculation. The time line throughout this book for discrete-time models should be understood as depicted in the following exhibit.

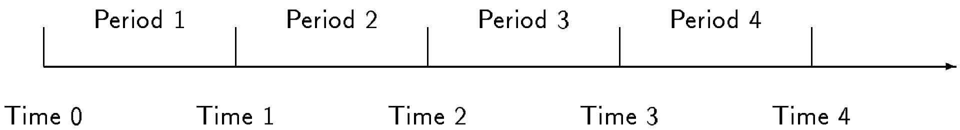

3.1 Time Value of Money: Compounding and Discounting

Interest is the cost of borrowing money, as the lender is giving up the present consumption for future purchasing power[1] [688]. Let r be the annual **nominal** interest rate. If the interest is **compounded** once per year, the **future value** of investing an amount P after n years is $\text{FV} = P(1+r)^n$. To look at it from another perspective, an amount FV that is available n years from now is worth only $P = \text{FV} \times (1+r)^{-n}$, its **present value.**[2] The process of obtaining the present value is called **discounting**. In general, if interest is compounded m times per annum, the future value is

$$\text{FV} = P\left(1 + \frac{r}{m}\right)^{nm}. \tag{3.1}$$

Hence, $\left(1 + \frac{r}{m}\right)^m - 1$ is the equivalent annual rate compounded once per annum, or the **effective** interest rate.

We have **annual** compounding with $m = 1$, **semiannual** compounding with $m = 2$, **quarterly** compounding with $m = 4$, **monthly** compounding with $m = 12$, **weekly** compounding with $m = 52$, and **daily** compounding with $m = 365$. The **bond equivalent yield** or **BEY** is the annualized yield with semiannual compounding, and the **mortgage equivalent yield** or **MEY** is the annualized yield with monthly compounding.

Example 3.1.1 With a nominal annual interest rate of 10% compounded twice per annum, each dollar will grow to be $(1 + (0.1/2))^2 = 1.1025$. The quoted annual rate is therefore equivalent to an interest rate of 10.25% compounded once per annum; 10.25% is the effective annual rate. □

Example 3.1.2 An insurance company has to pay 20 million dollars four years from now to pensioners. Suppose it can invest money at an annual rate of 7% compounded semiannually. How much should it invest today? Since the effective annual rate is $(1 + (0.07/2))^2 - 1 = 7.1225\%$, the present value of 20 million is $20,000,000 \times (1.071225)^{-4} = 15,188,231$. □

[1] For an interesting investigation of the sources of interest, see [685, Chapter V]. Written by Schumpeter at the age of 28, this book is one of the most important works in economics.

[2] The idea of present value is due to Irving Fisher (1867–1947) in 1896 [557].

Example 3.1.3 A person starts with P dollars. She is able to invest in a security that pays an effective interest rate of r_1 for n_1 years. At the end of the n_1th year, she expects to reinvest the proceeds at an interest rate of r_2 for n_2 more years. So she expects $P(1+r_1)^{n_1}(1+r_2)^{n_2}$ dollars after n_1+n_2 years. Note that r_1 and r_2 are expressed in decimal, not percentage. □

An interest rate of r compounded m times a year is equivalent to an interest rate of r/m per $1/m$ year by definition. If a security offers a return of 1% per month, then the nominal annual interest rate is 12% with monthly compounding. Usually, people use this number as a good approximation to the effective rate, which is harder to compute mentally. In the above example, the effective rate is 12.6825%. The disparity increases as the nominal rate increases.

In the limit as m approaches infinity and $\left(1+\frac{r}{m}\right)^m \to e^r$, we obtain **continuous compounding**,

$$\text{FV} = Pe^{rn},$$

where $e = 2.71828\ldots$. We shall call the compounding scheme of (3.1) **periodic** to differentiate it from continuous compounding. Numerically, continuous compounding is very close to daily compounding. For example, a rate of 8% compounded daily has an effective rate of 8.3278%. The effective rate becomes 8.3287% if the same rate is compounded continuously. Continuous compounding is also easy to work with algebraically. For instance, if both interest rates in Example 3.1.3 are compounded continuously, then the final amount would be $Pe^{r_1n_1+r_2n_2}$.

3.1.1 Efficient algorithms for present and future values

The present value of a cash flow, $C_1, \ldots, C_n$, is

$$\frac{C_1}{1+y} + \frac{C_2}{(1+y)^2} + \cdots + \frac{C_n}{(1+y)^n},$$

which can be computed by the algorithm in Fig. 3.1. One can easily verify that the variable d is equal to $(1+y)^{-i}$ at the beginning of the `for`-loop. As a result, the variable x becomes the partial sum, $\sum_{t=1}^{i} C_t(1+y)^{-t}$, at the end of each loop. This proves the correctness of the algorithm.

Algorithm for evaluating present value:

```
input: y, n, C_t (1 ≤ t ≤ n);
real   x, d;
x := 0;
d := 1 + y;
for i = 1 to n {
        1. x := x + (C_i/d);
        2. d := d × (1 + y);
}
return x;
```

Figure 3.1: ALGORITHM FOR PRESENT VALUE. C_t are the cash flows, y is the yield or interest rate, and n is the term of the investment. Typically, C_t are stored in an n-element array.

The computational complexity is $O(n)$ since the bulk of the computation lies in the four arithmetic operations during each execution of the loop, which is executed n times. One can save one arithmetic operation within the loop by creating a new variable, say z, and assigning $1+y$ to it before the loop. The statement $d := d \times (1+y)$ can then be replaced by $d := d \times z$. Since modern compilers often perform such optimization automatically, we will value readability over obvious optimizations in presenting algorithms, relying on the compiler to eliminate redundancies behind the scene. This observation also lends support to the argument for asymptotic analysis; in a complex environment where a lot of manipulations are done without our knowing them, the best we can do is often the trend, not specific numbers. We therefore will not take pains to optimize constant factors in the complexity.

One further simplification is to replace the loop with the single statement,

```
for i = n down to 1
        x := (x + C_i)/d;
```

The above tight loop computes the present value in the following way,

$$\left(\cdots\left(\left(\frac{C_n}{1+y}+C_{n-1}\right)\frac{1}{1+y}+C_{n-2}\right)\frac{1}{1+y}+\cdots\right)\frac{1}{1+y}.$$

This scheme, due to Horner in 1819 [510], is the most efficient possible in terms of the absolute number of arithmetic operations [93].

Computing future value is almost identical to the algorithm in Fig. 3.1. One simply uses the same algorithm with the following changes: (1) d is initialized to 1 instead of $1+y$, (2) i should start from n and run down to 1, and (3) $x := x + (C_i/d)$ is replaced by $x := x + (C_i \times d)$.

3.1.2 Conversion between compounding methods

Interest rates with different compounding methods can be compared by converting one into the other. Suppose r_1 is the interest rate with continuous compounding and r_2 is the equivalent rate compounded m times per annum. Then,

$$\left(1+\frac{r_2}{m}\right)^m = e^{r_1}.$$

Therefore,

$$r_1 = m \ln\left(1+\frac{r_2}{m}\right) \tag{3.2}$$

$$r_2 = m\left(e^{r_1/m}-1\right) \tag{3.3}$$

Example 3.1.4 Consider an interest rate at 10% with quarterly compounding. The equivalent rate with continuous compounding is

$$4 \times \ln\left(1+\frac{0.1}{4}\right) = 0.09877, \quad \text{or} \quad 9.877\%,$$

derived from (3.2) with $m=4$ and $r_2 = 0.1$. □

For a set of n compounding methods, there is a total of $n(n-1)$ possible pairwise conversions. Such huge numbers of cases usually invite programming errors and oversight. To make that number manageable, one can fix a ground case, say, continuous compounding, and then convert rates to their continuously compounded equivalents before comparison. This cuts down the number of possible conversions to $2(n-1)$, a more enviable situation.

3.1.3 Simple compounding

Besides periodic compounding and continuous compounding (hence **compound** interest), there is a different scheme for computing interest called **simple compounding** (hence **simple** interest). Under this scheme, interest is computed on the original principal; in other words, there is no interest on interest. Suppose P is borrowed at an annual rate of r. The simple total interest on P for a period of n years is Prn. Compare this with $P(1+r)^n - P$ under compound interest. The two numbers agree only if $n=1$ or $r=0$. To convert a simple interest rate r_{s} into an equivalent compound interest rate r_{c}, solve

$$P + Pr_{\mathsf{s}}n = P(1+r_{\mathsf{c}})^n$$

for r_{c}.

3.2 Annuity

A contract that pays out C dollars at the end of each year for n years is called **ordinary annuity**. With a nominal rate of r, the future value of this annuity at the end of the nth year is

$$\sum_{i=0}^{n-1} C(1+r)^i = C\,\frac{(1+r)^n - 1}{r}. \tag{3.4}$$

For the **annuity due**, cash flows are received at the beginning of each year, and the future value is

$$\sum_{i=1}^{n} C(1+r)^i = C\,\frac{(1+r)^n - 1}{r}\,(1+r). \tag{3.5}$$

If m payments are received per year (the so-called **general annuity**), and each payment has the amount of C, then (3.4) and (3.5) can be generalized to

$$C\,\frac{\left(1+\frac{r}{m}\right)^{nm} - 1}{\frac{r}{m}} \quad \text{and} \quad C\,\frac{\left(1+\frac{r}{m}\right)^{nm} - 1}{\frac{r}{m}}\left(1+\frac{r}{m}\right),$$

respectively. Unless stated otherwise, we assume the payment from an annuity is made at the end of each compounding period.

The present value of an annuity can be derived easily. If the payments are received m times a year, the present value is

$$\text{PV} = \sum_{i=1}^{nm} C\left(1+\frac{r}{m}\right)^{-i} = C\,\frac{1-\left(1+\frac{r}{m}\right)^{-nm}}{\frac{r}{m}}. \tag{3.6}$$

Example 3.2.1 The present value of an annuity of \$100 per annum for five years at an annual interest rate of 6.25% is

$$100 \times \frac{1-(1.0625)^{-5}}{0.0625} = 418.387$$

based on (3.6) with $m = 1$. □

Example 3.2.2 Suppose an annuity pays \$5,000 per month for nine years with an interest rate of 7.125% compounded monthly. Its present value can be derived from (3.6) with $C = 5000$, $r = 0.07125$, $n = 9$, and $m = 12$. The result is \$397,783. □

An annuity that lasts forever is called **perpetual annuity**. Its present value can be derived from (3.6) by letting n go to infinity as

$$\text{PV} = \frac{mC}{r}. \tag{3.7}$$

The above formula is useful for valuing perpetual fixed-coupon debts [557].

Example 3.2.3 A financial instrument promises to pay \$100 once a year forever. If the interest rate is 10%, the present value from (3.7) is $\frac{100}{0.1} = 1000$. If the price of the instrument is more than \$1,000, then the return rate is less than 10%. □

3.3 Amortization

Amortization is a method of repaying an interest-bearing loan through regular payments of interest *and* principal. Each payment consists of two parts. The size of the loan, or the **original balance**, is reduced by the principal part of the payment. The interest part of the payment pays the interest incurred on the **remaining principal balance**. As the principal gets paid down[3] over the term of the loan, the interest part of the payment diminishes as a result.

Amortization is typically used by home mortgages. By paying down the principal consistently, it reduces the risk to the lender. When the borrower sells the house, the remaining principal is due the lender. So, if the house sells for less than that amount net the commission, the borrower would be taking a capital loss. We consider for the rest of this section mainly the equal-payment case, i.e., fixed-rate, **level-payment**, **fully amortized** mortgages. Such mortgages are also known as **traditional mortgages** [288].

Example 3.3.1 A home buyer takes out a 15-year, \$250,000 loan at 8.0% interest rate. Applying (3.6) with $\text{PV} = 250000$, $n = 15$, $m = 12$, and $r = 0.08$ gives a monthly payment of \$2,389.13. The amortization schedule is shown in Fig. 3.2. One can verify that, in any month, (1) the principal and interest parts of the payment add up to \$2,389.13; (2) the remaining principal is reduced by the amount indicated under the Principal heading; and (3) the interest is computed by multiplying the remaining balance of the previous month by 0.08/12. The trend is for the interest payment to go down and the principal payment to go up. □

Month	Payment	Interest	Principal	Remaining principal
				250,000.000
1	2,389.13	1,666.667	722.464	249,277.536
2	2,389.13	1,661.850	727.280	248,550.256
3	2,389.13	1,657.002	732.129	247,818.128
		...		
178	2,389.13	47.153	2,341.980	4,730.899
179	2,389.13	31.539	2,357.591	2,373.308
180	2,389.13	15.822	2,373.308	0.000
Total	430,043.438	180,043.438	250,000.000	

Figure 3.2: AN AMORTIZATION SCHEDULE. See Example 3.3.1.

The remaining principal at any time is the present value of all the future payments. Let us be more precise. With the same notation as in the case of annuity, the amortization schedule allows the lender to receive C dollars for n years with m payments a year. Assume the annual interest rate is r. Right after the kth payment, the remaining principal is the present value of all the future $nm - k$ cash flows,

$$\sum_{i=1}^{nm-k} C\left(1+\frac{r}{m}\right)^{-i} = C\,\frac{1-\left(1+\frac{r}{m}\right)^{-nm+k}}{\frac{r}{m}}. \tag{3.8}$$

Example 3.3.2 With $C = 2389.13$, $n = 15$, $m = 12$, $r = 0.08$, and $k = 3$, Eq. (3.8) generates the same remaining principal as in the amortization schedule of Example 3.3.1 for the third month. □

A popular mortgage is the **adjustable-rate mortgage** or **ARM** [46]. The interest rate is no longer fixed but tied to some publicly available index such as the CMT (constant maturity Treasury) rate or COFI (Cost of Funds Index). The attractiveness of ARMs arises from the typically lower initial rate, thus qualifying the home buyer for a bigger mortgage, and the fact that the interest rate adjustments are capped. Chapters 28–30 cover mortgages in more details.

3.4 Sinking Fund

Another common method of paying off a long-term loan is for the borrower to pay interest on the loan and to pay into a sinking fund so that the debt can be retired with proceeds from the fund. The sum of the interest payment and the sinking-fund deposit is called the **periodic expense** of the debt. In the case where the sinking fund stipulates equal periodic expense and retirement of the debt at the end of the term, if the interest on the sinking fund is the same as that on the principal, then amortization and sinking fund methods produce the same periodic payment. In practice, sinking fund provisions vary. Some start several

[3] There are arrangements whereby the remaining principal actually increases and then decreases over the term (maturity) of the loan. The same principle applies here, too. See Exercise 3.3.2.

years after the issuance of the debt, others allow a balloon payment at maturity, and still others use the fund to periodically purchase bonds in the market [666].

Example 3.4.1 A company borrows \$100,000 at a semiannual interest rate of 10%. If the company pays into a sinking fund earning 8% to retire the debt in seven years, the semiannual payment can be calculated by (3.6) as follows,

$$\frac{\mathrm{PV} \times \frac{r}{m}}{1 - \left(1 + \frac{r}{m}\right)^{-nm}}$$

with $\mathrm{PV} = 100000$, $r = 0.08$, $m = 2$, and $n = 7$. The result is \$9,466.9. Interest on the loan is $100000 \times (0.1/2) = 5000$ semiannually. The periodic expense is thus $5000 + 9466.9 = 14466.9$. □

Exercise 3.4.1 In Keynes's masterpiece, *The Economic Consequences of the Peace*, Germany's capacity to pay Reparation demanded by the Allies during the Peace Conference of 1918 is calculated as follows [483, p. 200].

> [Suppose Germany could be made to pay \$500,000,000 annually for 30 years.] Such a figure, allowing 5 per cent for interest, and 1 per cent for repayment of capital, represents a capital sum having a present value of about \$8,500,000,000.

The footnote adds,

> If the amount of the sinking fund be reduced, and the annual payment is continued over a greater number of years, the present value—so powerful is the operation of compound interest—cannot be materially increased. A payment of \$500,000,000 annually *in perpetuity*, assuming interest, as before, at 5 per cent, would only raise the present value to \$10,000,000,000.

Verify Keynes's numbers.[4] ◇

3.5 Yields

The term **yield** refers to the return of an investment and is typically expressed as a percentage [245]. It is quoted in many ways, and one has to be careful as to which yield is meant. On *The Wall Street Journal* on August 26, 1997, for instance, a corporate bond is quoted as follows.

Company	Cur Yld.	Vol	Close	Net Chg.
ATT85/831	8.1	162	1061/2	−3/8

This is a bond issued by AT&T (American Telephone & Telegraph), maturing in year 2031 with a **nominal yield** of $8\frac{5}{8}$%, which is the same as the coupon of the bond and is part of the identification of the bond. For example, in the same paper, one can find other AT&T bonds: ATT43/498, ATT6s00, ATT51/801, ATT63/404, etc.

[4]Keynes (1883–1946) is recognized as among the greatest economists in history. See [706, 707] for a passionate and masterfully written biography.

The **current yield** of a bond is the coupon rate of interest divided by the current market price. In the above case, the coupon is $8\frac{5}{8}\%$. Hence, the annual interest is $8\,\frac{5}{8} \times 1000/100 = 86.25$, assuming a par value of \$1,000. The closing price is $106\,\frac{1}{2} \times 1000/100 = 1065$. The market prices of corporate bonds are quoted as a percentage of par value. Using the closing price as the purchase price, we obtain $86.25/1065 \approx 8.1\%$ as its current yield.

The above two yield measures are of little use in comparing returns. For example, the nominal yield completely ignores the market condition, while the current yield does not take the future into account even though it does depend on the current market price. One possible exception to the latter criticism is when the bond can be treated as a perpetual annuity, which shares the same formula as the current yield.

Another yield measure is related to securities like the U.S. **Treasury bills** (or **T-bills**) that pay interest based on the **discount method** rather than the **add-on method** [88]. With such a method, interest is subtracted from the par value of a security to derive the purchase price, and the investor receives the par value at maturity. Such securities are said to be issued **on a discount basis**. The **discount yield** or **discount rate** is simply

$$\frac{\text{par value} - \text{purchase price}}{\text{par value}} \times \frac{360 \text{ days}}{\text{number of days to maturity}}. \tag{3.9}$$

This yield is also called the **yield on a bank discount basis**. For the purpose of calculating this yield for short-term securities, a year is assumed to have 360 days [603, 730]. We will have more to say about day count conventions later on.

Example 3.5.1 Treasury bills are a short-term debt instrument with maturities of three months, six months, and twelve months. Treasury bills are issued in US\$10,000 denominations. Suppose an investor buys a US\$10,000, six-month Treasury bill for US\$9,521.45 with 182 days remaining to maturity. The discount yield is

$$\left(\frac{10000 - 9521.45}{10000}\right) \times \left(\frac{360}{182}\right) = 0.0947,$$

or 9.47%. It is this annualized yield on a bank discount basis that will be quoted. □

Example 3.5.2 A Treasury bill selling for \$98,000 with 100 days to maturity and a face value of \$100,000 will be quoted at 7.2% because

$$\frac{100,000 - 98,000}{100,000} \times \frac{360}{100} = 0.072.$$

For comparison, its effective yield with continuous compounding is

$$\frac{365}{100} \times \ln\left(\frac{100000}{98000}\right) = 0.07374,$$

or 7.374%. □

One commonly used method to make the discount rate more comparable to yield quotes for other money market instruments is the **CD-equivalent yield**, also called the **money market-equivalent yield**. It is computed as

$$\frac{360 \times \text{discount yield}}{360 - (\text{number of days to maturity} \times \text{discount yield})}.$$

This equation can be derived as follows. The CD-equivalent yield is a simple annualized interest rate defined as

$$\frac{\text{par value} - \text{purchase price}}{\text{purchase price}} \times \frac{360}{\text{number of days to maturity}}.$$

Now, plug in the discount rate formula of (3.9) and simplify.

To make the discount rate more comparable to the bond equivalent yield, compute

$$\frac{\text{par value} - \text{purchase price}}{\text{purchase price}} \times \frac{365}{\text{number of days to maturity}}.$$

For example, the discount rate in Example 3.5.2 (7.2%) becomes

$$\frac{2}{98} \times \frac{365}{100} = 0.0745, \quad \text{or} \quad 7.45\%. \tag{3.10}$$

Treasury bill's **ask yield** is computed in precisely this way [456].

The yield measure we will be considering for the rest of this section, unless stated otherwise, is the so-called **internal rate of return** (**IRR**). It is the interest rate which equates the present value of future cash flows of an investment and its price [245]. Mathematically, we are referring to the following identity,

$$P = \frac{C_1}{(1+y)} + \frac{C_2}{(1+y)^2} + \frac{C_3}{(1+y)^3} + \cdots + \frac{C_n}{(1+y)^n}, \tag{3.11}$$

where C_t is the cash flow in (at the *end* of) year t, P is the price, and n is the number of years during which the investment generates cash flows. Note that the right-hand side is the present value of the cash flows discounted at the internal rate of return y. Simple as it is, (3.11) and its various generalizations form the foundation upon which pricing methodologies are built [281].

Example 3.5.3 Suppose a financial instrument promises to pay \$1,000 for the next three years and is selling for \$2,500. Its yield is 9.7%. This can be verified as follows. Using $r = 0.097$ as the discounting rate, the present values of the three cash flows are $1000/(1+0.097)^t$ for $t = 1, 2, 3$. These three numbers are \$911.577, \$830.973, and \$757.5, which sum to \$2,500. □

The above example shows it is easy to verify if a yield equates the price with a set of cash flows. Finding the yield, however, requires numerical techniques because closed-form formulae in general do not exist. This issue will be picked up in §3.5.2. One exception is when there is a single, known cash flow as in the following example.

Example 3.5.4 A financial instrument can be bought for \$1,000, and the investor will end up with \$2,000 five years from now. The yield is simply the y that equates 1000 with $2000 \times (1+y)^{-5}$, the present value of \$2,000 five years hence. The yield is thus simply

$$\left(\frac{1000}{2000}\right)^{-1/5} - 1,$$

which is 14.87%. □

Recall that yield is the interest rate that equates future cash flows with the present value. Given a set of cash flows, we can calculate their future value as

$$\mathrm{FV} = \sum_{t=1}^{n} C_t \,(1+y)^{n-t}. \tag{3.12}$$

By (3.11), yield is simply the y that makes

$$\mathrm{FV} = \mathrm{PV} \times (1+y)^n$$

hold. Hence, in principle, multiple cash flows can be reduced to a single cash flow $\mathrm{PV} \times (1+y)^n$ at maturity. In Example 3.5.4, for instance, we did not care how the investor ended up with \$2,000 at the end of the fifth year. He might simply receive that much money then, or he might receive periodical payments which, with suitable arrangements, grew to \$2,000 eventually. This brings us to the second point. Look at (3.11) and (3.12) again. The reason they say the same thing is because it is implicitly assumed that all cash flows before the final date are reinvested at the *same* rate as the internal rate of return, y. This point is clear in (3.12), in which each cash flow C_t is compounded at the rate of y.

Example 3.5.4 therefore enunciates a general principle for yield computation: Calculate the future value and then find the yield that equates it with the present value. It is no longer mandatory that all cash flows be compounded at the same internal rate of return; in fact, a different reinvestment rate for each cash flow is possible. The yield thus derived can be more realistic than the internal rate of return due to the explicit reinvestment assumptions. This idea is behind the **holding period return** (**HPR**) methodology.[5]

Suppose the reinvestment rate has been exogenously determined to be r_e. Then the future value of (3.12) becomes

$$\mathrm{FV} = \sum_{t=1}^{n} C_t \,(1+r_e)^{n-t}.$$

We can then solve for the internal rate of return y such that $\mathrm{FV} = P\,(1+y)^n$.

Of course, if the reinvestment assumptions turn out to be wrong, then the yield will not be realized. Incidentally, this shows financial instruments without intermediate cash flows do not have **reinvestment risks**. They are perfect instruments to meet future obligations with certainty. As it is clearly impossible to know future reinvestment opportunities for sure, the internal rate of return may well be a reasonable compromise.

Example 3.5.5 A financial instrument promises to pay \$1,000 for the next three years and is selling for \$2,500. The investor figures that each cash flow can be put into a bank account paying interest at an effective rate of 5%. The future value of such an investment is

$$\sum_{t=1}^{3} 1000 \times (1+0.05)^{3-t} = 3152.5.$$

[5] Other names with identical connotation include **total return**, **horizon return**, **realized compound yield**, **effective yield**, **holding period yield**, **horizon total return**, and **investment horizon return** [557].

The return on this investment is thus $\left(\frac{3152.5}{2500}\right)^{1/3} - 1 = 0.08037$, or 8.037%. This return rate is considerably lower than the 9.7% derived in Example 3.5.3. The difference, as we mentioned above, comes from the implicit assumption in the internal rate of return that cash flows can be reinvested at an annual rate of 9.7%. Our investor above considers such a rate as being too rosy and scales it down to 5%. □

Cash flows may not be due at intervals of exactly one year. Monthly, as with mortgages, or semiannual, as with most bonds, cash flows are quite common. Equation (3.11) can be generalized without difficulty to

$$P = \sum_{t=1}^{n} \frac{C_t}{(1+y)^t}, \tag{3.13}$$

where C_t is the cash flow at *time* t (at the end of period t), P is the price of the investment, and n is the number of periods during which the investment generates cash flows. Note that y is now the period yield. Equivalently, $m \times y$ is the annual yield, compounded m times per annum, where m is the number of payments per year. For instance, if the cash flows are semiannual, then the computed y is a semiannual yield. The annual yield in this case would be obtained by doubling y.

In Example 3.5.3, suppose we were told of the same figures except that the cash flows occurred semiannually instead of annually. Then, the yield 9.7% had to be interpreted as a semiannual yield. The annual yield was thus $2 \times 0.097 = 0.194$, or 19.4%.

Example 3.5.6 A bank just lent a borrower \$260,000 for 15 years to purchase a house. This 15-year mortgage has a monthly payment of \$2,000. The annul interest rate is 4.583% because

$$\sum_{i=1}^{12\times 15} 2000 \times \left(1 + \frac{0.04583}{12}\right)^{-i} \approx 260000$$

from (3.13) with $C_t = 2000$, $n = 12 \times 15$, and $y = 0.04583/12$. □

Some people questioned the validity of the conversion formulae between different compounding schemes in §3.1.2, arguing for instance that 10% compounded quarterly can't possibly be the same as 9.877% compounded continuously as Example 3.1.4 implied, because one pays interest quarterly, whereas the other does so continuously. This objection *is* well-founded. In fact, such equivalence assumes interests are reinvested at their respective rates. The correct attitude to those conversion formulae is to treat them as *definitions* with which yields of different compounding frequencies can be compared.

3.5.1 Net present value (NPV)

An investment with cash flows $C_1, C_2, \ldots, C_n$ and selling for P has an internal rate of return y that satisfies (3.11). Suppose an investor believes this security should have a return rate of y^*. For this person, the security is said to have a **net present value** of

$$\sum_{t=1}^{n} \frac{C_t}{(1+y^*)^t} - P. \tag{3.14}$$

The net present value vanishes when $y = y^*$; in other words, the internal rate of return is the return rate that nullifies the net present value. In general, net present value is the difference between the present values of cash inflows and cash outflows. Businesses are often assumed to maximize their assets' net present values [687].

Example 3.5.7 The management is presented with the following two proposals.

Proposal	Investment now	Net cash flow at end of		
		Year 1	Year 2	Year 3
A	9,500	4,500	2,000	6,000
B	6,000	2,500	1,000	5,000

It believes the company can earn 15% effective on projects of this kind. The net present value for Proposal A is

$$\frac{4500}{1.15} + \frac{2000}{(1.15)^2} + \frac{6000}{(1.15)^3} - 9500 = -129.57$$

and that for Proposal B is

$$\frac{2500}{1.15} + \frac{1000}{(1.15)^2} + \frac{5000}{(1.15)^3} - 6000 = 217.64.$$

Proposal A is therefore dropped in favor of Project B. □

3.5.2 Numerical methods for finding yields

Let

$$f(y) \equiv \sum_{t=1}^{n} \frac{C_t}{(1+y)^t}.$$

(The symbol $\equiv$ introduces definitions.) The yield problem amounts to solving

$$f(y) - P = 0. \tag{3.15}$$

It will be assumed that $y > -1$.

Function $f(y)$ is monotonic in y if the C_t's are all positive. In this case, a simple geometric argument in Fig. 3.3 shows that a unique solution exists. Even in the general case where all the C_t's are not of the same sign, usually only one value makes economic sense [479, §5.3]. That something exists is not enough; one has to find it or at least approximate it. We now turn to this algorithmic problem.

The bisection method

One of the simplest and failure-free methods to solve equations like (3.15) for any well-behaved function is the **bisection method** [191]. Observe that $f(a)f(b) < 0$ implies $f(\xi) = 0$ for some ξ between a and b written as $\xi \in [a, b]$. If we then evaluate f at the mid-point c between a and b, then either (1) $f(c) = 0$, (2) $f(a)f(c) < 0$, or (3) $f(c)f(b) < 0$. In the first case we are done, in the second case we continue the process with the new bracket $[a, c]$, and in the third case we continue with $[c, b]$. Note that the bracket

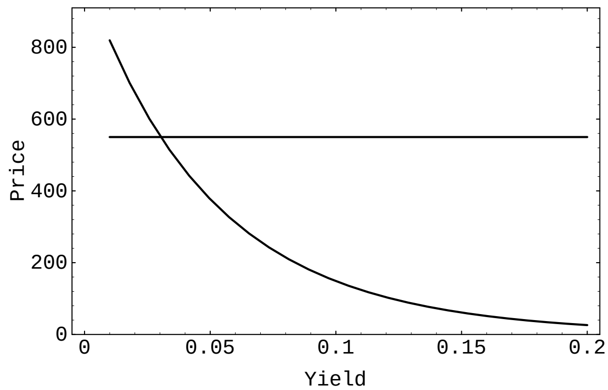

Figure 3.3: COMPUTING YIELDS. Yield is the number at which the current market price, represented by the horizontal line, equals the present value of future cash flows, represented by the downward-sloping curve. It is the value on the x-axis at which the two curves intersect.

The bisection method for solving equations:

```
input: ε, a, and b ( b > a and f(a)f(b) < 0 );
real   length, c;
length := b − a;
while [length > ε] {
        1. c := (b + a)/2;
        2. if [ f(c) = 0 ] return c;
        3. else if  [ f(a)f(c) < 0 ] b := c;
        4. else  a := c;
}
return c;
```

Figure 3.4: BISECTION METHOD. The number ϵ is an upper bound on the absolute error of the returned value c: $f(\xi) = 0$ for some ξ such that $|\xi - c| \leq \epsilon$. The bracket $[a, b]$ guarantees the existence of a root with the $f(a)f(b) < 0$ condition.

is halved. After n steps, we will have confined ξ within a bracket of length $(b - a)/2^n$. Figure 3.4 codes the idea.

The complexity of the bisection algorithm can be analyzed as follows. The `while`-loop is executed at most $1 + \log_2\left(\frac{b-a}{\epsilon}\right)$ times. Within the loop, the number of arithmetic operations is dominated by the evaluation of f. Denote this number by $\mathcal{C}_f$. The computational complexity is thus $O\left(\mathcal{C}_f \log_2\left(\frac{b-a}{\epsilon}\right)\right)$. In computing the internal rate of return, the complexity is $O\left(n \log_2\left(\frac{b-a}{\epsilon}\right)\right)$ since $\mathcal{C}_f = O(n)$ by the algorithm in Fig. 3.1.

The Newton-Raphson method

A method that converges faster than the bisection method is the **Newton-Raphson method** [400]. This method is iterative. In iterative methods, one starts with a first approximation, x_0, to a root. Successive approximations are then computed as

$$x_0, F(x_0), F(F(x_0)), \ldots$$

for some function F. In other words, if we let x_k denote the kth approximation, then $x_k = F^{(k)}(x_0)$, where

$$F^{(k)}(x) \equiv \overbrace{F(F(\cdots(F}^{k}(x))\cdots)).$$

The necessary condition for the convergence of such a procedure to a root ξ is

$$|F'(\xi)| \leq 1, \tag{3.16}$$

where F' denotes the derivative of F [400].

The Newton-Raphson method picks

$$F(x) \equiv x - \frac{f(x)}{f'(x)}.$$

In other words,

$$x_{k+1} \equiv x_k - \frac{f(x_k)}{f'(x_k)} \tag{3.17}$$

is the $(k+1)$st approximation. This is the method of choice when f' can be evaluated efficiently and is non-zero near the root [632]. See Fig. 3.5 for illustration and Fig. 3.6 for the algorithm. In the particular case of computing yields,

$$f'(x) = -\sum_{i=1}^{n} \frac{tC_t}{(1+x)^{t+1}}.$$

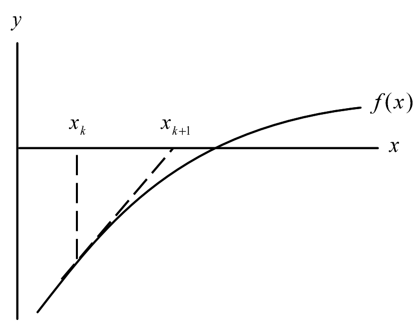

Figure 3.5: NEWTON-RAPHSON METHOD.

The Newton-Raphson method for solving equations:

```
input: ε, x_initial;
real   x_new, x_old;
x_old := x_initial;
x_new := ∞;
while [ | x_new − x_old | > ε ]
        x_new = x_old − f(x_old)/f'(x_old);
return x_new;
```

Figure 3.6: ALGORITHM FOR THE NEWTON-RAPHSON METHOD. One should start with a good initial guess [632].

Example 3.5.8 Suppose $f(x) = x^3 - x^2$. Then,

$$x_{k+1} = x_k - \frac{f(x_k)}{f'(x_k)} = x_k - \frac{x_k^3 - x_k^2}{3x_k^2 - 2x_k} = x_k - \frac{x_k^2 - x_k}{3x_k - 2}$$

by (3.17). □

Assume that we start with an x_0 near ξ. It can be shown that

$$\xi - x_{k+1} \approx -(\xi - x_k)^2 \frac{f''(\xi)}{2f'(\xi)}.$$

The method hence converges **quadratically**: Near the root, each iteration roughly doubles the number of significant digits. To achieve $|x_{k+1} - x_k| \leq \epsilon$ required by the algorithm, roughly $O(\log\log(1/\epsilon))$ iterations suffice. The computational complexity is thus

$$O\left((\mathcal{C}_f + \mathcal{C}_{f'}) \log\log(1/\epsilon) \right).$$

In particular, it implies a complexity of $O(n \log\log(1/\epsilon))$ for computing yields. This bound compares favorably with the $O\left(n \log\left((b-a)/\epsilon\right)\right)$ bound for the bisection method. In practice, one should put an upper bound on the number of iterations the routine is allowed to execute its loop. See [632, Chapter 9] for other practical algorithms and implementations.

A variant of the Newton-Raphson method that does not require differentiation is the **secant method** [29]. This method starts with two approximations, x_0 and x_1, and computes the $(k+1)$st approximation by replacing the derivative in (3.17) with

$$x_{k+1} = x_k - \frac{f(x_k)(x_k - x_{k-1})}{f(x_k) - f(x_{k-1})}.$$

This method may be preferred when the calculation of f' is to be avoided. Its convergence rate is slightly worse than the Newton-Raphson method.

3.5.3 Solving systems of nonlinear equations

The Newton-Raphson method can be extended to higher dimensions. Consider the two-dimensional case. Let (x_k, y_k) be the kth approximation to the solution of the two simultaneous equations,

$$f(x,y) = 0 \quad \text{and} \quad g(x,y) = 0.$$

The $(k+1)$st approximation (x_{k+1}, y_{k+1}) then satisfies the following linear equations,

$$\left[\begin{array}{cc} \partial f(x_k, y_k)/\partial x & \partial f(x_k, y_k)/\partial y \\ \partial g(x_k, y_k)/\partial x & \partial g(x_k, y_k)/\partial y \end{array}\right] \left[\begin{array}{c} \Delta x_{k+1} \\ \Delta y_{k+1} \end{array}\right] = -\left[\begin{array}{c} f(x_k, y_k) \\ g(x_k, y_k) \end{array}\right], \tag{3.18}$$

where $\Delta x_{k+1} \equiv x_{k+1} - x_k$ and $\Delta y_{k+1} \equiv y_{k+1} - y_k$. The above equations have a unique solution for $(\Delta x_{k+1}, \Delta y_{k+1})$ when the **Jacobian determinant** of f and g, defined as

$$J \equiv \left| \begin{array}{cc} \partial f/\partial x & \partial f/\partial y \\ \partial g/\partial x & \partial g/\partial y \end{array} \right|,$$

does not vanish at (x_k, y_k). The $(k+1)$st approximation is simply

$$(x_k + \Delta x_{k+1}, y_k + \Delta y_{k+1}).$$

Solving nonlinear equations has hence been reduced to solving a set of linear equations. Generalization to n dimensions is straightforward.

The bisection method may be applied where the Newton-Raphson method fails under $J = 0$. Let $y_f(x)$ be the y that makes $f(x,y) = 0$ given x. Similarly, let $y_g(x)$ be the y that makes $g(x,y) = 0$ given x. Assume y_f and y_g are continuous functions. Then the bisection method is applicable to the function $\phi(x) \equiv y_f(x) - y_g(x)$ if it starts with x_1 and x_2 such that $\phi(x_1)\phi(x_2) < 0$. In so doing, a two-dimensional problem is reduced to a one-dimensional problem. Unfortunately, this method has no obvious generalizations to $n > 2$ dimensions. Note that $y_f(x)$ and $y_g(x)$ may have to be numerically solved.

3.6 Bonds

A bond is a contract between the issuer (borrower) and the bondholder (lender). The issuer promises to pay the bondholder interest, if any, and principal on the remaining balance. Bonds usually refer to long-term debts. A bond has a **face value** (usually \$1,000 in the U.S), also called **denomination**, **par value**, **maturity value**, or **principal value**. The **redemption date** or **maturity (date)** specifies the date on which the loan will be repaid. A bond pays interest at the **coupon rate** on its par value at regular time intervals until the maturity date. The payment is usually made semiannually in the U.S. The **redemption value** is the amount to be paid at a redemption date. A bond is **redeemed at par**, which is usually the case for noncallable bonds, if the redemption value is the same as the par value. Redemption date may be different from maturity date [479].

There are several ways to redeem or retire a bond. A bond is **redeemed at maturity** if the principal is repaid at maturity. Most corporate bonds are **callable**, meaning the issuer can redeem some or all of the bonds before the stated maturity and usually at a price above the par value.[6] Since this provision gives the issuer the advantage of calling a bond when the prevailing interest rate is much lower than the coupon rate, the bondholders usually demand premium. A callable bond may also have **call protection** so that it is not callable for the first few years. **Refunding** involves using the proceeds from the issuance of new bonds to retire old ones. A corporation may deposit money into a sinking fund and use the funds to buy back some or all of the bonds at par value or from the open market [666]. A 80% sinker, for example, retires 80% of the debt before maturity. **Convertible bonds** can be converted into the issuer's common stocks. **Consols** are bonds that have no maturity; they pay interest forever. A consol can be analyzed as a perpetual annuity, and its value and yield satisfy the simple relationship due to (3.7),

$$P = \frac{c}{r}, \tag{3.19}$$

where c denotes the interest payout per annum.

The U.S. bond market is the largest in the world. It consists of U.S. Treasury securities, U.S. agency securities (excluding agency pass-through securities), corporate bonds, Yankee bonds, municipal securities, mortgages, and mortgage-backed securities. **Agency securities** refer to those issued by either the U.S. Federal government agencies or Federal government-sponsored organizations. The mortgage market is usually the largest (US\$5,260 billion in 1997), followed by the market for U.S. Treasury securities (US\$3,457 billion in 1997).

Treasury securities with maturities of one year or less are discount securities called Treasury bills. Discount securities do not pay interest. Treasury securities with original maturities between two and ten years are called **Treasury notes** or **T-notes**, while those with maturities greater than ten years are called **Treasury bonds** or **T-bonds**. Both Treasury notes and bonds are coupon securities, paying interest every six months.

[6] See [666] for reasons companies issue callable bonds. Callable bonds have not been issued by the U.S. Treasury since February 1985 [283].

3.6.1 Valuation of bonds

Start first with the **pure discount bonds**, also called **zero-coupon bonds** or simply **zeros**. They promise a single payment in the future and are sold at a discount from its par value. No interest is paid. The price of a zero-coupon bond that pays F in n periods is simply $F/(1+r)^n$, where r is the market interest rate or the required yield per period. The required yield is based on the yields on comparable investments in the market. Such bonds can be bought to meet future obligations without reinvestment risk. They are also an important theoretical tool in the analysis of coupon bonds, which can be thought of as a series of zero-coupon bonds. Although the U.S. Treasury does not issue such bonds with maturities over one year, there were companies which specialize in **coupon stripping** to create stripped Treasury securities. This innovation became redundant when the U.S. Treasury facilitated the creation of such securities via the STRIPS (Separate Trading of Registered Interest and Principal Securities) program in 1985 [699].

Example 3.6.1 Suppose the interest rate is 8% compounded semiannually. A zero-coupon bond that pays the par value 20 years from now will be priced at $1/(1.04)^{40}$, or 20.83%, of its par value; it will be quoted as 20.83. If the interest rate is 9% instead, the same bond will be priced at only 17.19. If the bond matures in 10 years instead of 20, its price would be 45.64 with an 8% interest rate. Clearly, both the maturity and the market interest rate have a strong impact on the price of a zero-coupon bond. □

A **level-coupon bond** pays interest based on the coupon rate and the par value. It also pays the par value at maturity but does not pay any principal until the end. Its price is

$$\mathrm{PV} = \sum_{i=1}^{n} \frac{C}{(1+r)^i} + \frac{F}{(1+r)^n} = C\, \frac{1-(1+r)^{-n}}{r} + \frac{F}{(1+r)^n}, \tag{3.20}$$

where n is the number of periods before maturity, r is the period interest rate or the required yield, F is the par value, and C is the coupon. An equivalent formulation is

$$\mathrm{PV} = \sum_{i=1}^{n} \frac{F\frac{c}{m}}{\left(1+\frac{r}{m}\right)^i} + \frac{F}{(1+\frac{r}{m})^n} = F\,\frac{c}{m}\, \frac{1-\left(1+\frac{r}{m}\right)^{-n}}{\frac{r}{m}} + \frac{F}{\left(1+\frac{r}{m}\right)^n}, \tag{3.21}$$

where c is the coupon rate, m is the number of payments per year, and r is the *annual* interest rate compounded m times per annum. Note that n continues to denote the number of periods and $C = Fc/m$. Unless stated otherwise, bonds should be understood as level-coupon bonds.

Example 3.6.2 Consider a 20-year, 9% bond with the coupon paid semiannually. This means a payment of $1000 \times 0.09/2 = 45$ will be made every six months until maturity, and \$1,000 will be made at maturity. The price of such a bond can be computed from (3.21) with $n = 2 \times 20$, $r = 0.08$, $m = 2$, $F = 1$, and $C = 0.09/2$. The result is 1.09896, or 109.896% of par value. In general, when the coupon rate is higher than the interest rate, as is the case here, the bond will be selling above its par value. □

The (annualized) **yield to maturity** of a level-coupon bond is its internal rate of return when the bond is held to maturity. In other words, it is the r that satisfies (3.21), where PV is the total price of the bond. This yield assumes coupons are reinvested at the same rate as the yield to maturity. For an investor with an 15% required bond equivalent yield to maturity, a ten-year, $1,000 bond with a coupon rate of 10% paid semiannually should sell for

$$50 \times \frac{1-(1+(0.15/2))^{-2\times 10}}{0.15/2} + \frac{1000}{(1+(0.15/2))^{2\times 10}} = 745.138,$$

for instance.

For a callable bond, the **yield to stated maturity** measures its yield to maturity as if it were not callable. Another common measure for return is **yield to call**. It is the yield to maturity satisfied by (3.21) with n denoting the number of remaining coupon payments until the first call date and F replaced by the **call price**, the price at which the bond will be called. Yet another variation is the **yield to effective maturity**. which replaces n with the **effective maturity date**, the redemption date when the bond is called. Of course, this uncertain date has to be estimated. Other return measures for callable bonds include **yield to par call** and **yield to worst**. The former assumes the call price is the par value, and the latter is defined as the minimum of all the yields to call under all possible call dates.

3.6.2 Price behaviors

The price of a bond goes in the opposite direction to interest rate movements. Bond prices fall when interest rate rises, and vice versa. This is because the present value decreases as the interest rate increases. It is a common mistake to reverse this basic relationship.[7]

Equation (3.21) can be used to show that a level-coupon bond would be selling **at a premium** (above its par value) when its coupon rate is above the market interest rate, **at par** (at its par value) when its coupon rate is equal to the market interest rate, and **at a discount** (below its par value) when its coupon rate is below the market interest rate. The table in Fig. 3.7 shows the relationship between the price of a bond and the required yield, assuming a 15-year, 9%-coupon bond. Bonds selling at par are also called **par bonds**.

The price/yield relationship for any noncallable bond has a **convex** shape as depicted in Fig. 3.8. Convexity is attractive for bondholders because the price decrease per percent rate increase is smaller than the price increase per percent rate decrease. This observation holds for bonds without **embedded options** such as the call option in callable bonds. The convexity property has far-reaching implications for bonds and will be explored later.

A bond selling at a discount will see its price move up toward par as the maturity date draws near. A bond selling at par will see its price remain at par as the maturity date draws near. A bond selling at a premium will see its price move down toward par as the maturity date draws near. This phenomenon is depicted in Fig. 3.9.

[7] See, for example, "If Japanese banks are hit by a liquidity problem, they may have to sell U.S. Treasury bonds. A strong sell-off could have the effect of pushing down bond yields and rattling Wall Street." Source: [607].

Yield (%)	Price (% of par)
7.5	113.37
8.0	108.65
8.5	104.19
9.0	100.00
9.5	96.04
10.0	92.31
10.5	88.79

Figure 3.7: SOME PRICE/YIELD RELATIONSHIPS.

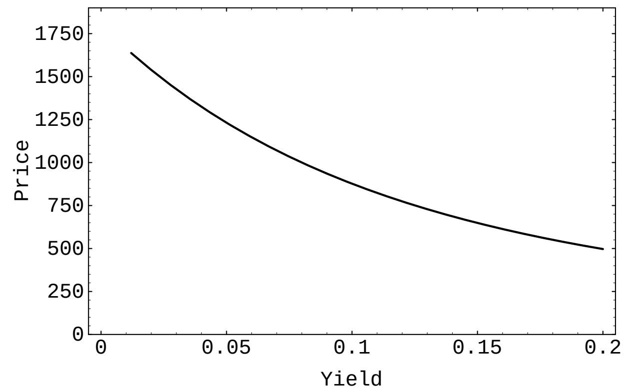

Figure 3.8: PRICE/YIELD RELATIONSHIP FOR BONDS. Plotted is a bond paying 8% interest on a par value of \$1,000, compounded annually. The term is 10 years.

We have cited two reasons for the change in the price of a bond: interest rate movements in the economy and a non-par bond moving toward maturity. Other reasons include (1) changes in the yield spread to Treasury bonds for non-Treasury bonds; (2) change in the perceived credit quality of the issuer; (3) change in the value of the embedded option.

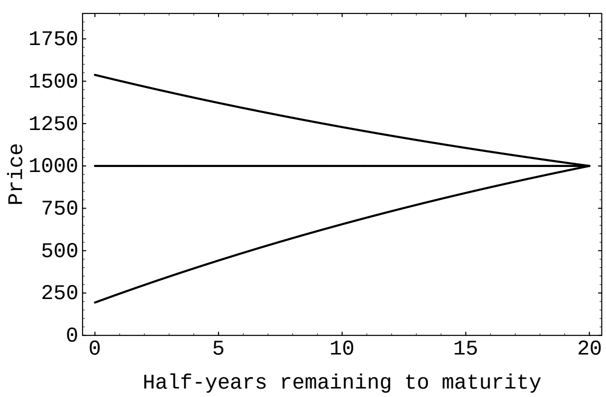

Figure 3.9: RELATIONSHIP BETWEEN PRICE AND TIME TO MATURITY FOR BONDS. Plotted are three curves for bonds, from top to bottom, selling at a premium, at par, and at a discount, with coupon rates of 12%, 6%, and 2%, respectively. The coupons are paid semiannually. The par value is \$1,000, and the required yield is 6%. The term is 10 years (the x-axis is measured in half-years).

3.6.3 Price quotes

Bonds are usually quoted as a percentage of par value. For Treasury notes and Treasury bonds, a quote of 100.05 means $100\frac{5}{32}\%$ of par value, not 100.05%. It is typically written as 100-05.

3.6.4 Day count conventions

Teach us to number our days aright,
that we may gain a heart of wisdom.
—Psalms 90:12

Right at the beginning of this chapter, we mentioned that one of the three essential ingredients in the evaluation of any investment is the times when cash flows occur. It is no exaggeration to say that correctly handling the issue of dating is often a major part of any financial software. We deal with day count conventions in this subsection.

In the so-called "**actual/actual**" day count convention, the first "actual" refers to the actual number of days in a month, and the second refers to the actual number of days in a coupon period. For example, for coupon-bearing Treasury securities, the number of days between June 17, 1992 and October 1, 1992 is 106 because there are 13 days in June, 31 days in July, 31 days in August, 30 days in September, and one day in October. As another example, the actual number of days in the coupon period between March 1, 1992 and September 1, 1992 is 184.

A convention popular with corporate and municipal bonds and agency securities is "**30/360**." Here, each month is assumed to have 30 days and each year 360 days. The number of days between June 17, 1992 and October 1, 1992 is now 104 because there are 13 days in June, 30 days in July, 30 days in August, 30 days in September, and one day in October. The number of days between two given dates under the "30/360" convention can be computed by

$$360 \times (y_2 - y_1) + 30 \times (m_2 - m_1) + (d_2 - d_1), \tag{3.22}$$

where y_i denote the years, m_i denote the months, and d_i denote the days [479].

Beneath the seemingly simplicity of (3.22) hide certain complications. What if one of the dates is the 31st? What about February, which has at most 29 days? One convention is this. If the beginning date $D_1 \equiv (y_1, m_1, d_1)$ is the last day of the month and $d_1 = 31$, change D_1 to be the first date of the following month. If the ending date $D_2 \equiv (y_2, m_2, d_2)$ is the last day of the month and $d_2 \neq 30$, change D_2 to $(y_2, m_2, 30)$. These two rules make sure a month has 30 days under the "30/360" convention. Such adjustments must be made before using (3.22). Consult Fig. 3.10 for concrete examples (dates so modified are in boldface).

Date		Date under 30/360	
From	To	From	To
July 30	March 15	July 30	March 15
July 31	March 15	**August 1**	March 15
August 1	March 15	August 1	March 15
February 28	March 15	February 28	March 15
February 29	March 15	February 29	March 15
July 30	August 31	July 30	**August 30**
July 31	August 31	**August 1**	**August 30**
August 1	August 31	August 1	**August 30**
February 1	February 29 (1996)	February 1	**February 30** (1996)
February 1	February 28 (1995)	February 1	**February 30** (1995)
February 1	February 28 (1996)	February 1	February 28 (1996)

Figure 3.10: "30/360" DAY COUNT EXAMPLES. Note: 1996 is a leap year.

3.6.5 Accrued interest

Up to now, we have assumed that the next coupon payment date is exactly one period (six months for bonds, for example) from now. In reality, the settlement date may fall on any day between two coupon payment dates, and yield measures have to be adjusted accordingly. Let

$$\omega \equiv \frac{\begin{array}{c}\text{number of days between the}\\ \text{settlement and the next coupon payment date}\end{array}}{\text{number of days in the coupon period}}, \tag{3.23}$$

where the day count is based on the convention applicable to the security in question. The price is now calculated as

$$\mathrm{PV} = \sum_{i=0}^{n-1} \frac{C}{\left(1+\frac{r}{m}\right)^{\omega+i}} + \frac{F}{\left(1+\frac{r}{m}\right)^{\omega+n-1}}, \tag{3.24}$$

where n is the number of remaining coupon payments, m is the number of coupon payments per year, r is the annual interest rate or the required yield, F is the par value, and C is the coupon [286]. This price is called the **full price**, **dirty price**, or **invoice price**. Note that the above equation reduces to (3.21) when $\omega = 1$.

The quoted price in the U.S. usually does not include the **accrued interest**, hence also called the **clean price** or **flat price**. As the issuer of the bond will not send the next coupon to the seller after the transaction, the buyer has to pay the seller part of the coupon during the time the bond was owned by the seller. More formally, the accrued interest is calculated by [281]

$$C \times \frac{\begin{array}{c}\text{number of days from the}\\ \text{last coupon payment to the settlement date}\end{array}}{\text{number of days in the coupon period}} = C \times (1-\omega).$$

The Street convention is that the buyer pays the quoted price of the bond plus the accrued interest. Yield to maturity of a bond is then the r in (3.24) with PV being the invoice price, i.e., the sum of the quoted price and the accrued interest. Yield calculation, day count, and accrued interest interact in quite complex ways. Consult [730] for a thorough treatment.

Example 3.6.3 Consider a bond with a coupon rate of 10% and paying interest semiannually. The maturity date is March 1, 1995, and the settlement date is July 1, 1993. The day count is "30/360." Since the number of days between July 1, 1993 and the next coupon date, September 1, 1993, is 60, the accrued interest is $(10/2) \times \frac{180-60}{180} = 3.3333$ per \$100 of par value. At the clean price of 111.2891, the yield to maturity is 3%. This can be verified from (3.24) with $\omega = 60/180$, $m = 2$, $C = 5$, PV$= 111.2891 + 3.3333$, and $r = 0.03$. □

The invoice price formula (3.24) is not as arbitrary as it might at first seem. Its alternative form below is perhaps more intuitive,

$$\frac{1}{\left(1+\frac{r}{m}\right)^{\omega}} \left[\sum_{i=0}^{n-1} \frac{C}{\left(1+\frac{r}{m}\right)^{i}} + \frac{F}{\left(1+\frac{r}{m}\right)^{n-1}}\right].$$

Also, if r' is the equivalent continuously compounded annual yield, then the formula above becomes

$$\sum_{i=0}^{n-1} \frac{C}{e^{r'(\omega+i)/m}} + \frac{F}{e^{r'(\omega+n-1)/m}},$$

which is as it should be.

3.6.6 Yield for a portfolio of bonds

Calculation for the yield to maturity for a portfolio of bonds is no different from that for a single bond. First, the cash flows of the individual bonds are combined. Then, the yield is calculated based on the combined cash flows as if they were from a single bond. The following example shows the fallacy of computing the yield as the average of the yields of individual bonds in the portfolio.

Example 3.6.4 A bond portfolio consists of two zero-coupon bonds. The bonds are selling at 50 and 20, respectively. The term is exactly three years from now. To calculate the yield the correct way, we solve

$$50 + 20 = \frac{100 + 100}{(1+y)^6}$$

for y. Since $y = 0.19121$, the annualized yield is 38.242%. On the other hand, the yields to maturity for the bonds are 24.4924% and 61.5321%. Neither a simple average (43.01225%) nor a weighted average (35.0752%) matches 38.242%. □

3.6.7 Components of return

Consider an investor who bought a bond and intended to hold it to maturity. Its price is

$$P = \sum_{i=1}^{n} \frac{Fc}{(1+y)^i} + \frac{F}{(1+y)^n} = Fc\,\frac{1-(1+y)^{-n}}{y} + \frac{F}{(1+y)^n},$$

where n is the number of periods before maturity, y is the period yield to maturity, and c is the period coupon rate. By definition, the **total monetary return** of this bond is $P(1+y)^n - P$, which is equal to

$$Fc\,\frac{(1+y)^n - 1}{y} + F - P = Fc\,\frac{(1+y)^n - 1}{y} + F - Fc\,\frac{1-(1+y)^{-n}}{y} - \frac{F}{(1+y)^n}.$$

This monetary return can be broken down into three components [281]:

- **Capital gain/loss:** $F - P$;
- **Coupon interest:** nFc;
- **Interest on interest:** It is equal to the total monetary return minus the above two items,

$$\begin{aligned}(P(1+y)^n - P) - (F - P) - nFc &= P(1+y)^n - F - nFc \\ &= Fc\,\frac{(1+y)^n - 1}{y} - nFc.\end{aligned}$$

We proceed to investigate the interest on interest's percentage of the total monetary return, or

$$\frac{c\,\frac{(1+y)^n-1}{y}-nc}{c\,\frac{(1+y)^n-1}{y}+1-c\,\frac{1-(1+y)^{-n}}{y}-\frac{1}{(1+y)^n}}.$$

Keeping y and n constant, the above can be seen to increase as c increases. This means the higher the coupon rate, the more dependent is the total monetary return upon the interest on interest. So bonds selling at a premium are more dependent upon the interest on interest component, given the same maturity and yield to maturity. It can be verified that, when the bond is selling at par ($c = y$), the longer the maturity, n, the higher the proportion of interest on interest among the total monetary return. The validity of the claim for bonds selling at a premium ($y < c$) or a discount ($y > c$) is tedious to show. Still, it is not hard to see the trend.

The above results reveal again the impact of reinvestment risk. Coupon bonds which obtain a higher percentage of their monetary return from the reinvestment of coupon interests are more vulnerable to changes in reinvestment opportunities. From this perspective, the inadequacy of the yield to maturity measure can be seen even more sharply. This yield assumes that all coupon payments can be reinvested at the same yield as the yield to maturity. This assumption is seldom true in a changing environment.

The holding period return measures the profit or loss incurred by holding the security until the **horizon date**. This period of time is called the **holding period** or the **investment horizon**. Holding period return is composed of (1) capital gain/loss on the horizon date, (2) cash flow income such as coupon and mortgage payments, and (3) reinvestment income from reinvesting cash flows received between the settlement date and the horizon date such as the interest on interest. Apparently, to derive the future value, one has to make explicit assumptions about the reinvestment rate during the holding period and the security's market price on the horizon date, its **horizon price**. **Scenario analysis** computes the holding period return for each assumption about the above two parameters. Each scenario may be weighted, and then an optimization algorithm is applied to find the best solution [791]. The **value at risk (VaR)** methodology is a refinement of scenario analysis. It constructs a confidence interval for the dollar return at horizon based on some stochastic model for price. VaR will be discussed in later chapters.

Example 3.6.5 Consider a five-year bond paying semiannual interest at a coupon rate of 10%. Assume the bond is bought for 90 and held to maturity with a reinvestment rate of 5%. The coupon interest plus interest on interest amounts to

$$\sum_{i=1}^{2\times 5}\frac{10}{2}\times\left(1+\frac{0.05}{2}\right)^{i-1}=56.017.$$

The capital gain is $100-90=10$. The holding period return is therefore $56.017+10=66.017$. The holding period yield is the y that satisfies

$$\left(1+\frac{y}{2}\right)^{2\times 5}=\frac{100+56.017}{90},$$

or 11.31%. As a comparison, its bond equivalent yield to maturity is 12.767%. Clearly, different holding period returns result under different reinvestment rate assumptions. If the security is to be sold before it matures, its horizon price needs to be figured out as well. □

Additional Reading

See [210, 281, 283, 286, 796] for more information regarding the materials in the chapter. Solving equations numerically is a field rich with techniques; consult [29, 191, 334, 369, 400, 632] for more information.