Optimal Monetary Institutions and Policy

ROBERT E. HALL

IN THE BEST of all possible worlds, the dollar would be a unit of purchasing power with the same stability as the inch and the gallon. Markets would clear instantly, so there would be no adverse consequences for real economic activity when monetary policy single-mindedly kept the price level on target. In the modern U.S. economy, however, there is a strong suspicion that purely nominal changes in monetary policy have real consequences in the short run. Moreover, the same suspicion exists for every other economy, both contemporary and historical. Though economics has not been notably successful in proving that monetary policy affects employment and output as well as the price level, and though economists differ acrimoniously about the strength of the real effects, almost nobody would assert that there are no real effects at all.

This chapter proposes a design of monetary policy for an economy where there is a suspicion of sluggish price movement and consequent real effects of monetary change. Its purpose is to recommend a practical policy that is as close as possible to the theoretical optimum. Though the recommendation is pratical in the sense that it considers all the issues that would arise if the policy were put in place in the modern U.S. economy, it does not consider the question of the political acceptability of the policy.

Three major objectives enter into the design of the policy:

1. *Microeconomic efficiency*. A monetary policy should avoid deadweight loss. Two major sources of loss in the current system are requirements that banks hold non-interest-bearing reserves and the prohibition of bearer securities in small denominations that would compete with Federal Reserve currency.

2. *Stability*. Every monetary system involves a unique asset—gold, silver, currency, or reserves—that unambiguously discharges a debt. In times of financial crisis, the demand for this asset rises sharply. A good monetary system will insulate the price level and real activity from these shifts in demand.

3. *Macroeconomic efficiency and robustness*. If markets do not clear instantly and monetary actions have real effects, policy faces a

Reprinted from *Alternative Monetary Regimes*, edited by Colin D. Campbell and William R. Dougan, with permission of The Johns Hopkins University Press.

tradeoff between price stability and real stability. An aggressive policy for price stabilization may bring sharp fluctuations in output and employment. Fundamentally, the choice between price and output stability involves subtle issues about the benefits of each. However, a very basic requirement is that policy minimize real fluctuations for whatever amount of price variability is chosen. Such a policy is efficient. Moreover, because of our ignorance about the structure of the economy, a policy should be robust; it should bring an acceptable result, both in an economy with perfectly flexible prices and in an economy with quite sticky prices.

All three of these objectives can be met—there is no tradeoff among them. Deadweight loss can be made inconsequential by basing the monetary system on interest-bearing reserves. Because the reserves are attractive financially, there is no need for reserve requirements or for the prohibition of private competitors. Further, by letting a substantial fraction of the federal debt serve as reserves, the problem of instability is solved. Sudden increases in demand during crises have much smaller consequences within a large market for hundreds of billions of dollars of federal debt than within a market for tens of billions of dollars of reserves.

A policy that is efficient from the macroeconomic point of view takes into account departures of the price level from its target and departures of real activity from its potential level. A remarkably simple characterization of efficient policies is this: all efficient policies make the price-level departures proportional to the unemployment departures. Another way to think of efficiency is by considering what I call an *elastic price target.* The price-level target is a constant plus an elasticity times the unemployment rate. The elasticity might be 8. Hawks (those who put heavy weight on price stability or believe that prices are highly flexible) would choose an elasticity below eight; doves might choose an elasticity even higher than eight. The simple consideration of efficiency, however, quite apart from the welfare value of price stability or the flexibility of prices, tells us to pursue a policy within the family of elastic price targets. All other policies, including especially the type of policy followed by the United States for the past few decades, are nonstarters. An elastic target policy can deliver better performance in terms of both price and employment stability.

Monetary Saturation

Milton Friedman (1969) has made the basic case for monetary saturation. Basically, if monetary instruments cost no more to issue and

service than do nonmonetary instruments, then simple minimization of deadweight loss requires that the yields of the two types of instruments be equal. In an economy where no interest is paid on monetary instruments, saturation can occur only when the nominal interest rate on nonmonetary instruments is zero. As Friedman pointed out, policy can achieve zero nominal rates by deflating at the negative of the real interest rate. But in a modern economy without legal or technological restrictions on paying interest on monetary instruments, saturation can be achieved with stable prices (or any other policy for the price level) by paying market interest rates on monetary instruments.

To make this discussion specific, I will consider the possibility that the Federal Reserve pays the three-month Treasury bill rate on reserves, minus a small differential. The actual magnitude of the differential is what much of the rest of this chapter is about; for now, I only stress that the diffferential is small. Because reserves pay a yield close to other short-term instruments, the demand for reserves would be substantial even though reserve requirements would no longer exist. Much of the wealth currently held by banks in the form of federal debt would be held as reserves instead. In 1983, commercial banks alone held $188 billion in Treasury securities, as against only $40 billion in reserves.

The portfolio of the Federal Reserve would swell to meet the added demand for reserves. Around half the $380 billion in short-term (less than one-year maturity) federal debt should be monetized to get close to saturation.

To achieve saturation in currency, a technological solution to the problem of paying interest is needed. Mere restoration of banks' rights to issue non-interest-bearing notes is probably a step backward because it trades tax revenue for the wasteful techniques that banks would use to keep their notes in circulation. The best step would be to grant all financial institutions the right to issue interest-bearing notes in small denominations. A note with a constant face value of $100 could earn interest for the holder, which would be credited to that person's Visa account. Each time a bank paid a depositor with a note, it would record the fact electronically with Visa so that interest could be credited to the current holder.

Deregulation of bearer notes would eliminate part of the deadweight loss associated with currency today, but I suspect that considerable demand for Federal Reserve notes would remain. The Federal Reserve would continue to earn substantial seigniorage on its currency issue.

Stability. Saturation of the economy in reserves should help to stabilize the economy by reducing the frequency of financial crises and by limiting their consequences when they do occur. In a crisis, debtors struggle to obtain the asset that underlies the monetary system—the asset that has the power to discharge a debt unambiguously. Under the gold standard, a crisis takes the form of a move out of paper assets and into gold. In a fiduciary monetary system, debtors move into the reserves at the central bank. Although the founding principle of the Federal Reserve System was that discretionary policy could accommodate such a move by issuing added reserves, monetary history suggests that the Federal Reserve can always think of a good reason not to follow through in any given crisis. In any case, it is difficult for the Federal Reserve to determine how much to expand reserves, especially when a crisis occurs during inflation.

When the economy is saturated in reserves, as it would be if those reserves paid virtually the market rate, it stands to reason that the demand for them will be more stable. In the existing system, there is outstanding at any time several trillion dollars' worth of promises to pay reserves, either on demand or at a specified term. In 1983, there was only $40 billion in reserves, and essentially all of it was tied up as required reserves. A little nervousness on the part of debtors adds tremendously to demand for reserves. As it stands, the economy can accommodate this demand only by sharply increasing interest rates. Were the economy saturated in several hundred billion dollars of reserves, a modest rearrangement of reserves would satisfy the new demand from nervous debtors.

The Interest-Rate Differential as the Instrument of Monetary Policy

Once the Federal Reserve has established the policy of paying interest on reserves so as to saturate the economy in reserves, it opens up the possibility of using, as a policy instrument, the differential between the reserve interest rate and other interest rates. Raising the differential—that is, lowering the reserve rate relative to other rates—stimulates the economy. It is equivalent to increasing the quantity of reserves because a larger differential decreases the demand for reserves. Similarly, reducing the differential constitutes monetary contraction.

Manipulation of the interest-rate differential offers an advantage over open-market operations as a technique for carrying out mon-

etary policy, by avoiding the brokerage costs of open-market operations. Under present policy, the Federal Reserve churns its portfolio of government securities in the process of trying to stay within its target ranges for the levels of monetary aggregates and interest rates. The process of crediting reserve accounts with interest, on the other hand, is purely a matter of making accounting entries and involves no brokerage.

Efficient and Robust Monetary Policy

Even though monetary saturation would improve macroeconomic performance by reducing the frequency and severity of financial crises, stabilization policy would retain many of its current problems. From time to time, unexpected shocks to aggregate demand would push unemployment above or below its normal level. These shocks would affect the price level as well, perhaps with a lag. A greater challenge to policy occurs when the price level jumps suddenly. In an economy with sticky prices, a sharp increase in the price of one factor such as oil has the initial impact of raising the general price level, though ultimately a price stabilization policy can effect a lowering of the prices of other factors of production as needed to keep the overall price level on target.

The success of monetary policies can be judged in terms of two basic outcomes in the economy: the variability of the price level and the variability of unemployment. The basic goal of monetary policy in the longer run is price stability. Every departure of the price level from a constant target is a shortcoming. Long-range financial planning by individuals and businesses is most effective if the future value of the dollar can be relied on, even many decades hence, to be close to its current value.

With respect to unemployment, there are many reasons to believe that the average level of unemployment is inefficiently high. In such a case, a reduction in unemployment below the rate at which the economy will tend toward unaided would be socially beneficial. It might seem that the level of unemployment, not just its variability, could be considered in judging monetary policies. However, as Milton Friedman (1968) argued persuasively, monetary policy is powerless to influence the average level of unemployment. It is reasonable to suppose that the marginal social cost of unemployment rises with the level of unemployment. The best that monetary policy can do is to limit fluctuations in unemployment, since each fluctuation has a net social cost when the upside outweighs the downside.

As a general matter, then, we can judge monetary policies by use

of a diagram pioneered by John Taylor (1980), where the horizontal axis is unemployment variability (measured by the standard deviation of the departure of unemployment from its normal or natural level) and the vertical axis is price variability (measured as the standard deviation of the percent deviation from a constant target). Figure 6–1 is such a diagram. Any point in the diagram is a combination of unemployment and price variability that might be brought by a particular monetary policy rule. In general, points close to the origin represent the best combinations. But structural characteristics of the economy limit the points that are attainable by even the best policy. There is a curve in Figure 6–1, labeled the *policy frontier*, made up of the points closest to the origin that can actually be achieved with a practical policy, given the degree of price stickiness in the economy.

Each point on the policy frontier corresponds to a policy that is efficient in the sense that no other policy gives better performance in terms of both unemployment variability and price variability. The policy frontier has a critical role in deriving optimal policy. Without knowing anything about social preferences for unemployment versus price stability, we can make the strong statement that any policy not on the frontier is irrational. Preferences turn out to have a sharply

Figure 6–1

The Policy Frontier for Unemployment and Price Variability

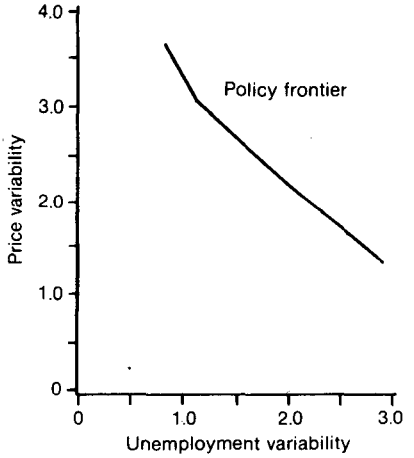

Note: Each point on the frontier is as close as any policy can get to small values of both unemployment and price variability. The axes are the standard deviation of the unemployment rate around its normal level of 6 percent and the standard deviation of the percent departure of the price level from a constant target level.

limited role in policy choice—the class of efficient monetary policies is quite small. The most important thing for policymakers to do is to get to the frontier. As I will show, actual policy since the mid-1960s put the U.S. economy at a point far above the frontier. Great improvements of the Pareto-superior type—reductions in both unemployment and price variability—could have been achieved without taking any position on the relative importance of unemployment and price stability.

Robustness. As I have stressed, economists don't really know how much influence monetary policy has on unemployment. The prevailing Keynesian model says that monetary policy moves affect unemployment during a transition period; during the same period, the price level is less sensitive to monetay policy than in the long run. The policy frontier in Figure 6–1 describes this case. But there are alternative models, not totally refuted by the data, where monetary policy cannot influence unemployment at all. Then the policy frontier is a vertical line.

A robust monetary policy gives reasonable performance under a wide variety of conditions. It gives a determinate price level, without too much variability, in the case where unemployment is unaffected by monetary policy. It gives a reasonable point on the policy frontier for the type of economy that the majority of practical macroeconomists believe we inhabit. Finally, it does not bring outrageous unemployment variability if prices prove less flexible than is generally thought.

Efficiency. Generally, the efficient monetary policy in a given macroeconomic model can be derived by minimizing the variance of the price level, given the objective of attaining a particular variance of unemployment. The structural equation of the model that matters for this calculation is the price adjustment equation, or Phillips curve. If that equation has a complicated form, the efficient policy may be correspondingly complicated. However, I avoid a detailed analysis of this type for two reasons. First, there is no professional consensus on the details of price adjustment. The only consensus is that prices move somewhat slowly to clear markets and that unemployment is one of the variables that might reasonably indicate the direction and magnitude of price adjustments. Second, as pointed out by Robert Lucas (1976) in his famous paper on econometric policy evaluation, the Phillips curve may change when a new policy comes into use. To estimate the shift requires a theory of price adjustment. Since

the whole point of this chapter is to avoid commitment to a particular theory, that avenue would be self-defeating.

What I will do instead is to examine the very simplest case, where the rate of inflation is governed, negatively, by the unemployment rate, without any expectational shifts or other modern complications. The absence of an expectational term is appropriate in considering alternative policies, all of which stabilize the price level. The shifting Phillips curve is a phenomenon of an economy with chronic upward drift in its price level. No policy considered here permits chronic drift.

In addition to a negative relation between the rates of unemployment and inflation, the Phillips curve in my analysis shifts randomly. Inflationary shocks from all sources other than aggregate demand—world oil and food markets, episodes of aggressive wage demands, and so forth—are wrapped into the shift. As it turns out, almost nothing is lost by treating these as a single composite rather than looking at them individually.

The problem, then, is to describe policies that bring efficient combinations of unemployment and price variability in an economy with a simple Phillips curve. I will assume that monetary policy has a single dimension that influences unemployment and prices. Specifically, when the interest-rate differential on reserves is raised, unemployment falls and prices rise in some combination; when it is lowered, unemployment rises and prices fall.

As a final step in setting the stage, let me assume that the inflationary shocks that perturb the Phillips curve are unpredictable from one year to the next. In dealing with this year's stabilization problems, monetary policy cannot anticipate what new shock will occur next year.

In this setup, efficient stabilization policies can be characterized in a particularly simple way:

> The efficient policy can be expressed as a requirement that the deviation of the price level from target be a fixed multiple, A, of the deviation of unemployment from normal.
>
> Any policy that keeps the price deviation in proportion to the unemployment deviation is efficient.

To carry out an efficient policy, the Federal Reserve simply pegs the price level at a multiple of the unemployment rate (when they are expressed as deviations from a constant long-run price target and the normal unemployment rate, respectively). Suppose the elasticity of the policy (the multiple, A) is chosen to be 8, which I consider a

reasonable choice. If the Federal Reserve finds that prices are on target but that unemployment is a percentage point too high, it will launch an expansion, which will bring down unemployment by one point, raise the price level by eight points, or have a combination of effects such that the new price level is 8 times the new unemployment rate.

An important feature of this type of policy is that it does not rest on any particular belief about the immediate impact of monetary policy. The opinion of a majority of economists is that the earliest effects are mostly real. In that case, policy will make adjustments that shift unemployment relative to an unresponsive price level in order to achieve the elastic target. Should the price level respond rapidly, the policy will work just as well. The elastic price target is a robust policy as a result.

The elastic price target achieves price stability in the long run without fail. No matter what happens to the economy, the average unemployment rate in the long run will equal the normal rate. Consequently, under this policy, the average price level must be equal to the target. Price-level drift, a major failing of actual U.S. policy since 1965, will not occur under the elastic target. If policymakers are unable to adjust the target formula to offset permanent changes in equilibrium unemployment, the level of prices may differ permanently from the target, but the average rate of change of prices will be zero.

The elasticity of the target provides the only necessary control over the choice between unemployment and price stability. A low elasticity keeps prices close to target at all times, at the cost of wide swings in unemployment. An elasticity of zero, which is an efficient policy, is strict price stabilization, as proposed by Knut Wicksell (1962). Under such a policy, the Federal Reserve is oblivious to unemployment and adjusts the reserve differential aggressively enough to keep prices right on target. If the majority of economists are right about the Phillips curve, then this policy involves hideous jumps in unemployment when inflationary shocks strike (e.g., the two oil price shocks of the 1970s).

An elasticity of 2.5 or 3 is a close approximation to nominal GNP targeting, a policy recommended by many economists in its own right. Under this policy, shocks are partially acccommodated in the short run. Prices are to rise and output is to fall by the same percentage, so that the product, nominal GNP, remains at a predetermined level. By Okun's law, the corresponding change in unemployment is ⅓ or 1/2.5 percentage points for each percentage of output, so the elasticity of the target should be 2.5 or 3. Stating the policy in terms of an

elastic price target with respect to unemployment avoids the inconvenience of having to prescribe a target path for nominal GNP.

Nominal GNP targeting, or its equivalent—an elastic price target with an elasticity around 3—turns out to be a fairly harsh policy because it calls for aggressive contraction when adverse price shocks occur. A more forgiving policy, with smaller unemployment fluctuations and correspondingly larger and longer departures of the price level from target, is obtained with an elasticity of 8.

Simulation of three variants of the elastic price target for the period 1952–83 for the U.S. economy, with a Phillips-curve tradeoff of 0.5 percent lower inflation for each extra percentage point of unemployment, gave the following three points on the policy frontier (for details, see Hall 1984):

Elasticity	Standard deviation of price	Standard deviation of unemployment
0	0	2.79
3	2.64	0.88
8	3.50	0.44

Pure price stabilization—an elasticity of zero—requires powerful expansion and contraction of real activity to offset price shocks. The standard deviation of unemployment of almost 3 percentage points means that unemployment rates of 9 percent are common and that rates of 12 percent occur about one time in twenty.

The next line shows that nominal GNP targeting—an elasticity of 3—gives much better real performance at the cost of some deviations of the price level from target. Unemployment is between 5 and 7 percent for more than two-thirds of the time. Rates above 8 percent are rare. The price level spends more than two-thirds of its time in a band between 97 percent and 103 percent of the target level.

The third line, corresponding to an elasticity of 8, gives even better real stability. The standard deviation of 0.44 percentage points for unemployment means that the rate is almost never above 7 percent or below 5 percent. The cost is a standard deviation of the price level of 3.5 percent.

What is most instructive is to contrast these three points on the policy frontier with the actual behavior of the price level and the unemployment rate over the same period, 1952–83. The actual standard deviation of unemployment was 1.74 percentage points, about double what it would have been under nominal GNP targeting. The high variability of unemployment should have given us an extremely stable price level, with a standard deviation of about 1 percent around a constant level. Instead, the price level had an extraordinary degree

of variability because it rose so much. Obviously, there was no constant target for the price level. One way of describing the failure of the policy is in terms of the standard deviation of the price level around its average for the period; this standard deviation was a staggering 38.4 percent. Actual policy was outside the policy frontier by miles. This illustrates my most basic point: it is much more important to have any efficient policy than to have an inefficient one. The actual choice of a point on the frontier is a subsidiary matter.

Properties of a Robust Policy. A monetary policy based on an elastic price target is robust. It is always efficient, regardless of the characteristics of the economy. It delivers a degree of unemployment stability that can be improved only by accepting more price instability. It has two other robustness properties.

Although an arbitrary choice of the elasticity always gives an efficient policy, the best policy is the one on the frontier that touches the highest social indifference curve between price and unemployment variability. Generally, the position of the policy frontier depends on the slope of the Phillips curve, so the tangency occurs at different points in the diagram for different slopes. The choice of the optimal policy rests on knowledge of the slope of the Phillips curve and on preferences about inflation and unemployment variability.

The interesting feature of the optimal choice, however, is that the optimal elasticity of the price target is very insensitive to the slope of the Phillips curve. It is true that an economy with a more responsive inflation rate—that is, a steeper Phillips curve—will have a more favorable policy frontier, closer to the origin. But the policy frontiers for different slopes are roughly concentric—one is more or less like another magnified. Mathematically, the policy frontier is approximately homothetic. The slopes of different frontiers corresponding to different Phillips curve parameters are the same along a ray from the origin. If the social indifference curves are roughly homothetic as well, then the expansion path showing the alternative optima for different Phillips curves' slopes is a ray from the origin. But all policies along a ray from the origin involve the same elasticity; the slope of the ray is the elasticity. A policy of making price deviations proportional to unemployment deviations makes the standard deviations of the two variables stand in the prescribed proportion as well.

The conclusion that emerges is the following: for a given set of social preferences, the choice of the elasticity of the elastic price target is roughly independent of the slope of the Phillips curve. If the best elasticity is 8 for an economy with a Phillips curve slope of

0.5, 8 will be a good choice for slopes of 0.1 and 2.0 as well. Because we don't really know the slope of the Phillips curve, this type of robustness is an important feature of the policy. It will not go badly astray even if the existing econometric evidence about price adjustment is severely biased, as some critics have charged.

The final robustness property of the policy concerns its performance in an economy with perfectly flexible prices. Certainly it would be an embarrassment to propose a policy that left the price level indeterminate in such an economy. A robust policy should make the price level determinate and not too unstable in an economy with instant market clearing. By the criteria of this chapter, only one policy is optimal for that economy: exact price stability. An elastic price target is a little suboptimal because it makes the price level vary as unemployment varies, even though there is no social gain from the elasticity. If we were completely confident that prices were fully flexible, we would set the elasticity at zero. But the cost of making the elasticity positive is not very high. The price level is determinate because the policy is one of pegging the price level to the unemployment rate, which is exogenous to price determination in an economy with perfectly flexible prices.

Taken together, these considerations suggest that an elastic price target with an elasticity of 8 is a desirable one, given our limited knowledge of the operating characteristics of the economy. It seems to give a reasonable simulated performance over the post–World War II period under the assumption that the slope of the Phillips curve is 0.5. Its performance would be much better if the Phillips curve were significantly steeper, or much worse if it were flatter, but it is not clear that we would choose a different elasticity in either of those cases. Even if the professional consensus on the transitory real effects of monetary policy is wrong, and the optimal policy is complete price stability, the cost of using the elastic target policy is not too high.

Automatic Execution of the Efficient Policy

Proponents of nominal GNP targeting have generally offered the target as a guide to making monetary policy in its present form. Their hope is that the Federal Reserve would trade government securities as necessary to keep nominal GNP on its target path. The chairman of the Federal Reserve Board would report to Congress periodically on the success of policy; the principal criterion for judging this success would be how close nominal GNP came to its prescribed level. Some type of penalty would be imposed if nominal GNP went too far astray.

That type of policymaking would be a big improvement over the current policy. It could not help but move us closer to the policy frontier. But the two basic ideas of this chapter can be linked to give a completely automatic policy. We can make the reserve-interest differential the instrument of monetary policy and then change it by formula so as to achieve the elastic price target. The result is a policy that is close to ideal on both microeconomic and macroeconomic grounds.

The basic idea of the automatic policy is to raise the reserve differential whenever the economy drops below the elastic target and to lower it when the economy is above target. However, some subtleties need to be handled in order to make the policy work smoothly. If policy responds just to the current state of the economy, and there is some lag before a policy change has much effect on either unemployment or the price level, then there is a potential for unstable feedback. Even if the system were not unstable under such a policy rule, a response this month to economic data from last month could be seriously disruptive if it is not self-limiting.

Both the price level and the unemployment rate are measured with a certain amount of unavoidable error. If the price level were high and unemployment low in a particular month, an aggressive policy might call for a reserve differential that would be so low as to attract a large amount of wealth into reserves. The immediate impact of this monetary contraction would be high interest rates and financial disruption. If the reserve differential were kept low for the whole month, without responding to anything that happened during the month, there could be a month-long economic crisis. Even if unemployment and prices responded a bit during the month, the crisis could easily extend into the second month as well. The response in later months could be so strong as to call for a later expansionary move to a high reserve differential.

All of these problems can be avoided by linking the reserve differential not to the most recent data but to future data, say for the forthcoming year. In other words, the elastic target is to be achieved over the average of the forthcoming twelve months, not for the current month alone. The advantage of imposing the target on expectations rather than on historical data is that expectations are instantly responsive to policy. The month-long crisis I just described could not occur if the reserve differential were linked to the near future instead of the immediate past. If a situation threatened where a higher differential might raise the demand for reserves and contract the economy, the expectation of that contraction would show imme-

diately in higher expected unemployment and lower expected prices. Quickly, a new equilibrium would be reached, where the amount of the differential was just enough to put the economy on the elastic price target in terms of expectations for the coming twelve months.

Many other economists have pointed out the virtue of guiding monetary policy by expectations or forecasts instead of the most recent actual behavior of the economy. Some economists have proposed that the Federal Reserve look at futures markets for commodities or for the cost of living index or, in the case of nominal GNP targeting, at nominal GNP forecasts from reliable outside forecasters. These all are good ideas. But when the reserve differential is used as the instrument of monetary policy, there is a particularly simple way to link policy to expectations. All that needs to be done is to pay interest later, once the actual performance of the economy becomes known.

The specific operating rule I have in mind is the following: the Federal Reserve keeps track of the average balance over the past twelve months for each reserve account. Each month, it credits each account for 1 percentage point of interest on its average balance for each point by which the price level exceeds the elastic target. The interest is credited toward the end of each month, when the price level and unemployment rate for the previous month are announced. In addition, interest is credited daily for each account on that day's balance at the three-month Treasury bill rate less 0.5 percent.

Under this system, it is impossible for the public consensus about the price level and the unemployment rate over the next twelve months to differ significantly from the elastic target. Suppose, for example, that the public believed that the price level would be 1 percent above the elastic target on average over the next twelve months. They would also believe that reserves were going to pay 0.5 percent more than Treasury bills. Such a situation would be highly contractionary. But this only shows that their original belief was incorrect—the price level could not be so far above target in the face of such a negative policy.

The policy of retrospective payment of interest based on departures from the target is close to the ideal implementation of the elastic target. It avoids the central problem of earlier proposals for nominal GNP targeting: What do we do when the Federal Reserve misses the target? Monetary policy would be in the hands of a perfectly reliable automaton. Unlike other proposals to put monetary policy on automatic pilot, such as the constant money growth rule, this one is guaranteed to stay out of trouble. Its sensitivity to unemployment

as well as to the price level means that it cannot bring the sharp recession or uncontrolled boom that is the bugaboo of constant money growth.

The proposed policy has two parameters that might benefit from fine tuning. One is the twelve-month period over which expectations are relevant. It might work a little better to use twenty-four months or six months. I have chosen twelve months because it appears that a substantial fraction of the total impact of monetary policy on the real variables takes place within that span. The other parameter is the number of percentage points of interest per point of departure from the target. One for one seems reasonable, but the system might work more smoothly with a smaller or larger multiplier.

Conclusion

I have recommended a completely practical general approach to monetary policy with a number of important properties of optimality. The approach has no political appeal whatsoever; it seems complicated and arcane and would require that the Federal Reserve give up all of its revenue and all of its responsibility for monetary policy. Nevertheless, the basic ideas of monetary saturation and an elastic price standard should be investigated by economists.

Let me conclude by restating the policy I have in mind and by listing its virtues. The Federal Reserve would monetize about half of the short-term federal debt; the actual amount is relatively unimportant. To do this, it would pay interest on reserves at about 50 basis points below the Treasury bill rate. Deadweight loss in financial markets would be further reduced by eliminating reserve requirements and by lifting restrictions on interest-bearing competitors to currency.

In addition to paying daily interest linked to Treasury bills, the Federal Reserve would credit each account monthly in proportion to its average balance over the preceding twelve months. The extra interest would be 1 percent (at annual rates) for each point by which the price level exceeded 8 times the unemployment rate. The price level would be stated as a deviation from a fixed target, and the unemployment rate as a deviation from the normal rate of 6 percent.

The effect of this policy would be to saturate the economy in reserves. Saturation has chronic microeconomic benefits in the form of eliminated deadweight loss. In addition, saturation would help stabilize financial markets and the economy as a whole. Episodes of economic crisis when the public shifts portfolio demands toward

reserves, are less disruptive in an economy in which large volumes of reserves are held for portfolio motives in the first place. Saturation provides the needed elasticity of reserves that the Federal Reserve was created to provide, but that it has failed to provide too many times.

By linking the reserve differential to the prospective behavior of the economy, the policy also achieves the efficient stabilization policy. It correctly balances price stability against unemployment stability. When an adverse inflationary shock hits the economy, the policy absorbs the shock first in the form of a higher price level and then gradually works the price level back to its fixed long-run target. Such a response is optimal.

The proposed policy is robust, in the sense that it functions well in economies with very different degrees of price flexibility and real responsiveness to monetary change. If prices are completely flexible, the policy simply pegs the price level to an exogenous variable, the unemployment rate. If prices and unemployment are related in the way suggested by mainstream macroeconomics, then the elastic price target is the efficient stabilization policy. That property holds for any degree of price stickiness. Moveover, the choice of the optimal policy from among the efficient ones—that is, the choice of the optimal elasticity of the target—is almost independent of the amount of price stickiness. An elasticity of 8 seem a good choice whether the slope of the Phillips curve is 0.1 points of inflation per point of unemployment or 1.0.

REFERENCES

Friedman, Milton. "The Role of Monetary Policy." *American Economic Review* 58 (March 1968), 1–17.

———. "The Optimum Quantity of Money." In *The Optimum Quantity of Money and Other Essays*. Chicago: Aldine Publishing Co., 1969. Pp. 1–50.

Hall, Robert E. "Monetary Strategy with an Elastic Price Standard." In *Price Stability and Public Policy*. Kansas City: Federal Reserve Bank of Kansas City, August 2–3, 1984. Pp. 137–59.

Lucas, Robert E. "Econometric Policy Evaluation: A Critique." In *The Phillips Curve and Labor Markets*, Karl Brunner and Allen H. Meltzer, eds. Carnegie-Rochester Conference Series on Public Policy, vol. 1. Amsterdam: North-Holland, 1976. Pp. 19–46.

Taylor, John B. "Output and Price Stability: An International Comparison." *Journal of Economic Dynamics and Control* 2 (1980), 109–32.

Wicksell, Knut. *Interest and Prices*, R.F. Kahn, trans. New York: Augustus M. Kelley, 1962.