Chapter 19

Optimal Monetary Policy

As we have discussed, expansionary monetary policies include decreases in the Fed funds rate and unexpected growth in the money supply. In the U.S., such expansionary monetary policies have tended to produce real expansions in output and increases in inflation. Conversely, contractionary monetary policies have tended to produce real contractions in output and decreases in inflation. In Chapter 18 Barro claims that these effects have been quite moderate, but recent empirical work lends support to the opposite view, that monetary shocks can have large effects on real variables in the short run.

Everyone agrees that expansionary monetary policies tend to lead to increases in inflation, while contractionary policies produce decreases in inflation. At this broad level, monetary policy would appear to be a matter of trading off inflation and output. Since unemployment tends to decrease as output increases, this is often cast a choice between inflation and unemployment. The empirical relationship between the two is called the *Phillips curve.*

In the U.S., as in most countries, monetary policy is under the control of the government. This immediately raises the question of how best to conduct monetary policy. As we shall see, this not so much a question of when and how to time expansions and contractions of the money supply, as economists used to think, as it is a question of what the private sector predicts the government will do and how the government can influence those predictions.

Before we can think fruitfully about monetary policy, we will have to have a reasonable model of how monetary shocks can influence the real economy. Our model will be a simplification of the seminal paper by Robert E. Lucas, Jr, "Expectations and the Neutrality of Money". In that model, the private sector is divided into different industries (called "islands") which observe only the price for their own product. This price is made up of a general price level (unobserved) and an industry-specific shock (also unobserved). The private sector has some forecast about inflation (never mind for the moment its origin) and uses this to derive an estimate of the industry specific shock it faces. If the estimated shock is high, the private sector increases production. If it is low, the private sector decreases production. The government chooses an inflation rate. An unexpected monetary expansion will produce a temporary increase in output. Thus Lucas's model highlights the role of *expectations* in the conduct of monetary policy.

In Lucas's model, only unanticipated changes in the price level have real effects. If a monetary expansion is completely expected, it has no real effects. This points to something quite important in the real conduct of monetary policy: only surprises matter. Moreover, the private sector does not enjoy being surprised, even if the monetary surprise produced a temporary boom. An older tradition in macroeconomics holds that governments should try to manipulate the money supply to cushion supply and demand disruptions. The central lesson from Lucas's research is that governments should instead strive to minimize the uncertainty surrounding monetary policy.

We then move away from the specific form of the Phillips curve derived from Lucas's model and start using a simple generalization in which inflation, inflationary expectations and unemployment are all related by a very simple formula. The government will have some preferences (and thus indifference curves) over unemployment and inflation (both will be bad), and monetary policy, if we ignore how expectations are formed, can be seen as a simple choice of unemployment and inflation.

Once we begin modeling the formation of expectations, we will see that the ability of the government to commit credibly to a particular inflationary path is critical. We will model explicitly a two-person game between the private sector and the government. With a so-called commitment device, the government will be able to play the *Ramsey strategy* and realize the *Ramsey outcome.* Recall the Ramsey optimal tax problem from Chapter 14. In that chapter we assumed that the government could commit to a particular tax sequence, hence the term "Ramsey". We did not consider what would happen if the government could not commit to a particular tax sequence. In this chapter we will see that without a commitment device, the government and the private sector will play *Nash strategies* and achieve the *Nash outcome.* The fundamental result of this chapter is that Ramsey is better than Nash. Both the government and the private sector are better off in the Ramsey outcome than in the Nash outcome. Indeed, under certain circumstances, the Nash outcome involves (temporarily) high inflation and high unemployment, the so-called "stagflationary" episode of the 1970s. At the time, stagflation was blamed on an oil price shock. We have to reconsider, and say that possibly it was the result of a lack of credible commitment by the government.

The theory in this chapter will give us an explanation for the "pain" associated with fighting inflation. There is a powerful maintained assumption in the media that policies that are anti-inflationary require some sacrifice of real output. As we shall see, when the private sector has formed strong expectations about continued high inflation, confounding those expectations with sudden, unexpected, low inflation can have a severe cost in terms of real output. This is not a reason to oppose anti-inflationary policies, it is a reason to campaign for a credible commitment to low inflation.

Finally, it is worth noting that, in this chapter, we will ignore the government budget constraint. In Chapter 18 we were very concerned about the relationship between persistent government budget deficits and inflation. In this chapter we will assume the government budget is more or less in balance, and that the government does not particularly need the seignorage revenue generated by high inflation. This is a safe assumption when thinking about inflation in the U.S. In thinking about inflation across different countries, though, the analysis of Chapter 18 is probably more appropriate in countries, like Brazil, that experience persistent inflation and large budget deficits. The "pain" of fighting inflation in those countries is the pain of raising direct taxes and decreasing government spending.

Both this chapter and Chapter 18 highlight the importance of credible government policies. In Chapter 18, to stop hyperinflations the government had to credibly commit to balancing its fiscal books. In this chapter, to prevent milder inflations, the government will have to credibly commit to keep its hands off of the monetary spigot. In both cases there is a role for international institutions as commitment devices.

19.1 The Model of Lucas (1972)

In this section we consider a simplified version of the important model of Lucas. We are going to get a relationship between the anticipated price level, the actual price level and something that looks like unemployment. We will use this relationship to argue for a particular functional form for the Phillips curve. We will not derive precisely a Phillips curve since our model is going to be static, to keep the exposition simple. The dynamic generalization is very elegant, and the interested reader is referred directly to the Lucas paper.

This model turns on the decisions made by many separated industries in the private sector. These industries cannot communicate with one another about prices. They will hire labor according to their estimate of the true state of demand for their product.

Let Q_i be output in industry i. Assume that all industries use only one input, labor. Let L_i be the number of workers hired in industry i. Assume that all industries have the common production function:

$$Q_i = L_i^{\alpha},$$

where the technology parameter α satisfies $0 < \alpha < 1$. Assume that all workers are paid the common wage of unity for their unit of labor supplied. To produce an output Q_i therefore requires labor input (and total costs) of $L_i^{1/\alpha}$. Thus the cost function in industry i is:

$$\text{Total Cost}(Q_i) = L_i^{\frac{1}{\alpha}}.$$

In industry i there will be a price P_i for that industry's output. It is known that this price is made up of two parts: a general price level P, common to all industries, and a shock term Z_i specific to industry i. These terms are related by the price equation:

$$P_i = PZ_i. \tag{19.1}$$

The shock term Z_i gives the *real price* of output in industry i. The general price level P will not be revealed until the end of the period, since the industries are on islands and cannot communicate during production.

All private-sector industries begin the period with a common *forecast* of P, which we denote by P^e. Thus an industry i's best estimate of its real price Z_i is:

$$Z_i^e = \frac{P_i}{P^e}. \tag{19.2}$$

Recall that industry i only observes P_i.

Equilibrium in industry i, assuming that it is competitive, requires that marginal cost equal estimated real price Z_i^e. Since we know the total cost curve, marginal cost must just be its derivative with respect to output Q_i. That is, equilibrium requires:

$$\frac{1}{\alpha} Q_i^{\frac{1}{\alpha}-1} = Z_i^e.$$

We can solve this to produce the equilibrium demand for labor conditional on the estimated shock Z_i^e:

$$L_i = (\alpha Z_i^e)^{\frac{1}{1-\alpha}}. \tag{19.3}$$

As expected, industries will demand more labor if they estimate that demand for their product is unusually strong (if Z_i^e is large).

The estimated shock Z_i^e is comprised of two parts: the known estimate of the price level P^e and industry-specific price level P_i, related by equation (19.2). Thus we can substitute from that equation into equation (19.3) to find the industry-specific demand for labor conditional on P^e and P_i:

$$L_i = \left(\alpha \frac{P_i}{P^e}\right)^{\frac{1}{1-\alpha}}.$$

Now we take logarithms of both sides. From now on, let lower-case variables denote logarithms. Thus $\ell_i = \ln(L_i)$ is given by:

$$\ell_i = \frac{1}{1-\alpha} \ln\left(\alpha \frac{P_i}{P^e}\right) = \frac{1}{1-\alpha} \ln(\alpha) + \frac{1}{1-\alpha}(p_i - p^e),$$

substitute $z_i + p$ for p_i from equation (19.1) above, and let $A = [1/(1-\alpha)]\ln(\alpha)$ to produce:

$$\ell_i = A + \frac{1}{1-\alpha}(z_i + p - p^e). \tag{19.4}$$

Equation (19.4) captures the log of labor demand as a function of the (log of the) shock, the common price level p and the common price forecast p^e.

Define u to be the "not employed rate" (not quite the unemployment rate, but something close).[1] If N is the total workforce, and $n = \ln(N)$, then define u as:

$$u = n - \sum_i \ell_i.$$

Assume for a moment that there are only two industries. Now:

$$u = n - 2A + \frac{2}{1-\alpha}(p^e - p) - \frac{2}{1-\alpha}(z_1 + z_2).$$

Define further:

$$u^\star = n - 2A,$$

where $u^\star$ is something like the natural rate of not-employment,

$$\begin{aligned} \varepsilon &= -\frac{2}{1-\alpha}(z_1 + z_2), \text{ and:} \\ \gamma &= \frac{2}{1-\alpha}. \end{aligned}$$

Now we can write the aggregate not-employment rate as:

$$u = u^\star + \gamma(p^e - p) + \varepsilon. \tag{19.5}$$

We will use some version of this equation throughout this chapter.

From the point of view of the government, the common price level p is a control variable. The government picks a level for p with monetary policy. Notice what equation (19.5) says about the relationship of unemployment (or not-employment), the price level and the forecast price level: unemployment is decreasing in the price level p but increasing in the forecast price level p^e. From the point of view of private industry, if the actual price level exceeds the forecast price level, $p > p^e$, the industry has produced too much and suffers losses as a result. From the point of view of the government, if $p > p^e$, it can stimulate a one-period boom in which unemployment is below its natural level.

19.2 Monetary Policy and the Phillips Curve

For the rest of this chapter we will be using a modified version of equation (19.5). Assume that:

$$u = u^\star + \gamma(\pi^e - \pi). \tag{19.6}$$

[1]The unemployment rate is $1 - (1/N)\sum_i L_i$ which doesn't translate well into logarithms.

Here u is the unemployment rate, $u^\star$ is the "natural rate" of unemployment, π^e is the *expected inflation rate* and π is the *actual inflation rate.* The natural rate of unemployment is the level of unemployment when inflation is perfectly anticipated, so no industries are fooled into thinking that relative demand is unusually high or low. The slope of this Phillips curve is $-\gamma$ where we assume $\gamma > 0$ (monetary expansions reduce unemployment). If we think that there is uncertainty about the state of the real economy, we can add a mean zero shock term, ε, to produce:

$$u = u^\star + \gamma(\pi^e - \pi) + \varepsilon.$$

For the most part we will assume that the monetary authority knows the state of the real economy with certainty.

In Figure (19.1) we plot Phillips curves with two different values of expected inflation π^e, a low value in which the expected inflation rate is zero, and a high value, in which the expected inflation rate is 8.3%. The dotted line gives the natural rate of unemployment (here $u^\star = 5\%$), and $\gamma = 0.3$. Notice that when inflationary expectations are high, to achieve any given unemployment rate requires a higher inflation rate, and to achieve zero inflation requires an unemployment rate well above the natural rate.

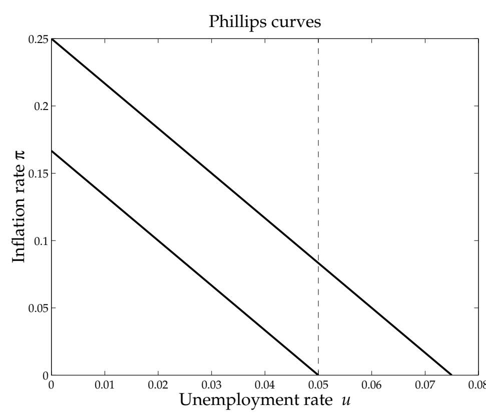

Figure 19.1: Phillips curves under two different expectations about inflation. The bottom curve assumes $\pi^e = 0$ and the top curve assumes $\pi^e = 0.0833$. The dotted line gives the natural rate of unemployment.

Monetary Policy with Fixed Expectations

Assume that the government (or which ever arm of the government controls monetary policy) has a utility function over unemployment and inflation of $V^g(u, \pi)$ given by:

$$V^g(u, \pi) = -u^2 - \pi^2.$$

That is, the government dislikes unemployment and inflation equally. We will assume this form for V for the rest of the chapter, so it's worth mentioning that the Federal Reserve Board is, by law, supposed to balance the twin goals of full employment and price stability. Thus this utility function seems to be written in law.

If we assume that π^e is given exogenously and fixed, we can substitute the Phillips curve in equation (19.6) into the government's utility function above to produce a maximization problem. Thus if the private sector has fixed expectations about the inflation rate given by π^e, then the government's optimal choice of inflation π is given by:

$$\max_{\pi} \left\{ -[u^\star + \gamma(\pi^e - \pi)]^2 - \pi^2 \right\}.$$

The first-order condition with respect to inflation π is:

$$2\gamma[u^\star + \gamma(\pi^e - \pi)] - 2\pi = 0.$$

We can solve this for π to get the optimal inflation choice when expected inflation is fixed at π^e and the natural rate is $u^\star$ (call it $\pi^\star(\pi^e)$):

$$\pi^\star(\pi^e) = \frac{\gamma}{1+\gamma^2}(u^\star + \gamma\pi^e).$$

We can plug $\pi^\star(\pi^e)$ into the Phillips curve in equation (19.6) to produce the associated unemployment rate, $u_0(\pi^e)$:

$$u_0(\pi^e) = \frac{1}{1+\gamma^2}u^\star + \frac{\gamma}{1+\gamma^2}\pi^e.$$

Notice that if π^e is "small" that $u_0(\pi^e)$ will lie below $u^\star$. The government trades off some inflation for a lower unemployment rate.

We plot $\pi^\star(\pi^e)$ in Figure (19.2) below. Notice that for low values of expected inflation, π^e, the government chooses inflation rates above expectations and for high values of π^e, the government chooses inflation rates below expectations. At one unique expected inflation rate, the government's best response is to choose an actual inflation rate exactly equal to the expected inflation rate. This will play a special role, as we shall see.

Two Stories About Inflationary Expectations

We are not yet ready to discuss the strategic interactions between the private sector and the government that determine inflationary expectations. However, we can study the outcomes under two different stories about inflationary expectations. These will help us to

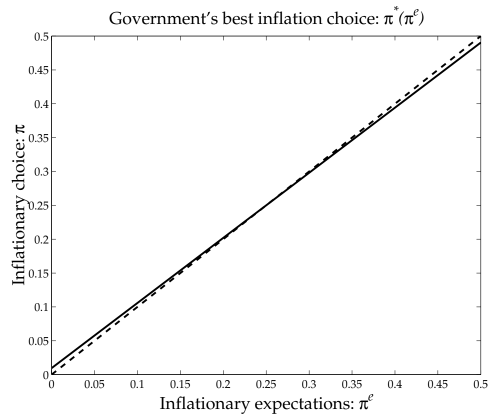

Figure 19.2: Government's optimal choice of inflation $\pi^\star$ as a function of different expectations of the private sector (solid) and the 45-degree line (dotted). Where the two lines cross is the inflationary expectation at which the government's best response is to meet expected inflation. In other words: the Nash inflation level.

think about the government's problem. First, we will assume that expectations are fixed, but that the private sector knows the government's maximization problem. If this is the case, then the private sector will set expectations to a unique value such that the government chooses to set inflation at exactly the same value the private sector anticipated. Second, we will assume that expected inflation exactly equals actual inflation in all cases. The private sector has a crystal ball (or a spy) which informs it precisely of the government's inflationary plan, no matter what the government picks.

Imagine for a moment that the private sector understands the government's maximization problem and correctly anticipates inflation. That is, assume that inflationary expectations satisfy:

$$\pi^e = \pi^\star(\pi^e).$$

From Figure (19.2) below, we see that there is exactly one such expected inflation rate. Expanding produces:

$$\pi^e = \frac{\gamma}{1+\gamma^2}(u^\star + \gamma\pi^e).$$

We can solve for this special value of π^e, call it π_1, to get:

$$\pi_1 = \gamma u^\star, \tag{19.7}$$

where π_1 is unique inflation rate such that when expectations satisfy $\pi^e = \pi_1$, the government's inflation target is also π_1. The associated unemployment rate is:

$$u_1 = u^\star,$$

since $\pi^e = \pi$. Thus the unemployment rate is at the natural rate $u^\star$ and inflation is relatively high at π_1. This will be the *Nash equilibrium* in inflation (as we shall see below).

Now imagine that the government is forced by law to correctly announce its inflation target each period. The private sector anticipates this and sets $\pi^e = \pi$. Thus the Phillips curve in equation (19.6) becomes:

$$u = u^\star + \gamma(\pi - \pi) = u^\star.$$

In other words, inflation does not affect output. If this is the case, the government chooses an inflation rate of zero (since inflation is costly and now provides no benefit), and the unemployment rate again goes to the natural rate. This will turn out to be the *Ramsey equilibrium* as we shall see below.

Contrast the Ramsey and the Nash equilibria. Both produced the natural rate of unemployment, but the Nash equilibrium also had a high inflation rate. Thus the government and the population are better off if the government is able to announce the inflation rate and be believed. As we shall see below, unfortunately, when the private sector expects inflation to be low, there is a temptation for the government to inflate.

Ramsey Monetary Policy

This last example was the Ramsey problem. If the government can credibly commit to a particular inflation rate, the private sector responds by setting inflationary expectations to the announced inflation target. As a result, the government announces an inflation target of zero, and the result is the natural rate of unemployment. What are some commitment devices? By making the monetary authority completely independent of the fiscal authority it can be insulated from political pressure. Further, if the central banker has a reputation for being an unpleasant misanthrope who cares only about defeating inflation, the private sector can become convinced over time that in fact the central bank will set $\pi = 0$ for all time.

Indeed, one reading of the deeply unpleasant recession in the early 1980s is that the private sector had to be convinced of the new central banker's commitment to low inflation. Paul Volcker arrived as Chairman of Federal Reserve Board at a time of high inflation and high unemployment. He announced that there would be low inflation in the future. The private sector did not adjust its expectations, but Volcker followed through on his promise. The result was the unusual case in which inflationary expectations exceeded actual inflation, that is $\pi^e > \pi$. As a result, unemployment shot above its natural rate in one of the deeper recessions of the century. After two years of this treatment, the private sector adjusted its expectations, convinced that Mr. Volcker was committed to low inflation.

Other countries, without the benefit of the tradition of anti-inflationary policies of the Fed to reassure the private sector, will completely let go of the reins of monetary policy. In Hong Kong, for example, the local currency is pegged to the U.S. dollar in an arrangement known as a *currency board*. For every 7.8 Hong Kong dollars issued, one U.S. dollar must be placed on deposit, so the currency is fully backed. The Hong Kong government cannot print money. Thus the exchange rate is immutably fixed, and there can be no depreciation of the local currency against the U.S. dollar. Countries will go to great lengths to convince the private sector that they are really committed to low inflation. They have to work so hard at it, we shall see, precisely because, if expectations are low, there is always a temptation to inflate.

19.3 Optimal Monetary Policy without Commitment: The Nash Problem

In this section we will explicitly model the strategic interaction between the private sector and the government when forming inflationary expectations. We will force the government to choose from only two possible inflation levels, and the private sector to pick from only two possible inflationary expectations. The results we derive here generalize to the case in which both choose from continuous distributions.

Inflation π can only take on one of two values: $\{0, \pi_1\}$. That is, inflation can be zero or the high level we derived in equation (19.7). The private sector expects π^e which can also only take on the values $\{0, \pi_1\}$, since it wouldn't make sense for the private sector to anticipate inflation rates that the government can't pick.

There are four possible combinations of expected and actual inflation, $\{\pi^e, \pi\}$. At each one of these four combinations we will specify the payoff to the private sector and to the government. These payoffs will be known by both players. We will look for a *Nash equilibrium*, which is simply a pair of choices (one for the private sector, one for the government) such that, given the other player's choice, no player can do better.

We now consider each of the four possible combinations. Let $V^g(\pi^e, \pi)$ be the payoff to the government and $V^p(\pi^e, \pi)$ be the payoff to the private sector at each possible $\{\pi^e, \pi\}$ combination. We will assume that the private sector suffers a penalty of -1 if it does not correctly forecast the inflation rate and gets a payoff of zero otherwise (this is just a normalization). We assume that the baseline government payoff (at zero inflation and the natural rate of unemployment) is 0, and that otherwise the government dislikes inflation and unemployment. At each of the four possible outcomes, the payoffs of the two players are:

Private Sector	Government	
	$\pi = 0$	$\pi = \pi_1$
$\pi^e = 0$	$V^g = 0, V^p = 0$	$V^g = 1, V^p = -1$
$\pi^e = \pi_1$	$V^g = -1, V^p = -1$	$V^g = -0.5, V^p = 0$

Notice that the government really dislikes $\{\pi^e = \pi_1, \pi = 0\}$; this corresponds to the Volcker play of low inflation when expectations are high. The result is unemployment above the natural rate. Also, the government dislikes (but not as much) $\{\pi^e = \pi_1, \pi = \pi_1\}$; here inflation is high, but unemployment is at the natural rate. The government would prefer to be at $\{\pi^e = 0, \pi = \pi_1\}$; here inflation is unexpectedly high, so unemployment is below the natural rate.

Now let us work through these payoffs to find the Nash equilibrium. If the household plays $\pi^e = 0$, the best response of the government is to set $\pi = \pi_1$. If the household plays $\pi^e = \pi_1$, the best response of the government is to play $\pi = \pi_1$. If the government plays $\pi = 0$ the best response of the household is $\pi^e = 0$, but this is not a Nash equilibrium since, if the household does play $\pi^e = 0$, we saw that the government will want to deviate to $\pi = \pi_1$. If the government plays $\pi = \pi_1$ then the household's best response is to play $\pi^e = \pi$. Since $\pi = \pi_1$ is the government's best response to a household play of $\pi^e = \pi_1$, this is the only Nash equilibrium in this example.

The Nash equilibrium then is unemployment at the natural rate combined with high inflation. Compare this to the Ramsey outcome of unemployment at the natural rate and inflation of zero.

19.4 Optimal Nominal Interest Rate Targets

In this section we will consider the government's optimal choice of nominal interest rates. We will consider the real cost of inflation, whereas previously we had simply taken it as given that the government disliked inflation. We will use the simple inventory model of cash holdings from Chapter 4 to show that households are best off when the nominal interest rate is zero. This is a form of what is known as the *Friedman rule.* It appears frequently in monetary economics.

Recall that the nominal interest rate R, the real interest rate r and the expected inflation rate π^e are related by the *Fisher formula*: $R = r + \pi^e$. For this discussion we will take the real interest rate as fixed and beyond the control of the government. Furthermore, we will assume that the government cannot directly manipulate inflationary expectations, and that the private sector correctly forecasts inflation. That is: $\pi^e = \pi$. Thus the government influences the nominal interest rate only through its choice of the actual inflation rate, π. The intuition behind the Fisher formula is quite compelling: households demand a premium of π for holding assets denominated in money, which is losing value at the rate of inflation.

In our model there will be no production. Households own a stock of interest bearing assets, which earn a nominal rate of return of R, and a stock of zero interest money. Money must be used for transactions. There is a fixed cost of μ of converting the interest bearing assets into money, which must be paid every time the household goes to the bank to replenish its cash inventory. The household has real consumption at a rate c per period which it does not vary.

The household goes to the bank x times in one year, so it goes $1/x$ of a year between trips to the bank. To have enough cash on hand to meet its consumption requirement c per period over those $1/x$ periods, the household has to withdraw a real amount c/x at every trip. Thus average real cash holdings over the entire year are $c/(2x)$. Those cash balances could have been invested in interest bearing assets earning an amount R over the year. Thus the foregone interest cost is: $(Rc)/(2x)$. Each time the household goes to the bank to replenish its cash inventory, it incurs a real cost of μ. Thus, transactions costs are: μx. Total costs for a particular policy x are:

$$\mu x + R\frac{c}{2x}.$$

The household minimizes total costs. The minimization problem has a first order condition of:

$$-\frac{cR}{2}\frac{1}{x^2} + \mu = 0.$$

Solving for x produces Baumol and Tobin's famous square-root rule for trips to the bank:

$$x = \sqrt{\frac{cR}{2\mu}}.$$

We can plug the household's decision x back into its cost function to determine the household's annual cash management costs, $\omega(\mu, c, R)$:

$$\omega(\mu, c, R) = \sqrt{2\mu cR}.$$

It is increasing in the fixed charge of going to the bank, μ, the rate of consumption, c and the nominal interest rate R.

A benevolent government that wishes to minimize the household's costs by choice of R would clearly choose to set $R = 0$. At this interest rate, the household goes to the bank only once in its lifetime and incurs no interest penalty for holding money. This is because money also earns in a real interest rate of r. This can only be the case if inflation is negative. From the Fisher formula $R = r + \pi$, we see that $R = 0$ implies that $\pi = -r$. So if $R = 0$, money is a perfect substitute for bonds. Holding a dollar isn't so bad, because next year the household will be able to purchase more with that dollar than it can now.

Although this model is quite limited, it points to one of the important real costs of inflation. Inflation causes households to engage in privately useful but socially useless activities. In times of high inflation, households find it in their interest to spend time and real resources economizing on cash balances. Notice that this is the first indication we have had this chapter that perfectly anticipated inflation is harmful.

If the idea of negative inflation rates seems outlandish, think of the Friedman rule instead as advocating paying interest on money. It is difficult (though not impossible) to pay interest on cash holdings (C. A. E. Goodhart, an English central banker, suggested having a lottery based on cash serial numbers), it is quite easy to pay interest on demand deposits.

Exercises

Exercise 19.1 (Easy)
True, False, or Uncertain (and explain):

1. The Consumer Price Index overstates increases in the "true" cost-of-living index.
2. Inflation is bad because to fight it the Fed increases interest rates, which hurts Americans.

Exercise 19.2 (Easy)
Do governments prefer Phillips curves that are relatively flat (low value of γ) or relatively steep (high values of γ)?

Exercise 19.3 (Moderate)
Assume that the government has a payoff over inflation π and unemployment u of:

$$V^g(u, \pi) = -\phi u^2 - \pi^2.$$

Here $\phi > 0$. The larger ϕ is, the nicer the central banker (that is, the more the central banker cares about the unemployed. Assume that there is a Phillips curve of the form in equation (19.6). Answer the following questions:

1. Assume that inflationary expectations are fixed at π^e. Find the optimal inflation rate choice of the government, $\pi_0(\phi)$.
2. For fixed inflationary expectations, find the corresponding choice of unemployment rate, $u_0(\phi)$.
3. Now assume that the private sector is aware of the government's maximization problem and knows ϕ perfectly. Find the inflation rate π_1 at which expectations are met. What is the associated unemployment rate, u_1?

Variable	Definition
Q_i, q_i	Output (and its log) in industry i.
L_i, ℓ_i	Employment (and its log) in industry i.
α	Production parameter common across industries.
P, p	Common but unobserved general price level (and its log).
P_i, p_i	Observed price (and its log) in industry i.
Z_i, z_i	Shock (and its log) specific to industry i, also industry i's relative price.
P^e, p^e	Common price forecast (and its log).
Z_i^e, z_i^e	Estimated industry-specific shock and its log.
N, n	Population and its log.
A	Parameters (used to make notation neat).
u	In Lucas model: the "not-employment" rate, elsewhere, the unemployment rate.
$u^\star$	The natural rate of unemployment, that is, the rate of unemployment when all industries correctly estimate their specific shocks.
γ	Slope of Phillips curve.
π	Actual inflation (chosen by government).
π^e	Expected inflation (chosen by private sector).
$\pi^\star(\pi^e)$	Government's optimal choice of inflation when inflationary expectations are π^e.
$V^g(u, \pi)$	Government's preferences over unemployment and inflation (it dislikes both equally).
V^p	Private sector payoff: industries dislike making errors in estimating the inflation rate.
π_1	Inflation rate at which expected and chosen inflation coincide, the Nash equilibrium.
u_1	Unemployment rate at Nash, just equal to $u^\star$.

Table 19.1: Notation for Chapter 19

4. Would you prefer to live in a country whose government has a high value of ϕ or a low value of ϕ?

Exercise 19.4 (Moderate)
For this exercise, we will consider what happens when the government and the private sector repeatedly interact. Unemployment in period t u_t, inflation π_t and inflationary ex-

pectations π_t^e are related by the simple Phillips curve:

$$u_t = u_t^\star + \gamma(\pi_t^e - \pi_t), \text{ for all } t = 0, 1, \ldots, \infty.$$

The parameter γ is fixed over time. The government knows about the Phillips curve, but the private sector does not. The government has preferences over unemployment and inflation in period t of:

$$V_t^g(\pi_t, u_t) = -{u_t}^2 - {\pi_t}^2, \text{ for all } t = 0, 1, \ldots, \infty.$$

The private sector sets inflationary expectations based on last period's inflation. This is known as *adaptive expectations.* As a result, π_t^e is given by:

$$\pi_t^e = \pi_{t-1}, \text{ for all } t = 1, 2, \ldots, \infty.$$

Assume that $\pi_0^e = 0$, that is, the private sector begins by believing that inflation will be zero. Answer the following questions:

1. Assume that the government takes as given expectations in a period π_t^e and picks the inflation rate π_t which gives it the highest payoff in period t. Find the government's choice rule $\pi_t^\star(\pi_t^e)$.
2. If the government sets inflation $\pi_t = \pi^\star(\pi_t^e)$, how do expectations evolve over time? Thus right down a law of motion for inflation, $\pi_t(\pi_{t-1})$.
3. What do the trajectories of inflation and unemployment look like over time? Are they rising or falling? Do they settle down? If so, where?
4. How would your answer have been different if, instead of the initial expected inflation being zero, it had been some very large number instead?
5. Now assume that the Phillips curve is augmented with a mean zero shock term, ε, so:

$$u_t = u^\star + \gamma(\pi_t^e - \pi_t) + \varepsilon_t.$$

   Assume that the government knows the value of ε_t and reacts appropriately. Now what happens?

Exercise 19.5 (Easy)
To answer this exercise, you need to answer Exercise 19.4 above. Imagine that the private sector has adaptive expectations about the government's inflationary policy over time, but that part of expected inflation is the government's *announced* inflation target. This announced inflation target is merely an announcement and has nothing to do with reality. If π_t^a is the announced target for period-t inflation, expectations satisfy:

$$\pi_t^e = \delta\pi_{t-1} + (1 - \delta)\pi_t^a, \text{ for all } t = 1, 2, \ldots, \infty.$$

Here $0 < \delta < 1$ is a parameter indexing how much weight the private sector puts on the government's announced inflation rate target. Assume that the government lies constantly, and announces $\pi_t^a = 0$ always. Assume that the government, as in Exercise 19.4, always chooses the inflation rate that maximizes its one-period payoff. Find the steady-state levels of inflation and unemployment.