We saw how the effects of monetary policy on output depend crucially on how expectations respond to monetary policy.

- In Chapter 19 we looked at the effects of monetary policy in an economy open in both goods markets and financial markets.

  We saw how, in an open economy, monetary policy affects spending and output not only through the interest rate, but also through the exchange rate. An increase in money leads both to a decrease in the interest rate and a depreciation, both of which increase spending and output. We saw how, under fixed exchange rates, the central bank gives up monetary policy as a policy instrument.
- In Chapter 20 we discussed the pros and cons of different monetary policy regimes, namely flexible exchange rates versus fixed exchange rates.

  We saw how, under flexible exchange rates, interest rate movements can lead to large changes in exchange rates. We saw how, under fixed exchange rates, speculation can lead to an exchange rate crisis and a sharp devaluation. We discussed the pros and cons of adopting a common currency such as the euro, or even giving up monetary policy altogether through the adoption of a currency board or dollarization.
- In Chapter 21 we looked at the problems facing macroeconomic policy in general, and monetary policy in particular.

  We saw that uncertainty about the effects of policy should lead to more cautious policies. We saw that even well-intentioned policy makers may sometimes not do what is best, and that there is a case to be made for putting restraints on policy makers. We also looked at the benefits of having an independent central bank and appointing a conservative central banker.

In this chapter we extend the analysis to look first at the inflation targeting framework in place before the crisis, and then at the challenges to monetary policy raised by the crisis.

23-2 From Money Targeting to Inflation Targeting

One can think of the goals of monetary policy as twofold: First, to maintain low and stable inflation. Second, to stabilize output around potential—to avoid or at least limit recessions or booms.

Money Targeting

Until the 1980s, the strategy was to choose a target rate of money growth and to allow for deviations from that target rate as a function of activity. The rationale was simple. A low target rate of money growth implied a low average rate of inflation. In recessions, the central bank could increase money growth, leading to a decrease in interest rates and an increase in output. In booms, the central bank could decrease money growth, leading to an increase in interest rates and a slowdown in output.

That strategy did not work well.

First, the relation between money growth and inflation turned out to be far from tight, even *in the medium run*. This is shown in Figure 23-1, which plots 10-year averages of the U.S. inflation rate against 10-year averages of the growth rate of money from 1970 up to the crisis (the way to read the figure: The numbers for inflation and for money growth for 2000 for example are the average inflation rate and the average growth rate of money from 1991 to 2000). The inflation rate is constructed using the consumer price index (CPI) as the price index. The growth rate of nominal money is constructed

Figure 23-1

***M1* Growth and Inflation: 10-Year Averages, 1970 to the crisis**

There is no tight relation between *M1* growth and inflation, even in the medium run.

Source: Series CPIAUSL and M1SL Federal Reserve Economic Data (FRED) http://research.stlouisfed.org/fred2/.

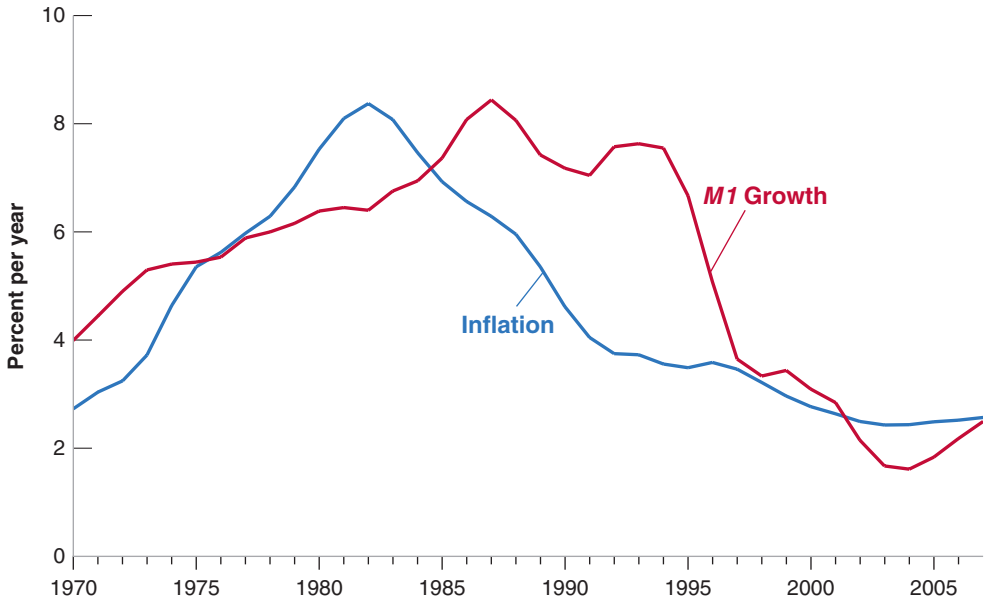

using the sum of currency and checkable deposits, known as ***M1***, as the measure for the money stock. The reason for using 10-year averages should be clear. In the short run, changes in nominal money growth affect mostly output rather than inflation. It is only in the medium run that a relation between nominal money growth and inflation should emerge. Taking 10-year averages of both nominal money growth and inflation is a way of detecting such a medium-run relation. The reason for stopping at the crisis is that, as we saw in Chapter 4, when an economy hits the zero lower bound (which the U.S. economy did at the end of 2008), increases in the money supply no longer have an effect on the policy rate, and by implication, the central bank is no longer able to affect output and inflation; so we want to exclude the period during which the U.S. economy was stuck at the zero lower bound.

Figure 23-1 shows that, for the United States, the relation between *M*1 growth and inflation was not tight. True, both went up in the 1970s, and both came down later. But note how inflation started declining in the early 1980s, whereas nominal money growth remained high for another decade and came down only in the 1990s. Average inflation from 1981 to 1990 was down to 4%, and average money growth over the same period was still running at 7.5%.

Second, the relation between the money supply and the interest rate *in the short run* also turned out also to be unreliable. A given decrease in money growth in response for example to low activity might lead to different effects on the interest rate, making money growth an unreliable instrument to affect demand and output.

Both problems, namely the poor relation between money growth and inflation in the medium run, and the poor relation of the interest rate to the money supply in the short run, had the same origin, namely *shifts in the demand for money*. An example will help here. Suppose, as the result of the introduction of credit cards, people decide to hold only half the amount of money they held before; in other words, the real demand for money decreases by half. In the short run, at a given price level, this large decrease in the demand for money will lead to a large decrease in the interest rate. In other words, we will see a large decrease in the interest rate with no change in the money supply. In the medium run, at a given interest rate, the price level will adjust, and the real money stock will eventually decrease by half. For a given nominal money stock, the price level

will eventually double. So, even if the nominal money stock remains constant, there will still be a period of inflation as the price level doubles. During this period, there will be no tight relation between nominal money growth (which is zero) and inflation (which would be positive).

Throughout the 1970s and the 1980s, these frequent and large shifts in money demand created serious problems for central banks. They found themselves torn between trying to keep a stable target for money growth and staying within announced bands (to maintain credibility), or adjusting to shifts in money demand (to stabilize output in the short run and inflation in the medium run). Starting in the early 1990s, a dramatic rethinking of monetary policy took place based on targeting inflation rather than money growth, and the use of an interest rate rule. Let's look at it more closely.

◀ From Chapter 5, equation (5.3): The real money supply (the left side) must be equal to the real demand for money (the right side):

$$\frac{M}{P} = YL(i)$$

If, as a result of the introduction of credit cards, the real demand for money halves, then

$$\frac{M}{P} = \frac{1}{2}YL(i)$$

In the short run, P does not move, and so the interest rate must adjust. In the medium run, P adjusts. For a given level of output and a given interest rate, M/P must halve. Given M, this implies P must double.

Inflation Targeting

If one of the main goals of the central bank is to achieve low and stable inflation, why not target inflation directly rather than money growth? And if the way to affect activity in the short run is to rely on the effect of the interest rate on spending, why not focus directly on the interest rate rather than on money growth? This is the reasoning which led to the elaboration of inflation targeting. Central banks committed to achieving a target inflation rate. And they decided to use the interest rate as the instrument to achieve it. Let's look at both parts of the strategy:

Committing to a given inflation target *in the medium run* is hardly controversial. Trying to achieve a given inflation target *in the short run* would appear to be much more controversial. Focusing exclusively on inflation would seem to eliminate any role monetary policy could play in reducing output fluctuations. But in fact, this is not the case.

To see why, return to the Phillips curve relation between inflation, π_t, expected inflation, π_t^e, and the deviation of the unemployment rate, u_t, from the natural rate of unemployment, u_n (equation (8.9)):

$$\pi_t = \pi_t^e - \alpha(u_t - u_n)$$

Let the inflation target be π^*. Assume that, thanks to the central bank's reputation, this target is credible, so that people expect inflation to be equal to the target. The relation becomes:

$$\pi_t = \pi^* - \alpha(u_t - u_n)$$

Note that, if the central bank is able to hit its inflation target exactly, so $\pi_t = \pi^*$, unemployment will be equal to its natural rate. By targeting and achieving a constant rate of inflation in line with inflation expectations, the central bank also keeps unemployment at the natural rate, and by implication keeps output at potential.

◀ $0 = -\alpha(u_t - u_n) \Rightarrow u_t = u_n$.

Put strongly: Even if policy makers did not care about inflation per se (they do) but cared only about output, inflation targeting would still make sense. Keeping inflation stable is a way of keeping output at potential. This result has been dubbed the **divine coincidence**. With a Phillips curve of the form given in equation (8.9), there is no conflict between keeping inflation constant and keeping output at potential. A focus on keeping stable inflation is thus the right approach to monetary policy, both in the short and the medium run.

This result is a useful benchmark, but it is too strong. Life is not that nice. The main objection is that, as we saw in Chapter 8, the Phillips curve relation is far from an exact relation. There are times when inflation may be above target and output below potential, reintroducing a trade-off between the two goals. The central bank then has to decide whether to focus on decreasing inflation and adopt a tighter monetary policy, or to focus on increasing output and adopt a more expansionary monetary policy. So, although

some central banks have been given a single mandate, namely stable and low inflation, others, such as the U.S. Fed, have a dual mandate, achieving both stable and low inflation and maintaining output close to potential. Also, all central banks have adopted what is called **flexible inflation targeting**. For the reasons we discussed in Chapter 21, uncertainty about the effects of the interest rate on output and in turn on inflation, central banks do not try to return to target inflation right away. Rather they adjust the interest rate to return to the target inflation rate over time. We now turn to the interest rate rule associated with inflation targeting.

The Interest Rate Rule

Inflation is not under the direct control of the central bank. The policy rate is. Thus, the question is how to set the policy rate so as to achieve the target rate of inflation. The answer is a simple one. When inflation is higher than the target, increase the policy rate to decrease the pressure on prices; when it is below the target rate of inflation, decrease the policy rate. With this in mind, in the 1990s, John Taylor, from Stanford University, suggested the following rule for the policy rate, a rule now known as the **Taylor rule**:

- Let π_t be the rate of inflation and π^* be the target rate of inflation.
- Let i_t be the policy rate, that is, the nominal interest rate controlled by the central bank, and i^* be the target nominal interest rate—the nominal interest rate associated with the neutral rate of interest, r_n, and the target rate of inflation, π^*, so $i^* = r_n + \pi^*$.
- Let u_t be the unemployment rate and u_n be the natural unemployment rate.

Think of the central bank as choosing the nominal interest rate, i. (Recall, from Chapter 4, that, through open market operations, and ignoring the liquidity trap, the central bank can achieve any short-term nominal interest rate that it wants.) Then Taylor argued, the central bank should use the following rule:

$$i_t = i^* + a(\pi_t - \pi^*) - b(u_t - u_n)$$

where a and b are positive coefficients chosen by the central bank.

Let's look at what the rule says:

- If inflation is equal to target inflation ($\pi_t = \pi^*$) and the unemployment rate is equal to the natural rate of unemployment ($u_t = u_n$), then the central bank should set the nominal interest rate, i_t, equal to its target value, i^*. This way, the economy can stay on the same path, with inflation equal to the target inflation rate and unemployment equal to the natural rate of unemployment.
- If inflation is higher than the target ($\pi_t > \pi^*$), the central bank should increase the nominal interest rate, i_t, above i^*. This higher interest rate will lead to an increase unemployment, and this increase in unemployment will lead to a decrease in inflation. The coefficient a should therefore reflect how much the central bank cares about inflation. The higher a, the more the central bank will increase the interest rate in response to inflation, the more the economy will slow down, the more unemployment will increase, and the faster inflation will return to the target inflation rate.

  In any case, as Taylor pointed out, a should be larger than one. Why? Because what matters for spending is the real interest rate, not the nominal interest rate. When inflation increases, the central bank, if it wants to decrease spending and output, must increase the *real* interest rate. In other words, it must increase the nominal interest rate more than one-for-one with inflation.
- If unemployment is higher than the natural rate of unemployment ($u_t > u_n$), the central bank should decrease the nominal interest rate. The lower nominal interest rate will lead to an increase output, leading to a decrease in unemployment.

Some economists argue that the increase in U.S. inflation in the 1970s was due to the fact that the Fed increased the nominal interest rate less than one-for-one with inflation. The result, they argue, was that an increase in inflation led to a decrease in the real interest rate, which led to higher demand, lower unemployment, more inflation, a further decrease in the real interest rate, and so on. ▶

The coefficient b should reflect how much the central bank cares about unemployment. The higher b, the more the central bank will be willing to deviate from target inflation to keep unemployment close to the natural rate of unemployment.

In stating this rule, Taylor did not argue that it should be followed blindly. Many other events, such as an exchange rate crisis or the need to change the composition of spending on goods, and thus the mix between monetary policy and fiscal policy, justify changing the nominal interest rate for other reasons than those included in the rule. But he argued, the rule provided a useful way of thinking about monetary policy. Once the central bank has chosen a target rate of inflation, it should try to achieve it by adjusting the nominal interest rate. The rule it should follow should take into account not only current inflation but also current unemployment.

The logic of the rule was convincing, and, by the mid-2000s, in advanced economies, most central banks had adopted some form of inflation targeting, that is, the choice of an inflation target together with the use of an interest rule.

Then the crisis came and raised many questions, from the choice of the inflation target, to what to do when the interest rate suggested by the interest rule reaches the zero lower bound, to whether and how the central bank should worry about financial stability in addition to inflation and activity. The next section discusses the choice of the inflation target, and the following sections discuss other questions raised by the crisis.

23-3 The Optimal Inflation Rate

Table 23-1 shows how inflation steadily decreased in advanced economies from the early 1980s. In 1981, average inflation in the OECD was 10.5%; in 2014, it was down to 1.7%. In 1981, only two countries (out of the 24 OECD members at the time) had an inflation rate below 5%; in 2014, the number had increased to 33 out of 34.

◀ The country with inflation above 5% was Turkey (8.8%).

Before the crisis, most central banks had aimed for an inflation rate of about 2%. Was this the right goal? The answer depends on the costs and benefits of inflation.

The Costs of Inflation

We saw in Chapter 22 how very high inflation, say a rate of 30% per month or more, can disrupt economic activity. The debate in advanced economies today, however, is not about the costs of inflation rates of 30% or more per month. Rather, it centers on the advantages of, say, 0% versus, say, 4% inflation per year. Within that range, economists identify four main costs of inflation: (1) shoe-leather costs, (2) tax distortions, (3) money illusion, and (4) inflation variability.

Shoe-Leather Costs

Recall that in the medium run, a higher inflation rate leads to a higher nominal interest rate, and so to a higher opportunity cost of holding money. As a result, people decrease

Table 23-1 Inflation Rates in the OECD, 1981–2014

Year	1981	1990	2000	2010	2014
OECD average*	10.5%	6.2%	2.8%	1.2%	1.7%
Number of countries with inflation below 5%**	2/24	15/24	24/27	27/30	33/34

*Average of GDP deflator inflation rates, using relative GDPs measured at PPP prices as weights.
**The second number denotes the number of member countries at the time.

their money balances by making more trips to the bank—thus the expression **shoe-leather costs**. These trips would be avoided if inflation were lower and people could be doing other things instead, such as working more or enjoying leisure.

During hyperinflations, shoe-leather costs become indeed quite large. But their importance in times of moderate inflation is limited. If an inflation rate of 4% leads people to go to the bank, say, one more time every month, or to do one more transaction between their money market fund and their checking account each month, this hardly qualifies as a major cost of inflation.

Tax Distortions

The second cost of inflation comes from the interaction between the tax system and inflation.

Consider, for example, the taxation of capital gains. Taxes on capital gains are typically based on the change in the price in dollars of the asset between the time it was purchased and the time it is sold. This implies that the higher the rate of inflation, the higher the tax. An example will make this clear:

- Suppose inflation has been running at $\pi\%$ a year for the last 10 years.
- Suppose also that you bought your house for \$50,000 10 years ago, and you are selling it today for \$50,000 times $(1 + \pi\%)^{10}$; so its real value is unchanged.
- If the capital-gains tax is 30%, the *effective tax rate* on the sale of your house—defined as the ratio of the tax you pay to the price for which you sell your house—is

The numerator of the fraction equals the sale price minus the purchase price. The denominator is the sale price. ▶

$$(30\%)\,\frac{50{,}000(1 + \pi\%)^{10} - 50{,}000}{50{,}000(1 + \pi\%)^{10}}$$

- Because you are selling your house for the same real price at which you bought it, your real capital gain is zero, so you should not be paying any tax. Indeed, if $\pi = 0$—if there has been no inflation—then the effective tax rate is 0. But if, for example, $\pi = 4\%$, then the effective tax rate is 9.7%: Despite the fact that your real capital gain is zero, you end up paying a high tax.

The problems created by the interactions between taxation and inflation extend beyond capital-gains taxes. Although we know that the real rate of return on an asset is the real interest rate, not the nominal interest rate, income for the purpose of income taxation includes nominal interest payments, not real interest payments. Or to take yet another example, until the early 1980s in the United States, the income levels corresponding to different income-tax rates were not increased automatically with inflation. As a result, people were pushed into higher tax brackets as their nominal income—but not necessarily their real income—increased over time, an effect known as **bracket creep**.

Some economists argue that the costs of bracket creep were much larger. As tax revenues steadily increased, there was little pressure on the government to control spending. The result, they argue, was an increase in the overall size of the government in the 1960s and 1970s far beyond what would have been desirable. ▶

You might argue this cost is not a cost of inflation per se, but rather the result of a badly designed tax system. In the house example we just discussed, the government could eliminate the problem if it *indexed* the purchase price to the price level—that is, it adjusted the purchase price for inflation since the time of purchase—and computed the tax on the difference between the sale price and the adjusted purchase price. Under this computation, there would be no capital gains and therefore no capital-gains tax to pay. But because tax codes around the world rarely define the tax base in real terms, the inflation rate matters and leads to distortions.

Money Illusion

The third cost comes from **money illusion**—the notion that people appear to make systematic mistakes in assessing nominal versus real changes in incomes and interest rates. A number of computations that would be simple when prices are stable become