1 The demand side

1.1 Overview

In the late 2000s, households and firms across the world cut back their spending and the global economy went into recession. There is little disagreement among economists that the dramatic fall in global GDP in 2008–09, which is now known as the global financial crisis, was a negative shock to aggregate demand.

This chapter begins the task of building a macroeconomic model with the demand side of the economy, which concerns the spending decisions of households, firms and government. As we shall see in Chapter 2, the supply side refers to the production activities in the economy. Dramatic changes in aggregate demand—that is, spending in the economy as a whole—as occurred in 2008–09 are unusual. In more normal times, shifts in aggregate spending decisions as well as shifts on the supply side of the economy are sources of the irregular fluctuations from recession to boom that form the business cycle.

If a period of depressed spending is forecast, there is much discussion in the financial press about whether the central bank will intervene and offset the likely recession by cutting the interest rate. The central bank would cut the interest rate because it expects this to encourage spending and help return the economy to stability. Similarly, if a boom in spending were forecast, the central bank would try to dampen it down by raising the interest rate. As well as the central bank, the government also needs to know how spending patterns are likely to evolve and affect output. A recession will depress tax revenue and increase spending on unemployment benefits. Forecasting patterns of spending is therefore a priority, not only for businesses, but also for the monetary and fiscal policy makers.

Aggregate demand and spending decisions

In this chapter, we focus on what lies behind the spending decisions in the economy and how they influence the level of economic activity. By the level of activity, we mean output or income. When output changes, employment also changes. For example, when output rises, more workers or longer hours are needed to produce the higher level of output. With more hours worked, the wage bill is larger and with higher sales, total profits are higher. This is why changes in economic activity are thought of as both changes in output and in income (i.e. wages and profits). When we think of a real-world economy, output is normally growing and recessions and booms produce fluctuations around a trend growth rate. However, we will often simplify by working with levels of output rather than its rate of growth.

Spending decisions are complex. For consumers they involve both a static component (what shall I buy today given my current income and the prices of the goods and services that are available?) and an intertemporal one (how do I allocate my spending over time given my expectations about how my income will evolve in the future?). Decision making for firms and the government also involves an intertemporal aspect. Firms make decisions to purchase machinery, equipment and premises based on a business plan that includes forecasts about how input costs and demand for their products will evolve over time. The government must also forecast demographic trends when making plans for building schools and hospitals.

Decisions by the entities that make up the economy—firms, families and government bodies—lie behind the demand side of the economy. However, macroeconomics is concerned with the aggregate sum of spending decisions of these groups, and the consequences of those decisions for economy-wide outcomes such as the rate of unemployment or inflation. Household spending decisions add up to aggregate consumption, C; firms' investment decisions add up to aggregate investment, I (note that I refers to spending on machinery, equipment and new houses and other buildings); and government spending on different goods and services adds up to a single number, G.

Including the purchase of new housing as part of investment highlights the difference between on the one hand, aggregate demand, which refers to spending on goods and services and on the other hand, the purchase of assets such as company shares or second-hand property. The purchase of a second-hand house does not contribute to aggregate demand. It is the transfer of the ownership of an asset from one household to another. In contrast, the building and selling of a new house uses resources and produces income, hence it influences aggregate demand.

Although the first chapters in the book relate to the closed economy (a single nation without links with others) as noted above, in the fuller model (set out in Chapters 9 and 10) the demand side includes foreign spending on home goods and services, exports (X), and domestic spending on foreign goods and services, imports (M). Aggregate demand is real expenditure on goods and services produced in the home economy. This can be summarized by an equation relating real expenditure, which is called y^D, to its individual components:

$$y^D = C + I + G + (X - M), \qquad \text{(aggregate demand)}$$

where we add to the trio of C, I and G, the expenditure of foreigners on home output, X and subtract the spending of home agents on output produced abroad, M.

Aggregate demand and government policy

Two of the major tools of macroeconomic policy, monetary and fiscal policy, work by influencing different elements on the demand side. Policy makers worry about fluctuations in aggregate demand because they affect unemployment and inflation.

Monetary policy seeks to stabilize aggregate demand by changing interest rates, which affect the investment decisions of firms and the purchase of durable goods like new cars and houses by households. A rise in the interest rate increases the cost of financing investment projects, and projects that would have gone ahead with lower interest rates are postponed or cancelled. Monetary policy also has indirect effects because the interest rate affects incentives to save and therefore shifts spending decisions over time. For example, by making new borrowing more expensive and increasing the return on saving, a higher interest rate will encourage households to postpone consumption.

Through changes in government spending on goods and services (G), fiscal policy affects aggregate demand directly. Fiscal policy can also be used to affect demand indirectly through its influence on household incomes and through that channel, on household spending. This is how changes in taxation and in the transfers made by the government to households in the form of pensions, disability and unemployment benefits feed through to affect aggregate demand. We will denote taxes minus transfers by T.

An important reason to study the demand side is to construct a model of the transmission mechanism by which monetary and fiscal policy, via the spending decisions of households, firms and the government, affect the economy. Before explaining the first building block of the model, which is the *IS* curve, we provide some facts about the demand side.

1.1.1 Facts about the demand side and business cycles

Shares of GDP

Table 1.1 sets out the average composition of gross domestic product (GDP) in five major economies between 2000 and 2005. GDP is the national accounts measure of national output. We discuss the calculation of GDP next subsection. Table 1.1 shows that consumption makes up the largest proportion of GDP in all the five economies. However, the importance of consumption in GDP ranges from 43% in China to nearly 70% in the US. A large part of this cross-country variation can be explained by differences in the contribution of investment. For example, an exceptionally high proportion of Chinese GDP arises from investment spending, which crowds out consumption spending. The table also highlights the variation across countries in the contribution to GDP from net exports. On average, between 2000 and 2005, the UK and the US ran current account deficits with exports in excess of imports, whereas the other three countries were in surplus (see Chapter 10 for further discussion and definitions).

The upper panel of Fig. 1.1 shows how the shares of consumption, investment and government spending in GDP changed over time for the UK from 1948 to 2010. The share of consumption follows a shallow U shape over the period; investment is mildly hump shaped and the share of government spending drifts upward throughout the period. Government spending is dominated by spending on services (e.g. health and education) where

Table 1.1 Shares of GDP (in %), current prices, average 2000–2005.

	Consumption	Investment	Government spending	Net exports
China	43.0	37.5	15.1	2.9
Germany	58.6	18.7	19.0	3.5
Japan	57.1	23.7	17.7	1.4
United Kingdom	65.3	16.7	20.0	−2.5
United States	69.9	19.0	15.3	−4.5

Note: GDP in current prices is a measure of nominal GDP; shares might not add up to 100% due to stockbuilding and statistical discrepancies.
Source: OECD National Accounts (data accessed November 2011).

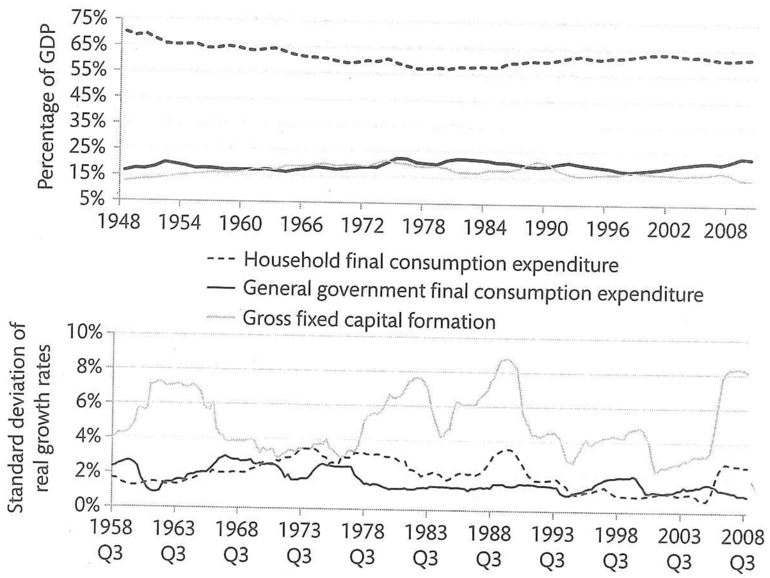

Figure 1.1 Components of GDP in the UK: as a percentage of GDP between 1948 and 2010 (upper panel) and volatility of growth rates between 1958 Q3 and 2008 Q4 (lower panel).

Source: UK Office for National Statistics (data accessed October 2011).

Note: Volatility has been calculated as the standard deviation of the GDP growth rate over a rolling 21-quarter period. The upper graph uses data from the series 'Gross Domestic Product by category of expenditure, current prices' and the lower graph uses data from the series 'Gross Domestic Product by category of expenditure, chained volume measures'.

productivity growth is slower than in industry. In sectors with comparatively slow productivity growth, prices rise relative to the rest of the economy. This is an important structural reason for the tendency for government expenditure, when measured in current prices, to rise as a proportion of GDP.

National accounts

The national accounts are used to measure the output of an economy. The most commonly used measure for calculating national output is gross domestic product, or GDP. GDP can be measured in three different ways. All three equations for calculating GDP are identities, which means that because of the way the variables are defined, the left hand side must always equal the right hand side. This special feature of an identity is signalled by the use of the equals sign with three bars.

Firstly, the *expenditure* method, which is one we use in the model of the demand side. This method measures GDP as the total expenditure on the economy's output of goods and services:

$$y \equiv C + I + G + (X - M), \qquad \text{(GDP identity, expenditure method)}$$

where $y \equiv$ output (or GDP), $C \equiv$ consumption, $I \equiv$ investment (including changes in stocks of raw material and finished goods), $G \equiv$ government spending and $(X - M) \equiv$ net exports. This is an identity, as it simply breaks down GDP into its constituent components. A fall in one of the components on the right hand side, for example consumption, will therefore always result in an equivalent reduction in measured GDP on the left hand side.

When a *new* house is bought, just like a new piece of machinery, this is an investment decision—the house provides a flow of services to the household over many years—and is included in the national accounts under *I*. The house and the piece of machinery form part of the economy's capital stock. National accountants have to make many tricky classification decisions: although households buy a variety of durable goods that provide services over many years, by convention it is only housing that is considered as investment in the national accounts. Cars and furniture, for example, are treated in the national accounts as consumption.

It is important to note that only government spending on goods and services is part of aggregate demand. *G* does not include government expenditure on transfers (e.g. pensions or social security payments). When the recipients of transfers spend their income from benefits or pensions, this is then recorded as consumption.

The expenditure method only includes *final* goods. For example, during the production process, firms buy raw materials and *intermediate* goods to make their final goods. The only purchases that are counted in GDP are when the firm sells the finished goods to the consumer. This avoids any double counting. In this way, GDP captures only the *value added* created in the economy. This leads onto the second approach for calculating GDP; the *value added* method. This method measures GDP as the value added created in all sectors of the economy, such that:

$$y \equiv \text{value of output sold} - \text{costs of raw materials and intermediate goods.}$$
$$\text{(GDP identity, value added method)}$$

The third and final approach for calculating GDP is the *income* method. This method measures GDP as the total income of all agents in the economy, such that:

$$y \equiv \text{salaries of workers} + \text{profits of the owners of capital.}$$
$$\text{(GDP identity, income method)}$$

The three methods of calculating GDP are all identities and each holds at each point in time. It makes intuitive sense that the total income in the economy is equal to the total expenditure, because for every transaction there is both a buyer and a seller. What is expenditure for the buyer is income to the seller. In practice, however, GDP calculated using the three methods may differ. This is a matter of *measurement error*, which can arise for several reasons. Measurement error may be more serious for one rather than another of the three methods. For example, tax evasion or the existence of a substantial black market are likely to produce a greater problem of underestimation of GDP when measured by incomes than for the other methods. If the underlying components could be measured completely accurately, then each of the three approaches would yield exactly the same estimate of GDP.

Relative volatility

The lower panel of Fig. 1.1 shows the relative volatility of growth rates for the three components of aggregate demand: a higher standard deviation means higher volatility. We can see that investment is much more volatile than consumption and government spending:

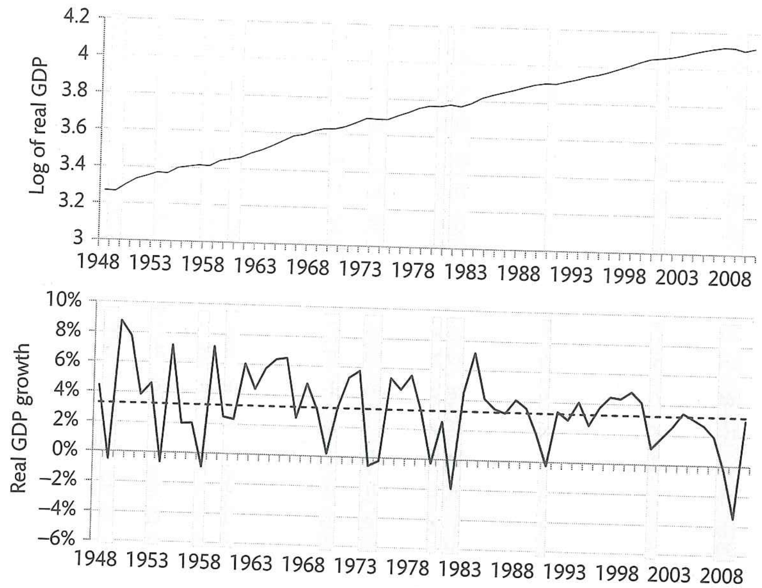

Figure 1.2 Business cycles in the United States between 1948 and 2010–Log of real GDP (upper panel) and real GDP growth (lower panel).

Source: US Bureau of Economic Analysis (data accessed October 2011).

Note: The shaded grey areas represent recessions as defined by NBER. Both graphs use data from the series 'Real GDP, chained dollars, billions of chained (2005) dollars'.

the line showing the standard deviation for investment lies above that for consumption and government spending. Investment depends on expected post-tax profits and is very dependent on how optimistic firms are, so it tends to flourish in boom periods and collapse in recessions, making it more volatile than the other components of GDP, or GDP itself. In addition to this, investment can also be postponed in recessions, whereas government spending and consumption cannot be as easily delayed. For example, a household still has to spend money on food and drink in a recession, whereas a firm may choose to wait until the economy has recovered before it undertakes investment. We shall also see—in Section 1.2.6—that fixed investment decisions are often bunched, further accounting for the greater volatility of investment.

Growth and cycles

The idea that economies fluctuate between phases of boom and recession is confirmed by looking at the data for the USA. Figure 1.2 illustrates the long-run growth and business cycles in the US economy. In the top panel, the log of GDP in constant prices is plotted and shows the rather steady growth rate of GDP over the long run. Using the log scale, a straight line would represent a constant growth rate from 1948 to 2010.[1] In fact, annual growth rates in each decade from 1948 were: 3.3%, 4.8%, 3.2%, 3.0%, 3.1% and 2.5%.

[1] See Section 3 of Chapter 8 for a detailed discussion of growth concepts and the use and interpretation of natural logs.

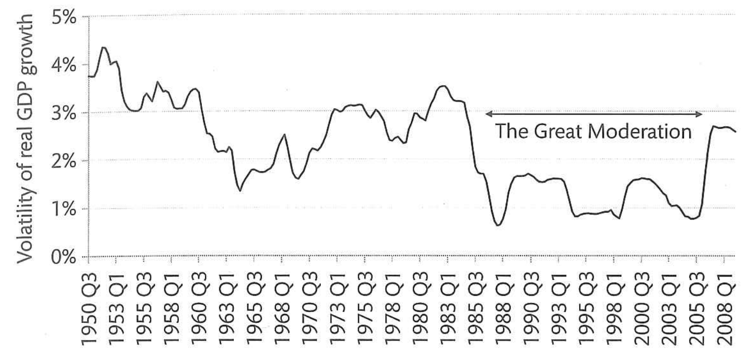

Figure 1.3 Volatility of real GDP growth in the United States: 1950 Q3–2009 Q1.

Source: US Bureau of Economic Analysis (data accessed October 2011).

Note: Volatility has been calculated as the standard deviation of the GDP growth rate over a rolling 21-quarter period. Graph uses data from the series 'Real GDP, chained dollars, billions of chained (2005) dollars'.

By plotting the *annual* growth rate of GDP in the lower panel, the fluctuations of the economy around the long-run trend are highlighted. The dotted line represents the average growth rate between 1948 and 2010 of 3.3%. Peaks and troughs of growth occur about twice a decade. In the US, an independent body called the National Bureau of Economic Research (NBER) establishes the dates of US business cycles and states there have been eleven separate contractions (i.e. recessions) over this period, which are represented by the shaded grey areas on the graphs. Recessions start at the peak of a business cycle and end at the trough: during recessions, economic activity is contracting. In the period we are analysing, the recessions range from 6 to 18 months in duration. This is why the *annual* GDP growth rate shown in Fig. 1.2 is not negative in all of the recessions (e.g. the 2001 recession). The longest of these recessions was the 2007–09 global financial crisis (GFC), which spanned a year and a half.

Before the global financial crisis, many macroeconomists claimed that better policy making had produced greater macroeconomic stability. This led to the use of the term 'the Great Moderation' to describe the calmer macroeconomic conditions from the mid-1980s (fewer and shorter recessions as illustrated in Fig. 1.2).

A more precise way to document the Great Moderation is to look at the volatility of GDP growth. Figure 1.3 shows the volatility of US GDP growth since the late 1950s. By plotting the standard deviation of the growth rate over a 21-quarter rolling period, changes in volatility can be seen easily. There is a noticeable fall in the volatility of GDP growth from the mid 1980s. However, the spike in volatility experienced during the recent recession brought the era of the Great Moderation to an end.

1.1.2 **Introducing the *IS* curve**

In this chapter, we begin to build up a way of capturing essential elements of the complex real-world macro-economy in a model. The model is called the 3-equation model because there are three core elements, representing the demand side, the supply side and the policy maker. The *IS* is the curve that represents the demand side. It is called the *IS* curve because it refers to planned investment and savings decisions; hence the *I* and the *S*. It was originally formulated as a simple shorthand version of John Maynard Keynes's description of the demand side in his General Theory.[2]

Models are important for providing a framework in which to think about complex problems more easily and more carefully. By modelling the goods market we can answer interesting questions about the demand side that we might otherwise struggle with, such as:

1. If you give an individual a $100 bonus, how much of it will they spend?
2. How much will output increase following a rise in government spending or a boost in the confidence of households and firms about their future prospects?
3. To what extent will a rise in the interest rate curtail investment in new housing, machinery and equipment?

The *IS* curve is a way of summarizing the way that aggregate output in the economy is affected by changes in the spending decisions of families, firms and government bodies. For example, when firms increase their spending on new equipment, this triggers increased production in the capital goods producing sector of the economy. More people are employed to produce the extra capital goods and as they spend their wages, demand for consumer goods goes up and employment and output expand in those sectors as well. As a result of this process, the economy will move to a higher level of output and employment. We shall see that in response to the initial expansion of spending, output will expand until the extra *S*aving households want to make just balances the extra spending on *I*nvestment. At that point, the process by which the impact of the initial increase in aggregate demand is multiplied through the economy in subsequent rounds of higher employment, spending by newly employed workers, higher demand for consumer goods, etc. comes to an end.

Aggregate demand in the private sector

Focusing attention on decisions by the private sector, the level of aggregate demand will be affected by current income and by the following factors.

1. *Expectations about the future*: the plans of firms to invest in new equipment and premises depend on their expectations of *future post-tax profits*. If firms anticipate high levels of capacity utilization and strong order books, they will increase investment in new capacity. Households prefer to have smooth rather than fluctuating consumption, which means they need to save and borrow in order to spread their consumption more evenly over time. To make their saving and borrowing decisions, they must form a view about the future growth of their income. The *life-cycle motive* for saving refers to the planning of saving taking into account the projected pattern of income during a person's working life and retirement. Households will revise upward their estimate of how much they can spend each period if they have new information that leads them to expect their income to grow more strongly.

[2] See Keynes (1936). The *IS* curve also underlies the famous *IS/LM* model, which was introduced by John Hicks in 1937 (Hicks, 1937). For a concise and interesting discussion of its origins and impact, see Durlauf and Hester's article entitled *IS-LM* in the New Palgrave Dictionary of Economics: 2nd edition (2008).

Firms and households form their expectations in the face of uncertainty. For example, an increase in the unemployment rate in the economy could be a signal to households that *uncertainty* about their future income has risen, which would trigger an increase in saving for *precautionary* reasons.

2. The extent of *credit constraints*: these arise because of problems faced by banks in assessing the creditworthiness of households and firms. It is impossible for banks to have full information about borrowers' projects and actions. Borrowing by households and by small and medium-sized firms is therefore often restricted by banks. Households and business that cannot borrow as much money as they would like are said to be *credit constrained*. The information problems inherent in bank lending mean that access to credit is often highly dependent on the amount of *collateral* the borrower has with which to secure the loan. The most common form of collateral for households is the value of their house. This means that changes in the value of collateral, which occur for example when house prices change, affect consumption and investment because they either tighten or relax credit constraints.

3. The *interest rate*: there are a number of channels through which the interest rate will affect aggregate demand. When interest rates go up, households find it more expensive to get a mortgage. This reduces the demand for new houses and reduces the demand for furnishings and other consumer durables that go along with moving house. A higher interest rate will lead firms to rein in their spending plans on new capital equipment and buildings. Households will tend to postpone consumption spending because of the improved returns from saving. However, creditor and debtor households would be expected to react differently. For example, a creditor household will find their income has gone up when the interest rate rises and this will boost spending through the income effect. For debtor households, the effect will be the opposite. The first and third effects are normally stronger than the second one and a higher interest rate is associated with lower aggregate demand.

The IS curve

The *IS* curve is a way of summarizing in a diagram the demand side in the macroeconomic model. It shows the combinations of the interest rate and output at which aggregate spending in the economy is equal to output.

Figure 1.4 shows the *IS* curve. This is a downward-sloping relationship. To see this, think of the combination of a high interest rate and low output. When the interest rate is high, spending on housing, consumer durables, machinery and equipment will be low. This means aggregate demand is low and a low level of output will satisfy the low demand. Hence, we have the first point on the *IS* curve. Now, take a combination of a low interest rate and high output. Here the situation is the opposite: buoyant spending on new houses, consumer durables and investment goods generates a high level of output and high incomes for households. This is the second point on the downward-sloping *IS* curve.

To show how changes in profit or income growth expectations, uncertainty and the value of collateral can be captured in the diagram, we hold the interest rate constant and look at shifts in the *IS* curve. As an example, in a situation of depressed profit expectations, we would expect firms to postpone new investment. The result is lower investment spending at any interest rate. The *IS* curve shifts to the left (shown by the curve labelled *IS* (pessimistic expectations)).

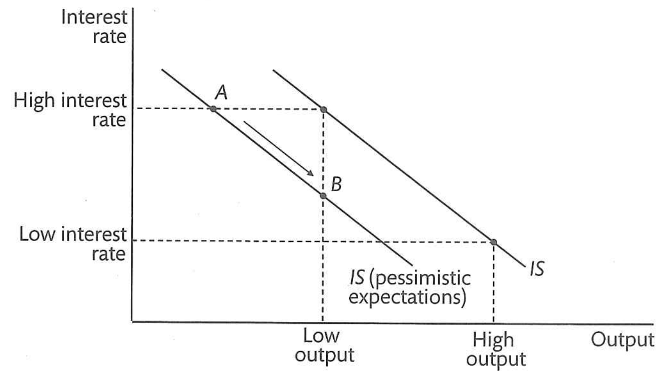

Figure 1.4 The *IS* curve—the effects of changes in optimism and economic policy.

Using the *IS* diagram, we can show how the central bank or the government can affect the demand side. Following up the situation where the business environment becomes more pessimistic, the central bank could lower the interest rate to stimulate investment. This would be shown as a move *along* the *IS* curve from point *A* to point *B*. As an example, there was much discussion about the decision by the US central bank, the Federal Reserve, to reduce the policy interest rate and keep it low as a way of stimulating investment in the aftermath of the collapse of the Dotcom bubble in 2001. The long period of low interest rates stimulated investment in new house construction, and we shall see in Chapter 7 the role that this played in the background to the global financial crisis of 2008.

Another response to a leftward shift of the *IS* curve due to pessimistic profit expectations could be action by the government rather than central bank. If the government decides to launch a major expenditure programme, such as to improve healthcare provision or to install an information super-highway, this will *shift* the *IS* curve to the right: at a given interest rate, the government purchases a larger amount of goods and services. Under our assumption that suppliers will respond to the higher demand, the economy moves to a higher level of output and employment. In the diagram, this would shift the *IS* curve to the right back towards its initial position.

A much debated issue in macroeconomics is by how much a government expenditure programme would increase output in the economy. Would extra government spending increase output one-for-one, by more or by less? This debate is frequently couched in terms of the size of the 'multiplier': i.e. by how much would expenditure of an extra dollar raise GDP? As we shall see in Section 1.2, in the *IS* diagram, the larger the multiplier,

1. the larger is the rightward shift of the *IS* curve associated with any given amount of additional government spending and
2. the flatter is the *IS* curve, since with a larger multiplier, a given cut in the interest rate has a larger effect on output.

The demand side is only part of the macroeconomic model. To understand how the spending decisions of households, firms and the government fit into the bigger picture, we need to include the supply side and the motivations and behaviour of the policy makers.

Chapter 2 addresses the supply side and Chapter 3 brings in the policy maker. Before moving on, we set out the modelling of the demand side in more detail.

1.2 Modelling

The Modelling sections of the chapters in the book provide more details of the models and their components introduced in the Overview sections. Detailed derivations of some of the results are available for download from the book's website.

In order to better understand the demand side of the macro-economy, we build a simple model, which is captured by the *IS* curve. This section focuses on the closed economy *IS* curve and we will come back to the open economy *IS* curve in Chapters 9 and 10. The closed economy *IS* curve splits demand into three components; consumption, investment and government spending.

As a first step in this chapter, we introduce the concept of goods market equilibrium and the definition and mechanics of the multiplier process. Throughout this section, we assume that firms are willing to meet the higher demand for their goods and services, and workers are willing to take the extra jobs or work the extra hours that are offered. In simple terms, we assume that the supply of output adjusts to meet the demand for goods, services and labour. The next step is to bring in additional determinants of aggregate spending listed under points (1) to (3) in the previous section and to modify the *IS* equation to take these into account. Along the way, we refer to empirical evidence to support the richer model.

1.2.1 Goods market equilibrium

We begin our discussion and modelling of the goods market by considering each of the components of aggregate demand in the closed economy. The aggregate—or economy-wide—demand for goods and services consists of:

Consumption demand: expenditure by individuals on goods and services. Spending is on both *durable* products such as a car, laptop or sofa and on *non-durable* products such as a theatre show or groceries.

Investment demand: expenditure on capital goods (machinery, equipment, and buildings). Spending is by households on new houses, by firms on new capital goods, including structures, and building up inventories of materials or finished goods, and by government (on public infrastructure projects, such as building a new high speed rail line).

Demand stemming from government purchases: government expenditure on salaries, goods and services. Spending includes public sector wages (e.g. civil servants, doctors), purchases of goods (e.g. educational supplies, ammunition for the army) and purchases of services (e.g. waste disposal and contract cleaning).

The model of goods market equilibrium

We now set out a model of the goods market. In the closed economy, aggregate demand, y^D, is given by

$$y^D = C + I + G, \tag{1.1}$$

and we ignore imports and exports. All variables are in real terms, which is also referred to as constant price terms. Equilibrium in the goods market requires that planned real expenditure on goods and services (i.e. aggregate demand) is equal to real output

$$y^D = y, \qquad \text{(goods market equilibrium)}$$

where y is output. y is also income: spending on the output of the economy in turn becomes the income of those producing it (wages and profits). This circular flow of income to expenditure and output, and back to the incomes of producers means we can use the terms output, income and expenditure interchangeably. We shall come back to the goods market equilibrium condition after providing some more detail about consumption and investment behaviour.

To start, we assume that aggregate consumption is a simple linear function of after-tax or disposable aggregate income

$$C = c_0 + c_1(y - T), \tag{1.2}$$

where T is total taxes net of transfers.

To get to the consumption function we use as the core of the model of the demand side, we make the additional assumption that taxes are a fixed proportion of income—i.e. $T = ty$, where $0 < t < 1$. The consumption function then becomes

$$C = c_0 + c_1(1 - t)y. \qquad \text{(Keynesian consumption function)}$$

This is referred to as a Keynesian consumption function. It consists of a constant term, c_0, which is often referred to as autonomous consumption, and $c_1(1 - t)y$, which shows households spending a fixed proportion of their disposable income. c_1 is referred to as the marginal propensity to consume and lies between zero and one ($0 < c_1 < 1$).

Note that the marginal propensity to consume (MPC) shows the change in consumption as the result of a change in post-tax or disposable income:

$$MPC \equiv \frac{\Delta C}{\Delta y^{\text{disp}}} = c_1, \qquad \text{(marginal propensity to consume)}$$

where $y^{\text{disp}} = (1 - t)y$.

In this model of consumption, saving is $S = y - C$ and adds to one's assets. Income is the amount that can be consumed without reducing one's assets. The corresponding marginal propensity to save is s_1 and $c_1 + s_1 = 1$ (i.e. disposable income can be either saved or spent).

If disposable income is zero, the consumption function predicts consumption equal to c_0. For this to be the case, households must have some savings and or they must be able to borrow. A more satisfactory model of consumption must therefore include forward-looking behaviour in order to account for savings, and we build this in to the model of consumption later in the chapter.

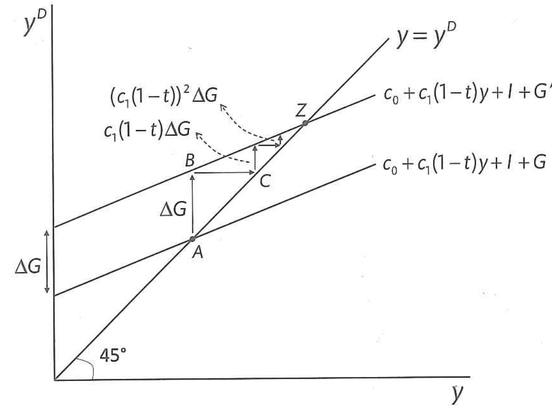

Note: $\Delta G = G' - G$

Figure 1.5 Keynesian cross—Increase in government spending.

1.2.2 The multiplier

We begin with the Keynesian consumption function in order to introduce in the simplest way the concept of how equilibrium in the goods market is determined and how the multiplier process works. If we substitute the consumption function into the equation for aggregate demand, then we get a relation between aggregate demand,

$$y^D = c_0 + c_1(1 - t)y + I + G, \qquad \text{(aggregate demand)}$$

and output, y, which we can draw on a diagram with aggregate demand on the y-axis and output on the x-axis (Fig. 1.5). The intercept of the curve will be $(I + G + c_0)$ and the slope $c_1(1 - t)$ which, given the assumptions that both c_1 and t are between 0 and 1, will itself be between zero and one. We can also draw the goods market equilibrium condition on the same graph, which will be a 45° line, because in equilibrium

$$y = y^D.$$

The point at which the two curves intersect shows the level of output where planned real expenditure on goods (and services) by firms, households and government is exactly equal to the level of goods being supplied in the economy. To the left of the 45° line, the demand for goods is greater than the supply.

To understand how the model works, we disturb the initial equilibrium at point A. Suppose there is an increase in government spending. From the aggregate demand equation, we can see that this shifts the aggregate demand curve upwards by the change in government spending, ΔG. Aggregate demand now exceeds output (Point B). As the government increases its purchases of goods (e.g. of office equipment) the stocks of these goods in warehouses decline. The inventory management software records the fall in stocks and triggers an increase in production: output rises. This is the move from B to C, where once again $y = y^D$. The higher output in turn raises incomes (in the form of wages of the additional workers employed and the profits of the owners of the firms making higher sales) and according to the consumption function, some proportion of the higher incomes are spent on goods and services in the economy, raising aggregate demand further. The process continues until the new goods market equilibrium is reached at point Z, where output and aggregate demand are equal.

Adjustment to shifts in aggregate demand via inventory accumulation and decumulation shows how a purely quantity-based adjustment to a new goods market equilibrium can take place. For the goods market to be in equilibrium it must be the case that $y = y^D$. If we substitute y into the aggregate demand equation we can rearrange to define equilibrium output in terms of the *exogenous variables*. Exogenous variables are those that are determined outside the model. In the model of goods market equilibrium, the exogenous variables are autonomous consumption (c_0), investment (I) and government spending (G), the marginal propensity to consume (c_1) and the tax rate (t). If we know the value of these five variables, we can work out the level of equilibrium output,

$$y = c_0 + c_1(1-t)y + I + G \tag{1.3}$$

$$y - c_1(1-t)y = c_0 + I + G$$
$$y = \underbrace{\frac{1}{1-c_1(1-t)}}_{\text{multiplier}} \underbrace{(c_0 + I + G)}_{\text{autonomous demand}}. \tag{1.4}$$

Because c_1 and t are between 0 and 1, this simplified model implies that the *multiplier* is greater than 1. This means that a 1% increase in autonomous demand would be predicted to lead to an increase in output of more than 1%. This is referred to as the short-run multiplier as we hold the interest rate and all other policy responses constant—i.e. output is the only variable allowed to change. Note that if consumption was not a function of current income, i.e. if the marginal propensity to consume, c_1, is equal to zero, then the aggregate demand line is horizontal and there is no multiplier process amplifying the impact on the economy of the rise in government spending.

In the consumption function, c_0 is an exogenous variable, meaning that it is not influenced by any other variables in the equation (i.e. it is determined outside the model). However, the second part of the equation, $c_1(1-t)y$, is *endogenous*. In order to calculate the level of aggregate consumption in equilibrium, we therefore need to first work out the equilibrium value of y and then compute the level of consumption.

We now combine algebra and geometry to deepen our understanding of the multiplier process. We return to the previous example shown in Fig. 1.5, where the government uses fiscal policy to increase activity in the economy. Since the aggregate demand curve is flatter than the market-clearing condition, equilibrium output will increase by more than the increase in government expenditure, i.e. the multiplier is greater than one. The initial boost in government spending increases aggregate demand by ΔG. This increases output, and income, and so aggregate consumption increases by $c_1(1-t)\Delta G$. Aggregate demand increases by the same amount, and, given that the goods market clears, output and income increase too. Aggregate consumption increases again by $[c_1(1-t)]^2 \Delta G$, and so on to infinity. The total increase in output is given by the sum of these changes:

$$\Delta y = \Delta G + [c_1(1-t)]\Delta G + [c_1(1-t)]^2 \Delta G.... \tag{1.5}$$
$$= \left(1 + [c_1(1-t)] + [c_1(1-t)]^2 + ...\right)\Delta G. \tag{1.6}$$

This is a geometric progression and its sum to infinity is given by[3]

$$\Delta y = \underbrace{\frac{1}{1-c_1(1-t)}}_{\text{multiplier}} \Delta G = k\Delta G, \tag{1.7}$$

where k is the multiplier.

Application: the paradox of thrift

Imagine the economy is in a recession and the following question is put to an economist: should a policy be introduced to encourage or to discourage saving in order to help the economy recover? On the one hand, the idea of encouraging more saving sounds helpful—if more is saved and invested in new capital stock, that would seem a good recipe for recovery. On the other hand, if more saving means less consumption and there is no increase in investment, won't aggregate demand fall and the recession get worse?

To answer this question requires spelling out how agents are assumed to behave. In an economic model, assumptions about how agents make their decisions are summarized by behavioural equations. In macroeconomics, it is aggregate behaviour that is captured by these equations. In the very simple model we have set out, the behavioural equation for consumption is the Keynesian consumption function. Investment and government spending are assumed to remain constant: in the model so far, they are not affected by other variables (i.e. they are exogenous). The behavioural equations are as follows:

$$C = c_0 + c_1(1-t)y\,;\, G = \overline{G}\,;\text{ and } I = \overline{I}, \tag{behavioural equations}$$

where we put a bar over the G and the I to emphasize that these are both exogenous. From the behavioural equations, we can write planned real expenditure (aggregate demand) as

$$y^D = c_0 + c_1(1-t)y + \overline{I} + \overline{G}. \tag{1.8}$$

As usual, the goods market equilibrium condition is $y^D = y$, and if we substitute Equation 1.8 into this, then in equilibrium, planned savings are equal to planned investment plus government spending:

$$\underbrace{(1-c_1(1-t))y - c_0}_{\text{Planned saving}} = \underbrace{\overline{I} + \overline{G}}_{\text{Planned } I \text{ and } G}, \tag{equilibrium condition}$$

[3] This is the sum of a geometric series. We want to find an expression for the series $1 + [c_1(1-t)] + [c_1(1-t)]^2 + [c_1(1-t)]^3 + \cdots$, which we call x. If we put

$$x = 1 + [c_1(1-t)] + [c_1(1-t)]^2 + [c_1(1-t)]^3 + \cdots,$$

and multiply both sides by the common factor, $[c_1(1-t)]$, then

$$[c_1(1-t)] \cdot x = [c_1(1-t)] + [c_1(1-t)]^2 + [c_1(1-t)]^3 + \cdots$$

and if we subtract the bottom equation from the one above, we have

$$x(1 - [c_1(1-t)]) = 1$$
$$\Rightarrow x = \frac{1}{1-[c_1(1-t)]}.$$

We model the proposal to encourage savings by a fall in autonomous consumption from c_0 to c_0'. Aggregate demand falls initially by ($c_0 - c_0'$) and the multiplier process works in the downwards direction. Using the same logic as illustrated in Fig. 1.5 in reverse, this process continues until we get to the new equilibrium at lower output. At the new lower output level, y', planned savings on the left hand side of Equation 1.9 is equal to planned investment plus government spending on the right hand side, which have not changed:

$$(1 - c_1(1 - t))y' - c_0' = \bar{I} + \bar{G}\,. \tag{1.9}$$

An important insight emerges from this example. The initial equilibrium was disturbed by a fall in autonomous consumption as households sought to increase their savings. However, the intention to save more did not lead to higher aggregate savings because *income fell* ($\Delta y = k(c_0 - c_0')$). This must be the case for Equation 1.9 to hold.[4] This is called the *paradox of thrift*, because if greater thriftiness is not matched by higher investment in fixed capital, income will fall and there will be no overall increase in savings in the economy.

To summarize, the answer to the question of whether saving should be encouraged in a recession depends on the model of the economy the economist is using–and of course, on how well that model matches the real economy under study. Using the model we have developed so far in this chapter, the answer is clear: encouragement to save more will not help the economy to exit recession. The reason is that there is no mechanism in *this* model through which higher saving is translated into higher investment. Investment remains at $\bar{I}$ throughout. Hence the result of a higher propensity to save is that aggregate demand falls, output falls and the recession is deepened.

By contrast, if the model includes a central bank, then the recession could be averted by the central bank cutting the interest rate and boosting investment to offset the fall in savings. In the next section, we remove the simplifying assumption of exogenous investment to show the impact on the demand side of investment being responsive to changes in the interest rate.

1.2.3 The *IS* curve

Deriving the IS equation

Throughout this book, we use the *IS* curve to represent the demand side of the economy. In this subsection, we derive the *IS* curve and use it as a starting point to think about how monetary and fiscal policy work. As we discussed in Section 1.1.2, the *IS* curve shows the combinations of output and the real interest rate at which the goods market is in equilibrium.

We first set out the relationship between the real interest rate, r, and the nominal interest rate, i. This relationship is shown by the Fisher equation:

$$r = i - \pi^E. \qquad \text{(Fisher equation)}$$

4 We can confirm that the change in aggregate saving is equal to zero. Using Equation 1.7:

$$\Delta y = \frac{1}{1-c_1(1-t)}\Delta c_0,$$

$$\Delta S = \left[(1 - c_1(1 - t))\frac{1}{1-c_1(1-t)}\Delta c_0\right] - \Delta c_0 = 0.$$

This equation says that the real interest rate is simply the nominal interest adjusted for expected inflation. It is the real interest rate that is most important for investment and saving decisions, as it represents the true cost of borrowing (and the true return on saving). This interest rate is therefore the one used in the investment equation and in the *IS* curve. When the central bank sets the nominal interest rate, it does this with the intention of achieving a particular real interest rate since it aims to affect interest-sensitive spending. In the Appendix, we show how the Fisher equation is derived.

At this stage, we assume there is just one interest rate in the economy that applies to all borrowing and saving. We make the assumption to keep the maths as simple as possible. In reality, there is a spectrum of interest rates. For example, the interest rate on bank lending will typically be higher than the interest rate set by the central bank. We discuss the banking mark-up in Chapter 5 and the difference between short-term and long-term interest rates in Chapter 7.

Although we shall amend this later in the chapter, in the model to this point, consumption is independent of interest rates. This means that the transmission of monetary policy will work through its effect on investment spending, including on new houses. Up to this point in the chapter, we have assumed that investment is determined by expected future profits, which we have assumed to be exogenous. We now incorporate the interest rate into the investment function as follows:

$$I = a_0 - a_1 r, \qquad \text{(investment function)}$$

where r is the real interest rate, a_0 and a_1 are constants and $a_1 > 0$. The main determinant of investment is expected future post-tax profits, which is captured by the term, a_0.

We derive the *IS* curve as follows. Substituting the investment function into Equation 1.4 gives us a relationship between the real interest rate and output, which is the *IS* curve. We can then use k to denote the multiplier and simplify to achieve a simple equation for the *IS* curve:

$$\begin{aligned}
y &= \underbrace{\frac{1}{1 - c_1(1-t)}}_{\text{multiplier}}[c_0 + (a_0 - a_1 r) + G] && (1.10)\\
&= k[c_0 + (a_0 - a_1 r) + G] && (1.11)\\
&= k(c_0 + a_0 + G) - k a_1 r && (1.12)\\
&= A - ar, && (\textit{IS}\ \text{curve})
\end{aligned}$$

where $A \equiv k(c_0 + a_0 + G)$ and $a \equiv k a_1$.

The derivation of the equation for the *IS* curve highlights the fact that given r, equilibrium output, y, is found by multiplying autonomous consumption and investment, and government spending by the multiplier $k = \frac{1}{1-c_1(1-t)}$. This fixes an r-y combination on the *IS* curve. It is clear from Equation 1.12 that a higher multiplier, k, or higher interest-sensitivity of investment to the interest rate, a_1, increases the effect of a change in the interest rate on output, making the *IS* curve flatter.

The *IS* curve is derived graphically in Fig. 1.6 using the *IS* curve equation shown above. Three curves are shown on the figure: the investment function, the investment function plus autonomous consumption and government spending and the *IS* curve. Two real interest

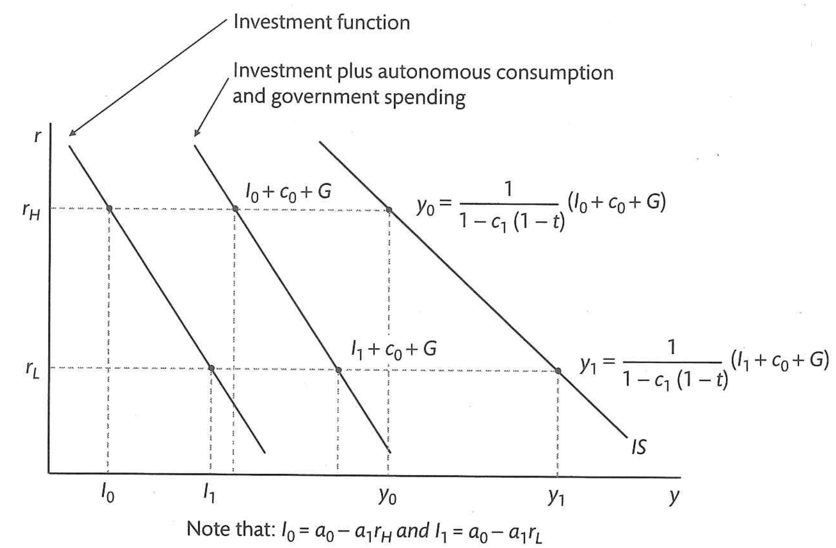

Figure 1.6 Deriving the IS curve.

rates are shown in the figure: a high real interest rate, r_H, and a low real interest rate, r_L. At r_H, investment is more costly, so the level of planned investment is low, whereas the opposite is true at r_L. This leads to a downward-sloping investment function. Autonomous consumption and government spending are unaffected by the interest rate, so when added on, the curve shifts out parallel. The last step to derive the *IS* curve is to multiply $I_0 + c_0 + G$ by the multiplier. This gives output equal to planned expenditure, y, on the *IS* curve. The *IS* curve shows all the combinations of output and real interest rate where the goods market is in equilibrium. It is flatter than the other two curves because the multiplier effect of investment is higher when interest rates are low.

Using the *IS* curve equation and the diagram, we can summarize its properties.

1. The *IS* curve is downward sloping because a low interest rate generates high investment, which will be associated with high output. By contrast, when the interest rate is high, investment and consequently equilibrium output is low.
2. The slope of the *IS*:
   (a) Any change in the size of the *multiplier* will change the slope of the *IS* curve. For example, a rise in the marginal propensity to consume, c_1, or a fall in the tax rate, t, will increase the multiplier, making the *IS* flatter: it rotates counter clockwise from the intercept on the vertical axis.
   (b) Any change in the *interest sensitivity of investment* (a_1) will lead to a consequential change in the slope of the *IS* curve: a more interest-sensitive investment function (i.e. $\uparrow a_1$) will be reflected in a flatter *IS* curve.
3. Shifts in the *IS* curve: any change in *autonomous consumption, autonomous investment* or *government expenditure* (c_0, a_0, G) will cause the *IS* curve to shift by the change in autonomous spending times the multiplier.

1.2.4 Forward-looking behaviour

The spending decisions of agents in the present are influenced by their expectations of the future. This means there is an intertemporal component to both consumption and investment.

1. Households adjust their current spending based on their expected income in the future. For example, a final year economics undergraduate who secures a lucrative job contract (which is set to commence after university), has low current income, but high expected future income. Having got the job, this individual may borrow in the present, consume more and pay back the money when they start their job. This behaviour is referred to as *consumption smoothing*.

2. Firms make decisions (e.g. purchasing machinery, equipment and premises) based on a business plan that includes forecasts about future demand for their products and input costs. For example, a firm selling cars to China might choose to build a new factory (i.e. undertake investment) based on forecasts that Chinese incomes (and therefore their demand for cars) will continue to rise rapidly over the next 20 years. Investment is intrinsically forward looking, as it incurs a cost today, but the stream of benefits occurs in the future.

Present value calculation

We can calculate the present value of a flow of income or profits received in future periods in the following way. We take the example of a firm making an investment decision.

We assume that firms aim to maximize profits, so they undertake investment projects if these offer a return that is higher than costs. The outlay on investment typically precedes the returns, which may be lumpy and spread over a number of years. The way to deal with this is to calculate the 'Present Value' (V) of the expected flow of profits Π, where Π_t, Π_{t+1}, ... is the flow of profits in period t, $t+1$ etc. Suppose the interest rate is 10%. Then if I save €100 today I will have €110 in a year's time (100 (1 + 10%)). Expressed the other way around, we could say that €110 in a year's time has the same value as €100 today—its present value is €100.

More generally, if interest rates are constant at r, the value of X in n years' time is the same as that of $X/(1+r)^n$ today. We can calculate the present value of our stream of expected profits, Π^E, from an investment project with profits over T years by:

$$V_t^E = \Pi_t^E + \frac{\Pi_{t+1}^E}{(1+r)} + \frac{\Pi_{t+2}^E}{(1+r)^2} + \cdots = \sum_{i=0}^{T} \frac{1}{(1+r)^i}\Pi_{t+i}^E.$$

(expected present value calculation)

We use the superscript E to denote the expected value of a variable.

If the cost of the machine is greater than the present value of the flow of profits from the machine, then it would be more profitable not to buy the machine but instead put the money in the bank or in bonds (which earn interest of r). Similarly, if the money to purchase the machine is being borrowed, then if the cost of the machine is greater than the present value, buying the machine will be unprofitable. On the other hand, if the present value is greater than the cost, then this investment is profitable, and a profit-maximizing firm will go ahead with it.

We can also apply the same logic to modelling consumption decisions. We can account for the fact that the future will affect consumption decisions by calculating the present value of the stream of expected income over a person's lifetime. If we assume that individuals live forever, we can use the formula for calculating present value to calculate the expected present value of lifetime wealth (Ψ^E, pronounced 'sigh'), which is defined at time t as follows,

$$\Psi_t^E = (1+r)A_{t-1} + \sum_{i=0}^{\infty} \frac{1}{(1+r)^i} y_{t+i}^E,$$

(expected present value of lifetime wealth)

where $\sum_{i=0}^{\infty} \frac{1}{(1+r)^i} y_{t+i}^E$ is the present value of expected post-tax lifetime labour income and $(1+r)A_{t-1}$ are the resources available in period t from the assets the individual held at the end of the previous period.[5]

1.2.5 Consumption

People's income fluctuates over their life cycle; it also fluctuates when they lose their job, move to a different job or get a promotion. Given that people prefer to smooth out the fluctuations in their income in their spending behaviour, they must take account of the future and they must be able to save and borrow. The modelling of consumption should include how households form expectations about the future and how they borrow and save.

The desire to smooth consumption in the face of fluctuations in income is captured by the assumption of diminishing marginal utility of consumption. A simple way of visualizing this is to consider the choice of consumption in a world of just two periods. If we know income will be higher next period, how does that affect consumption now? Consider the choice between either low consumption equal to income in the first period and high consumption equal to income in the second period, or having the average of the two consumption levels in each period. If there is diminishing marginal utility of consumption, more consumption always increases utility, but successive increases in consumption deliver smaller and smaller benefits. Therefore, households will make the second choice, because consuming the average in both periods offers higher utility than the first choice.

The permanent income hypothesis (PIH)

The permanent income hypothesis (PIH) states that individuals optimally choose how much to consume by allocating their resources across their lifetimes.[6] Their resources include their

[5] The assumption that individuals 'live forever' can be thought of as assuming that households have children or heirs and that they incorporate the utility of their children or heirs into their consumption function. In other words, households behave 'as if' they last forever. The assumption is clearly unrealistic, but it greatly simplifies the mathematical derivation of the permanent income hypothesis (PIH). From this point onward, however, we will refer to lifetime income and earnings in the text, as it is more intuitive to think about forward-looking behaviour in terms of a household maximising utility over their lifetime.

[6] This view of the consumption function was first developed by Modigliani and Brumberg (1954). See also Friedman (1957). For a review of Friedman's theory, its influence on modern economics and the relevant empirical literature, see Meghir (2004).

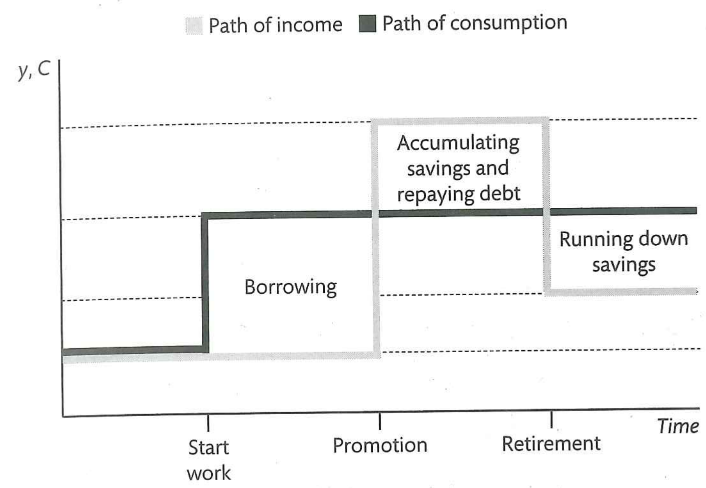

Figure 1.7 The permanent income hypothesis over the life cycle.

assets and their current and future income. This is a forward looking decision and will depend on interest rates, asset values, expectations of future income and expectations of future taxes. The PIH predicts that optimal consumption is smooth as compared to income. For example, when individuals start work, their income is low and they will borrow to consume more; when income increases, they keep consumption constant and use the excess income to pay off debts and save for retirement; then at retirement, their income falls and they draw down their savings. Figure 1.7 shows how consumption and income change over the life cycle in this simplified example. The important point is that the PIH predicts that households will borrow and save in order to smooth consumption over their lifetimes. Likewise, over the business cycle, if an individual becomes unemployed, the model predicts that they will borrow in order to sustain consumption during the spell of unemployment. As we shall see in Chapter 14 on fiscal policy, the government plays an important role in smoothing consumption through the provision of unemployment benefits.

The PIH model of consumption provides a stark contrast with the Keynesian consumption function where there is no explicit consideration of the future. Aggregate consumption of households is modelled there as the consumption of a fixed amount, c_0 and a fixed proportion of the current period's disposable income.

The PIH model is derived in more detail in Section 1.4.2 of the Appendix, but it is useful to set out the intuition and predictions of the model here, as it provides a framework for thinking about forward-looking consumption decisions. Given that income will fluctuate over a person's lifetime, the starting point of the PIH is their desire to smooth out fluctuations in consumption and their ability to save and borrow in order to achieve this.

The next question that arises is whether an individual prefers their smooth consumption path to be one of constant consumption in each period or of rising or falling consumption. This will depend on the relationship between the interest rate on saving and borrowing and the rate at which the individual trades off consumption in the future for consumption in the present. The latter is the subjective discount rate ρ, called rho which is a measure of an individual's impatience.

The household chooses a path of consumption to maximize its lifetime utility subject to its lifetime budget constraint, which is

$$\Psi_t^E = (1+r)A_{t-1} + \sum_{i=0}^{\infty} \frac{1}{(1+r)^i} y_{t+i}^E. \qquad \text{(lifetime budget constraint)}$$

Note that when considering the PIH, income, y is defined as post-tax income. The key to solving this problem can be thought about over just two periods, since the same considerations will apply to every subsequent pair of periods.[7] The answer is a simple relationship between consumption this period and next period, which is called an Euler equation (the derivation is in Section 1.4.2 of the Appendix):

$$C_t = \frac{1+\rho}{1+r} C_{t+1}^E. \qquad \text{(Euler equation)}$$

The intuition is clearest if we take the case where the interest rate, r, and the subjective discount rate, ρ, are the same. This means the household gets the same (objective) return from saving, r, as their (subjective) willingness to trade off consumption in the present for consumption in the future, ρ. In this case, the agent prefers a constant level of consumption each period, $C_t = C_{t+1}^E$.

We can see that if the discount rate is above the interest rate, $C_t > C_{t+1}^E$ then consumption is falling over time, reflecting the 'impatience' of the household. The opposite is the case if the discount rate is below the interest rate; a patient household will have a path in which consumption rises over time, i.e. $C_t < C_{t+1}^E$. The crucial point to note is that whichever of these consumption patterns is chosen, it is independent of the period by period changes in income.

From now on we focus on the PIH where the subjective discount rate is equal to the real rate of interest, i.e. ($\rho = r$). To implement a life-time consumption plan like this, given each period's income the household must do whatever saving and borrowing is called for in order to deliver the constant level of consumption each period. If current income is above permanent income, individuals save and if it is below, they borrow.

The next step is to use the Euler equation to find out how much individuals consume each period (the derivation is in Section 1.4.2 of the Appendix). When $\rho = r$, consumption is

$$C_t = \frac{r}{1+r} \Psi_t^E. \qquad \text{(PIH consumption function)}$$

The intuition is easy to understand: an individual with this consumption function will borrow and save to deliver a perfectly smooth consumption path (in expectation). The amount they consume each period is equal to the annuity value of expected lifetime wealth and is called 'permanent income'. The individual consumes their permanent income and the formula ensures that in expectation, they will be able to do this forever. Note that consumption remains constant unless Ψ_t^E changes.

[7] See Hall (1978) for a full derivation of the PIH and an empirical test of the main implications of the theory.

Predictions and empirical evidence on the PIH

The predictions of the PIH can be tested. How does consumption react when a change in income occurs? The answer depends on the nature of the change in income.

1. *Anticipated* or foreseen changes in income should have no effect on consumption when they occur. The reason is that anticipated changes will already have been incorporated into consumption through the recalculation of permanent income. Hence, when the change in current income is recorded, the marginal propensity to consume is predicted to be zero (the multiplier is equal to one). A finding of *'excess sensitivity'* to anticipated income changes would contradict the full smoothing behaviour predicted by the PIH.
2. *News or unanticipated* changes in income should affect consumption because they require the recalculation of future lifetime wealth, Ψ_t^E.
   (a) News of a temporary increase in income. If current income y_t increases unexpectedly by one unit, consumption increases by the extent to which this raises permanent income. Since the increase in one unit will be spread over the entire future, the PIH consumption function tells us that permanent income and hence consumption rise very little, just by $\frac{r}{1+r}$ times the increase in lifetime wealth. The marginal propensity to consume out of temporary income is $\frac{r}{1+r}$. (For example, if the real interest rate is 4%, the MPC is 3.8%, which implies a multiplier barely above one.)
   (b) News of a permanent increase in income. If there is news that income y_t is higher from now and for every future period by one unit, then permanent income and hence consumption rise by the full one unit. This means the marginal propensity to consume out of post-tax permanent income is one. A finding of *'excess smoothness'* of consumption in response to news of permanent income changes would contradict the simple PIH.

Excess sensitivity to anticipated *changes in income*

The first testable hypothesis suggests that there should be no change in consumption at the time income changes, if the change in income was known in advance. This is because consumption should already have adjusted as soon as the news arrived of the change in future income. An influential study by Campbell and Mankiw (1989) tested this hypothesis econometrically using aggregate data on consumption and income from the G7 countries. The study rejected a model in which all consumers were following the PIH but could not reject a model in which half of all consumers were simply following the 'rule of thumb' of spending their current income. In other words, they found that current consumption was overly sensitive to expected changes in income across the G7.

A recent study used the 2001 federal income tax rebates in the US as a testing ground: according to the PIH, the one-off (temporary) tax rebate should have very little effect on spending and if there is any effect, it should occur when the announcement was made and not when the cheques arrived. Johnson, Parker and Souleles (2006) were able to identify the causal effect of the tax rebate by using household data and the fact that the timing of the sending of cheques was based on the taxpayer's social security number, which is random. They found that the average household spent 20–40% of the (predictable) tax rebate in the three month period when the cheque was received rather than when the programme was announced. They also found that it was households with low income and wealth that responded most, underlining the likely role of limits on the ability of households to borrow, i.e. credit constraints. This is referred to as the *excess sensitivity* of consumption and is evidence against the strong predictions of the PIH.[8]

Is consumption excessively sensitive to the predictable fall in income at retirement? What is observed is that when income falls in a predictable way on retirement, consumption falls. Although this is a violation of the simple PIH, explanations consistent with consumption smoothing in a more detailed model have been proposed. These include the possibility that consumption falls because spending was related to being at work (i.e. complementary to working, such as the purchase of business suits) and that in the light of the increased leisure of the household, home production substitutes for consumption spending (e.g. cooking replaces the purchase of more expensive ready-made meals). Jappelli and Pistaferri (2010) conclude that evidence on the so-called consumption–retirement puzzle is not clearly in contradiction with the spirit of the PIH. Putting this together with the above discussion, the fact that there is excess sensitivity to anticipated rises in income but not so much to falls in income underlines the important role of credit constraints, which are ignored in the PIH.

Excess sensitivity to news *about temporary income and excess smoothness to* news *about permanent income*

Turning to the hypotheses that relate to news about income, studies have found that consumption over-responds to temporary income shocks. In the case of positive income shocks, this strongly violates the PIH. An early example of this was the large consumption response of US veterans after the Second World War to an unexpected windfall payout of the National Service Life Insurance. This kind of evidence questions the assumption in the PIH that households make optimizing decisions over a time horizon stretching a long way into the future and use a low discount rate (equal to the interest rate) to compare consumption at near and far horizons. The fact that people respond to a windfall by raising spending suggests that discount rates are higher than assumed in the simple PIH: people appear to be more *impatient* than the hypothesis assumes.

It is also likely that uncertainty about whether observed income changes are temporary or permanent prevents households from acting exactly as PIH would predict. For example, if a household mistakenly thought a temporary change in their current income was permanent, then they would consume more of the income change than would be consistent with PIH behaviour.

Jappelli and Pistaferri (2010) review the recent evidence and suggest that it shows that consistent with the PIH, consumption responds more to permanent than transitory income shocks, but that households do not revise their consumption fully into line with permanent shocks.

[8] For a detailed summary of the early literature on excess sensitivity see Deaton (1992). A recent excellent overview of the empirical testing of the excess sensitivity and excess smoothness hypotheses is Jappelli and Pistaferri (2010).

Credit constraints, impatience and uncertainty

The evidence suggests there are three reasons for the failure of the simple PIH to provide an adequate model of aggregate consumption:

1. The presence of credit constraints, which prevent borrowing by households who lack wealth or collateral.
2. Impatience, which prevents some households from saving as would be indicated by a permanent income view.
3. Uncertainty about future income, which explains precautionary saving over and above the level predicted by the PIH.

Credit constraints

If people prefer to smooth their consumption, but are prevented from doing so because they cannot borrow to bring forward consumption when their current income is below their expected permanent income, they are said to face credit constraints. The prevalence of credit constraints is an important factor that stops the world from operating as the simple PIH theory would imply. Because of information problems facing banks in assessing creditworthiness, banks are not always willing to lend to households without the wealth or collateral to secure a loan in order that they can smooth their consumption (see Chapter 5). Figure 1.8 shows the consumption response for PIH and credit-constrained households to an anticipated rise in income. The PIH households borrow to raise consumption as soon as the news of the future income increase is received, whereas the credit-constrained households do not have

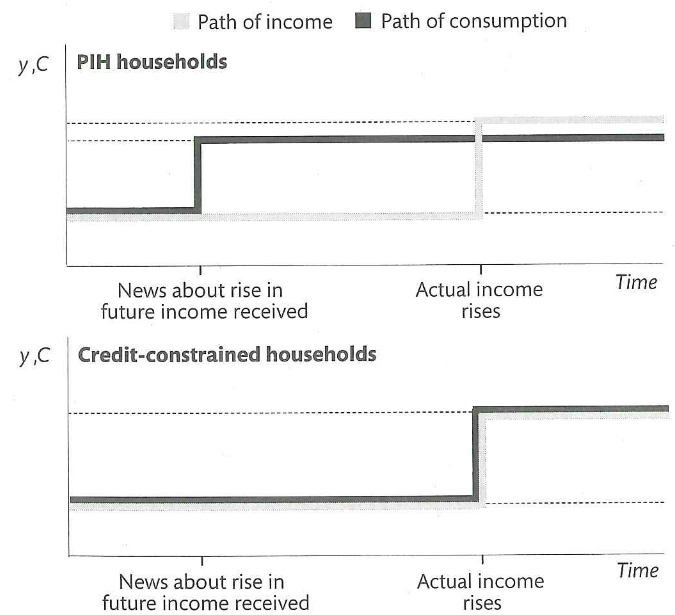

Figure 1.8 The consumption response of PIH and credit-constrained households to an anticipated rise in income.

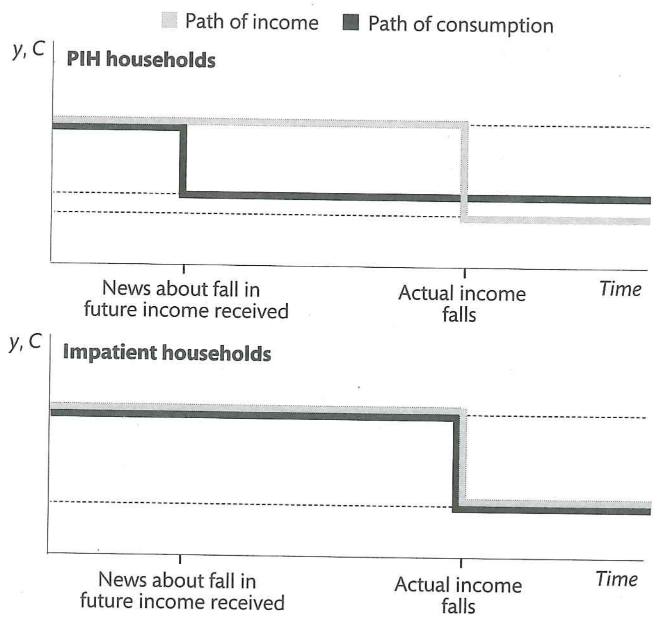

Figure 1.9 The consumption response of PIH and impatient households to an anticipated fall in income.

this option, so consumption can only rises when income does. Estimates from studies using household data have found that around 20% of the population in the United States is credit constrained.[9] A much higher estimate of closer to two-thirds comes from a study of the way people respond to automatic increases in the limits on their credit cards: if spending responds to such automatic increases, it suggests the card-owner was credit constrained.[10]

Impatience

The observation that consumption of a substantial fraction of households changes one-for-one with their current income (i.e. no borrowing or saving) reflects two features of the real world. The most important is the presence of credit constraints, which can prevent households from borrowing to smooth their consumption. This seems to especially affect young and low income households. Faced with a spell of unemployment, for example, constraints on borrowing will prevent a household from maintaining its level of consumption.

Whereas credit constraints can explain the inability of households to borrow to smooth consumption, it cannot explain the failure of some households to save, which would allow for some consumption smoothing without relying on bank loans. It seems that there are households who are not only credit constrained, but also impatient. Figure 1.9 shows the consumption response for PIH and impatient households to an anticipated fall in income. The PIH households start saving as soon as the news of the fall in future income is received, which alllows them to smooth consumption over the two periods. In contrast, the impa-

[9] See Japelli (1990), Mariger (1986) and Hall and Mishkin (1982).
[10] See Gross and Souleles (2002).

tient households fail to reduce current consumption upon receipt of the news, meaning consumption falls dramatically when income actually falls.

Experimental evidence on impatience

The impatience of households is represented in a model of consumption by a higher discount rate for the short run than the longer run. A simple illustration of what this means is to consider the following experiment: if you were choosing today for eating *next week*, would you choose fruit or chocolate? About three-quarters of experimental subjects choose fruit. When asked whether they would choose fruit or chocolate to eat *today*, 70% choose chocolate. From this, we can see that preferences are not consistent over time, because if we run time forward to next week, the three quarters who said they would prefer fruit would have shrunk to 30%. When next week comes, they would in fact, choose chocolate.[11]

Impatience is also documented in experiments related to consumption and saving behaviour. For example, subjects have a budget and can choose between two accounts paying the same interest rate: a 'freedom account' with no constraints on withdrawal or a 'commitment account' where there are restrictions on withdrawal before a date chosen by the participant. From a purely economic point of view, the freedom account is preferable because it imposes no constraints on withdrawal. In spite of this, more than half of investment was put in the commitment account (no withdrawal before the goal date), which reveals the desire of the experimental subjects to commit themselves to saving.[12]

Implications of impatience for consumption and saving

Very large differences between a high short-run and a low long-run discount rate help explain why people simultaneously save in illiquid assets such as housing (where the interest rate is low) and borrow on credit cards, paying a very high interest rate. Households with behaviour of this kind will have a marginal propensity to consume of one. Although a person can agree that their expected well-being is higher if they save more now, they are unable to resist consuming and blowing a hole in their saving plan when they get an unexpected increase in their income. One way of committing to saving is to take out a mortgage. Another is for the government to automatically enrol people in schemes for saving for retirement (from which they can opt out) such as pensions.

Uncertainty and precautionary saving

If there is uncertainty about future job opportunities or about health, then the household may save as insurance for future contingencies. For example, having a 'buffer stock of savings' could be important to sustain consumption in the event of job loss for a credit-constrained household.[13]

In the face of uncertainty, households would tend to save more early in life than the PIH predicts. Instead of the average propensity to consume falling as income rises over the working life as in the simple PIH, the desire for precautionary savings leads to saving early on with result that consumption rises with income (not by more than income) early in the

[11] See Read and van Leeuwen (1998).
[12] See Ashraf et al. (2006).
[13] See Carroll (1997) for an exposition of the buffer-stock savings model.

life cycle and the average propensity to consume rises later on. If saving for a rainy day is important, then the utility provided by having assets to tide the household over if need be outweighs the utility that would have come from higher consumption in the early years.

An empirical consumption function

To conclude the discussion of consumption, we consider a consumption function estimated on macroeconomic data for the UK, US and Japan.[14] Although the empirical model does not impose the straitjacket of the simple PIH, it nevertheless allows for permanent income considerations to affect consumption; households will have a notion of sustainable consumption constrained by their budget, which includes their income forecasts. They will therefore know that they cannot spend their existing assets now without impairing future consumption and that future income affects their sustainable consumption.

In addition to permanent income considerations based on forecasts, the consumption function that is estimated includes the following factors.

1. Consumption is predicted to respond more strongly to current income than the PIH predicts due to the presence of credit constraints.

2. The effects of income uncertainty on the desire of households to have precautionary savings. Income uncertainty is measured by the change in the unemployment rate and an increase is predicted to reduce consumption.

3. The housing collateral effect through higher house prices is expected to interact with increased credit availability to increase consumption (in the presence of credit constraints). The presence of this effect will depend on the institutional arrangements for mortgage borrowing in the country in question. For example, the availability or otherwise of home equity loans is a crucial distinction. A home equity loan is a loan secured against the equity in an individual's home. Home equity is the difference between the current market value of the house and the remaining mortgage, hence it increases as house prices rise. The collateral channel works through credit constraints being relaxed, but if it is not possible to get a home equity loan, as in France, Germany, Japan and China, this channel does not operate. In fact, if home equity loans are unavailable and high down-payments are required to get a mortgage (as in Japan), then higher house prices can actually dampen consumer spending, as young people have to save more to accumulate the required down-payment.

4. The current real interest rate. There are many interest rate channels, some of which point to a negative and some to a positive effect of a higher real interest rate on consumption.

5. Increases in households' wealth to income ratio are expected to increase consumption.

The consumption function includes the housing and financial variables that are at the centre of the discussion in Chapters 5-7, where the banking system, financial cycles and financial crises are brought into the macro model.

[14] See Aron et al. (2012).

For the UK and US, the consumption functions estimated using macroeconomic data suggest that roughly 50% of household consumption is accounted for by permanent income considerations. Current income has an important positive effect on consumption and the real interest rate has a negative effect. The wealth to income ratio has a positive effect but different types of wealth have different effects. For example, cash in a bank or savings account is more spendable than pension wealth. Also, the evidence is that for a given degree of access to new credit, higher debt of \$100 reduces consumer spending more than higher pension or stock market wealth of \$100 increases spending.

For these two economies, where home equity loans are important, increased credit availability and house prices interact to raise consumption through the collateral effect. A house price boom in conditions where access to credit is easier boosts consumption. We return to look at the implications of the richer analysis of consumption for the *IS* curve in Section 1.2.7.

1.2.6 **Investment**

The decision a firm makes over whether to invest or not will depend on their expectations about future after-tax profits. The role of expectations helps explain the excess volatility of investment over the other components of GDP visible in Fig. 1.1. In developing the *IS* equation, we used a simple linear relation to model investment in Section 1.2, where the amount of investment depended positively on expected future profits (captured by the constant term, a_0) and negatively on the real interest rate (a measure of firms' borrowing costs). In this subsection we introduce a more sophisticated forward-looking modelling approach—the q theory of investment.

Tobin's *q* theory of investment

This theory was developed by Nobel Prize winner James Tobin.[15] Like the PIH, Tobin's q theory of investment is forward looking. Firms choose the amount of investment to undertake with a view to maximising the expected discounted profits over the lifetime of the project. The q theory amounts to comparing the benefits from investment in an increase in the capital stock with the costs of doing so: if the expected benefits exceed the costs, investment should take place.

Marginal and average q models

The theory is derived explicitly in the Web Appendix to this chapter, but it is useful to set out the equation for *marginal q* here, as it provides a framework for thinking about forward-looking investment decisions:

$$q = \frac{\text{MB of investment}}{\text{MC of investment}} = \frac{Pf_K}{\delta + r}. \qquad \text{(marginal } q \text{ model of investment)}$$

The model is intrinsically forward looking since the firm is considering the future benefits from investment spending now. Before looking at the expression on the right, we note that the optimal amount of investment occurs when $q = 1$. This is where the expected marginal

[15] See Tobin (1969) and Tobin and Brainard (1977).

benefits of investment are equal to the expected marginal cost. Thus, if $q > 1$, the marginal benefit of investment exceeds the marginal cost, so firms should invest to increase the capital stock until $q = 1$. If $q < 1$ the firms should reduce their capital stock.

It is important to note that the q theory of investment is a model in which firms are assumed to make their investment decisions (a flow) by pinning down the optimal capital stock (a stock). Small changes in the desired capital stock can therefore translate into large changes in the flow of investment. This is one reason for the excess sensitivity of investment over GDP (see Fig. 1.1).

The marginal benefits from investment are higher when the price of output is higher and when the investment is more productive. The equation shows that firms should undertake more investment if there is

1. an increase in the price of output, P,
2. an increase in the marginal productivity of capital, f_K, which indicates the increase in output the new capital equipment will produce,
3. a reduction in the rate of interest, r, or
4. a reduction in the rate of depreciation, δ. For example, a rise in the expected rate of depreciation (e.g. as a result of uncertainty about future legislation banning fuel-inefficient cars) would lead to a reduction in the level of current investment because it would reduce the expected benefits to the firm from additional investment in the auto industry.

Is q-theory a good way to represent real-world investment behaviour? Marginal q implies firms adjust their level of investment in each period to equate the marginal benefits and marginal costs of investment. This does not fit with the real-world data on investment. When economists looked at very detailed microeconomic data for very large numbers of plants, they found that, contrary to the idea of a smooth investment flow, investment was in fact very bunched.[16] For example, a growing manufacturing firm might build a new factory once every 10 years.

The share price as indicator of expected profits

Testing q theory directly is difficult because q itself is difficult to measure. This is because it depends on the marginal product of capital and a measure of technology, which cannot be observed. In the real world, firms do not know their production functions and do not think explicitly in terms of marginal products. To operationalize the theory, the market value of the firm as reflected in its stock market valuation is compared with the replacement cost of the capital stock. If the market value is higher then this signals that the firm should increase investment. On the other hand, if the market value is below the replacement cost, then the firm would not want to build a *new* factory because it could buy an existing one more cheaply.

We can define average Q as follows:

$$Q = \frac{\text{Market value of firm}}{\text{Replacement cost of capital}}. \qquad \text{(average } Q \text{ model of investment)}$$

16 See the survey chapter on 'Investment' by Caballero (1999).

This is often called *average* Q to distinguish it from marginal q. Whereas q is the ratio of the marginal benefit of a unit of investment to its marginal cost, Q depends on the total expected return from the firm's capital divided by the total cost. For publicly listed companies, the stock market provides a forward-looking measure of the market value of the firm, which can be used in the numerator of the Q equation: when the firm's market value rises relative to the replacement cost of the firm, as reflected in a rise in the price of its shares, the model suggests that investment should go up. A higher interest rate and depreciation rate will raise the replacement cost of capital.

The idea is that the market value incorporates information about how well the firm is expected to be able to implement the investment, whether new competitors will enter the market, whether there are new technological innovations likely to affect the firm's value, the state of the macro-economy and the labour market and the future path of the interest rate. Since investing in a firm is a wager on this uncertain future, investors continuously evaluate these factors and, under certain conditions, the share price and hence the market value of the firm will reflect all the information available.

The stock market is notoriously volatile. For example, in 2008, the UK's leading share index—the FTSE 100—fell 31% and GDP contracted by just 1.1%. The following year, the FTSE 100 rose 22%, whilst GDP continued to contract, but this time by 4.4%.[17] In times of extreme uncertainty, the share price may not reflect the fundamental value of a firm. Large movements in share prices heavily affect market capitalization, which is the measure used to proxy for the market value of the firm in the calculation of average Q. As the stock market is volatile and subject to bubbles, fads and herd behaviour, it may not be a good indicator of the firm's prospects.

However, a study using US micro data (over 1,000 firms from 1982 to 1999) shows that a Q model of investment can be successfully estimated when analysts' forecasts of a firm's profits are used instead of the stock market price to measure the numerator in the Q equation.[18]

Empirical evidence on investment

The empirical evidence on the q theory is that Q helps to explain investment but it is not the only influence. In particular, just as was the case for consumption, credit constraints play an important role in explaining investment spending.

Credit constraints—the role of cash flow

A testable prediction of q theory is that current cash flow should have no impact on investment. The reason is that forward-looking firms should take into account any credit constraints that they face: these should already be incorporated into the stock market valuation, Q. However, the role of cash flow in empirical studies of investment suggests that *capital market imperfections* are important. The importance of cash flow variables in estimated investment equations is reminiscent of the importance of current income in estimated

17 GDP data from IMF World Economic Outlook database, April 2012. Stock market data from Bloomberg News articles.

18 See Bond and Cummins (2001).

consumption functions. As we have seen, the presence of credit-constrained households causes excess sensitivity of consumption to predictable changes in income. Similarly, the presence of cash flow terms in estimated investment functions strongly suggests that firms are credit constrained. Although some firms are able to borrow as much or sell as much equity as they would like in order to finance their investment plans, for others, their investment will be limited by their internal funds.

This is sometimes referred to as the *excess sensitivity of investment to internal funds*. For firms like this, cash flow would be expected to be an important determinant of investment.[19] The reason why firms may be rationed in their access to bank loans is explored in Chapter 5.

A study using UK company data for the period 1975–86 found that company investment was significantly influenced by Q and by credit constraints. However, the impact of Q was found to be very small: a 10% rise in the stock market value of a company was associated with an immediate rise in the investment rate of only 2.5%. Cash flow, on the other hand, was very important. The sample period was divided into two, and it was found that the impact of Q was lower and the impact of cash flow higher in the first part of the sample, during which the UK economy was in a deep recession. This is consistent with the idea of credit constraints biting especially hard in a recession.

Uncertainty—the option value of waiting

The decision of firms about whether to invest or not is also influenced by uncertainty about the future. Under certain circumstances, there can be a value in waiting to undertake an investment project, since with the passing of time more information arrives, the costs of delay (lost profits from the project) may be outweighed by more secure information on the balance between the costs (including those sunk in the project) and the benefits of undertaking it. The upshot of such considerations is that an expected rate of return considerably higher than the cost of capital will be required to trigger investment. Dixit (1992) gives examples to show that including the so-called 'option value of waiting' can double the hurdle rate (i.e. the return required to trigger investment) for an investment project to proceed.

1.2.7 **Consumption, investment and the *IS* curve**

In the previous subsection, we introduced models of consumption and investment that took account of forward-looking behaviour. Expectations of future income and of future profits play a key role in the spending decisions of households and firms. We also reported empirical evidence on aggregate consumption and investment, which highlighted the role of credit constraints for households and firms. What are the implications for the *IS* curve?

The basic form of the *IS* curve was presented in Section 1.2.3 and we shall see that we can bring together the additional insights about consumption and investment by focusing on the factors that affect the slope of the *IS* curve and which shift it.

[19] See, for example, Chirinko (1993).

[20] See Blundell et al. (1992).

The *IS* equation is:

$$\begin{aligned} y &= k(c_0 + a_0 + G) - ka_1 r \qquad &(1.13)\\ &= A - ar, \qquad &\text{(IS curve)} \end{aligned}$$

where $A \equiv k(c_0 + a_0 + G)$ and $a \equiv ka_1$; k is the multiplier. A larger multiplier makes the *IS* curve flatter, which increases the impact on output of a change in the interest rate; it also increases the impact on output of a shift in c_0, a_0 or G for a given interest rate. This makes it clear that the size of the multiplier is important for understanding the effectiveness of both fiscal and monetary policy.

As we shall see in Chapter 14 on Fiscal Policy, there has been a great deal of debate about the empirical size of the multiplier in the context of the use of fiscal stimulus to respond to the Great Recession following the global financial crisis. We shall see that that there is no simple answer to the question: what is the size of the multiplier? It depends on characteristics of the economy and on the context in which, for example, a fiscal stimulus is being applied.

We can summarize the implications of the analyses of consumption and investment from Sections 1.2.5 and 1.2.6 by thinking about how they affect the multiplier and the *IS* curve.

Factors affecting the multiplier

The PIH predicts that the size of the multiplier is dependent on the permanency of the income shock. The hypothesis suggests the marginal propensity to consume from temporary unanticipated changes in income is close to the size of the interest rate, i.e. only $\left(\frac{r}{1+r}\right)$. The multiplier is therefore predicted to be very close to one for temporary income shocks. On the other hand, permanent income shocks will result in a multiplier larger than one as households update their consumption to reflect their new, higher permanent income.

The presence of credit-constrained and impatient households also impacts the size of the multiplier. For these households, the multiplier effect will be larger than one even for temporary income shocks, as their marginal propensity to consume out of changes in income is equal to one. The larger is the proportion of these 'hand to mouth' or 'rule of thumb' households in the economy, the larger will be the multiplier.

Uncertainty over the permanency of income shocks is another factor that will influence the multiplier. If a substantial proportion of observed income fluctuations are perceived as permanent, then this will result in a multiplier larger than one. In this case, households could be following the PIH, but uncertainty leads their behaviour to be consistent with a Keynesian consumption function.

A change in the multiplier will shift the *IS* curve and change its slope (a higher multiplier makes the *IS* flatter).

Other factors affecting the slope of the IS *curve*

In addition to the effect of the size of the multiplier, the slope of the *IS* curve is affected by the interest sensitivity of investment and of consumption. The theoretical prediction of the impact of the interest rate on consumption is ambiguous; a reduction in interest rates boosts consumption through some channels and dampens it through others. The empirical evidence on the consumption function suggests that national institutional structures, particularly in relation to the financial system, are important for determining the strength and direction of the relationship. Aron et al. (2012) find a negative relationship between consumption and the real interest rate in the financially liberalized UK and US economies, but they find a positive relationship in Japan, where households' huge liquid deposits far outweigh household debt.

Other factors shifting the IS *curve*

In addition to the effect of the size of the multiplier, the *IS* curve is shifted by a number of other factors. Looking first at consumption, the PIH predicts that anything that changes expected lifetime wealth, Ψ_t^E, such as changes in asset prices or news about a future promotion, shifts the *IS*. The empirical findings about consumption highlight three other factors that can shift the *IS* curve:

1. The role of uncertainty: a rise in the rate of unemployment will raise savings for precautionary purposes, shifting the *IS* to the left.
2. A house price boom will boost consumption in a country with home equity loans by loosening credit constraints, shifting the *IS* to the right. However, a house price boom could also shift the *IS* to the left in countries where home equity loans are unobtainable and large down-payments are needed to get a mortgage.
3. A shift in credit market architecture that increases household access to credit, such as financial innovation or deregulation, will shift the *IS* to the right (at least until the accumulation of debt eventualy cancels out some of the shift).

Turning our attention to investment, Tobin's marginal q predicts that the following factors will shift the *IS* curve to the right:

1. an increase in the price of output, P,
2. an increase in the marginal productivity of capital, f_K; and
3. a reduction in the rate of depreciation, δ.

Lastly, the average Q equation highlights the role of expectations of future profits as a shift factor for the *IS* curve: a rise in the stock market tends to boost fixed investment as it signals a rise in the value of companies relative to their replacement cost.

1.3 **Conclusions**

This chapter has provided the first building block in the model of the macro-economy. We have set out the *IS* curve, which is used to model the demand side of the economy and is one of the three equations that underpin the 3-equation model that will form the core model of this textbook.

The *IS* curve allows us to think systematically about how changes in the spending behaviour of firms, households and governments can influence output and drive business cycles. The *IS* curve shows the combinations of the real interest rate and output where the goods market is in equilibrium. It slopes downwards to account for the fact that households' consumption

and new housing decisions respond negatively to the interest rate; firms will also undertake fewer investment projects as the cost of borrowing increases.

We return to the questions at the start of the modelling section (1.2) and see how the model we have developed helps us work towards an answer. This acts as an exercise to summarize what we have learnt in this chapter.

1. If you give an individual a \$100 bonus, how much of it will they spend? This depends on the model of household spending we are using. If consumer behaviour is well modelled by a Keynesian consumption function (i.e. $C_t = c_0 + c_1(1-t)y_t$) then households will spend a fixed proportion of any extra income they receive. In the case of a \$100 bonus, they will spend everything left over once they have paid tax and put some money into savings. However, if consumer behaviour is better modelled by the permanent income hypothesis (PIH), then individuals' current consumption is a function of their expected lifetime wealth (i.e. $C_t = \left(\frac{r}{1+r}\right)\Psi_t^E$). As the bonus is a one-off, it will only increase expected lifetime wealth by a small amount, resulting in a small increase in consumption in this and all future periods. This leads the individual to save the majority of the bonus in the current period. In Section 1.2.7, we showed how the proportions of rule of thumb and forward-looking households in the economy affect the multiplier and the *IS* curve. We also discussed how the extent to which observed changes in income are perceived as permanent can help to reconcile the Keynesian and PIH consumption functions.

2. How much will output increase following a rise in autonomous demand, such as higher government spending or a boost in the confidence of households and firms? The multiplier will determine the extent to which an €x increase in autonomous demand will increase output. In the simple model presented in Section 1.2.1, the multiplier is always greater than one, as both c_1 and t are between zero and one. This means that in the short run, an injection of autonomous demand will always boost output by more than one for one. The extent of the boost of output will depend on the size of the multiplier. The multiplier will be larger if the tax rate (t) is low and households have a high marginal propensity to consume (c_1). In Section 1.2.7, we saw that if all households are described by the PIH, the permanency of the income change is important for determining the size of the multiplier. When income changes are permanent, the multiplier is greater than one as households update their consumption to reflect the increase in their permanent income. When the income change is temporary, the marginal propensity to consume out of current income is essentially zero as it has a negligible effect on permanent income, and the multiplier is therefore equal to one. Temporary and unanticipated income shocks can, however, result in a multiplier greater than one if some proportion of households are credit constrained or consume according to a rule of thumb. We discuss the multiplier further, with a particular focus on government spending, in Chapters 4, 7 and 14.

3. To what extent will a rise in the interest rate curtail investment? In the simple model in Section 1.2.3, where $I = a_0 - a_1 r$, a rise in the interest rate of Δr will reduce investment by $a_1\Delta r$. In the marginal q theory of investment (with $q = \frac{PBf_K}{\delta+r}$), an increase in the interest rate will increase the marginal cost of investment. If we assume that q is initially equal to one, (i.e. the marginal cost and marginal benefit of carrying out the investment project are equal) then an increase in r will lead firms to reduce investment. The extent to which investment

will fall depends on how quickly the marginal productivity of capital (f_K) rises as investment falls to bring q back toward a value of 1.

This chapter has provided a model of the demand side which can shed light on macro-economic questions. We have discussed how the models of consumption and investment relate to characteristics of the real world. For example, the greater volatility of investment as compared with consumption observed in the data can be partly explained by the factors that influence investment spending decisions and by consumers borrowing and saving to smooth their consumption over the economic cycle. Households vary—some are impatient and find it difficult to save, whereas others are prudent and save for precautionary reasons. The government contributes to consumption smoothing through provision of unemployment benefits. Credit constraints facing both households and firms help better align the models of consumption and investment with the empirical data.

Although the demand side is a major influence on economic activity, it is only part of the story of how the macro-economy works. To develop our understanding of economic fluctuations and longer-run trends, we need to introduce the supply side. This will provide a framework for thinking about how wages and prices are set and what determines the unemployment rate. The supply side is the subject of the next chapter.

1.4 Appendix

1.4.1 Real and nominal interest rates and the Fisher equation

To clarify why it is the real rather than the nominal interest rate that affects real expenditure decisions in the economy, think about a firm considering an investment project. A higher money or nominal rate of interest will not impose a greater real burden on the firm if it is balanced by correspondingly higher inflation, because the expected profits from the investment project will be higher in money terms and the balance between the real cost and the real return on the project will not have changed.

The real interest rate is defined in terms of *goods* and the nominal interest rate, in terms of *money*. Thinking of a consumer good, the real rate of interest, r, is how much extra in terms of units of this good - namely (1+r) units - would have to be paid in the future in order to borrow one unit of the goods today. The nominal rate of interest is how much extra in euros would have to be paid in the future in order to borrow one euro today. If goods prices remain constant then it is clear that the real and nominal interest rates are the same: if you lent one euro today, you would be able to buy $(1+r)$ goods in the future. In general,

$$1+r=(1+i)\cdot\frac{P}{P^E_{t+1}},$$

where it is the expected price level in the future (P^E_{t+1}) that comes into play since at time t, we do not know what the price level will be at $t+1$. If we use the following definition of expected inflation:

$$\pi^E=\frac{P^E_{t+1}-P}{P},$$

then

$$\frac{P}{P^E_{t+1}}=\frac{1}{1+\pi^E}.$$

By rearranging the above expression, it follows that

$$(1+r)=\frac{(1+i)}{(1+\pi^E)},$$

and therefore that

$$r=\frac{i-\pi^E}{1+\pi^E}.$$

When expected inflation is low, the denominator of this expression is close to one and we have the standard approximation for the relationship between the real and the nominal rate of interest:

$$i\approx r+\pi^E. \qquad \text{(Fisher equation)}$$

Inflation expectations will drive the divergence between the real and nominal interest rates. It should be noted that only one of these three terms is observable: the nominal interest rate, i. The real interest rate can be estimated from historical data on the nominal interest rate and the rate of inflation: this gives a measure of the so-called *ex post* real rate of interest. Alternatively, an *ex ante* measure can be derived from a model that is able to predict inflation. Finally, if bonds have been issued in the economy that are protected against inflation because the face value is indexed by the rate of inflation, then the yield on such a bond is a real rate of interest and can provide a third measure. But there are only a few countries that have issued index-linked or inflation-proof bonds (UK in 1981, the USA in 1997, France in 1998).

1.4.2 Deriving the Euler equation and the PIH consumption function

Deriving the Euler equation

The first step in deriving the Euler equation is to set out an expression for the present value of a consumer's utility:

$$V^E_t=U(C_t)+\frac{U(C^E_{t+1})}{1+\rho}+\frac{U(C^E_{t+2})}{(1+\rho)^2}+\ldots \qquad \text{(expected present value of consumption)}$$

The present value equation uses the expectations operator (i.e. C^E_{t+i} for all $i>0$) to show that future consumption is uncertain in the present period. In addition, the utility of consumption is discounted to the present period by the consumer's rate of time preference. Future consumption will be worth less in the current period to more impatient consumers.

To derive the Euler equation, we need to assume a specific functional form for the utility function. We choose $U(C_t)=\log C_t$ because it exhibits diminishing marginal returns and is easy to work with. We can take the first derivative of the utility function to find the marginal utility of consumption, which is $U'(C_t)=\frac{1}{C_t}$.

The Euler equation shows the optimal C_t in relation to C^E_{t+1}. If we just think about these two periods, then it is easy to derive the Euler equation. In period t the consumer must weigh up the gain from consuming more in this period, against the discounted loss of consuming less in the next period. The gain from consuming one unit more in this period is simply the marginal utility of consumption $\frac{1}{C_t}$. The subjective present value loss of utility in the next period is equal to $\frac{1}{C^E_{t+1}}\frac{1}{1+\rho}(1+r)$. This is because the consumer's income will have fallen by (1+r) next period, multiplied by the marginal utility of consumption next period multiplied by the impatience factor.

The consumer cannot gain any more utility when the gain from consuming one more unit in the current period is exactly the same as the subjective present value loss of consuming one unit less in the next period. This gives us the Euler equation:

$$\frac{1}{C_t} = \frac{1}{C^E_{t+1}}\frac{1}{1+\rho}(1+r)$$
$$C_t = \frac{1+\rho}{1+r}C^E_{t+1}. \qquad \text{(Euler equation)}$$

Deriving the permanent income hypothesis consumption function

The simplest case of the Euler equation is when the rate of interest is equal to the consumer's subjective discount rate $r = \rho$. This is the case we have focused on throughout this chapter. We can see from the Euler equation that this implies $C_t = C^E_{t+1}$. In other words, it implies that consumption is constant in expectation in all future periods.

In order to derive the permanent income hypothesis, we make the additional assumption that future consumption is known with certainty, so that $C_t = C_{t+1} = \cdots = C_{t+i} = C$. To calculate C, we need to set out the expression for the present value of consumption, $\overline{C_t}$:

$$\overline{C_t} = C_t + \frac{C_{t+1}}{1+r} + \ldots + \frac{C_{t+i}}{(1+r)^i} + \ldots = C\left(1 + \frac{1}{1+r} + \ldots + \frac{1}{(1+r)^i} + \ldots\right)$$
$$= C\left(\frac{1}{1-\frac{1}{1+r}}\right) = C\left(\frac{1+r}{r}\right). \qquad \text{(present value of consumption)}$$

The permanent income hypothesis implies that the present value of consumption has to be equal to the present value of income, Ψ^E_t, so that:

$$\overline{C_t} = C\left(\frac{1+r}{r}\right) = \Psi^E_t$$
$$C_t = C = \left(\frac{r}{1+r}\right)\Psi^E_t, \qquad \text{(PIH consumption function)}$$

where C_t is the amount the consumer consumes in each period. It is their permanent income and is equal to the annuity value of their expected lifetime wealth.

1.5 Questions

1.5.1 Checklist questions

1. What is the *IS* curve? Why does it slope downwards?
2. Use the equation for the *IS* curve shown in Section 1.2.3 to answer the following questions:
   (a) If we assume that $0 < c_1, t < 1$, then what can we say about the size of multiplier?
   (b) If there is a decrease in government spending of ΔG, by how much does this decrease output?
   (c) Describe the feedback process that means a decrease in government spending can decrease output by more than 1 : 1.
3. Use the Keynesian cross to show the effect of a decrease in autonomous investment on the economy. Discuss the path of the economy as it adjusts to the new medium-run equilibrium. Why does the economy not continue to contract?
4. Use the Keynesian cross to illustrate the paradox of thrift. Model the change in savings behaviour as an increase in the marginal propensity to save, s_1 (remember that $c_1 + s_1 = 1$). Show how a rise in investment can counteract the reduction in output associated with the rise in savings.
5. Use the equation for the *IS* curve shown in Section 1.2.3 and Fig. 1.6 to discuss what happens to the *IS* curve in response to the following shocks. In each case provide a real world example of what might cause the shock.
   (a) An increase in autonomous consumption (i.e. $\uparrow c_0$).
   (b) A reduction in the interest sensitivity of investment (i.e. $\downarrow a_1$).
   (c) An increase in the marginal rate of taxation (i.e. $\uparrow t$).
6. According to the permanent income hypothesis, how will the paths of borrowing and consumption change in response to:
   (a) A temporary decrease in income when it occurs.
   (b) A permanent decrease in income when it occurs.
   (c) Are the answers different if the changes in income are unanticipated, i.e. if they are 'news'? Comment on the size of the marginal propensity to consume and the size of the multiplier.
7. Assuming the real interest rate is 4%, calculate how, according to the PIH, consumption and borrowing would change in each of the following cases
   (a) A stock market crash permanently reduces the value of an individual's assets by 1,000.
   (b) Households are told that in a year's time, they will receive a one-off bonus of 1,000. Then in one year's time, it is not paid.
   (c) Comment briefly on your results.
8. Explain the concepts of excess sensitivity and excess smoothness that arise from the empirical literature on the permanent income hypothesis. What could explain these findings?
9. What does Tobin's q tell us firms' investment decisions depend upon? According to Tobin's q, when should a firm invest?

10. What is the key problem with measuring marginal q? Is there an alternative measurement that can be used instead? Is this alternative measurement likely to be an accurate proxy for marginal q?
11. Use Section 1.2.7 to discuss what is expected to happen to the IS curve in response to the following shocks:
    (a) A crash in the stock market.
    (b) An increase in the retirement age.
    (c) A decrease in the rate of depreciation.
    (d) An increase in the cost of oil.
    (e) An increase in the rate of technological progress.

1.5.2 Problems and questions for discussion

1. This question involves collecting data from national statistics agencies (e.g. the UK Office for National Statistics) and/or international organizations (e.g. the OECD or the IMF). First, select an emerging and a developed economy.
    (a) Collect annual data on real GDP as far back as it is available for both countries. Convert the data to the log scale and plot on a graph. Comment on your findings. Calculate the growth rates of GDP for the two countries over the period and plot them on a graph. How do the business cycles of the two countries compare? To what extent do they appear synchronized?
    (b) Collect current price data on GDP and its components for the two countries. Calculate the percentage of GDP according to each of the four main types of expenditure—i.e. household consumption, government consumption, gross fixed capital formation and net exports (i.e. exports minus imports) and plot the series. How does the composition of GDP and changes over time in the two countries compare? Discuss possible reasons for the differences.
2. We start this question with a simple version of the aggregate demand function, where, to keep the maths simple, we omit taxation and government spending:

$$y^D = c_0 + c_1 y + \bar{I}. \tag{1.14}$$

Now, assume that $c_0 = 200$ and $\bar{I} = 200$. In addition, assume that $c_1 = 0.8$, such that:

$$y^D = 200 + 0.8y + 200. \tag{1.15}$$

(a) What is the level of output in goods market equilibrium?
(b) Assume there is a fall in the marginal propensity to consume, c_1, to 0.6—i.e. there is a rise in the savings rate, as $s_1 + c_1 = 1$. If we assume that the rise in planned savings leads to the accumulation of unplanned inventories of goods and does not increase planned investment ($\bar{I}$), then what is the level of output that satisfies the new goods market equilibrium?
(c) Compare the level of savings in the new and old goods market equilibria. Note that in goods market equilibrium; $S = s_1 y - c_0$ and $I = \bar{I}$.
(d) Comment on your findings in (c) in response to the question of whether it is advisable for the policy maker to encourage households to increase their savings to help escape a recession. Describe in words a mechanism, which is not included in this model, that could provide a connection from a policy encouraging households to save more and exit from a recession.

3. Use Section 1.2.4 and appropriate readings from beyond this book to decide whether the following statements (S1 and S2) are both true or whether one of them is false. Justify your answer.

    S1: According to Tobin's q theory, the path of investment is independent of current cash flow (and profits).

    S2: The empirical evidence shows that current cash flow is an important determinant of investment.

4. Aggregate consumption varies less than GDP and aggregate investment varies more. Can you reconcile these observations with the assumption that consumption and investment decisions are taken by rational, forward-looking agents?