

Aggregate Supply

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

— Milton Friedman

Most economists analyze 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 9, we established that aggregate supply behaves differently in the short run than 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 rate. 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 9 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.

Unfortunately, one fact makes this task more difficult: economists disagree about how best to explain aggregate supply. As a result, this chapter begins by presenting three prominent 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 also 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 tradeoff between two measures of economic performance—inflation and unemployment. According to this tradeoff, to reduce the rate of inflation policymakers must temporarily raise unemployment, and to reduce unemployment they must accept higher inflation. As the quotation at the beginning of the chapter suggests, the tradeoff between inflation and unemployment is only temporary. One goal of this chapter is to explain why policymakers face such a tradeoff in the short run and, just as important, why they do not face it in the long run.

13-1 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, roughly in the order of their development. In all the models, some market imperfection (that is, some type of friction) causes the output of the economy to deviate from the classical benchmark. As a result, the short-run aggregate supply curve is upward sloping, rather than vertical, and shifts in the aggregate demand curve cause the level of output to deviate temporarily from the natural rate. These temporary deviations 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

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

where Y is output, \bar{Y} is the natural rate of output, P is the price level, and P^e is the expected price level. This equation states that output deviates from its natural rate 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 highlights a particular reason why unexpected movements in the price level are associated with fluctuations in aggregate output.

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 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 labor cheaper.
2. The lower real wage induces firms to hire more labor.
3. The additional labor 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 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 labor supply and demand. More likely, the target real wage is higher than the equilibrium real wage: as discussed in Chapter 6, union power and efficiency-wage considerations tend to keep real wages above the level that brings supply and 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

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

Nominal Wage = Target Real Wage \times Expected Price Level.

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

$$\frac{W}{P} = \omega \times \frac{(P^e/P)}{1}$$

Real Wage = 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 labor 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 labor as the firms wish to buy at the predetermined wage. We describe the firms' hiring decisions by the labor demand function

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

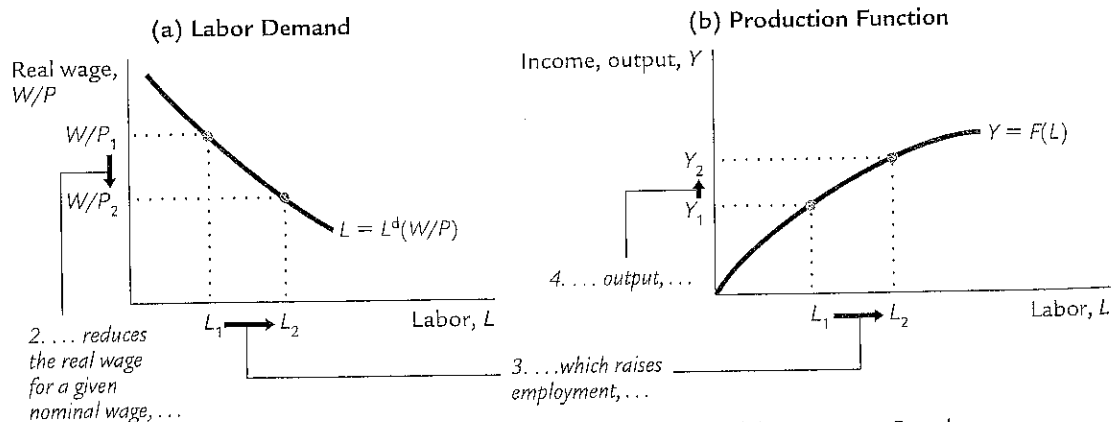
which states that the lower the real wage, the more labor firms hire. The labor demand curve is shown in panel (a) of Figure 13-1. Output is determined by the production function

$$Y = F(L),$$

which states that the more labor is hired, the more output is produced. This is shown in panel (b) of Figure 13-1.

Panel (c) of Figure 13-1 shows the resulting aggregate supply curve. Because the nominal wage is sticky, an unexpected change in the price level

FIGURE 13-1



2. ... reduces the real wage for a given nominal wage, ...

3. ... which raises employment, ...

1. An increase in the price level...

6. The aggregate supply curve summarizes these changes.

5. ... and income.

The Sticky-Wage Model Panel (a) shows the labor 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 labor demanded from L_1 to L_2 . Panel (b) shows the production function. An increase in the quantity of labor 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 .

moves the real wage away from the target real wage, and this change in the real wage influences the amounts of labor hired and output produced. The aggregate supply curve can be written as

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

Output deviates from its natural level when the price level deviates from the expected price level.¹

CASE STUDY

The Cyclical Behavior of the Real Wage

In any model with an unchanging labor 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 labor hired and the amount of output produced. Thus, the real wage should be *countercyclical*: 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 13-2 is a scatterplot of the percentage change in real compensation per hour and the percentage change in real GDP using annual data for the U.S. economy from 1960 to 2000. 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 the real wage and output, and it is the opposite of what Keynes predicted. That is, if the real wage is cyclical at all, it is slightly *procyclical*: the real wage tends to rise when output rises. Abnormally high labor 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 labor demand curve shifts over the business cycle. These shifts may arise because firms have sticky prices and cannot sell all they want at those prices; we discuss this possibility later. Alternatively, the labor demand curve may shift because of shocks to technology, which alter labor productivity. The theory we discuss in Chapter 19, called the theory of real business cycles, gives a prominent role to technology shocks as a source of economic fluctuations.²

¹ For more on the sticky-wage model, see Jo Anna Gray, “Wage Indexation: A Macroeconomic Approach,” *Journal of Monetary Economics* 2 (April 1976): 221–235; and Stanley Fischer, “Long-Term Contracts, Rational Expectations, and the Optimal Money Supply Rule,” *Journal of Political Economy* 85 (February 1977): 191–205.

² For some of the recent work on the cyclical behavior of the real wage, see Scott Sumner and Stephen Silver, “Real Wages, Employment, and the Phillips Curve,” *Journal of Political Economy* 97 (June 1989): 706–720; and Gary Solon, Robert Barsky, and Jonathan A. Parker, “Measuring the Cyclical Behavior of Real Wages: How Important Is Composition Bias?” *Quarterly Journal of Economics* 109 (February 1994): 1–25.

Figure 13-3



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

Source: U.S. Department of Commerce and U.S. Department of Labor.

The Imperfect-Information Model

The second explanation for the upward slope of the short-run aggregate supply curve is called the **imperfect-information model**. Unlike the sticky-wage model, 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 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 some of each has happened. 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 that is now familiar:

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

Output deviates from the natural rate when the price level deviates from the expected price level.³

The Sticky-Price Model

Our third 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

³ 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* 58 (March 1968): 1-17; and Robert E. Lucas, Jr., "Understanding Business Cycles," *Stabilization of the Domestic and International Economy*, vol. 5 of Carnegie-Rochester Conference on Public Policy (Amsterdam: North-Holland, 1977).

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 catalog 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 behavior 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.

We write the firm's desired price as

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

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 rate $Y - \bar{Y}$. The parameter a (which is greater than zero) measures how much the firm's desired price responds to the level of aggregate output.⁴

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

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

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 rate, so that the last term, $a(Y^e - \bar{Y}^e)$, is zero. Then these firms set the price

$$p = P^e.$$

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

⁴ *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 the natural rate.

⁵ For a
tic Pric

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 $1 - s$ the fraction with flexible prices, then the overall price level is

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

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

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

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

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

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:

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

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

Although the sticky-price model emphasizes the goods market, consider briefly what is happening in the labor market. If a firm's price is stuck in the short run, then a reduction in aggregate demand reduces the amount that the firm is able to sell. The firm responds to the drop in sales by reducing its production and its demand for labor. Note the contrast to the sticky-wage model: the firm here does not move along a fixed labor demand curve. Instead, fluctuations in output are associated with shifts in the labor demand curve. Because of these shifts in labor demand, employment, production, and the real wage can all move in the same direction. Thus, the real wage can be procyclical.⁵

⁵ For a more advanced development of the sticky-price model, see Julio Rotemberg, "Monopolistic Price Adjustment and Aggregate Output," *Review of Economic Studies* 49 (1982): 517-531.

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* 63 (June 1973): 326–334.

⁷ Laurence Ball, N. Gregory Mankiw, and David Romer, "The New Keynesian Economics and the Output-Inflation Tradeoff," *Brookings Papers on Economic Activity* (1988:1): 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 nominal wages are sticky; the second assumes information about prices is imperfect; the third assumes prices are sticky. 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 behavior 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

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

This equation states that deviations of output from the natural rate 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 rate. If the price level is lower than the expected price level, output falls short of its natural rate. Figure 13-3 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.

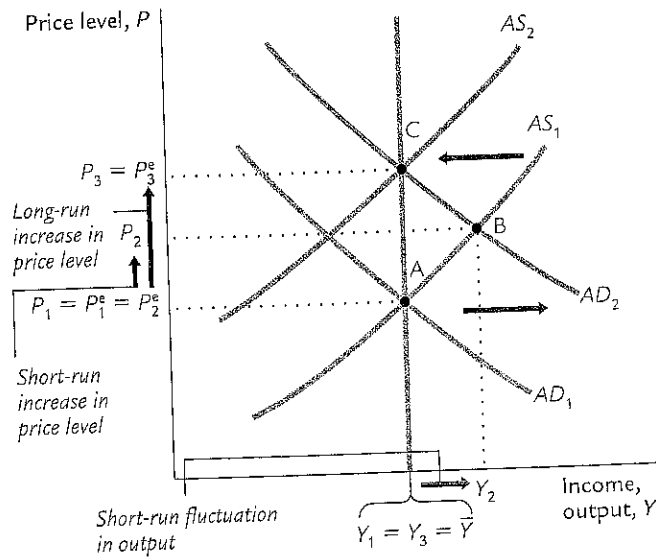
Now that we have a better understanding of aggregate supply, let's put aggregate supply and aggregate demand back together. Figure 13-4 uses our aggregate supply equation to show how the economy responds to an unexpected increase in aggregate demand attributable, say, to an unexpected monetary expansion. 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 rate \bar{Y} . Thus, the unexpected expansion in aggregate demand causes the economy to boom.

Figure 13-3



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

FIGURE 13-4



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 rate, 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 rate.

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 nonneutrality are perfectly compatible. Short-run nonneutrality 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.

13-2 Inflation, Unemployment, and the Phillips Curve

Two goals of economic policymakers are low inflation and low unemployment, but often these goals conflict. Suppose, for instance, that policymakers 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 13-4 shows this as the change from point A to point B.) Higher output means lower unemployment, because firms need more workers when they produce more. A higher price level, given

the previous year's price level, means higher inflation. Thus, when policymakers 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 tradeoff 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 policymakers 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:

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

Inflation = Expected Inflation - ($\beta \times$ Cyclical Unemployment) + 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: high unemployment tends to reduce inflation. This equation summarizes the relationship between inflation and unemployment.

From where does this equation for the Phillips curve come? 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

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

With one addition, one subtraction, and one substitution, we can manipulate this equation to yield a 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:

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

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

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

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 :

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

Third, to go from output to unemployment, recall from Chapter 2 that Okun's law gives a relationship between these two variables. One version of Okun's law states that the deviation of output from its natural rate is inversely related to the deviation of unemployment from its natural rate; that is, when output is higher than the natural rate of output, unemployment is lower than the natural rate of unemployment. We can write this as

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

Using this Okun's law relationship, we can substitute $-\beta(u - u^n)$ for $(1/\alpha)(Y - \bar{Y})$ 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 that 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 $dP = d(\log \text{ price level}) = d(\text{price level})/\text{price level}$.

The History of the Modern Phillips Curve

The Phillips curve is named after New Zealand-born economist A. W. Phillips. 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 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.

Adaptive Expectations and Inflation Inertia

To make the Phillips curve useful for analyzing the choices facing policymakers, 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 of unemployment is sometimes called the Non-Accelerating Inflation Rate of Unemployment, or *NAIRU*.

The first term in this form of the Phillips curve, π_{-1} , implies that inflation has inertia. That is, like an object moving through space, inflation keeps going unless something acts to stop it. In particular, if unemployment is at the NAIRU and if there are no supply shocks, the continued rise in price level neither speeds up nor slows down. This inertia arises because past inflation influences expectations of future inflation and because these expectations influence the wages and prices

⁹ A. W. Phillips, "The Relationship Between Unemployment and the Rate of Change of Money Wages in the United Kingdom, 1861-1957," *Economica* 25 (November 1958): 283-299.

that people set. Robert Solow captured the concept of inflation inertia well when, during the high inflation of the 1970s, he wrote, “Why is our money ever less valuable? Perhaps it is simply that we have inflation because we expect inflation, and we expect inflation because we’ve had it.”

In the model of aggregate supply and aggregate demand, inflation inertia is interpreted as persistent upward shifts in both the aggregate supply curve and the aggregate demand curve. Consider first aggregate supply. If prices have been rising quickly, people will expect them to continue to rise quickly. Because the position of the short-run aggregate supply curve depends on the expected price level, the short-run aggregate supply curve will shift upward over time. It will continue to shift upward until some event, such as a recession or a supply shock, changes inflation and thereby changes expectations of inflation.

The aggregate demand curve must also shift upward to confirm the expectations of inflation. Most often, the continued rise in aggregate demand is caused by persistent growth in the money supply. If the Fed suddenly halted money growth, aggregate demand would stabilize, and the upward shift in aggregate supply would cause a recession. The high unemployment in the recession would reduce inflation and expected inflation, causing inflation inertia to subside.

Two Causes of Rising and Falling Inflation

The second and third terms in the Phillips curve equation show the two forces that can change the rate of inflation.


The second term, $\beta(u - u^n)$, shows that cyclical unemployment—the deviation of unemployment from its natural rate—exerts upward or downward pressure on inflation. Low unemployment pulls the inflation rate up. This is called **demand-pull inflation** because high aggregate demand is responsible for this type of inflation. High unemployment pulls the inflation rate down. The parameter β measures how responsive inflation is to cyclical unemployment.

The third term, v , shows that inflation also rises and falls because of supply shocks. An adverse supply shock, such as the rise in world oil prices in the 1970s, implies a positive value of v and causes inflation to rise. This is called **cost-push inflation** because adverse supply shocks are typically events that push up the costs of production. A beneficial supply shock, such as the oil glut that led to a fall in oil prices in the 1980s, makes v negative and causes inflation to fall.

Inflation and Unemployment in the United States

Because inflation and unemployment are such important measures of economic performance, macroeconomic developments are often viewed through the lens of the Phillips curve. Figure 13-5 displays the history of inflation and unemployment in the United States since 1961. These four decades of data illustrate some of the causes of rising or falling inflation.

Figure



Inflation and Unemployment in the United States Since 1961 This figure uses annual data on the unemployment rate and the inflation rate (percentage change in the GDP deflator) to illustrate macroeconomic developments over the past four decades.

Source: U.S. Department of Commerce and U.S. Department of Labor.

The 1960s showed how policymakers can, in the short run, lower unemployment at the cost of higher inflation. The tax cut of 1964, together with expansionary monetary policy, expanded aggregate demand and pushed the unemployment rate below 5 percent. This expansion of aggregate demand continued in the late 1960s largely as a by-product of government spending for the Vietnam War. Unemployment fell lower and inflation rose higher than policymakers intended.

The 1970s were a period of economic turmoil. The decade began with policymakers trying to lower the inflation inherited from the 1960s. President Nixon imposed temporary controls on wages and prices, and the Federal Reserve engineered a recession through contractionary monetary policy, but the inflation rate fell only slightly. The effects of wage and price controls ended when the controls were lifted, and the recession was too small to counteract the inflationary impact of the boom that had preceded it. By 1972 the unemployment rate was the same as a decade earlier, whereas inflation was 3 percentage points higher.

Beginning in 1973 policymakers had to cope with the large supply shocks caused by the Organization of Petroleum Exporting Countries (OPEC). OPEC first raised oil prices in the mid-1970s, pushing the inflation rate up to about 10 percent. This adverse supply shock, together with temporarily tight monetary policy, led to a recession in 1975. High unemployment during the recession reduced inflation somewhat, but further OPEC price hikes pushed inflation up again in the late 1970s.

The 1980s began with high inflation and high expectations of inflation. Under the leadership of Chairman Paul Volcker, the Federal Reserve doggedly pursued monetary policies aimed at reducing inflation. In 1982 and 1983 the unemployment rate reached its highest level in 40 years. High unemployment, aided by a fall in oil prices in 1986, pulled the inflation rate down from about 10 percent to about 3 percent. By 1987 the unemployment rate of about 6 percent was close to most estimates of the natural rate. Unemployment continued to fall through the 1980s, however, reaching a low of 5.2 percent in 1989 and beginning a new round of demand-pull inflation.

Compared to the previous 30 years, the 1990s were relatively quiet. The decade began with a recession caused by several contractionary shocks to aggregate demand: tight monetary policy, the savings-and-loan crisis, and a fall in consumer confidence coinciding with the Gulf War. The unemployment rate rose to 7.3 percent in 1992. Inflation fell, but only slightly. Unlike in the 1982 recession, unemployment in the 1990 recession was never far above the natural rate, so the effect on inflation was small.

By the late 1990s, inflation and unemployment both reached their lowest levels in many years. Some economists explain this fortunate development by claiming that the economy's natural rate of unemployment fell (for reasons discussed in Chapter 6). Others argue that various temporary factors (such as a strong U.S. dollar attributable to a financial crisis in Asia) yielded favorable supply shocks. Most likely, a combination of events helped keep inflation in check, despite low unemployment. In 2000, however, inflation did begin to creep up.


Thus, U.S. macroeconomic history exhibits the many causes of inflation. The 1960s and the 1980s show the two sides of demand-pull inflation: in the 1960s low unemployment pulled inflation up, and in the 1980s high unemployment pulled inflation down. The 1970s with their oil-price hikes show the effects of cost-push inflation.

The Short-Run Tradeoff Between Inflation and Unemployment

Consider the options the Phillips curve gives to a policymaker who can influence aggregate demand with monetary or fiscal policy. At any moment, expected inflation and supply shocks are beyond the policymaker's immediate control. Yet, by changing aggregate demand, the policymaker can alter output, unemployment, and inflation. The policymaker can expand aggregate demand to lower unemployment and raise inflation. Or the policymaker can depress aggregate demand to raise unemployment and lower inflation.

Figure 13-6 plots the Phillips curve equation and shows the short-run tradeoff between inflation and unemployment. When unemployment is at its natural rate ($u = u^n$), inflation depends on expected inflation and the supply shock ($\pi = \pi^e + v$). The parameter β determines the slope of the tradeoff between inflation and unemployment. In the short run, for a given level of expected inflation, policymakers can manipulate aggregate demand to choose a

Figure 13-6




The Short-Run Tradeoff Between Inflation and Unemployment In the short run, inflation and unemployment are negatively related. At any point in time, a policymaker who controls aggregate demand can choose a combination of inflation and unemployment on this short-run Phillips curve.

combination of inflation and unemployment on this curve, called the *short-run Phillips curve*.

Notice that the position of the short-run Phillips curve depends on the expected rate of inflation. If expected inflation rises, the curve shifts upward, and the policymaker's tradeoff becomes less favorable: inflation is higher for any level of unemployment. Figure 13-7 shows how the tradeoff depends on expected inflation.

Because people adjust their expectations of inflation over time, the tradeoff between inflation and unemployment holds only in the short run. The policymaker cannot keep inflation above expected inflation (and thus unemployment below its natural rate) forever. Eventually, expectations adapt to whatever inflation rate the

Figure 13-7



Shifts in the Short-Run Tradeoff The short-run tradeoff between inflation and unemployment depends on expected inflation. The curve is higher when expected inflation is higher.

How Precise Are Estimates of the Natural Rate of Unemployment?

If you ask an astronomer how far a particular star is from our sun, he'll give you a number, but it won't be accurate. Man's ability to measure astronomical distances is still limited. An astronomer might well take better measurements and conclude that a star is really twice or half as far away as he previously thought.

Estimates of the natural rate of unemployment, or NAIRU, are also far from precise. One problem is supply shocks. Shocks to oil supplies, farm harvests, or technological progress can cause inflation to rise or fall in the short run. When we observe rising inflation, therefore, we cannot be sure if it is evidence that the unemployment rate is below the natural rate or evidence that the economy is experiencing an adverse supply shock.

A second problem is that the natural rate changes over time. Demographic changes (such as the aging of the baby-boom generation), policy changes (such as minimum-wage laws), and institutional changes (such as the declining role

of unions) all influence the economy's normal level of unemployment. Estimating the natural rate is like hitting a moving target.

Economists deal with these problems using statistical techniques that yield a best guess about the natural rate and allow them to gauge the uncertainty associated with their estimates. In one such study, Douglas Staiger, James Stock, and Mark Watson estimated the natural rate to be 6.2 percent in 1990, with a 95-percent confidence interval from 5.1 to 7.7 percent. A 95-percent confidence interval is a range such that the statistician is 95-percent confident that the true value falls in that range. The large confidence interval here of 2.6 percentage points shows that estimates of the natural rate are not at all precise.

This conclusion has profound implications. Policymakers may want to keep unemployment close to its natural rate, but their ability to do so is limited by the fact that we cannot be sure what that natural rate is.¹⁰

policymaker has chosen. In the long run, the classical dichotomy holds, unemployment returns to its natural rate, and there is no tradeoff between inflation and unemployment.

Disinflation and the Sacrifice Ratio

Imagine an economy in which unemployment is at its natural rate and inflation is running at 6 percent. What would happen to unemployment and output if the central bank pursued a policy to reduce inflation from 6 to 2 percent?

The Phillips curve shows that in the absence of a beneficial supply shock, lowering inflation requires a period of high unemployment and reduced output. But by how much and for how long would unemployment need to rise above the natural rate? Before deciding whether to reduce inflation, policymakers must know how much output would be lost during the transition to lower inflation. This cost can then be compared with the benefits of lower inflation.

¹⁰ Douglas Staiger, James H. Stock, and Mark W. Watson, "How Precise Are Estimates of the Natural Rate of Unemployment?" in Christina D. Romer and David H. Romer, eds., *Reducing Inflation: Motivation and Strategy* (Chicago: University of Chicago Press, 1997).

Much research has used the available data to examine the Phillips curve quantitatively. The results of these studies are often summarized in a number called the **sacrifice ratio**, the percentage of a year's real GDP that must be forgone to reduce inflation by 1 percentage point. Although estimates of the sacrifice ratio vary substantially, a typical estimate is about 5: for every percentage point that inflation is to fall, 5 percent of one year's GDP must be sacrificed.¹¹

We can also express the sacrifice ratio in terms of unemployment. Okun's law says that a change of 1 percentage point in the unemployment rate translates into a change of 2 percentage points in GDP. Therefore, reducing inflation by 1 percentage point requires about 2.5 percentage points of cyclical unemployment.

We can use the sacrifice ratio to estimate by how much and for how long unemployment must rise to reduce inflation. If reducing inflation by 1 percentage point requires a sacrifice of 5 percent of a year's GDP, reducing inflation by 4 percentage points requires a sacrifice of 20 percent of a year's GDP. Equivalently, this reduction in inflation requires a sacrifice of 10 percentage points of cyclical unemployment.

This disinflation could take various forms, each totaling the same sacrifice of 20 percent of a year's GDP. For example, a rapid disinflation would lower output by 10 percent for 2 years: this is sometimes called the *cold-turkey* solution to inflation. A moderate disinflation would lower output by 5 percent for 4 years. An even more gradual disinflation would depress output by 2 percent for a decade.

Rational Expectations and the Possibility of Painless Disinflation

Because the expectation of inflation influences the short-run tradeoff between inflation and unemployment, it is crucial to understand how people form expectations. So far, we have been assuming that expected inflation depends on recently observed inflation. Although this assumption of adaptive expectations is plausible, it is probably too simple to apply in all circumstances.

An alternative approach is to assume that people have **rational expectations**. That is, we might assume that people optimally use all the available information, including information about current government policies, to forecast the future. Because monetary and fiscal policies influence inflation, expected inflation should also depend on the monetary and fiscal policies in effect. According to the theory of rational expectations, a change in monetary or fiscal policy will change expectations, and an evaluation of any policy change must incorporate this effect on expectations. If people do form their expectations rationally, then inflation may have less inertia than it first appears.

¹¹ Arthur M. Okun, "Efficient Disinflationary Policies," *American Economic Review* 68 (May 1978): 348-352; and Robert J. Gordon and Stephen R. King, "The Output Cost of Disinflation in Traditional and Vector Autoregressive Models," *Brookings Papers on Economic Activity* (1982:1): 205-245.

Here is how Thomas Sargent, a prominent advocate of rational expectations, describes its implications for the Phillips curve:

An alternative "rational expectations" view denies that there is any inherent momentum to the present process of inflation. This view maintains that firms and workers have now come to expect high rates of inflation in the future and that they strike inflationary bargains in light of these expectations. However, it is held that people expect high rates of inflation in the future precisely because the government's current and prospective monetary and fiscal policies warrant those expectations. . . . Thus inflation only seems to have a momentum of its own; it is actually the long-term government policy of persistently running large deficits and creating money at high rates which imparts the momentum to the inflation rate. An implication of this view is that inflation can be stopped much more quickly than advocates of the "momentum" view have indicated and that their estimates of the length of time and the costs of stopping inflation in terms of foregone output are erroneous. . . . [Stopping inflation] would require a change in the policy regime: there must be an abrupt change in the continuing government policy, or strategy, for setting deficits now and in the future that is sufficiently binding as to be widely believed. . . . How costly such a move would be in terms of foregone output and how long it would be in taking effect would depend partly on how resolute and evident the government's commitment was.¹²

Thus, advocates of rational expectations argue that the short-run Phillips curve does not accurately represent the options that policymakers have available. They believe that if policymakers are credibly committed to reducing inflation, rational people will understand the commitment and will quickly lower their expectations of inflation. Inflation can then come down without a rise in unemployment and fall in output. According to the theory of rational expectations, traditional estimates of the sacrifice ratio are not useful for evaluating the impact of alternative policies. Under a credible policy, the costs of reducing inflation may be much lower than estimates of the sacrifice ratio suggest.

In the most extreme case, one can imagine reducing the rate of inflation without causing any recession at all. A painless disinflation has two requirements. First, the plan to reduce inflation must be announced before the workers and firms who set wages and prices have formed their expectations. Second, the workers and firms must believe the announcement; otherwise, they will not reduce their expectations of inflation. If both requirements are met, the announcement will immediately shift the short-run tradeoff between inflation and unemployment downward, permitting a lower rate of inflation without higher unemployment.

Although the rational-expectations approach remains controversial, almost all economists agree that expectations of inflation influence the short-run tradeoff between inflation and unemployment. The credibility of a policy to reduce inflation is therefore one determinant of how costly the policy will be. Unfortunately, it is often difficult to predict whether the public will view the announcement of a new policy as credible. The central role of expectations makes forecasting the results of alternative policies far more difficult.

¹² Thomas J. Sargent, "The Ends of Four Big Inflations," in Robert E. Hall, ed., *Inflation: Causes and Effects* (Chicago: University of Chicago Press, 1982).

CASE STUDY

The Sacrifice Ratio in Practice

The Phillips curve with adaptive expectations implies that reducing inflation requires a period of high unemployment and low output. By contrast, the rational-expectations approach suggests that reducing inflation can be much less costly. What happens during actual disinflations?

Consider the U.S. disinflation in the early 1980s. This decade began with some of the highest rates of inflation in U.S. history. Yet because of the tight monetary policies the Fed pursued under Chairman Paul Volcker, the rate of inflation fell substantially in the first few years of the decade. This episode provides a natural experiment with which to estimate how much output is lost during the process of disinflation.

The first question is, how much did inflation fall? As measured by the GDP deflator, inflation reached a peak of 9.7 percent in 1981. It is natural to end the episode in 1985 because oil prices plunged in 1986—a large, beneficial supply shock unrelated to Fed policy. In 1985, inflation was 3.0 percent, so we can estimate that the Fed engineered a reduction in inflation of 6.7 percentage points over four years.

The second question is, how much output was lost during this period? Table 13-1 shows the unemployment rate from 1982 to 1985. Assuming that the natural rate of unemployment was 6 percent, we can compute the amount of cyclical unemployment in each year. In total over this period, there were 9.5 percentage points of cyclical unemployment. Okun's law says that 1 percentage point of unemployment translates into 2 percentage points of GDP. Therefore, 19.0 percentage points of annual GDP were lost during the disinflation.

table 13-1

Unemployment During the Volcker Disinflation

Year	Unemployment Rate, u	Natural Rate, u^n	Cyclical Unemployment, $u - u^n$
1982	9.5%	6.0%	3.5%
1983	9.5	6.0	3.5
1984	7.4	6.0	1.4
1985	7.1	6.0	1.1
			Total 9.5%

Now we can compute the sacrifice ratio for this episode. We know that 19.0 percentage points of GDP were lost and that inflation fell by 6.7 percentage points. Hence, $19.0/6.7$, or 2.8, percentage points of GDP were lost for each percentage-point reduction in inflation. The estimate of the sacrifice ratio from the Volcker disinflation is 2.8.

This estimate of the sacrifice ratio is smaller than the estimates made before Volcker was appointed Fed chairman. In other words, Volcker reduced inflation at a smaller cost than many economists had predicted. One explanation is that Volcker's tough stand was credible enough to influence expectations of inflation directly. Yet the change in expectations was not large enough to make the disinflation painless: in 1982 unemployment reached its highest level since the Great Depression.

Although the Volcker disinflation is only one historical episode, this kind of analysis can be applied to other disinflations. A recent study documented the results of 65 disinflations in 19 countries. In almost all cases, the reduction in inflation came at the cost of temporarily lower output. Yet the size of the output loss varied from episode to episode. Rapid disinflations usually had smaller sacrifice ratios than slower ones. That is, in contrast to what the Phillips curve with adaptive expectations suggests, a cold-turkey approach appears less costly than a gradual one. Moreover, countries with more flexible wage-setting institutions, such as shorter labor contracts, had smaller sacrifice ratios. These findings indicate that reducing inflation always has some cost, but that policies and institutions can affect its magnitude.¹³

Hysteresis and the Challenge to the Natural-Rate Hypothesis

Our discussion of the cost of disinflation—and indeed our entire discussion of economic fluctuations in the past four chapters—has been based on an assumption called the **natural-rate hypothesis**. This hypothesis is summarized in the following statement:

Fluctuations in aggregate demand affect output and employment only in the short run. In the long run, the economy returns to the levels of output, employment, and unemployment described by the classical model.

The natural-rate hypothesis allows macroeconomists to study separately short-run and long-run developments in the economy. It is one expression of the classical dichotomy.

Recently, some economists have challenged the natural-rate hypothesis by suggesting that aggregate demand may affect output and employment even in the long run. They have pointed out a number of mechanisms through which recessions might leave permanent scars on the economy by altering the natural rate of unemployment. **Hysteresis** is the term used to describe the long-lasting influence of history on the natural rate.

A recession can have permanent effects if it changes the people who become unemployed. For instance, workers might lose valuable job skills when unemployed, lowering their ability to find a job even after the recession ends.

¹³ Laurence Ball, "What Determines the Sacrifice Ratio?" in N. Gregory Mankiw, ed., *Monetary Policy* (Chicago: University of Chicago Press, 1994).

Alternatively, a long period of unemployment may change an individual's attitude toward work and reduce his desire to find employment. In either case, the recession permanently inhibits the process of job search and raises the amount of frictional unemployment.

Another way in which a recession can permanently affect the economy is by changing the process that determines wages. Those who become unemployed may lose their influence on the wage-setting process. Unemployed workers may lose their status as union members, for example. More generally, some of the *insiders* in the wage-setting process become *outsiders*. If the smaller group of insiders cares more about high real wages and less about high employment, then the recession may permanently push real wages further above the equilibrium level and raise the amount of structural unemployment.

Hysteresis remains a controversial theory. Some economists believe the theory helps explain persistently high unemployment in Europe, because the rise in European unemployment starting in the early 1980s coincided with disinflation but continued after inflation stabilized. Moreover, the increase in unemployment tended to be larger for those countries that experienced the greatest reductions in inflations, such as Ireland, Italy, and Spain. Yet there is still no consensus whether the hysteresis phenomenon is significant, or why it might be more pronounced in some countries than in others. (Other explanations of high European unemployment, discussed in Chapter 6, give little role to the disinflation.) If it is true, however, the theory is important, because hysteresis greatly increases the cost of recessions. Put another way, hysteresis raises the sacrifice ratio, because output is lost even after the period of disinflation is over.¹⁴

13-3 Conclusion

We began this chapter by discussing three models of aggregate supply, each of which focuses on a different reason why the short-run aggregate supply curve is upward sloping. The three models have similar predictions for the aggregate economy, and all of them yield a short-run tradeoff between inflation and unemployment. A convenient way to express and analyze that tradeoff is with the Phillips curve equation, according to which inflation depends on expected inflation, cyclical unemployment, and supply shocks.

Keep in mind that not all economists endorse all the ideas discussed here. There is widespread disagreement, for instance, about the practical importance of rational expectations and the relevance of hysteresis. If you find it difficult to fit all the pieces together, you are not alone. The study of aggregate supply remains one of the most unsettled—and therefore one of the most exciting—research areas in macroeconomics.

¹⁴ Olivier J. Blanchard and Lawrence H. Summers, "Beyond the Natural Rate Hypothesis," *American Economic Review* 78 (May 1988): 182–187; Laurence Ball, "Disinflation and the NAIRU," in Christina D. Romer and David H. Romer, eds., *Reducing Inflation: Motivation and Strategy* (Chicago: University of Chicago Press, 1997): 167–185.

Summary

1. The three theories of aggregate supply—the sticky-wage, imperfect-information, and sticky-price models—attribute deviations of output and employment from the natural rate to various market imperfections. According to all three theories, output rises above the natural rate when the price level exceeds the expected price level, and output falls below the natural rate when the price level is less than the expected price level.
2. Economists often express aggregate supply in a relationship called the Phillips curve. The Phillips curve says that inflation depends on expected inflation, the deviation of unemployment from its natural rate, and supply shocks. According to the Phillips curve, policymakers who control aggregate demand face a short-run tradeoff between inflation and unemployment.
3. If expected inflation depends on recently observed inflation, then inflation has inertia, which means that reducing inflation requires either a beneficial supply shock or a period of high unemployment and reduced output. If people have rational expectations, however, then a credible announcement of a change in policy might be able to influence expectations directly and, therefore, reduce inflation without causing a recession.
4. Most economists accept the natural-rate hypothesis, according to which fluctuations in aggregate demand have only short-run effects on output and unemployment. Yet some economists have suggested ways in which recessions can leave permanent scars on the economy by raising the natural rate of unemployment.

KEY CONCEPTS

Sticky-wage model	Adaptive expectations	Rational expectations
Imperfect-information model	Demand-pull inflation	Natural-rate hypothesis
Sticky-price model	Cost-push inflation	Hysteresis
Phillips curve	Sacrifice ratio	

QUESTIONS FOR REVIEW

1. Explain the three theories of aggregate supply. On what market imperfection does each theory rely? What do the theories have in common?
2. How is the Phillips curve related to aggregate supply?
3. Why might inflation be inertial?
4. Explain the differences between demand-pull inflation and cost-push inflation.
5. Under what circumstances might it be possible to reduce inflation without causing a recession?
6. Explain two ways in which a recession might raise the natural rate of unemployment.

PROBLEMS AND APPLICATIONS

1. Consider the following changes in the sticky-wage model.

a. Suppose that labor contracts specify that the nominal wage be fully indexed for inflation. That is, the nominal wage is to be adjusted to fully compensate for changes in the consumer price index. How does full indexation alter the aggregate supply curve in this model?

b. Suppose now that indexation is only partial. That is, for every increase in the CPI, the nominal wage rises, but by a smaller percentage. How does partial indexation alter the aggregate supply curve in this model?

2. In the sticky-price model, describe the aggregate supply curve in the following special cases. How do these cases compare to the short-run aggregate supply curve we discussed in Chapter 9?

a. No firms have flexible prices ($s = 1$).

b. The desired price does not depend on aggregate output ($a = 0$).

3. Suppose that an economy has the Phillips curve

$$\pi = \pi_{-1} - 0.5(u - 0.06).$$

a. What is the natural rate of unemployment?

b. Graph the short-run and long-run relationships between inflation and unemployment.

c. How much cyclical unemployment is necessary to reduce inflation by 5 percentage points? Using Okun's law, compute the sacrifice ratio.

d. Inflation is running at 10 percent. The Fed wants to reduce it to 5 percent. Give two scenarios that will achieve that goal.

4. According to the rational-expectations approach, if everyone believes that policymakers are committed to reducing inflation, the cost of reducing inflation—the sacrifice ratio—will be lower than if the public is skeptical about the policymakers' intentions. Why might this be true? How might credibility be achieved?

5. Assume that people have rational expectations and that the economy is described by the sticky-wage or sticky-price model. Explain why each of the following propositions is true:

a. Only unanticipated changes in the money supply affect real GDP. Changes in the money supply that were anticipated when wages and prices were set do not have any real effects.

b. If the Fed chooses the money supply at the same time as people are setting wages and prices, so that everyone has the same information about the state of the economy, then monetary policy cannot be used systematically to stabilize output. Hence, a policy of keeping the money supply constant will have the same real effects as a policy of adjusting the money supply in response to the state of the economy. (This is called the *policy irrelevance proposition*.)

c. If the Fed sets the money supply well after people have set wages and prices, so the Fed has collected more information about the state of the economy, then monetary policy can be used systematically to stabilize output.

6. Suppose that an economy has the Phillips curve

$$\pi = \pi_{-1} - 0.5(u - u^n),$$

and that the natural rate of unemployment is given by an average of the past two years' unemployment:

$$u^n = 0.5(u_{-1} + u_{-2}).$$

a. Why might the natural rate of unemployment depend on recent unemployment (as is assumed in the preceding equation)?

b. Suppose that the Fed follows a policy to reduce permanently the inflation rate by 1 percentage point. What effect will that policy have on the unemployment rate over time?

c. What is the sacrifice ratio in this economy? Explain.

d. What do these equations imply about the short-run and long-run tradeoffs between inflation and unemployment?

7. Some economists believe that taxes have an important effect on labor supply. They argue that higher taxes cause people to want to work less and that lower taxes cause them to want to work more. Consider how this effect alters the macroeconomic analysis of tax changes.

- a. If this view is correct, how does a tax cut affect the natural rate of output?
 - b. How does a tax cut affect the aggregate demand curve? The long-run aggregate supply curve? The short-run aggregate supply curve?
 - c. What is the short-run impact of a tax cut on output and the price level? How does your answer differ from the case without the labor-supply effect?
 - d. What is the long-run impact of a tax cut on output and the price level? How does your answer differ from the case without the labor-supply effect?
8. Economist Alan Blinder, whom Bill Clinton appointed to be Vice Chairman of the Federal Reserve, once wrote the following:

The costs that attend the low and moderate inflation rates experienced in the United States and in other

industrial countries appear to be quite modest—more like a bad cold than a cancer on society. . . . As rational individuals, we do not volunteer for a lobotomy to cure a head cold. Yet, as a collectivity, we routinely prescribe the economic equivalent of lobotomy (high unemployment) as a cure for the inflationary cold.¹⁵

What do you think Blinder meant by this? What are the policy implications of the viewpoint Blinder is advocating? Do you agree? Why or why not?

9. Go to the Web site of the Bureau of Labor Statistics (www.bls.gov). For each of the past five years, find the inflation rate as measured by the consumer price index (all items) and as measured by the CPI excluding food and energy. Compare these two measures of inflation. Why might they be different? What might the difference tell you about shifts in the aggregate supply curve and in the short-run Phillips curve?

¹⁵ Alan Blinder, *Hard Heads, Soft Hearts: Tough-Minded Economics for a Just Society* (Reading, Mass.: Addison-Wesley, 1987), 51.

APPENDIX

A Big, Comprehensive Model

In the previous chapters, we have seen many models of how the economy works. When learning these models, it can be hard to see how they are related. Now that we have finished developing the model of aggregate demand and aggregate supply, this is a good time to look back at what we have learned. This appendix sketches a large model that incorporates much of the theory we have already seen, including the classical theory presented in Part II and the business cycle theory presented in Part IV. The notation and equations should be familiar from previous chapters.

The model has seven equations:

$Y = C(Y - T) + I(r) + G + NX(\epsilon)$	IS: Goods Market Equilibrium
$M/P = L(i, Y)$	LM: Money Market Equilibrium
$NX(\epsilon) = CF(r - r^*)$	Foreign Exchange Market Equilibrium
$i = r + \pi^e$	Relationship Between Real and Nominal Interest Rates
$\epsilon = eP/P^*$	Relationship Between Real and Nominal Exchange Rates
$Y = \bar{Y} + \alpha(P - P^e)$	Aggregate Supply
$\bar{Y} = F(\bar{K}, \bar{L})$	Natural Rate of Output

These seven equations determine the equilibrium values of seven endogenous variables: output Y , the natural rate of output \bar{Y} , the real interest rate r , the nominal interest rate i , the real exchange rate ϵ , the nominal exchange rate e , and the price level P .

There are many exogenous variables that influence these endogenous variables. They include the money supply M , government purchases G , taxes T , the capital stock K , the labor force L , the world price level P^* , and the world real interest rate r^* . In addition, there are two expectational variables: the expectation of future inflation π^e and the expectation of the current price level formed in the past P^e . As written, the model takes these expectational variables as exogenous, although additional equations could be added to make them endogenous.

Although mathematical techniques are available to analyze this seven-equation model, they are beyond the scope of this book. But this large model is still useful, because we can use it to see how the smaller models we have examined are related to one another. In particular, *many of the models we have been studying are special cases of this large model*. Let's consider six special cases.

Special Case 1: The Classical Closed Economy Suppose that $P^e = P$, $L(i, Y) = (1/V)Y$, and $CF(r - r^*) = 0$. In words, this means that expectations of the price

level adjust so that expectations are correct, that money demand is proportional to income, and that there are no international capital flows. In this case, output is always at its natural rate, the real interest rate adjusts to equilibrate the goods market, the price level moves parallel with the money supply, and the nominal interest rate adjusts one-for-one with expected inflation. This special case corresponds to the economy analyzed in Chapters 3 and 4.

Special Case 2: The Classical Small Open Economy Suppose that $P^e = P$, $L(i, Y) = (1/V)Y$, and $CF(r - r^*)$ is infinitely elastic. Now we are examining the special case when international capital flows respond greatly to any differences between the domestic and world interest rates. This means that $r = r^*$ and that the trade balance NX equals the difference between saving and investment at the world interest rate. This special case corresponds to the economy analyzed in Chapter 5.

Special Case 3: The Basic Model of Aggregate Demand and Aggregate Supply Suppose that α is infinite and $L(i, Y) = (1/V)Y$. In this case, the short-run aggregate supply curve is horizontal, and the aggregate demand curve is determined only by the quantity equation. This special case corresponds to the economy analyzed in Chapter 9.

Special Case 4: The IS-LM Model Suppose that α is infinite and $CF(r - r^*) = 0$. In this case, the short-run aggregate supply curve is horizontal, and there are no international capital flows. For any given level of expected inflation π^e , the level of income and interest rate must adjust to equilibrate the goods market and the money market. This special case corresponds to the economy analyzed in Chapter 10 and 11.

Special Case 5: The Mundell-Fleming Model with a Floating Exchange Rate Suppose that α is infinite and $CF(r - r^*)$ is infinitely elastic. In this case, the short-run aggregate supply curve is horizontal, and international capital flows are so great as to ensure that $r = r^*$. The exchange rate floats freely to reach its equilibrium level. This special case corresponds to the first economy analyzed in Chapter 12.

Special Case 6: The Mundell-Fleming Model with a Fixed Exchange Rate Suppose that α is infinite, $CF(r - r^*)$ is infinitely elastic, and e is fixed. In this case, the short-run aggregate supply curve is horizontal, huge international capital flows ensure that $r = r^*$, but the exchange rate is set by the central bank. The exchange rate is now an exogenous policy variable, but the money supply M is an endogenous variable that must adjust to ensure the exchange rate hits the fixed level. This special case corresponds to the second economy analyzed in Chapter 12.

You should now see the value in this big model. Even though the model is too large to be useful in developing an intuitive understanding of how the economy works, it shows that the different models we have been studying are closely related. Each model shows a different facet of the larger and more realistic model

presented here. In each chapter, we made some simplifying assumptions to make the big model smaller and easier to understand. When thinking about the real world, it is important to keep the simplifying assumptions in mind and to draw on the insights learned in each of the chapters.

MORE PROBLEMS AND APPLICATIONS

1. Let's consider some more special cases of this large model. Starting with the large model, what extra assumptions would you need to yield each of the following models:
 - a. The model of the classical large open economy in the appendix to Chapter 5.
 - b. The Keynesian cross in the first half of Chapter 10.
 - c. The *IS-LM* model for the large open economy in the appendix to Chapter 12.