Part I

Expected Returns on Financial Assets

1 The cost of capital under certainty

The purpose of Part I is to examine how the expected returns on financial assets are determined in simple theoretical settings. We explain in Section 1.1 that a project can be thought of as a financial asset, and that its cost of capital is the expected return on the asset at market value. So we are really looking at how the cost of capital is determined. The settings are rather abstract, and may seem odd at first sight. The reason for the abstraction is to help in understanding the economic processes that determine a project's cost of capital in the real world.

The current chapter begins with a brief account of what the terms 'capital' and 'cost of capital' mean. It then considers the interest rate in a world in which the future is known with certainty. The assumption of certainty, though unrealistic, provides a relatively easy starting point. The analysis serves to establish several ideas that will be useful when uncertainty about the future is introduced in Chapter 2.

1.1 Concepts

1.1.1 Capital and investment

It is not as easy as one might expect to say what is meant by 'capital', and by the related terms 'investment', 'saving', 'income' and the 'cost of capital'. In fact, there is a sizable literature in economics on these questions (e.g. Parker and Harcourt, 1969; Hirshleifer, 1970). We offer a brief discussion based on everyday usage.

There are at least three meanings of *capital* in common use.

(i) Many readers will think first of capital in the context of personal finance, in which it means the same as savings. Most people would consider the value of their savings (their wealth or non-human capital) to be the value of their financial assets, such as money in a bank account or shares in a mutual fund, plus the value of their house, net of their personal borrowing.

(ii) In the context of company finance, capital means money tied up in the company in the form of equity and debt. Equity and debt are financial assets that have been issued by the company at some time in exchange for cash. There are two ways of measuring the amount of money tied up that are in general use. The value of the equity and debt on the balance sheet – the accounting or book value of the capital – measures the amount of cash raised by the company and yet to be repaid, including retained cash flow, net of depreciation. At least, this is what book value measures under the pure historic-cost convention of preparing accounts, under which the cost of an asset is adjusted only for depreciation. In contrast, the market value of the money tied up measures the amount of cash that could be raised by the current holders of the equity and debt if they chose to sell their holdings.

It is normal to speak of projects or operations or divisions as having capital, regardless of whether they are actually constituted as companies. Such business units are implicitly being thought of as discrete entities, which might – in principle – exist as companies, able to issue equity and debt.

(iii) In economics, the central meaning of capital is probably assets used in the production process: tangible assets such as machinery, and intangible assets such as know-how. Such real (as opposed to financial) assets are used to provide goods and services consumed by individuals, or to produce real assets used as inputs in further production processes. Thus, capital in the economist's sense refers to resources created or nurtured as a means to an end, the ultimate end being consumption by people. The resources exist in physical form, or in the form of knowledge and skills possessed by individuals and organisations.

The common thread across the three meanings is that capital is money tied up or invested in something that provides income in the future. The 'something' is a financial asset, or property, in the context of personal finance; it is a company or project in corporate finance; and it is any real asset in economics.

The term *investment* means, in general, addition to capital. In personal finance, investment means the same as saving – that is, addition to one's financial assets or property. In company finance, investment is the injection of new cash in a company through an increase in the cash tied up as equity, via the retention of net cash flow or the issue of shares, or through an increase in the amount of debt outstanding. Investment is also the commitment of cash to a project. In economics, investment is expenditure to create or purchase real assets, which therefore increases the owner's stock of real assets.

1.1.2 A project and its cost of capital

We are concerned with the cost of capital for projects. We shall leave aside the distinction between equity and debt until Part II, and shall assume that all capital is in the form of equity. *A project* is a discrete undertaking that requires capital, in the corporate finance sense of cash tied up as equity, and that is expected to provide a positive real rate of return on its capital, net of operating costs (otherwise it would not be undertaken). A project's *cost of capital* is the minimum expected rate of return needed to attract the required capital. The minimum expected rate is an opportunity cost, not a cash cost. It is the rate that could be obtained by investing the cash in the next-best alternative. We shall see that the next-best alternative is to invest in other projects, or in financial assets, of the same risk as the project in question. If the asset market is efficient, assets of the same risk will offer the same expected rate of return, so there will be one expected rate for each level of risk. Any project that offers an expected rate above the market rate for its level of risk is a project that will increase the wealth of the owner. Hence, the cost of capital can also be thought of as the project's *hurdle rate*, the minimum expected rate of return that leaves investors (the providers of capital) at least as well off as they would have been had they invested in another asset of the same risk. These ideas will become clearer in the context of a model of project valuation (below and Section 2.1).

We have, in talking about an *expected* rate of return, implicitly assumed that the future is uncertain. If the future is assumed to be certain, there is no risk, and the minimum rate of return needed to attract the required capital will be the same for all projects. The cost of capital can then be called the *interest rate*. We shall assume certainty for the rest of the present chapter.

A word on terminology

A return is a percentage gain or loss in a single period. A rate of return or interest rate is a percentage gain or loss per period. They mean the same thing given a single period, or given identical returns per period. But some care is needed if there are $T > 1$ periods and the returns per period differ. A return over a particular single period will be different from the rate of return over the T periods. For the moment, we assume either a single period or identical returns per period.

1.1.3 Capital budgeting and valuation

The core application of the cost of capital is in capital budgeting, which is the process by which a company decides which projects to undertake. Capital budgeting involves the estimation of the market value of projects that have yet to be undertaken, and that do not have an observable market value. This is because there is no market for would-be projects, and financial assets are not issued that provide claims to the future net cash flows from would-be projects. The estimated market value of a project is then compared with the initial cash cost. If the value exceeds the cost, then the project adds value and should be undertaken, assuming that the aim of the investment is to increase shareholder wealth. The estimated market value of an untraded asset, whether it is a project (a real asset) or a financial asset, is an estimate of the price that investors would pay for it were it tradable. A valuation model is a model that provides an estimate of the market value; it explains the price that investors would pay for an asset.

Single-period setting

A project is something that provides one or more net cash flows or pay-offs in the future. Consider a project that will provide a single pay-off of Y_1 in real terms after one period, at date 1. The present date is date 0. Let the capital, or cash investment, required at date 0 to undertake the project be I_0. The percentage real return on the cash investment, R_{inv}, is

$$R_{\text{inv}} = Y_1/I_0 - 1 \tag{1.1}$$

The market value V_0 of the shares issued to raise I_0 is determined by Y_1 and by the interest rate R:

$$V_0 = Y_1/(1 + R)$$

or

$$V_0(1 + R) = Y_1 \tag{1.2}$$

Equation 1.2 is a model of the market value of the project, given Y_1 and R. Investors are willing to pay for the project because ownership gives entitlement to the future pay-off. Equation 1.2 says that the price in equilibrium that investors would be willing to pay, V_0, is such that, given the future pay-off, the return on the purchase is R. At this price, investors are indifferent between investing in the shares in the project and investing in the next-best alternative, which is another asset that provides a percentage return of R.

The project will increase shareholder wealth only if $V_0 > I_0$ (i.e. if the *net present value* (NPV)) is positive, and it will attract the required capital only if $V_0 \geq I_0$. It can be seen by comparing equations (1.1) and (1.2) that $V_0 \geq I_0$ if $R_{\text{inv}} \geq R$, since Y_1 is given. Since the cost of capital is defined as the minimum return on the cash needed to attract the amount I_0, and no one will invest unless $V_0 \geq I_0$, the cost of capital is given by the interest rate R.

Example

Suppose the investment required to undertake a project is £100 and the pay-off at date 1 is £120. The interest rate is 5 per cent. The market value of the project is then

$$V_0 = £120/1.05 = £114$$

from equation (1.2). Someone buying the project at this price obtains a return of 5 per cent. The net present value is positive, at £114 − £100 = £14, and the return on the cash investment of £100, which is 20 per cent, exceeds the cost of capital, which is 5 per cent.

Multiperiod setting

Things are a little more complicated with more than one period, even assuming certainty. The normal measure of the return on investment over more than one period is the *internal rate of return* (IRR). This is defined as the discount rate at which the present value of the cash flows is equal to the cash investment required at date 0. That is, we solve

$$-I_0 = Y_1/(1 + x) + Y_2/(1 + x)^2 + \cdots Y_T/(1 + x)^T \tag{1.3}$$

for x, given V_0 and the cash flows for the T periods; and the value found for x is the IRR. The virtue of the IRR is that, if it exceeds the interest rate, then the NPV will be positive. Thus, for the purpose of working out whether the NPV is positive, the IRR can be compared with the interest rate, assuming that the interest rate is the same in all periods.

Despite this, there are several familiar problems with using the IRR in capital budgeting. IRR takes no account of the size of the project. If one or more of the cash flows is negative, there will be more than one value of x, or no value of x, that satisfies equation (1.3). It is debatable whether IRR is the best measure of the rate of return on the cash invested in the project. The amount of cash invested changes over time. The IRR is the rate of return that will be received assuming that all the cash received before the end of the project (date T) is invested until T at the IRR. It would, arguably, be more sensible to assume that cash received can be invested at the interest rate.

The problems with IRR can be avoided by using the rule that says: accept projects with positive NPV. The important point for our purposes is that the positive-NPV rule involves the concept of the cost of capital. It is equivalent to saying: accept projects with a rate of return on investment that exceeds the interest rate. We can see this by thinking of the claim to each (positive) cash flow in equation (1.3) as a separate asset or 'mini-project', each of which provides the cash flow in question. The next-best alternative to a claim to Y_t is another asset providing the same pay-off at the same date t periods from date 0. The rate of return on this asset at market value must be R per period, so its price must be $Y_t/(1+R)^t$. Thus the market value of the package of claims to $Y_1, Y_2 \ldots Y_T$ is

$$V_0 = Y_1/(1+R) + Y_2/(1+R)^2 + \cdots Y_T/(1+R)^T$$

If $V_0 > I_0$, it must be the case that the rate of return on the package – i.e. the project – exceeds the interest rate R.

1.2 The interest rate under certainty

1.2.1 Choice in a timeless world

The question to which we turn in the rest of the chapter is how the interest rate is determined assuming certainty. The model we present originates with Fisher (1930), and our presentation is based on Hirshleifer (1970) and Fama and Miller (1972).

The 'classical' microeconomic theory of interest is an extension of the traditional theory of demand for goods in a timeless world of perfect markets. We start with a sketch of this theory. The main assumptions are as follows.

(i) A perfect market is one in which there are no transactions costs, there is no taxation or other interference from outside the market, information is freely available and all the participants are price takers, unable to affect the market price by their trades.

(ii) Goods and services are traded in units that are 'small' in relation to the budgets of individuals. This ensures that marginal utility per pound spent can be equated across goods. We abstract from the facts that some goods are much more expensive than others and that goods are not 'infinitely divisible'.

(iii) Individuals are assumed to be rational and to seek to maximise their utility, or material welfare. Utility derives from the consumption of goods and services by individuals.

The following assumptions are normally made about the rationality and utility functions of individuals.

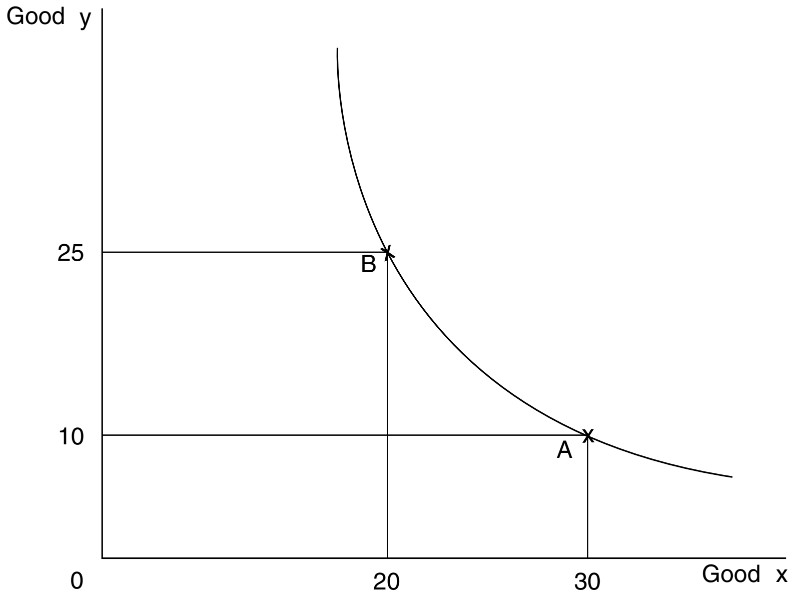

Figure **1.1** An indifference curve

1. *Comparability between goods*, or *completeness of preferences*. For any two bundles of goods *A* and *B*, the individual either prefers *A* to *B* or *B* to *A*, or is indifferent between them.
2. *Transitivity of preferences*, or *consistency*. If the individual prefers *A* to *B* and prefers *B* to a third bundle *C*, he or she prefers *A* to *C*. If he/she is indifferent between *A* and *B* and between *B* and *C*, he/she is indifferent between *A* and *C*.
3. *Non-satiation*. The individual prefers or is indifferent to having more of a given good, assuming that this does not mean having less of any other good.
4. *Diminishing marginal utility*. Starting from any combination of two individual goods, *x* and *y*, for each extra unit of one good given up the individual requires an increasing amount of the other in order to maintain the same level of total utility.

Diminishing marginal utility implies that indifference curves are convex to the origin (the curve lies below a straight line between any two points), as illustrated in Figure 1.1 for the case of two goods. Imagine two points, *A* and *B*, that lie on the same indifference curve. These points represent different bundles of the two goods for a given individual. At point *A* the individual has thirty units of *x* and ten units of *y*, and he requires one unit of *y* for every one unit of *x* given up. In other words, the slope of the curve at that point, or *marginal rate of substitution* of *y* for *x*, is -1. At point *B* the individual has twenty units of *x* and twenty-five units of *y* and requires two units of *y* for every one of *x* given up, so the slope at this point is -2. The individual has more of *y* at *B* than *A*, so his marginal utility from *y* – his utility from an extra unit of *y* – is smaller at *B*. The total utility from each of these combinations is the same, which is why they are on the same indifference curve.

It can easily be shown that individuals maximise their utility by allocating their budget in such a way that the marginal utility, the utility from spending an extra £1 on any good, is the same. This means that, for any two goods x and y, a given individual will buy amounts of each such that the ratio of the marginal utilities for each good, and hence the marginal rate of substitution between the two, is equal to the ratio of their prices:

$$MU_x/MU_y = P_x/P_y \tag{1.4}$$

where MU_x is the marginal utility of good x and P_x is the price of one unit of x. The market for goods and services is in equilibrium if equation (1.4) holds across all goods and all individuals. It is not necessary for individuals to have the same preferences (utility functions) for equilibrium to arise; two people with the same budget but different preferences will maximise their respective utilities by buying goods in different amounts. Also, the analysis does not require that explicit numerical (cardinal) values be assigned to different levels of utility. An ordinal utility function is sufficient. This means that the individual can rank combinations of goods in consistent orders of preference (assumptions 1 and 2 above). This is enough for different levels of utility to be identified, for indifference curves not to cross and for an individual to say how much of a good he wants in exchange for one unit of another good, so that the marginal rates of substitution are known.

1.2.2 Choice in a two-date world

There is, obviously, no role for an interest rate, or for capital, unless the analytical framework is extended to incorporate time. A simple way of introducing time is to suppose that there are just two dates, the present and a future date, with a period of time in between. The individual decides how to allocate his current budget between consumption today and saving, which enables him to consume more at the future date. At the heart of the theory of interest under certainty is the view that the utility from a unit of consumption today is not necessarily the same as the utility today from a unit of consumption in the future. The individual chooses between 'bundles' of differing amounts of present and future consumption in a way that is analogous to the choice between bundles of goods in a timeless world.

Under certainty, there is no chance that the individual will face an unforeseen shortfall in income or an unforeseen requirement for expenditure, so there is no precautionary motive for saving. There is also no opportunity to seek a higher return by taking risk. The sole reason for saving is to have more to spend in the future. We shall see that utility is maximised when the loss in utility from forgoing one unit of consumption today equals the gain in utility from being able to spend $(1 + R)$ units more in the future, where R is the real rate of interest.

Suppose for the moment that all an individual can do with his money saved is to put it under the mattress. Then £1 not spent today provides £1 more to spend at date 1. Let c_t be one unit of consumption, or bundle of goods and services, at date t. The price per unit of current consumption is $P_0 = £1$. c_1 is exactly the same unit except that it will be consumed at date 1, and its price is P_1. The real value of £1 saved – the value in terms of the consumption units that £1 buys – depends on P_1.

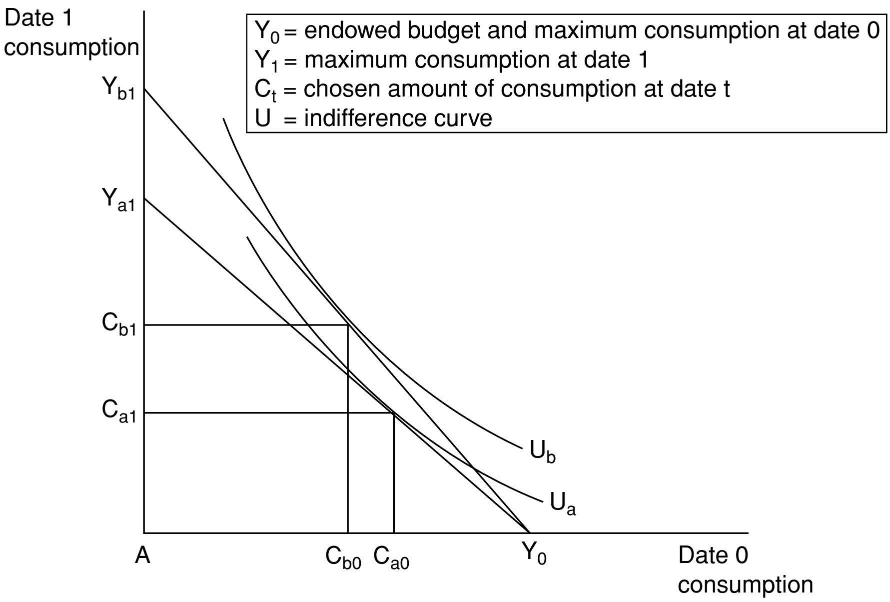

Figure **1.2** Consumption across two dates

In Figure 1.2 an individual has funds (or income or endowment) of Y_0 pounds at date 0. His budget constraint says that the total amount he can spend on consumption at dates 0 and 1 cannot exceed Y_0. It can be written

$$P_1C_1 = Y_0 - P_0C_0$$

where C_t is the number of units consumed at date t. The line between Y_{a1} and Y_0 represents the budget constraint assuming the price $P_{a1} = P_0 = £1$, so its slope is -1. The interest rate is zero; the real value of the funds saved stays the same.

U_a is an indifference curve showing combinations of units consumed at the two dates between which the individual is indifferent. It is convex because diminishing marginal utility of consumption at each date is assumed. Curves further away from the origin represent higher levels of 'lifetime utility' – that is, utility at date 0 from consumption at date 0 and from consumption-to-come at date 1. The shape of the individual's indifference curves depends upon his *time preference*. Greater time preference means greater utility from a given amount of consumption now in relation to utility from the same amount of consumption in the future; the indifference curves have a steeper slope, so the point of tangency occurs nearer to point Y_0, and the individual saves less. The indifference curve tangential to the budget constraint represents the highest level of lifetime utility the individual can attain, given his budget. At the point of tangency, the slope of the indifference curve equals the slope of the budget constraint, so the loss in utility from forgoing one unit of consumption at date 0 equals the gain in utility from spending the extra £1 at date 1. The individual consumes AC_{a0} and saves $C_{a0}Y_0$ at date 0, and consumes AC_{a1} at date 1.

Now suppose we have $P_{b1} < £1$, which means that £1 buys more units at date 1 than at date 0 (there is price deflation). This is the same as saying that the real value of funds saved increases at the rate R_b:

$$1/(1 + R_b) = P_{b1}/P_0$$

therefore

$$R_b = P_0/P_{b1} - 1 = £1/P_{b1} - 1$$

For example, if $P_{b1} = £0.95$, $R_b = 5.26$ per cent. Each £1 saved still sits under the mattress until it is spent; the interest rate is determined entirely by the price of c_1 in relation to the price of c_0. A positive interest rate enables an individual with a given endowment at date 0 to consume more in total, across both dates, and to reach an indifference curve U_b representing greater utility than is possible with an interest rate of zero.

In Figure 1.2 the lower price of c_1 is shown as resulting not only in an increase in the quantity demanded of c_1 but also in the proportion of the budget spent on c_1. This is not an inevitable outcome. An individual could choose to spend less on future units as their price falls. For such a person, the income effect, caused by the rise in real income considered across both dates, is negative for consumption at date 1 and outweighs the positive substitution effect, caused by the fact that date 1 units have become cheaper in relation to date 0 units. However, it is plausible to assume that individuals in aggregate (society) will save more as the real interest rate rises.

More formal exposition

The preceding discussion includes the notion of the maximisation of lifetime utility by an individual. Since this is an important notion, it is worthwhile to provide a more formal exposition. Each individual seeks to maximise his lifetime utility subject to his budget constraint. In the single-period setting, his lifetime utility at date 0, U, is a function of the amounts he consumes at dates 0 and 1:

$$U = U(C_0, C_1) \tag{1.5}$$

where $U(.)$ indicates an unspecified function that expresses the relation between the variables in the brackets and the amount of lifetime utility. The budget constraint assuming no endowment at date 1 is

$$P_1C_1 = Y_0 - P_0C_0 \tag{1.6}$$

where Y_0 is the individual's initial endowment at date 0 and P_0 and P_1 are given. The problem is to maximise (1.5), by choosing the appropriate amounts C_0 and C_1, subject to (1.6). This is the same as finding the point of tangency of the indifference curves with the budget constraint. The problem can be solved using the Lagrangian multiplier technique. We form a Lagrangian expression to be maximised:

$$L(C_0, C_1, \lambda) = U(C_0, C_1) - \lambda[P_0C_0 + P_1C_1 - Y_0]$$