1

Asset markets and asset prices

Overview

Financial markets encompass a broad, continually evolving and not altogether clearly delimited collection of institutions, formal and informal, that serve to facilitate the exchange of assets. More to the point, the concept of an 'asset' is open to a variety of interpretations.[1] Rather than get bogged down in arbitrary classifications – and in ultimately fruitless distinctions – the nature of 'assets' and the markets in which they are traded is allowed to emerge from examples. To place the examples in context, the chapter begins by reviewing, in section 1.1, the fundamental properties of financial systems, and identifies various sorts of capital market, several of which receive attention later in the book.

The main objective of this chapter is to outline the ideas that underpin explanations of asset prices and hence rates of return. Sections 1.2, 1.3 and 1.4 describe a framework for modelling asset price determination and comment on alternative approaches.

Central to an understanding of finance is the process of arbitrage. Arbitrage trading policies seek, essentially, to exploit price discrepancies among assets. Of more interest than the policies themselves are their unintended consequences, namely the implications they have for tying asset prices together in predictable patterns. The examples in section 1.5 serve to introduce arbitrage. Its consequences emerge in several places throughout the book.

Observers and analysts of capital markets frequently seek ways to appraise the performance of the markets. The concepts of 'efficiency' introduced in section 1.7 show that different criteria can be applied in making judgements about how well the markets function.

[1] Perhaps it would be more accurate to use the clumsier term 'financial instrument', or possibly 'security', instead of 'asset'. But, for the purposes of this book, 'asset' is simpler and should not cause confusion.

1.1 Capital markets

Financial innovations are to the financial system what technological advances are to the economy as a whole. They embrace changes in the methods of doing business as well the assets traded in markets. In the broadest terms, financial innovations refer to development in the institutions of finance made in response to changes in the environment in which the institutions exist. The process of financial innovation involves institutional adaptation and evolution even when the functions of the system remain the same.

Merton and Bodie (chap. 1 in Crane et al., 1995) argue that the *functions* of financial systems change more slowly than their institutions. They propose a sixfold classification of functions.

1. *Clearing and settling payments*. Financial systems provide mechanisms that facilitate exchanges of goods and services, as well as assets, followed by settlement, transferring ownership in return for the agreed remuneration.
2. *Pooling resources and subdividing shares*. Financial systems enable multiple investors to contribute to projects that no one of them alone could afford. Also, even if a single investor could afford to fund a project, there may be incentives for diversification, each investor contributing a small portion of the project's cost and bearing a small portion of its risks.
3. *Transferring resources across time and space*. A fundamental purpose of investing is to delay consumption, for example as households accumulate wealth for retirement or for the benefit of future generations. Firms in one industry, or in one location, may seek to invest surplus funds in other industries or at other locations. Financial systems enable the assignment of these funds from households and firms with surplus resources to others that seek to acquire resources for investment and (intended) future return.
4. *Managing risk*. Financial systems provide ways for investors to exchange, and thereby to control, risks. For example, insurance enables the pooling of risks, hedging enables the transfer of risk to speculators, diversification exploits low correlations that may exist among risky projects.
5. *Providing information*. Financial systems enable *price discovery* – that is, for those who wish to trade to observe the prices (rates of exchange) at which agreements can be made. Other information, for example about expectations of future asset price volatility, can be inferred from market prices. (Chapter 19 explains how observed option prices enable inferences about the magnitude of expected asset price fluctuations in the future.)
6. *Dealing with incentive problems*. It is reasonable to suppose that contractual obligations can never stipulate the actions to be taken in every eventuality, even if every contingency could be imagined. Financial systems can help individuals to construct the sorts of contracts that fulfil their needs and to cope with the contingencies that the contracts do not explicitly take into account. For instance, the shareholders of a firm may finance its operations partly with debt, the contractual obligations for which are designed to provide incentives for the firm's managers to act in the interests of the shareholders.

What explains financial innovation (i.e. what accounts for institutional change)? There are many possible causes, including (a) technological change – e.g. advances in information technology; (b) changes in the 'real' economy – e.g. the growth of new industries and markets in South-East Asia; (c) changes in the demand for assets – e.g. ageing populations saving for retirement; and (d) changes in government regulation – e.g. the liberalization of trading rules, creating new opportunities, or new regulations providing incentives to avoid, bypass or otherwise profit from their introduction.

This book explores the operation of mature financial systems as of the early twenty-first century. While there are hints about the pattern of financial innovation, this is not a main focus of analysis. Also, the relationships between the functions of the financial system and the institutions that currently perform them remain implicit, though they should be straightforward enough to infer.

The following list of capital markets, although not comprehensive, identifies the differences among markets (differences relevant for this book, anyway) and the assets traded in them.

1. *Equity, or stock, markets*. The stock exchange is the main 'secondary' market for shares in corporations – i.e. limited liability companies.[2] It is a secondary market in the sense that the shares are already in existence, so that trade takes place between investors and need not directly involve the corporations themselves. The 'primary' market involves the issue of new shares by corporations. There are various categories of shares (e.g. ordinary shares, preference shares) but the distinctions among them are neglected here, being peripheral to the basic principles of price determination. The pattern of share prices is normally summarized by reference to particular well-known stock price averages or indexes, such as the Dow-Jones Industrial Average, Standard and Poor's 500 index, or the Financial Times Stock Exchange 100 index (see appendix 1.1).
2. *Bond markets*. These are markets for long-term securities such as government debt (known as gilt-edged securities in Britain) or corporate bonds.

   Bonds are usually regarded as less risky than shares because bonds normally oblige the issuer to promise to take specific actions at definite dates in the future. The distinction is not quite as clear as it might first seem because bond contracts can include clauses that provide for different actions in a multitude of different contingencies. Also, it is possible that the issuer of the bond will default with respect to some clause in the agreement. Even so, a typical bond is a promise to pay (a) a sequence of *coupons* (commonly twice a year) and (b) a lump sum *maturity value* (or face value) at a specified date in the future.

[2] If there is any distinction between 'stocks' and 'shares', it is not one of any significance here. A company's 'stock' could refer to the whole value of its equity, while 'shares' could refer to the ownership of a portion of that stock.

Bonds are commonly traded on stock exchanges in much the same way as shares. A feature of *medium-term* and *long-term* bonds is that, like shares, much of the trade is amongst investors, without the direct involvement of the issuer (government or company).

3. *Money markets*. Money markets exist to facilitate the exchange of securities such as treasury bills (commonly, three-month or six-month government debt) or other loans with a short time to maturity. Although such securities are traded in markets, any holder does not have to wait long before the issuer is obliged to redeem the debt in compliance with the terms of the contract.
4. *Commodity markets*. Markets of some form exist for almost every commodity, though financial studies are usually confined to highly organized markets for a fairly narrow range of commodities, including precious metals (gold, silver, platinum), industrial metals (such as lead, tin and copper), petrochemicals or agricultural commodities (such as cereals, soya beans, sugar and coffee). This list is not exhaustive but it does suggest that the commodities in question need to have certain physical characteristics: namely, that they can be graded according to well-defined attributes, that they are divisible into precisely defined units, and that they are storable (though often subject to deterioration over time). As will be described later, most organized commodity markets involve trading in contracts for the delivery of the stated commodity at a future date, though perhaps one very near to the present.
5. *Physical asset markets*, such as for real estate. In this case, the relevant asset for financial analysis is often a security (e.g. a mortgage) constructed to have a well-defined relationship with the physical asset (e.g. a mortgage being a loan secured against the equity of the property). It is not uncommon for mortgages to be *securitized* by financial intermediaries that issue bonds backed by (and with payoffs defined by) bundles of mortgages.
6. *Foreign exchange markets* – ‘FOREX’ or ‘FX’ markets. These are markets for one currency against another. Governments often intervene in such markets – not infrequently with disastrous consequences – to fix, or at least influence, exchange rates among currencies. Two notable features of FX markets are (a) the vast turnover of funds (often about \$1.5 trillion each day in mid-2001) and (b) round-the-clock trading.
7. *Derivatives markets*. Corresponding to most of the above categories are derivative, or synthetic, securities. They are ‘derivative’ in the sense that their payoffs are defined in terms of the payoffs on an underlying asset or assets. The underlying asset could itself be a derivative, so that a whole hierarchy of such instruments emerges. Almost all derivatives are variants of two generic contracts.

   (a) *Forward agreements*. These are contracts in which the parties agree to execute an action (typically, the exchange of a specified amount of money for a specified amount of some ‘good’) at a stipulated location and date in the future. For example, a forward contract might specify the delivery of 5000 bushels of domestic feed wheat to a grain elevator in Chicago, six months from the date of the agreement, at a price equal to \$3.50 per bushel. A *futures* contract is a special type of forward contract designed to allow for trading in the contract itself. *Repo* contracts are combinations of loans and forward agreements. *Swaps* are sequences of forward contracts packaged together.

(b) *Options*. Options are contracts for which the holder has the right, but not the obligation, to execute a specified action at an agreed date, or over a range of dates. For example, an option might stipulate that its owner can purchase 100 IBM ordinary shares for $220 per share at any time prior to the following 30 September. Many sorts of option contracts are traded. For example, *options on futures* are options to purchase or sell futures contracts; *swaptions* are options on swap contracts. *Exotic* options encompass a variety of contracts involving non-standard terms for their execution.

1.2 Asset price determination: an introduction

1.2.1 A single asset market

The simplest economic theory of price determination applied to asset markets is that of 'supply and demand'. The prices of many assets are highly flexible, with rates of change that are rapid compared with the rates of change in the total volume of the asset in existence. At each instant of time the total stock of the asset is assumed fixed. The market price is allowed to adjust so that wealth holders, in the aggregate, are just prepared to hold the existing stock – the demand to hold the asset equals the stock in existence. Figure 1.1 depicts an equilibrium price of p^* that equates demand with the given stock denoted by $\overline{Q}$.

In some cases, it makes sense to treat the *total* stock of the asset in existence as *zero*. For example, corresponding to every futures contract there must be exactly the same volume of purchases ('long' positions) as sales ('short' positions): they net out to zero. The stock of outstanding purchases (or sales) – known as 'open interest' – will, of course, change over time, but at each instant the total of purchases and the total of sales each equals the open interest.

From this perspective, the relevant question is: what determines the demand to hold the asset? An immediate but superficial response is that the demand for an asset is determined by the same things as the demand for any good: (a) *preferences*, (b) the *price* of this and other assets, and (c) *income* (here the *stock* of wealth, not the flow of income, forms the relevant constraint). A more complete and satisfactory response involves delving beneath the surface to analyse the role of each of these elements.

1.2.2 Multiple asset markets: a more formal approach

What are the forces that determine the market prices for different assets? As a start, consider a world with many market participants – *investors* – each of whom has an initial amount of wealth available for investment.

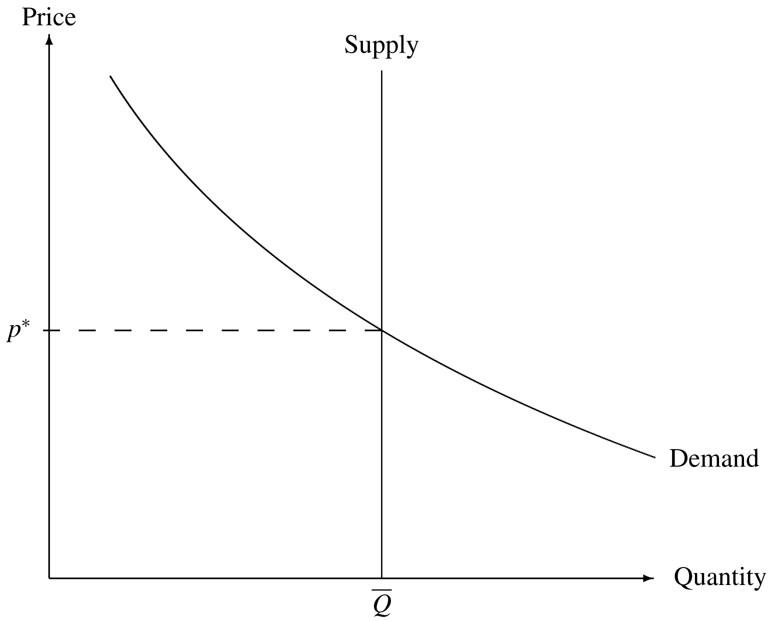

Fig. 1.1. Market equilibrium for a single asset

At each instant of time the total stock of the asset is fixed, say at $\overline{Q}$. The demand to hold the asset is depicted by the negatively sloped curve. At price p^* the market is in equilibrium – i.e. the demand to hold the asset equals the stock available to be held.

In the presence of a large number of investors, it is plausible to assume that each investor is a price taker, in the sense that no one investor has enough market power to influence prices. Each investor thus treats asset prices as parametric, though not necessarily constant over time. Initial wealth is also parametric, being equal to the sum of each asset's price multiplied by the quantity of the asset that the investor starts out with (i.e. holds as a consequence of past decisions).

Faced with given asset prices and with given initial wealth, each investor selects a portfolio in accordance with a *decision rule*. The decision rule – which can be unique to each investor – determines the number of units of each asset to hold as a function of the observed prices and initial wealth. Theories of *decision making under uncertainty* provide the necessary foundation from which each investor's decision rule is derived (see chapters 4, 5 and 11).

The *market equilibrium* at each date is defined by a set of asset prices and an allocation (portfolio) of assets among investors that, together, satisfy the following conditions.

1. Each investor's portfolio is determined according to the investor's decision rule. In particular, the chosen portfolio is optimal subject to the investor's preferences (i.e. willingness to bear risk), beliefs (about assets' payoffs) and constraints (the given level of initial wealth and, perhaps, institutional limits on permissible trades).
2. Demand equals supply; that is, the total stock of each asset equals the total demand aggregated over all investors.

Note that, in principle, some or all investors may be allowed to hold assets in negative amounts – investors may be able to 'short-sell' assets (see section 1.4.2).

The main components of the approach so far are as follows.

1. At each instant of time total asset stocks (netting out assets and liabilities) are given.
2. Asset prices adjust so that existing stocks are willingly held.
3. With the passage of time asset stocks change (e.g. because companies issue new shares and debt, or repurchase shares and redeem existing debt). Also, investors revise their portfolios in response to changes in their circumstances or their beliefs about the future. As a consequence, prices change.

This is merely the skeleton of a framework and makes no definite, testable predictions. Even so, it is a useful way of viewing asset markets because most of the models in the remainder of the book emerge as special cases, each of which fits within the framework. The capital asset pricing model (see chapters 6 and 11), for instance, is perhaps the most notorious special case. It would be wrong, however, to conclude that the approach outlined above is the *only* way to model asset prices; an alternative framework, based on asset flows rather than stocks, is explored in chapter 2.

1.2.3 Rates of return

Assets are typically held because they yield – or, at least, are expected to yield – a rate of return. A general way of writing the rate of return on an asset is

$$\text{rate of return} \equiv \frac{\text{payoff} \quad \textit{minus} \quad \text{price}}{\text{price}} \tag{1.1}$$

where 'price' is the observed market price (or outlay on the asset) as of today, date t, and 'payoff' is the value of the asset at the next relevant point of time, date $t+1$ (where $t+1$ could be tomorrow, next month, next year or whenever). The *gross* rate of return on an asset is commonly defined as $\frac{\text{payoff}}{\text{price}}$. Thus, while the rate of return might be a number such as 0.064 (6.4 per cent), the gross rate of return would be 1.064.

An asset's payoff may have several components according to the type of asset. For a bond, the payoff is its market price at $t+1$, plus any coupons received between t and $t+1$. For a bank deposit, the payoff is the principal at t plus the interest accumulated between t and $t+1$ minus bank charges. For a company's shares, the payoff is the share's market price at $t+1$ plus the dividends, if any, paid between t and $t+1$.

Let the asset's price at t be denoted by p_t and its payoff at $t+1$ by v_{t+1}. Then the asset's rate of return between t and $t+1$, y_{t+1}, is defined by

$$y_{t+1} \equiv \frac{v_{t+1} - p_t}{p_t} \tag{1.2}$$

where y is intended to stand for 'yield'. It is often convenient to interpret the price at $t+1$, p_{t+1}, to include any dividends or coupons received between t and $t+1$. With this interpretation, $v_{t+1} = p_{t+1}$. In words: the rate of return is the proportional rate of change of the asset's market price. Slightly more generally, the rate of return is measured by the proportional rate of change of the asset's market value (i.e. it includes flows such as dividends or coupons as well as the market price).

The *real* rate of return on an asset is defined as the rate of return measured not in units of account, 'money', as in expression (1.1), but in terms of aggregate 'real' output.[3] Call the rate of return in (1.1) the *nominal* rate of return. Then the relationship between real and nominal rates of return – often attributed to the eminent American economist Irving Fisher (1867–1947), of Yale University – can be written as

real rate of return = nominal rate of return *minus* rate of inflation

(See appendix 1.2 for a derivation.) More substantively, the Fisher hypothesis is commonly interpreted as the prediction that the real rate of interest is constant – that fluctuations in the nominal rate and inflation tend to offset one another.

The distinction between nominal and real rates of return is important in many branches of economics, especially monetary economics and macroeconomics (where another distinction – between actual and expected inflation – is particularly relevant). In this book the distinction between nominal and real rates of return is not prominent. Where necessary, an adjustment from nominal to real rates can be made by subtracting the rate of inflation from the nominal rate. This simple-minded approach is not intended to underrate the importance of the difference between nominal and real rates. Rather, it serves to emphasize that the determination of expected and actual rates of inflation is not studied here.

[3] In principle, the rate of return can be defined in the units of any commodity, service or asset. In practice, an index of aggregate output is used in an attempt to measure output as a whole.

1.2.4 The roles of prices and rates of return

The most important aspect of rates of return for decision making is that they are *forward-looking*: they depend on future payoffs. For almost all assets, the payoff is, at least in part, *uncertain* when viewed from the present, date t. For example, the prices of stocks and shares at date t can be observed at date t, but their prices at date $t+1$ are matters of conjecture.

The current, observed market price for an asset plays two distinct roles in financial economics.

1. The price represents an *opportunity cost*. An asset's price appears in the wealth constraint as the amount that has to be paid, or is received, per unit of the asset. This is the conventional role for prices in economic analysis.
2. The price conveys *information*. Today's asset price reveals information about prices in the future.

The information conveyed by prices affects investors' *beliefs* and hence their actions (portfolios selected). Investors' actions determine the demand to hold assets in the aggregate and hence influence the assets' market prices.

1.3 The role of expectations

A famous passage in John Maynard Keynes's *General Theory* illustrates the role of expectations formation in financial markets (Keynes, 1936, p. 156).

> ... professional investment may be likened to those newspaper competitions in which the competitors have to pick out the six prettiest faces from a hundred photographs, the prize being awarded to the competitor whose choice most nearly corresponds to the average preferences of the competitors as a whole; so that each competitor has to pick, not those faces which he himself finds prettiest, but those which he thinks likeliest to catch the fancy of the other competitors, all of whom are looking at the problem from the same point of view. It is not a case of choosing those which, to the best of one's judgement, are really the prettiest, nor even those which average opinion genuinely thinks the prettiest. We have reached the third degree where we devote our intelligences to anticipating what average opinion expects average opinion to be. And there are some, I believe, who practise the fourth, fifth and higher degrees.

Here Keynes is posing a conundrum without proposing how to resolve it. Keynes's example may seem to involve circular reasoning: asset prices affect expectations, expectations affect decisions, decisions affect prices, and so on. Regardless of whether this is circular reasoning, the puzzle pinpoints the simultaneous interactions that occur between observed prices in the present and beliefs about prices in the future.

One implication is that the demand curve drawn in figure 1.1 should be treated with the utmost caution; when a price conveys information (as well as representing an opportunity cost) a simple downward-sloping demand curve may be difficult to justify – for a higher price today could lead investors to infer that the price will be even higher tomorrow, thus encouraging a greater demand to hold the asset in anticipation of a capital gain. In the presence of such 'extrapolative expectations', the demand curve could display a *positive* slope, at least for some prices.

It is common to assume that investors have 'rational expectations'; that is, their expectations are formed with an awareness of the forces that determine market prices. Moreover, in a rational-expectations equilibrium, the forces that determine prices include the decisions made by investors. This does not imply that investors are blessed with perfect foresight, but, at least, it does exclude expectations that are systematically wrong.

The rational-expectations hypothesis, on its own, is not much help in explaining asset prices. Firstly, rational expectations make sense only in the context of a model of price determination, including assumptions about investors' preferences and the information they possess. Secondly, investors may differ in the information they can bring to bear on their decisions – there may be *asymmetric information*. Thirdly, the information available changes over time as investors learn from their experience, or forget.

It is hardly surprising, in view of all these considerations, that building expectations formation into asset-pricing theories is both (a) central to any explanation of prices and (b) fraught with complications.

In an attempt to account for some of the imponderable features of price fluctuations, Fischer Black (1986) has introduced the concept of *noise* to financial analysis. From this perspective, some investors are assumed to act in arbitrary ways that are difficult – perhaps impossible – to explain as the outcome of consistent behaviour. These investors are called *noise traders*. *Rational traders* (sometimes called 'information traders' or 'smart-money investors'), on the other hand, are assumed to behave according to more coherent precepts, or to have better information, or better ways of processing the available information, than noise traders. (Asset price determination in the presence of noise traders is examined in more detail in chapters 2 and 10.)

The noise-trader approach falls with the broader framework of *behavioural finance*, which exploits ideas from outside conventional economics, including psychology. Behavioural finance can be understood as a modelling strategy that seeks to explain many otherwise puzzling phenomena – for example, empirical evidence that appears to be incompatible with the so-called *efficient markets hypothesis* (see below, section 1.7, and chapter 3). Whether behavioural finance can do a better job than orthodox theories in this regard remains an open question. At present, behavioural finance has succeeded more as a critique of conventional models than as a constructive alternative. Consequently, orthodoxy is likely to