This rule of thumb makes the simple models all the more valuable. Although they do not describe perfectly the world in which we live, they do provide a useful guide to the effects of economic policy.

MORE PROBLEMS AND APPLICATIONS

1. Imagine that you run the central bank in a large open economy with a floating exchange rate. Your goal is to stabilize income, and you adjust the money supply accordingly. Under your policy, what happens to the money supply, the interest rate, the exchange rate, and the trade balance in response to each of the following shocks?

a. Taxes are raised to reduce the budget deficit.

b. The import of foreign cars is restricted.

2. Since the 1980s, the economies of the world have become more financially integrated. That is, investors in all nations have become more willing and able to take advantage of financial opportunities abroad. Consider how this development affects the ability of monetary policy to influence the economy.

a. If investors become more willing and able to substitute foreign and domestic assets, what happens to the slope of the *CF* function?

b. If the *CF* function changes in this way, what happens to the slope of the *IS* curve?

c. How does this change in the *IS* curve affect the Fed's ability to control the interest rate?

d. How does this change in the *IS* curve affect the Fed's ability to control national income?

3. Suppose that policy makers in a large open economy want to raise the level of investment without changing aggregate income or the exchange rate.

a. Is there any combination of domestic monetary and fiscal policies that would achieve this goal?

b. Is there any combination of domestic monetary, fiscal, and trade policies that would achieve this goal?

c. Is there any combination of monetary and fiscal policies at home and abroad that would achieve this goal?

4. The appendix considers the case of a large open economy with a floating exchange, but suppose instead that a large open economy has a fixed exchange rate. That is, the central bank announces a target for the exchange rate and commits itself to adjusting the money supply to ensure that the equilibrium exchange rate equals the target.

a. Describe what happens to income, the interest rate, and the trade balance in response to a fiscal expansion, such as an increase in government purchases. Compare your answer with the case of a small open economy with a fixed exchange rate.

b. Describe what happens to income, the interest rate, and the trade balance if the central bank expands the money supply by buying bonds from the public. Compare your answer with the case of a small open economy with a fixed exchange rate.

CHAPTER 14

Aggregate Supply and the Short-Run Trade-Off between Inflation and Unemployment

Probably the single most important macroeconomic relationship is the Phillips Curve.

– George Akerlof

There is always a temporary trade-off between inflation and unemployment; there is no permanent trade-off. The temporary trade-off comes not from inflation per se, but from unanticipated inflation, which generally means, from a rising rate of inflation.

– Milton Friedman

Most economists analyse short-run fluctuations in aggregate income and the price level using the model of aggregate demand and aggregate supply. In the previous three chapters, we examined aggregate demand in some detail. The *IS-LM* model – together with its open-economy cousin, the Mundell–Fleming model – shows how changes in monetary and fiscal policy and shocks to the money and goods markets shift the aggregate demand curve. In this chapter, we turn our attention to aggregate supply and develop theories that explain the position and slope of the aggregate supply curve.

When we introduced the aggregate supply curve in Chapter 10, we established that aggregate supply behaves differently in the short run to how it does in the long run. In the long run, prices are flexible, and the aggregate supply curve is vertical. When the aggregate supply curve is vertical, shifts in the aggregate demand curve affect the price level, but the output of the economy remains at its natural level. By contrast, in the short run, prices are sticky and the aggregate supply curve is not vertical. In this case, shifts in aggregate demand do cause fluctuations in output. In Chapter 10 we took a simplified view of price stickiness by drawing the short-run aggregate supply curve as a horizontal line, representing the extreme situation in which all prices are fixed. Our task now is to refine this understanding of short-run aggregate supply to better reflect the real world in which some prices are sticky and others are not.

Unfortunately, one fact makes our task more difficult: economists disagree about how best to explain aggregate supply. As a result, this chapter begins by presenting three models of the short-run aggregate supply curve. Among

economists, each of these models has some prominent adherents (as well as some prominent critics), and you can decide for yourself which you find most plausible. Although these models differ in some significant details, they are related in an important way: they share a common theme about what makes the short-run and long-run aggregate supply curves differ, and a common conclusion that the short-run aggregate supply curve is upward sloping.

After examining the models, we examine an implication of the short-run aggregate supply curve. We show that this curve implies a trade-off between two measures of economic performance – inflation and unemployment. This trade-off, called the *Phillips curve,* tells us that to reduce the rate of inflation, policy makers must temporarily raise unemployment, and to reduce unemployment, they must accept higher inflation. As the quotation from Milton Friedman at the beginning of the chapter suggests, the trade-off between inflation and unemployment is only temporary. One goal of this chapter is to explain why policy makers face such a trade-off in the short run and, just as important, why they do not face it in the long run.

14-1 Three Models of Aggregate Supply

When classes in physics study balls rolling down inclined planes, they often begin by assuming away the existence of friction. This assumption makes the problem simpler and is useful in many circumstances, but no good engineer would ever take this assumption as a literal description of how the world works. Similarly, this book began with classical macroeconomic theory, but it would be a mistake to assume that this model is always true. Our job now is to look more deeply into the 'frictions' of macroeconomics.

We do this by examining three prominent models of aggregate supply. In all the models, some market imperfection (that is, some type of friction) causes the output of the economy to deviate from its natural level. As a result, the short-run aggregate supply curve is upward sloping, rather than vertical, and shifts in the aggregate demand curve cause output to fluctuate. These temporary deviations of output from its natural level represent the booms and busts of the business cycle.

Although each of the three models takes us down a different theoretical route, each route ends up in the same place. That final destination is a short-run aggregate supply equation of the form

$$Y = \overline{Y} + \alpha(P - P^e), \quad \alpha > 0,$$

where Y is output, $\overline{Y}$ is the natural level of output, P is the price level and P^e is the expected price level. This equation states that output deviates from its natural level when the price level deviates from the expected price level. The parameter α indicates how much output responds to unexpected changes in the price level; $1/\alpha$ is the slope of the aggregate supply curve.

Each of the three models tells a different story about what lies behind this short-run aggregate supply equation. In other words, each model highlights a particular reason why unexpected movements in the price level are associated with fluctuations in aggregate output.

The Sticky-Price Model

Our first explanation for the upward-sloping short-run aggregate supply curve is called the **sticky-price model**. This model emphasizes that firms do not instantly adjust the prices they charge in response to changes in demand. Sometimes prices are set by long-term contracts between firms and customers. Even without formal agreements, firms may hold prices steady in order not to annoy their regular customers with frequent price changes. Some prices are sticky because of the way markets are structured: once a firm has printed and distributed its catalogue or price list, it is costly to alter prices.

To see how sticky prices can help explain an upward-sloping aggregate supply curve, we first consider the pricing decisions of individual firms, and then add together the decisions of many firms to explain the behaviour of the economy as a whole. Notice that this model encourages us to depart from the assumption of perfect competition, which we have used since Chapter 3. Perfectly competitive firms are price takers rather than price setters. If we want to consider how firms set prices, it is natural to assume that these firms have at least some monopoly control over the prices they charge.

Consider the pricing decision facing a typical firm. The firm's desired price p depends on two macroeconomic variables:

- The overall level of prices P: a higher price level implies that the firm's costs are higher – hence, the higher the overall price level, the more the firm would like to charge for its product.
- The level of aggregate income Y: a higher level of income raises the demand for the firm's product – because marginal cost increases at higher levels of production, the greater the demand, the higher the firm's desired price.

$$p = P + a(Y - \overline{Y}).$$

This equation says that the desired price p depends on the overall level of prices P and on the level of aggregate output relative to the natural level $Y - \overline{Y}$. The parameter a (which is greater than zero) measures how much the firm's desired price responds to the level of aggregate output.[1]

Now assume that there are two types of firms. Some have flexible prices: they always set their prices according to this equation. Others have sticky prices: they announce their prices in advance, based on what they expect economic conditions to be. Firms with sticky prices set prices according to

$$p = P^e + a(Y^e - \overline{Y}^e),$$

where, as before, a superscript 'e' represents the expected value of a variable. For simplicity, assume that these firms expect output to be at its natural level, so that the last term, $a(Y^e - \overline{Y}^e)$, is zero. Then these firms set the price

$$p = P^e.$$

[1] *Mathematical note:* The firm cares most about its relative price, which is the ratio of its nominal price to the overall price level. If we interpret p and P as the logarithms of the firm's price and the price level, then this equation states that the desired relative price depends on the deviation of output from its natural level.

That is, firms with sticky prices set their prices based on what they expect other firms to charge.

We can use the pricing rules of the two groups of firms to derive the aggregate supply equation. To do this, we find the overall price level in the economy, which is the weighted average of the prices set by the two groups. If s is the fraction of firms with sticky prices, and $1 - s$ is the fraction with flexible prices, then the overall price level is

$$P = sP^e + (1 - s)[P + a(Y - \overline{Y})].$$

The first term is the price of the sticky-price firms weighted by their fraction in the economy, and the second term is the price of the flexible-price firms weighted by their fraction. Now subtract $(1 - s)P$ from both sides of this equation to obtain

$$sP = sP^e + (1 - s)[a(Y - \overline{Y})].$$

Divide both sides by s to solve for the overall price level:

$$P = P^e + [(1 - s)a/s](Y - \overline{Y}).$$

The two terms in this equation are explained as follows:

- When firms expect a high price level, they expect high costs. Those firms that fix prices in advance set their prices high. These high prices cause the other firms to set high prices also. Hence, a high expected price level P^e leads to a high actual price level P.
- When output is high, the demand for goods is high. Those firms with flexible prices set their prices high, which leads to a high price level. The effect of output on the price level depends on the proportion of firms with flexible prices.

Hence, the overall price level depends on the expected price level and on the level of output.

Algebraic rearrangement puts this aggregate pricing equation into a more familiar form:

$$Y = \overline{Y} + \alpha(P - P^e),$$

where $\alpha = s/[(1 - s)a]$. The sticky-price model says that the deviation of output from the natural level is positively associated with the deviation of the price level from the expected price level.[2]

The Sticky-Wage Model

To explain why the short-run aggregate supply curve is upward sloping, many economists stress the sluggish adjustment of nominal wages. In many industries, nominal wages are set by long-term contracts, so wages cannot adjust quickly

[2] For a more advanced development of the sticky-price model, see Julio Rotemberg, 'Monopolistic Price Adjustment and Aggregate Output', *Review of Economic Studies*, 1982, vol. 49, pp. 517–531.

when economic conditions change. Even in industries not covered by formal contracts, implicit agreements between workers and firms may limit wage changes. Wages may also depend on social norms and notions of fairness that evolve slowly. For these reasons, many economists believe that nominal wages are sticky in the short run.

The **sticky-wage model** shows what a sticky nominal wage implies for aggregate supply. To preview the model, consider what happens to the amount of output produced when the price level rises:

1. When the nominal wage is stuck, a rise in the price level lowers the real wage, making labour cheaper.
2. The lower real wage induces firms to hire more labour.
3. The additional labour hired produces more output.

This positive relationship between the price level and the amount of output means that the aggregate supply curve slopes upward during the time when the nominal wage cannot adjust to a change in the price level.

To develop this story of aggregate supply more formally, assume that workers and firms bargain over and agree on the nominal wage before they know what the price level will be when their agreement takes effect. The bargaining parties – the workers and the firms – have in mind a target real wage. The target may be the real wage that equilibrates labour supply and demand. More likely, the target real wage is higher than the equilibrium real wage: as discussed in Chapter 7, union power and efficiency-wage considerations tend to keep real wages above the level that brings labour supply and labour demand into balance.

The workers and firms set the nominal wage W based on the target real wage ω and on their expectation of the price level P^e. The nominal wage they set is

$$W = \omega \times P^e,$$

$$\text{Nominal Wage} = \text{Target Real Wage} \times \text{Expected Price Level}.$$

After the nominal wage has been set and before labour has been hired, firms learn the actual price level P. The real wage turns out to be

$$W/P = \omega \times (P^e/P),$$

$$\text{Real Wage} = \text{Target Real Wage} \times \frac{\text{Expected Price Level}}{\text{Actual Price Level}}.$$

This equation shows that the real wage deviates from its target if the actual price level differs from the expected price level. When the actual price level is greater than expected, the real wage is less than its target; when the actual price level is less than expected, the real wage is greater than its target.

The final assumption of the sticky-wage model is that employment is determined by the quantity of labour that firms demand. In other words, the bargain between the workers and the firms does not determine the level of employment in advance; instead, the workers agree to provide as much labour as the firms

wish to buy at the predetermined wage. We describe the firms' hiring decisions by the labour demand function

$$L = L^d(W/P),$$

which states that the lower the real wage, the more labour firms hire. The labour demand curve is shown in panel (a) of Figure 14-1. Output is determined by the production function

$$Y = F(L),$$

FIGURE 14-1

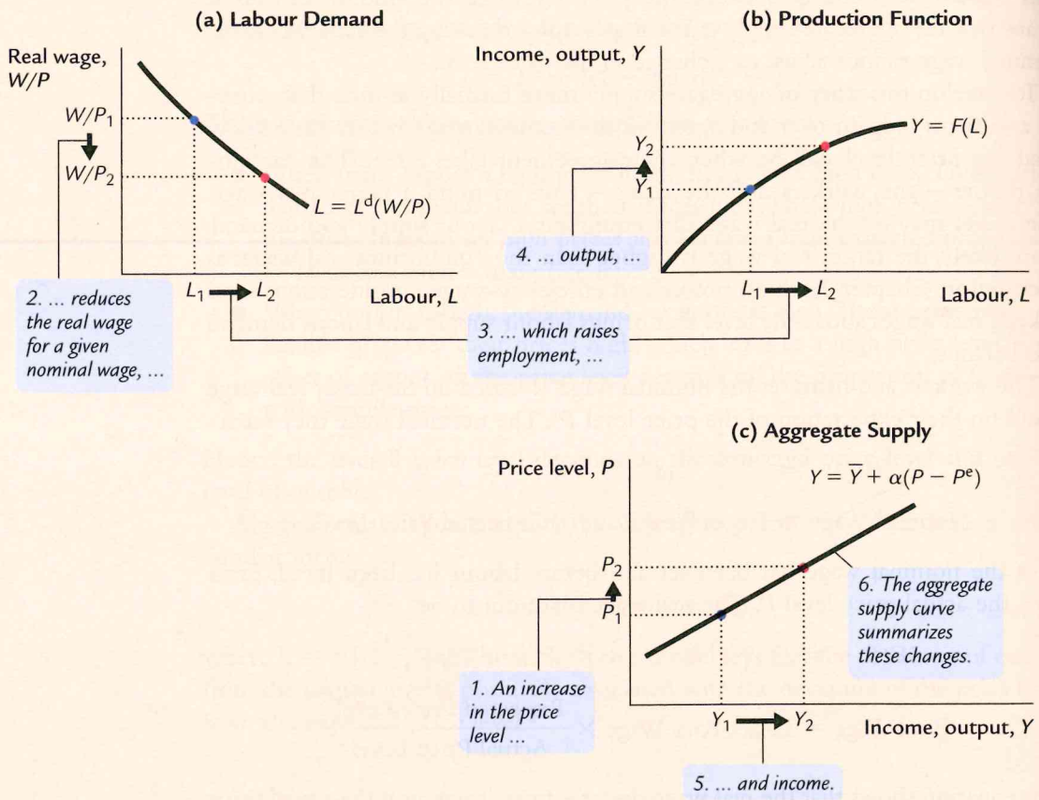

The Sticky-Wage Model Panel (a) shows the labour demand curve. Because the nominal wage W is stuck, an increase in the price level from P_1 to P_2 reduces the real wage from W/P_1 to W/P_2. The lower real wage raises the quantity of labour demanded from L_1 to L_2. Panel (b) shows the production function. An increase in the quantity of labour from L_1 to L_2 raises output from Y_1 to Y_2. Panel (c) shows the aggregate supply curve summarizing this relationship between the price level and output. An increase in the price level from P_1 to P_2 raises output from Y_1 to Y_2.

which states that the more labour is hired, the more output is produced. This is shown in panel (b) of Figure 14-1.

Panel (c) of Figure 14-1 shows the resulting aggregate supply curve. Because the nominal wage is sticky, an unexpected change in the price level moves the real wage away from the target real wage, and this change in the real wage influences the amounts of labour hired and output produced. The aggregate supply curve can be written as

$$Y = \overline{Y} + \alpha(P - P^e).$$

Output deviates from its natural level when the price level deviates from the expected price level.[3]

CASE STUDY

The Cyclical Behaviour of the Real Wage

In any model with an unchanging labour demand curve, such as the model we just discussed, employment rises when the real wage falls. In the sticky-wage model, an unexpected rise in the price level lowers the real wage and thereby raises the quantity of labour hired and the amount of output produced. Thus, the real wage should be *counter-cyclical:* it should fluctuate in the opposite direction from employment and output. Keynes himself wrote in *The General Theory* that 'an increase in employment can only occur to the accompaniment of a decline in the rate of real wages'.

The earliest attacks on *The General Theory* came from economists challenging Keynes's prediction. Figure 14-2 is a scatterplot of the percentage change in average real earnings and the percentage change in real GDP, using annual data for the UK economy from 1960 to 2011. If Keynes's prediction were correct, the dots in this figure would show a downward-sloping pattern, indicating a negative relationship. Yet the figure shows only a weak correlation between real earnings and output, and it is the opposite of what Keynes predicted. That is, if the real wage is cyclical at all, it is slightly *pro-cyclical:* the real wage tends to rise when output rises. Abnormally high labour costs cannot explain the low employment and output observed in recessions.

How should we interpret this evidence? Most economists conclude that the sticky-wage model cannot fully explain aggregate supply. They advocate models in which the labour demand curve shifts over the business cycle. These shifts may arise because firms have sticky prices: when prices are stuck too high, firms sell less of their output and reduce their demand for labour. Alternatively, the labour demand curve may shift because shocks to technology alter labour productivity. Economists have developed this idea, known as the theory of real

[3] For more on the sticky-wage model, see Jo Anna Gray, 'Wage Indexation: A Macroeconomic Approach', *Journal of Monetary Economics*, April 1976, vol. 2, pp. 221–235; and Stanley Fischer, 'Long-term Contracts, Rational Expectations, and the Optimal Money Supply Rule', *Journal of Political Economy*, February 1977, vol. 85, pp. 191–205.

FIGURE 14-2

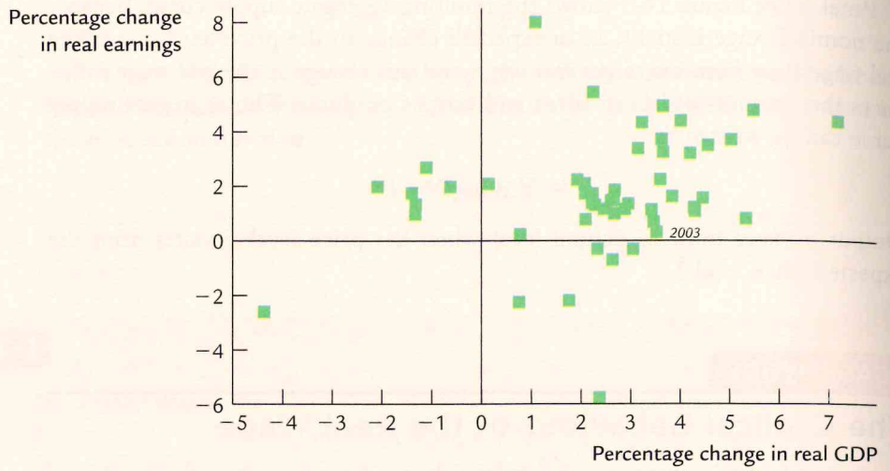

The Cyclical Behaviour of the Real Wage This scatterplot shows the percentage change in real UK GDP and the percentage change in the UK real wage (measured here as average real earnings). As output fluctuates, the real wage typically moves in the same direction. That is, the real wage is somewhat pro-cyclical. This observation is inconsistent with the sticky-wage model.

Source: Lawrence H. Officer and Samuel H. Williamson, 'What Was the UK GDP Then?', MeasuringWorth, 2013 (www.measuringworth.com/ukgdp).

business cycles, which gives a prominent role to technology shocks as a source of economic fluctuations.[4] ■

The Imperfect-Information Model

The third explanation for the upward slope of the short-run aggregate supply curve is called the **imperfect-information model**. Unlike the previous two models, this model assumes that markets clear – that is, all wages and prices are free to adjust to balance supply and demand. In this model, the short-run and long-run aggregate supply curves differ because of temporary misperceptions about prices.

The imperfect-information model assumes that each supplier in the economy produces a single good and consumes many goods. Because the number of goods is so large, suppliers cannot observe all prices at all times. They monitor closely the prices of what they produce, but less closely the prices of all the goods they

[4] For some of the recent work on the cyclical behaviour of the real wage, see Scott Sumner and Stephen Silver, 'Real Wages, Employment, and the Phillips Curve', *Journal of Political Economy*, June 1989, vol. 97, pp. 706–720; and Gary Solon, Robert Barsky and Jonathan A. Parker, 'Measuring the Cyclicality of Real Wages: How Important is Composition Bias?', *Quarterly Journal of Economics*, February 1994, vol. 109, pp. 1–25.

consume. Because of imperfect information, they sometimes confuse changes in the overall level of prices with changes in relative prices. This confusion influences decisions about how much to supply, and it leads to a positive relationship between the price level and output in the short run.

Consider the decision facing a single supplier – a wheat farmer, for instance. Because the farmer earns income from selling wheat, and uses this income to buy goods and services, the amount of wheat she chooses to produce depends on the price of wheat relative to the prices of other goods and services in the economy. If the relative price of wheat is high, the farmer is motivated to work hard and produce more wheat, because the reward is great. If the relative price of wheat is low, she prefers to enjoy more leisure and produce less wheat.

Unfortunately, when the farmer makes her production decision, she does not know the relative price of wheat. As a wheat producer, she monitors the wheat market closely and always knows the nominal price of wheat. But she does not know the prices of all the other goods in the economy. She must therefore estimate the relative price of wheat using the nominal price of wheat and her expectation of the overall price level.

Consider how the farmer responds if all prices in the economy, including the price of wheat, increase. One possibility is that she expected this change in prices. When she observes an increase in the price of wheat, her estimate of its relative price is unchanged. She does not work any harder.

The other possibility is that the farmer did not expect the price level to increase (or to increase by this much). When she observes the increase in the price of wheat, she is not sure whether other prices have risen (in which case wheat's relative price is unchanged) or whether only the price of wheat has risen (in which case its relative price is higher). The rational inference is that both have increased to some degree. In other words, the farmer infers from the increase in the nominal price of wheat that its relative price has risen somewhat. She works harder and produces more.

Our wheat farmer is not unique. When the price level rises unexpectedly, all suppliers in the economy observe increases in the prices of the goods they produce. They all infer, rationally but mistakenly, that the relative prices of the goods they produce have risen. They work harder and produce more.

To sum up, the imperfect-information model says that when actual prices exceed expected prices, suppliers raise their output. The model implies an aggregate supply curve with the familiar form:

$$Y = \overline{Y} + \alpha(P - P^e).$$

Output deviates from the natural level when the price level deviates from the expected price level.[5]

[5] Two economists who have emphasized the role of imperfect information for understanding the short-run effects of monetary policy are the Nobel Prize winners Milton Friedman and Robert Lucas. See Milton Friedman, 'The Role of Monetary Policy', *American Economic Review*, March 1968, vol. 58, pp. 1–17; and Robert E. Lucas, Jr., 'Understanding Business Cycles', *Carnegie-Rochester Conference on Public Policy*, vol. 5: *Stabilization of the Domestic and International Economy*, Amsterdam: North-Holland, 1999, pp. 7–29.

CASE STUDY

International Differences in the Aggregate Supply Curve

Although all countries experience economic fluctuations, these fluctuations are not exactly the same everywhere. International differences are intriguing puzzles in themselves, and they often provide a way to test alternative economic theories. Examining international differences has been especially fruitful in research on aggregate supply.

When economist Robert Lucas proposed the imperfect-information model, he derived a surprising interaction between aggregate demand and aggregate supply: according to his model, the slope of the aggregate supply curve should depend on the volatility of aggregate demand. In countries where aggregate demand fluctuates widely, the aggregate price level fluctuates widely as well. Because most movements in prices in these countries do not represent movements in relative prices, suppliers should have learned not to respond much to unexpected changes in the price level. Therefore, the aggregate supply curve should be relatively steep (that is, α will be small). Conversely, in countries where aggregate demand is relatively stable, suppliers should have learned that most price changes are relative price changes. Accordingly, in these countries, suppliers should be more responsive to unexpected price changes, making the aggregate supply curve relatively flat (that is, α will be large).

Lucas tested this prediction by examining international data on output and prices. He found that changes in aggregate demand have the biggest effect on output in those countries where aggregate demand and prices are most stable. Lucas concluded that the evidence supports the imperfect-information model.[6]

The sticky-price model also makes predictions about the slope of the short-run aggregate supply curve. In particular, it predicts that the average rate of inflation should influence the slope of the short-run aggregate supply curve. When the average rate of inflation is high, it is very costly for firms to keep prices fixed for long intervals. Thus, firms adjust prices more frequently. More frequent price adjustment, in turn, allows the overall price level to respond more quickly to shocks to aggregate demand. Hence, a high rate of inflation should make the short-run aggregate supply curve steeper.

International data support this prediction of the sticky-price model. In countries with low average inflation, the short-run aggregate supply curve is relatively flat: fluctuations in aggregate demand have large effects on output and are slowly reflected in prices. High-inflation countries have steep short-run aggregate supply curves. In other words, high inflation appears to erode the frictions that cause prices to be sticky.[7]

Note that the sticky-price model can also explain Lucas's finding that countries with variable aggregate demand have steep aggregate supply curves. If the price level is highly variable, few firms will commit to prices in advance (s will be small). Hence the aggregate supply curve will be steep (α will be small). ■

[6] Robert E. Lucas, Jr., 'Some International Evidence on Output–Inflation Tradeoffs', *American Economic Review*, June 1973, vol. 63, pp. 326–334.

[7] Laurence Ball, N. Gregory Mankiw and David Romer, 'The New Keynesian Economics and the Output–Inflation Trade-off', *Brookings Papers on Economic Activity*, 1988, vol. 1, pp. 1–65.

Summary and Implications

We have seen three models of aggregate supply and the market imperfection that each uses to explain why the short-run aggregate supply curve is upward sloping. One model assumes that the prices of some goods are sticky; the second assumes that nominal wages are sticky; the third assumes that information about prices is imperfect. Keep in mind that these models are not incompatible with one another. We need not accept one model and reject the others. The world may contain all three of these market imperfections, and all may contribute to the behaviour of short-run aggregate supply.

Although the three models of aggregate supply differ in their assumptions and emphases, their implications for aggregate output are similar. All can be summarized by the equation

$$Y = \overline{Y} + \alpha(P - P^e).$$

This equation states that deviations of output from the natural level are related to deviations of the price level from the expected price level. *If the price level is higher than the expected price level, output exceeds its natural level. If the price level is lower than the expected price level, output falls short of its natural level.* Figure 14-3 graphs this equation. Notice that the short-run aggregate supply curve is drawn for a given expectation P^e and that a change in P^e would shift the curve.

FIGURE 14-3

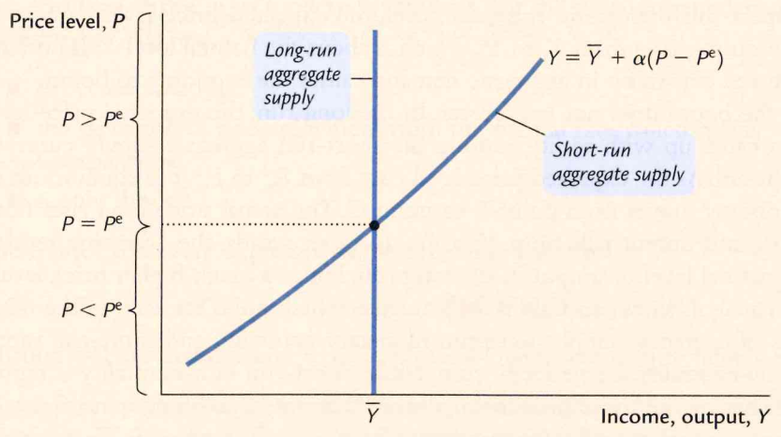

The Short-Run Aggregate Supply Curve Output deviates from its natural level $\overline{Y}$ if the price level P deviates from the expected price level P^e.

FIGURE 14-4

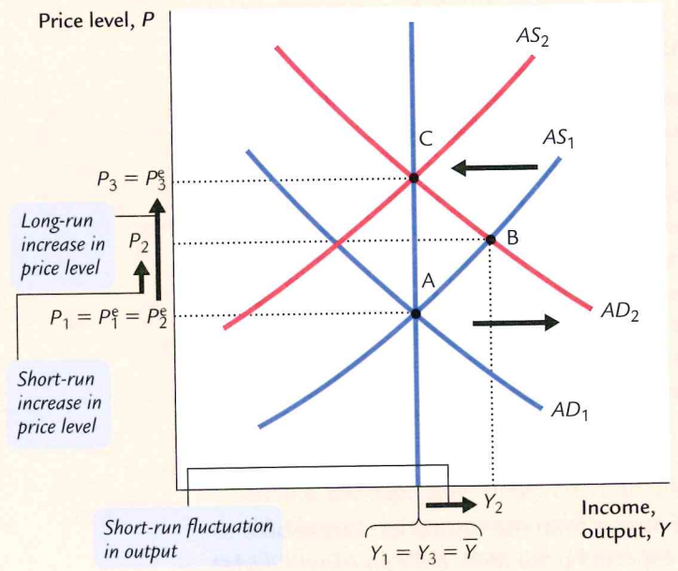

How Shifts in Aggregate Demand Lead to Short-Run Fluctuations Here the economy begins in a long-run equilibrium, point A. When aggregate demand increases unexpectedly, the price level rises from P_1 to P_2. Because the price level P_2 is above the expected price level P_2^e, output rises temporarily above the natural level, as the economy moves along the short-run aggregate supply curve from point A to point B. In the long run, the expected price level rises to P_3^e, causing the short-run aggregate supply curve to shift upward. The economy returns to a new long-run equilibrium, point C, where output is back at its natural level.

Now that we have a better understanding of aggregate supply, let's put aggregate supply and aggregate demand back together. Figure 14-4 uses our aggregate supply equation to show how the economy responds to an unexpected increase in aggregate demand attributable, say, to an unexpected monetary expansion. In the short run, the equilibrium moves from point A to point B. The increase in aggregate demand raises the actual price level from P_1 to P_2. Because people did not expect this increase in the price level, the expected price level remains at P_2^e, and output rises from Y_1 to Y_2, which is above the natural level $\overline{Y}$. Thus, the unexpected expansion in aggregate demand causes the economy to boom.

Yet the boom does not last forever. In the long run, the expected price level rises to catch up with reality, causing the short-run aggregate supply curve to shift upward. As the expected price level rises from P_2^e to P_3^e, the equilibrium of the economy moves from point B to point C. The actual price level rises from P_2 to P_3, and output falls from Y_2 to Y_3. In other words, the economy returns to the natural level of output in the long run, but at a much higher price level.

This analysis shows an important principle, which holds for each of the three models of aggregate supply: long-run monetary neutrality and short-run monetary *non*-neutrality are perfectly compatible. Short-run non-neutrality is represented here by the movement from point A to point B, and long-run monetary neutrality is represented by the movement from point A to point C. We reconcile the short-run and long-run effects of money by emphasizing the adjustment of expectations about the price level.

14-2 Inflation, Unemployment and the Phillips Curve

Two goals of economic policy makers are low inflation and low unemployment, but often these goals conflict. Suppose, for instance, that policy makers were to use monetary or fiscal policy to expand aggregate demand. This policy would move the economy along the short-run aggregate supply curve to a point of higher output and a higher price level. (Figure 14-4 shows this as the change from point A to point B.) Higher output means lower unemployment, because firms employ more workers when they produce more. A higher price level, given the previous year's price level, means higher inflation. Thus, when policy makers move the economy up along the short-run aggregate supply curve, they reduce the unemployment rate and raise the inflation rate. Conversely, when they contract aggregate demand and move the economy down the short-run aggregate supply curve, unemployment rises and inflation falls.

This trade-off between inflation and unemployment, called the *Phillips curve,* is our topic in this section. As we have just seen (and will derive more formally in a moment), the Phillips curve is a reflection of the short-run aggregate supply curve: as policy makers move the economy along the short-run aggregate supply curve, unemployment and inflation move in opposite directions. The Phillips curve is a useful way to express aggregate supply because inflation and unemployment are such important measures of economic performance.

Deriving the Phillips Curve from the Aggregate Supply Curve

The **Phillips curve** in its modern form states that the inflation rate depends on three forces:

- expected inflation;
- the deviation of unemployment from the natural rate, called *cyclical unemployment;*
- supply shocks.

These three forces are expressed in the following equation:

$$\pi = \pi^e - \beta(u - u^n) + v,$$

$$\text{Inflation} = \text{Expected Inflation} - (\beta \times \text{Cyclical Unemployment}) + \text{Supply Shock},$$

where β is a parameter measuring the response of inflation to cyclical unemployment. Notice that there is a minus sign before the cyclical unemployment term: other things equal, higher unemployment is associated with lower inflation.

Where does this equation for the Phillips curve come from? Although it may not seem familiar, we can derive it from our equation for aggregate supply. To see how, write the aggregate supply equation as

$$P = P^e + (1/\alpha)(Y - \overline{Y}).$$

With one addition, one subtraction and one substitution, we can transform this equation into the Phillips curve relationship between inflation and unemployment.

Here are the three steps. First, add to the right-hand side of the equation a supply shock v to represent exogenous events (such as a change in world oil prices) that alter the price level and shift the short-run aggregate supply curve:

$$P = P^e + (1/\overline{Y})(Y - \overline{Y}) + v.$$

Next, to go from the price level to inflation rates, subtract last year's price level P_{-1} from both sides of the equation to obtain

$$(P - P_{-1}) = (P^e - P_{-1}) + (1/\overline{Y})(Y - \overline{Y}) + v.$$

The term on the left-hand side, $P - P_{-1}$, is the difference between the current price level and last year's price level, which is inflation π.[8] The term on the right-hand side, $P^e - P_{-1}$, is the difference between the expected price level and last year's price level, which is expected inflation π^e. Therefore, we can replace $P - P_{-1}$ with π, and $P^e - P_{-1}$ with π^e:

$$\pi = \pi^e + (1/\overline{Y})(Y - \overline{Y}) + v.$$

Third, to go from output to unemployment, we simply note that output will be a positive function of the number of people employed, through the aggregate production function. But if output is positively related to the level of employment, it must in general be negatively related to the level of unemployment. In particular, the deviation of output from its natural level will tend to be inversely related to the deviation of unemployment from its natural rate; that is, when output is higher than the natural level of output, unemployment is lower than the natural rate of unemployment. We can write this as

$$(1/\alpha)(Y - \overline{Y}) = -\beta(u - u^n).$$

Using this relationship, we can substitute $-\beta(u - u^n)$ for $(1/\alpha)(Y - \overline{Y})$ in the previous equation to obtain

$$\pi = \pi^e - \beta(u - u^n) + v.$$

Thus, we can derive the Phillips curve equation from the aggregate supply equation.

All this algebra is meant to show one thing: the Phillips curve equation and the short-run aggregate supply equation represent essentially the same macroeconomic ideas. In particular, both equations show a link between real and nominal variables which causes the classical dichotomy (the theoretical separation of real and nominal variables) to break down in the short run. According to the short-run aggregate supply equation, output is related to unexpected movements in the price level. According to the Phillips curve equation, unemployment is related to unexpected movements in the inflation rate. The aggregate supply curve is more convenient when we are studying output and the price level, whereas the Phillips curve is more convenient when we are studying unemployment and inflation. But we should not lose sight of the fact that the Phillips curve and the aggregate supply curve are two sides of the same coin.

[8] *Mathematical note*: This statement is not precise, because inflation is really the *percentage* change in the price level. To make the statement more precise, interpret P as the logarithm of the price level. By the properties of logarithms, the change in P is roughly the inflation rate. The reason is that $dP = d(\text{log price level}) = d(\text{price level})/\text{price level}$.

FYI

The History of the Modern Phillips Curve

The Phillips curve is named after New Zealand-born economist A. W. Phillips, who worked at the London School of Economics. In 1958, Phillips observed a negative relationship between the unemployment rate and the rate of wage inflation in data for the United Kingdom.[9] The Phillips curve that economists use today differs in three ways from the relationship Phillips examined.

First, the modern Phillips curve substitutes price inflation for wage inflation. This difference is not crucial, because price inflation and wage inflation are closely related. In periods when wages are rising quickly, prices are rising quickly as well.

Second, the modern Phillips curve includes expected inflation. This addition is due to the work of Nobel Prize winners Milton Friedman and Edmund Phelps. In developing early versions of the imperfect-information model in the 1960s, these two economists emphasized the importance of expectations for aggregate supply.

Third, the modern Phillips curve includes supply shocks. Credit for this addition goes to OPEC, the Organization of Petroleum Exporting Countries. In the 1970s, OPEC caused large increases in the world price of oil, which made economists more aware of the importance of shocks to aggregate supply.

[9] A. W. Phillips, 'The Relationship between Unemployment and the Rate of Change of Money Wages in the United Kingdom, 1861–1957', *Economica*, November 1958, vol. 25, pp. 283–299.

Adaptive Expectations and Inflation Inertia

To make the Phillips curve useful for analysing the choices facing policy makers, we need to say what determines expected inflation. A simple and often plausible assumption is that people form their expectations of inflation based on recently observed inflation. This assumption is called **adaptive expectations**. For example, suppose that people expect prices to rise this year at the same rate as they did last year. Then expected inflation π^e equals last year's inflation π_{-1}:

$$\pi^e = \pi_{-1}.$$

In this case, we can write the Phillips curve as

$$\pi = \pi_{-1} - \beta(u - u^n) + v,$$

which states that inflation depends on past inflation, cyclical unemployment and a supply shock. When the Phillips curve is written in this form, the natural rate