

The Inflation/Output Variability Trade-off Revisited

*John B. Taylor**

Describing the nature of the trade-off between inflation and output or unemployment has long been difficult and controversial. The Friedman-Phelps hypothesis, that there is no long-run Phillips curve trade-off between inflation and unemployment, has clearly won over most macroeconomists, but the debate has continued over what, if any, trade-off remains. The subtle notion that an uncertain short-run trade-off, but no long-run trade-off, exists between inflation and output has proved more difficult to analyze and describe.

The debate over monetary policy tightening in the United States in 1994 illustrates some of these difficulties. The distinction between long-run and short-run trade-offs was again blurred as many commentators expressed concern that the Federal Reserve's goal of low inflation would reduce real GDP growth. Typical of much financial and political reporting was a *New York Times* article on the rise in interest rates in 1994, which concluded, "the balance between . . . more growth and less inflation, shifts again—toward a slower economy" (Uchitelle 1994). The article even quoted Paul Volcker for support: "If you have a weaker economy, you have lower [nominal interest] rates. That is not a great world but that is the way it is." But a long-term analysis of the output versus inflation or interest rate trade-off would be stated differently. A weaker economy does not imply a lower inflation rate or a lower interest rate: In 1978, the unemployment rate was 6 percent, while interest rates

*Professor of Economics, Stanford University. This research was supported by a grant from the National Science Foundation at the National Bureau of Economic Research. The author is grateful to Laurence Ball and other conference participants for helpful comments.

and inflation were in double digits. Today the unemployment rate is no higher and the economy is no weaker, but interest rates and inflation are well below double digits. Thus, there is no long-run trade-off between a strong economy and low inflation or low nominal interest rates.

Several years ago, in an effort to more clearly delineate the short-run versus the long-run trade-off, I estimated a different type of trade-off between inflation and output (Taylor 1979). Rather than a long-run trade-off between the *levels* of inflation and output, I defined and estimated a long-run trade-off between the *variability* of inflation and of output. Because of this trade-off, efforts to keep the inflation rate too stable would result in larger fluctuations in real GDP and unemployment. Conversely, efforts to smooth out the business cycle too much would result in a more volatile inflation rate.

Such a variability trade-off is consistent with rational expectations and sticky prices and implies no long-run trade-off between the levels of inflation and output. It can also be estimated with stochastic optimal control methods. In fact, recent estimates by Fuhrer and Moore (1993) using modern techniques have found the shape and positions of the trade-off curve to be very similar to the one I estimated earlier. However, although little technical criticism has been made of the idea of such a variability trade-off, it is safe to say that it has not become part of the popular debate on the subject. While technically useful, the trade-off has not helped to clarify the distinction between the short run and the long run in popular discussions. It certainly has not replaced the Phillips curve!

The idea of the trade-off between inflation and output has been made even more confusing to outside observers by more recent strands of research. One strand—Fischer (1993), Lucas (1994), and Motley (1994)—has found that inflation has a quantitatively significant long-run effect on real GDP, or real GDP growth. Another strand of research (see Caballero and Hammour 1991) has examined whether efforts to smooth out short-term business cycle fluctuations might reduce long-run economic growth.

The purpose of this paper is to revisit the trade-off between the variability of inflation and output in order to clarify in simple terms the nature of this trade-off, as implied by recent research and experience, and to compare it with other notions of a trade-off. Rather than estimate a new trade-off using stochastic optimal control techniques—a topic of much current research (see Bryant, Hooper, and Mann 1993; Taylor 1993b; and Fuhrer and Moore 1993)—this paper will take a different approach, developing a more intuitive analysis using a series of simple diagrams and graphs. Such an analysis complements the ongoing technical research and provides additional insights that can improve public discussion and perhaps even public policy.

A Stylized Macroeconomic Model

Consider the following simple three-equation summary of the relationships between real GDP, the nominal interest rate, and the inflation rate:

$$y_t = -\beta(i_t - \pi_t - r^*) + u_t \quad (1)$$

$$\pi_t = \pi_{t-1} + \alpha y_{t-1} + e_t \quad (2)$$

$$\dot{i}_t = \pi_t + g y_t + h(\pi_t - \pi^*) + r^f + v_t \quad (3)$$

where y_t is real GDP measured as a percentage deviation from potential GDP; \dot{i}_t is the short-term nominal interest rate measured in percentage points; π_t is the inflation rate measured in percentage points; and e_t , v_t , and u_t are shocks that equal zero on average. The parameters of the model are π^* , r^f , r^* , α , β , g , and h , and are all positive.

Equation (1) describes an inverse relationship between the real interest rate and the deviations of real GDP from potential GDP. The deviations of real GDP from potential GDP are assumed to be due to fluctuations in aggregate demand; each component of aggregate demand—consumption, investment, and net exports—is assumed to depend negatively on the real interest rate. (Net exports depend on the real interest rate through the positive relationship between the real exchange rate and the real interest rate.) Potential GDP is assumed to be described by a production function—not shown separately—in which increases in capital, labor, and total factor productivity cause potential GDP to grow. Potential GDP is thus the normal or natural level of real GDP, rather than an upper bound on real GDP. When real GDP equals potential GDP ($y = 0$), the ex post real interest rate equals r^* , which is, therefore, the equilibrium real interest rate in the economy. Greater accuracy might be achieved in equation (1) by using the ex ante expected real interest rate as well as the long-term interest rate—using rational expectations for the term structure. However, to keep the model simple, only the actual inflation rate is included in equation (1). The variable u_t in equation (1) could represent changes in government purchases or any other factor that shifts aggregate demand.

Equation (2) summarizes price adjustment in the economy. When real GDP rises above potential GDP, inflation increases, with a lag because of the stickiness of prices. When real GDP falls below potential GDP, inflation decreases, again with a lag. The random variable e_t represents price shocks. Staggered wage and price setting as well as limited information are possible rationales for the stickiness of prices.

Equation (3) summarizes monetary policy in terms of the interest rate reaction of the central bank to deviations of inflation from a target π^* and to the deviations of real GDP from potential GDP. When inflation

rises, the policy calls for the nominal interest rate to rise by more than the inflation rate; and when real GDP rises relative to potential GDP, the interest rate also rises. The intercept term r^f in this relationship is the implicit real interest rate in the central bank's reaction function. The central bank takes actions to affect the nominal interest rate by open market operations, and these have implications for the growth rate of the money supply. Although these open market operations and money supply growth are not stated explicitly in these equations, they play an important role in the setting of interest rates.¹

Long-Run Averages

The long-run average values of real GDP, inflation, and the nominal interest rate implied by the model can be found by setting the change in the inflation rate and all the shocks to zero in equations (1) through (3). This gives:

$$y = 0 \quad (4)$$

$$i = r^* + \pi \quad (5)$$

$$\pi = \pi^* + (r^* - r^f)/h. \quad (6)$$

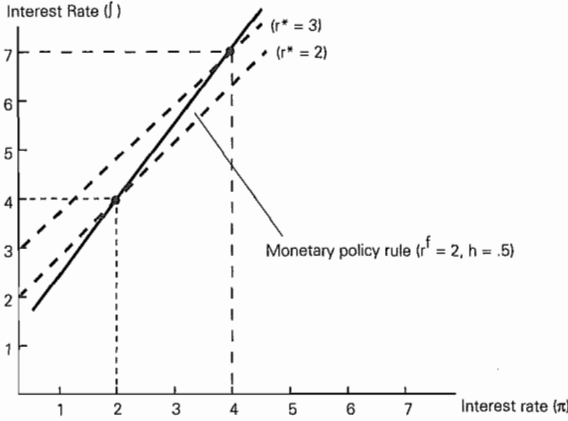
Equation (4) simply states that real GDP equals potential GDP in the long run; this equation follows immediately from the price adjustment equation (2). Equation (5) then follows from equation (1) with $y = 0$. Equation (6) follows from equation (3) with y and i given by equations (4) and (5). In addition, the growth rate of the money supply equals the growth rate of potential GDP plus the inflation rate π in equation (6) minus the growth rate of velocity.

It is obvious from equation (4) that no long-run trade-off exists between the inflation rate and the deviations of real GDP from potential GDP. With the deviations of unemployment from the natural rate proportional to the deviations of real GDP from potential (Okun's law), no long-run trade-off exists, therefore, between inflation and unemployment. Of course, the equations have been designed to capture these properties. It is certainly possible for either the natural rate of unem-

¹ Equation (3) could also be interpreted as the result of a monetary policy with a fixed growth rate of the money supply. The target inflation rate would then be the long-run average inflation rate implied by the quantity equation with the constant money growth rate. Then, when inflation rose above the target inflation rate, the demand for money would rise relative to the supply of money and interest rates would rise as shown in equation (3).

Figure 1

The Multiplier Effect of a Change in the Equilibrium Real Rate of Interest



ployment or potential GDP to be dependent on the inflation rate. However, strong evidence (summarized below) suggests no long-run trade-off.

Note that equation (6) implies that if the central bank chooses a monetary policy with an implicit real interest rate r^f , different from the equilibrium real interest rate in the economy r^* , then the steady state inflation rate π will not equal the target inflation rate π^* . If the equilibrium real interest rate r^* changes—perhaps because of a change in government spending policy—then the steady state inflation rate will change unless the central bank also adjusts its implicit real interest rate r^f . If the parameter h is less than 1, then equation (6) implies that the change in the equilibrium real interest rate has a multiplier effect on the inflation rate; that is, the inflation rate rises by more than the equilibrium real interest rate.

This multiplier effect is illustrated in Figure 1, which also shows how the long-run nominal interest rate and the long-run average inflation rate implied by the monetary policy rule combine with the given equilibrium real interest rate. The solid line showing the reaction function of the Federal Reserve is plotted in the case where $g = 0.5$, $h = 0.5$, $\pi^* = 2$ percent, and $r^f = 2$ percent. The dashed line shows the

relationship between inflation and the interest rate that must hold in the steady state if the equilibrium real interest rate is r^* . In the example in Figure 1, $r^* = 2$ percent, so that $r^* = r^f$ and the steady state inflation rate is 2 percent. However, if r^* rises from 2 percent to 3 percent, then the dashed line shifts up, the steady state nominal interest rate rises to 7 percent, and the steady state inflation rate rises from 2 percent to 4 percent, unless of course the central bank shifts up the policy rule so that r^f also equals 3 percent. Similarly, a decline in the equilibrium interest rate would lead to a decline in the steady state inflation rate unless the Fed adjusted its policy.

In reality, the central bank does not know the equilibrium real interest rate, so that we cannot expect it to accurately set r^f equal to r^* , and this is a disadvantage of a policy rule like equation (3) in contrast to money growth rules. If the central bank uses an incorrect estimate of the equilibrium real interest rate when using a monetary policy like equation (1), then an inflation rate higher or lower than targeted will result. However, such an error will not result in continuing increases or continuing decreases in inflation, as would a policy that tries to peg the real interest rate above or below the equilibrium real interest rate. Moreover, equation (6) shows that the impacts of the error on the long-run average inflation rate depend on the size of the response of monetary policy to the inflation rate. The larger the response parameter h , the smaller the impact of a change in the equilibrium real interest rate on the long-run average inflation rate. This is a reason not to choose a monetary policy with h too close to zero.

Because equation (3) may be less familiar than equations (1) and (2), Figure 2 is presented, showing how the equation describes actual Fed behavior in recent years. The actual federal funds rate and the federal funds rate implied by equation (2) are shown in Figure 2. After the Fed tightening moves early this year, the policy rule is back on track. (Figure 2 is an updated version of a similar plot from Taylor 1993a.)

Short-Run Fluctuations

Now, the fluctuations of real GDP and inflation will be considered. Substitute equation (3) into equation (1) to obtain:


$$y_t = -c(\pi_t - \pi^*) - (c/h)(r^f - r^*) + (u_t - \beta v_t)/(1 + \beta g) \quad (7)$$

where $c = \beta h/(1 + \beta g)$. If $r^f = r^*$, then the middle term on the right-hand side of the above expression drops out.

An easy way to derive the trade-off between the variability of inflation and the variability of the deviations of real GDP from potential is to substitute equation (7) into equation (2). This gives

Figure 2

Actual Federal Funds Rate and the Rate Implied by a Policy Rule (Equation 2)



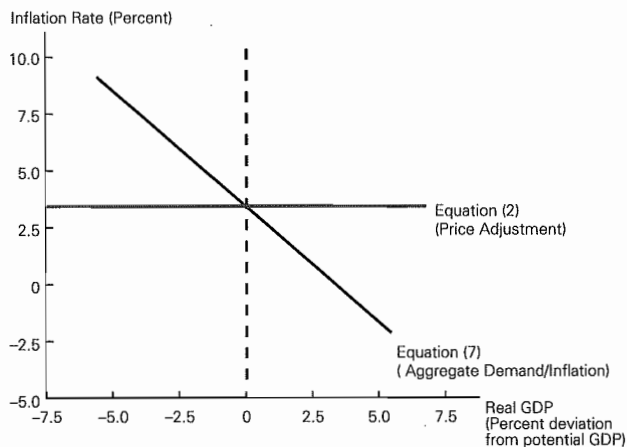
$$\pi_t - \pi^* = (1 - \alpha c)(\pi_{t-1} - \pi^*) - (\alpha c/h)(r^f - r^*) + \alpha(u_{t-1} - \beta v_{t-1})/(1 + \beta g) + e_t \tag{8}$$

which is simply a first order autoregression in $\pi_t - \pi^*$. The variance of $\pi_t - \pi^*$ can easily be obtained from equation (8), and from this the variance of y can be obtained using equation (7). For example, in the case with only price adjustment shocks (e_t), the standard deviation of $\pi - \pi^*$ is $\sigma/(1 - (1 - \alpha c)^2)^{1/2}$ and the standard deviation of y is $c\sigma/(1 - (1 - \alpha c)^2)^{1/2}$. In this case, a trade-off is traced out by varying c , which depends on the two policy parameters h and g .

However, the aim here is to provide an intuitive understanding of this trade-off. The two key relationships in the model describing inflation and the deviations of real GDP from potential are equations (2) and (7). Both describe a dynamic relationship between inflation and real GDP. These two relationships are graphed in Figure 3. The downward-sloping curve shows equation (7); it indicates how real GDP and inflation are negatively related. Recall that equation (7) combines the relationship between the interest rate and inflation with the central bank's policy rule. As inflation rises, the central bank raises the interest

Figure 3

Relationship between Inflation and Deviation of Real GNP from Potential, as Described in Equations (2) and (7)



rate and this lowers real GDP. The opposite occurs if inflation falls. The policy also calls for higher interest rates when real GDP rises above potential GDP, and that is also incorporated in equation (7) and the downward-sloping line in Figure 3. The downward-sloping line describes how aggregate demand depends on inflation, and therefore is called the aggregate demand/inflation (ADI) curve.


Equation (2) is shown as a flat line in Figure 3 because contemporaneous real GDP does not appear in the equation; only $y_t - 1$ appears. If real GDP rises above potential GDP, then inflation will start to rise and the flat line in Figure 3 will shift up over time. If real GDP falls below potential GDP, then the flat line will shift down. The intersection of the two lines determines real GDP at any particular time.

Now, fluctuations in real GDP and inflation occur if either of the two curves in Figure 3 shifts. The downward-sloping aggregate demand/inflation curve will shift to the right with a shift in monetary policy to a higher inflation target, a monetary policy mistake (v_t), or a shift (u_t) of equation (1). The price adjustment line will shift up if a price shock (e_t) to equation (2) occurs.

An example of how shifts in these two curves trace out fluctuations in inflation and output is shown in Figure 4. First, imagine that the Fed—either on purpose or by mistake—shifts monetary policy towards

Figure 4

Effects on Real GDP and Inflation of Shifts in the Price Adjustment or Aggregate Demand Equations



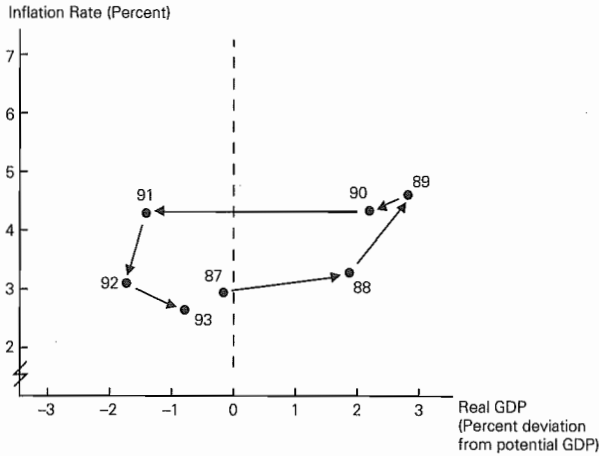
a higher inflation rate. Such a shift could be due to a shift to a higher rate of money growth, which would imply a higher inflation target. As seen in Figure 4, this causes an expansion of real GDP above potential GDP. In the short run, this has no effect on inflation, but over time inflation rises and real GDP moves back to potential GDP. If the Fed made no further changes in policy, then that would be the end of the story: higher inflation with real GDP back to potential GDP, consistent with the nonexistence of any long-run trade-off between inflation and real GDP.

On the other hand, the Fed could shift the policy back again—perhaps after learning that its policy mistake has increased inflation. Then the aggregate demand/inflation curve shifts back to where it was originally, causing a decline in real GDP below potential GDP, as shown in Figure 4. After a lag, inflation will start to decline; the price adjustment line shifts down gradually over time until real GDP returns to potential GDP.

The pattern of inflation and real GDP traced out in Figure 4 is a typical boom-bust cycle, with a boom and then rising inflation followed by a recession with subsequent falling inflation. Figure 5 shows that the actual pattern of real GDP and inflation in the past seven years looks similar to the points in Figure 4.

Figure 5

Relationship of the Deviation of Real GDP from Potential and the Inflation Rate, 1987 to 1993




Now consider the trade-off between the fluctuations in inflation and real GDP. Monetary policy determines the slope of the aggregate demand curve, because the slope of the curve is given by $-(1 + \beta g)/\beta h$, and g and h are the parameters of the monetary policy rule. Thus, the curve is flatter either if h is higher—the central bank responds more aggressively to inflation—or if g is lower—the central bank responds less actively to deviations of real GDP from potential GDP. A lower h or a higher g makes the curve steeper.

The effects of the different policy parameters are shown in Figure 6. The hypothetical shifts in the price adjustment line in the two left-hand panels represent a given size of shifts to equation (3); that is, a given size for the price shocks. If the aggregate demand curve is flatter, then output declines by a large amount when a price shock occurs. If aggregate demand is steeper, then output declines by a small amount with a shock to inflation. Clearly the variance of real GDP is much smaller when the aggregate demand curve is steep—which is the case where h is small and the central bank does not respond very much when inflation rises.

It may appear from Figure 6 that the variance of inflation is not affected by the slope of the aggregate demand/inflation curve; however,

Figure 6

Effects of the Policy Parameters



the variance of inflation is lower in the case where the variance of real output is higher, because inflation returns to the target level more quickly after a shock. After a price shock, the sum of squared deviations of inflation from target inflation is smaller in the case where output falls

(or rises, in the case of negative price shocks) by a larger amount. In other words, if price shocks were the only shocks affecting the economy, then the vertical spread in a diagram like Figure 5 should not depend on the policy rule; only the speed at which real GDP returns to potential GDP would depend on the policy rule. By affecting the speed, the policy rule affects the variance of inflation as well as the variance of real GDP.

Finally, the right-hand panel of Figure 6 illustrates how the pairs of fluctuations in real GDP and inflation trace out the trade-off curve. The variance of inflation (π) is on the vertical axis, while the variance of real GDP (y) is on the horizontal axis. The parameters of the monetary policy rule change so that as interest rates respond more to inflation (higher h) and less to real GDP (lower g), the aggregate demand/inflation curve flattens and the variance of real GDP rises while the variance of inflation falls.

Observe that the objective of policy should be to keep the aggregate demand/inflation curve stable. However, if for some reason—a policy mistake or an unavoidable shock to consumption—the aggregate demand/inflation curve shifts, then inflation will move away from the target and the central bank is faced with the same trade-off as in the case of shock to the price adjustment line.

Note also that although only the slope of the aggregate demand/inflation line is important for the trade-off between inflation and real output fluctuations, the absolute sizes of the parameters g and h affect fluctuations in the interest rate. Choosing a g that is very high, for example, could result in large fluctuations in interest rates. Although fluctuations in interest rates are not directly a cause of concern in this model, in a more realistic model with lags and expectations such fluctuations would likely lead to instrument instability. Thus, raising g and h very high would probably not be a good policy in reality.

The estimated trade-off curve bends very sharply at a point such as that designated by the open circle in Figure 6. In other words, the opportunity costs of reducing inflation variability below the level at the open circle are very high, in terms of higher output variability. Similarly, the opportunity costs of reducing output variability below the open circle are also very high, in terms of higher inflation variability. This suggests that the optimal choice for policy is likely to be near the open circle. Even with large changes in preferences over time—say, because of a change in political sentiment—a country would therefore not be likely to move far from the sharp curvature point.

How does one go about choosing points on such a trade-off curve? Which utility function to use is not obvious. Perhaps the best way to make a choice is to examine different scenarios, such as the one depicted in Figure 5. In this 1987–93 scenario, the standard deviation of inflation is considerably less than the standard deviation of output. If the Fed had been successful in achieving a soft landing in 1990–91 rather than a

recession, then the standard deviation of real output would have been lower. The policy rule indicates that the actual pattern was achieved with $h = 0.5$ and $g = 0.5$. Hence, to lower the output fluctuation and thereby raise the inflation fluctuation, one would have to raise the coefficient on real GDP in the policy rule or lower the coefficient on inflation. For the reasons mentioned above, lowering the coefficient on inflation probably would be unwise because of uncertainty about the real interest rate. Hence, raising the coefficient on real GDP—perhaps to 0.7 rather than 0.5—might be considered.

Empirical Evidence

The trade-off between the variability of inflation and that of real GDP can be better understood by examining some data on real GDP, unemployment, and inflation in addition to the data presented in Figure 5.

Inflation and Unemployment

That no long-term trade-off exists between inflation and unemployment, or between inflation and the deviations of real GDP from potential GDP, has been well established. For completeness in this graphic analysis, Figure 7 provides a simple picture that summarizes the relevant evidence. It shows four years during which the economy was operating where real GDP was close to potential GDP, neither in recession nor in boom. Whether inflation was high, as in the late 1970s, or low, as in the early 1960s or the early 1990s, the unemployment rate was close to 6 percent. Clearly, no long-term trade-off exists between the levels of unemployment and inflation. The assumption in the above stylized model that the equilibrium value of y is zero is thus a good one.


The Effects of Inflation on Long-Term Growth

Observations on economic growth in different countries indicate that inflation is negatively correlated with economic growth. In the stylized model this would mean that potential GDP depends on the inflation rate, but the assumption would still be maintained that the deviations of real GDP from potential GDP converge to zero in the long run.

How large are the long-run effects on potential GDP? Fischer (1993) and Motley (1994) provide a comprehensive set of estimates based on data in both developed and less developed countries. Motley finds that a reduction (increase) in the inflation rate of 1 percentage point would increase (reduce) the long-run productivity growth rate by 0.06 percentage points per year in the developed countries. For example, an increase

Figure 7

*Rates of Unemployment and Inflation in Selected Years
When Real GDP Close to Potential GDP*



in the inflation rate from 2 percent—close to its level in the early 1960s—to 12 percent—close to where it was in the late 1970s—would lower productivity growth by 0.6 percentage points.

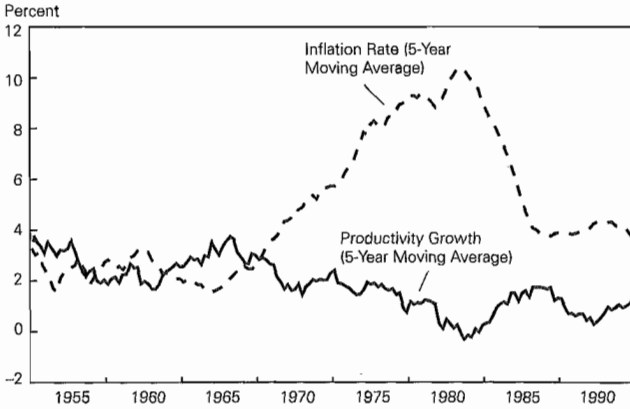
Figure 8a examines the pattern of inflation and labor productivity growth in the nonfarm business sector. (To abstract from the large cyclical productivity and inflation swings, Figure 8a reports five-year moving averages of both the inflation rate and the productivity growth rate.) Note that the start of the increase in inflation in the mid 1960s occurred at about the same time as the start of a slowdown in labor productivity growth. Moreover, the productivity growth slowdown ended at about the same time as the disinflation of the early 1980s, which ended the very high inflation period of the 1970s. Similar productivity growth slowdowns and inflation increases occurred in other countries.

Although the productivity growth slowdown has ended, the growth rate of productivity has not yet returned to the levels of the 1950s and 1960s. Figure 8b shows how much of a revival in productivity growth would be expected if the simple statistical relationship between productivity growth and inflation observed during those years persisted. According to Figure 8b, a rise in inflation of 1 percentage point

Figure 8


A

Inflation Rate and Rate of Growth in Labor Productivity, Nonfarm Business Sector



B

Predicted Increase in Productivity Growth



leads to a decline in productivity growth of 0.25 percent; this is a much larger effect of inflation on productivity growth than that reported in Motley (1994).

Figure 8b does not prove that inflation was the key factor in the great labor productivity growth slowdown. People have pointed to many other factors. Moreover, there is no reason to expect the 0.25 coefficient to be stable; most certainly it would not hold outside of the narrow range of observations in Figures 8a and 8b, but Figure 8b certainly suggests that inflation should be considered along with other reasons more commonly given for the productivity growth slowdown, such as lagging research and development, education, or public infrastructure investment.

Effects of Output Variability on Long-Term Growth

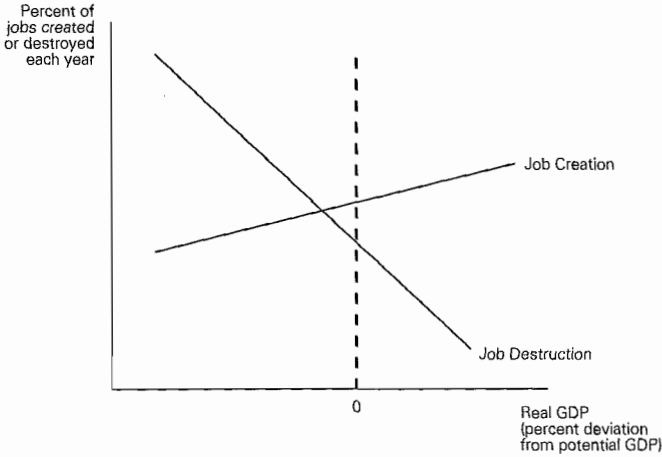
Schumpeter (1939) first pointed out the close link between economic growth and economic fluctuations. According to Schumpeter, booms are periods when inventions spread throughout the economy through innovation. Recessions are periods when the destruction of firms and jobs overtakes the creation of jobs. Schumpeter's analysis raised the possibility that recessions might enhance productivity growth, as firms take the opportunity of slack times to make structural adjustments.

Davis and Haltiwanger's (1990) recent studies of job creation and job destruction have been influential in assessing this "cleansing effect of recessions." Their finding that job creation is much less sensitive to the business cycle than job destruction has led Caballero and Hammour (1991) to argue that the cyclical fluctuations—in particular, recessions—are needed for the creative destruction described in Schumpeter's theory.

Figure 9 attempts to summarize the implications of Davis and Haltiwanger's research for the question about the effect of economic fluctuations on long-term growth. The figure presents a job creation curve and a job destruction curve, each showing the sensitivity to fluctuations in real GDP. As shown in Figure 9, when real GDP equals potential GDP ($y = 0$), job creation in the U.S. economy is greater than job destruction; thus, job creation is positive as the number of jobs in the economy grows. As real GDP falls below potential GDP, job destruction increases and job creation falls. Observe that the job destruction curve is steeper than the job creation curve, corresponding to Davis and Haltiwanger's findings. However, their finding that the slope of the job destruction curve is steeper than the slope of the job creation curve does not indicate that recessions are needed to increase productivity growth. Even with the steeper slope of the job destruction curve as shown in Figure 9, there appears to be little need for recessions to cleanse the economy. A considerable amount of job destruction occurs in normal years when real GDP equals potential GDP.

Figure 9

Response of Job Creation and Job Destruction to Deviations of Real GNP from Potential GDP



Moreover, as Figure 9 illustrates, the effect of output fluctuations on real GDP growth is related to the *size* of the fluctuations in real GDP—booms as well as recessions, because the level of unemployment cannot be affected by stabilization policy. Without non-linearities in the job creation and job destruction curves in Figure 9, larger fluctuations in real GDP around potential would not increase the amount of structural adjustment. More job destruction in recessions would cancel out with less job destruction in booms. The average would be the same, regardless of the size of the fluctuations. In any case, the evidence is mixed on the effects of recessions on long-term productivity growth.

Conclusion

This paper has explored the trade-off between the variability of inflation and that of output. The trade-off exists because of the slow adjustment of prices; monetary policy can determine where on the trade-off curve the economy lies. Although the trade-off is more abstract than the old Phillips curve trade-off, the simple graphs presented in the paper are meant to provide a better understanding of the trade-off and

why it exists. This approach is meant to complement ongoing econometric work on estimating such trade-offs.

The paper has also compared this variability trade-off with several other types of trade-offs relating to inflation and output. While no long-term relationship exists between inflation and the deviation of real GDP from potential GDP, inflation seems to have strong effects on productivity growth and therefore on the growth of potential GDP. Evidence was also presented that casts doubt on the idea that larger fluctuations in real GDP would increase the growth of potential GDP.

A useful extension of this paper would be to examine whether indirect evidence can be found for variability trade-offs; preliminary empirical work looking at different historical periods in the United States and other countries indicates that a negative trade-off may be difficult to find. Perhaps this is because, throughout history, countries have been far from the trade-off curve because of inefficient monetary policies.

References

- Bryant, Ralph, Peter Hooper, and Catherine Mann. Eds. 1993. *Evaluating Policy Regimes: New Research in Empirical Macroeconomics*. Washington, D.C. The Brookings Institution.
- Caballero, Ricardo and Mohamad L. Hammour. 1991. "The Cleansing Effect of Recessions." National Bureau of Economic Research Working Paper No. 3922.
- Davis, Steven and John Haltiwanger. 1990. "Gross Job Creation and Destruction: Microeconomic Evidence and Macroeconomic Implications." *NBER Macroeconomic Annual*, Volume 5, pp. 123-68.
- Fischer, Stanley. 1993. "The Role of Macroeconomic Factors in Growth." *Journal of Monetary Economics*, vol. 32 (December), pp. 485-512.
- Fuhrer, Jeffrey and George Moore. 1993. "Monetary Policy and the Behavior of Long-Term Real Interest Rates." FEDS working paper, Board of Governors of the Federal Reserve System.
- Lucas, Robert E., Jr. 1994. "On the Welfare Costs of Inflation." Paper presented at Center for Economic Policy Research-Federal Reserve Bank of San Francisco conference, March.
- Motley, Brian. 1994. "Growth and Inflation: A Cross-Country Study." Paper presented at Center for Economic Policy Research-Federal Reserve Bank of San Francisco conference, March.
- Schumpeter, Joseph A. 1939. *Business Cycles: A Theoretical, Historical, and Statistical Analysis of the Capitalist Process*. New York: McGraw Hill.
- Taylor, John B. 1979. "Estimation and Control of a Macroeconomic Model with Rational Expectations." *Econometrica*, vol. 47, no. 5 (September), pp. 1267-86.
- . 1993a. "Discretion versus Policy Rules in Practice." *Carnegie Rochester Conference Series on Public Policy* no. 39 (Autumn), pp. 195-214.
- . 1993b. *Macroeconomic Policy in the World Economy: From Econometric Design to Practical Operation*. New York: W.W. Norton.
- Uchitelle, Louis. 1994. "Why America Won't Boom." *The New York Times*, June 12, Section 4.

Discussion

Laurence M. Ball*

I have four sets of comments on John Taylor's paper.

First, the paper asks exactly the right question for policymakers: What are the trade-offs in choosing among policy rules? Policy discussions often focus on the short-run Phillips curve, which describes the trade-off facing policymakers at a given moment. When considering monetary strategy, however, we need to ask how different rules affect the stochastic behavior of the economy. And if we believe the natural rate hypothesis—so policy does not affect average output—then the variance of output and the variance of inflation are the right variables to focus on. We need to know how alternative policies affect these two variances.

Second, I very much like Taylor's methodology, his use of a simple, textbook-style model. The model consists of three linear equations that we can understand *fully* and use to build intuition. Current research tends to emphasize rigorous microeconomic foundations and quantitative accuracy, with the result that models are very messy. Often the models are too complicated to understand, and we lose track of the basic economic forces at work. Microfoundations and quantitative work are certainly desirable, but they should come *after* simple models, so we know what we are seeking foundations for or trying to quantify. We need more research in the style of Taylor's paper.

My third set of comments concerns the model itself. The first two equations of the model are conventional I-S and Phillips-curve equations. These equations are deservedly popular among applied researchers because they capture behavior that we see in actual economies. The third equation of the model is more novel. It is a description of monetary policy: The Fed varies the interest rate to offset deviations of output from

*Professor of Economics, The Johns Hopkins University.

its natural level and deviations of inflation from some target. Is this a good specification?

The answer to this question depends on whether we are considering normative or positive issues. Taylor's framework is nice for normative analysis: We can derive the optimal coefficients on output and inflation in the interest-rate rule. The model is less useful as a positive description of policy, because it assumes a constant inflation target π^* . This assumption implies that inflation reverts to a fixed mean, whereas actual inflation has had an important random walk component in recent decades. In the history of actual policymaking, much of the interesting action is changes in the Fed's inflation target. Paul Volcker, for example, reduced the target in the early 1980s, and Arthur Burns increased it by accommodating supply shocks in the 1970s. To explain monetary policy, we need to understand the reasons, both economic and political, for changes in inflation targets. Taylor's model does not address this issue.

Taylor shows that his interest-rate rule fits the data well for the period since 1987. However, this period happens to be one in which the Fed's inflation target was fairly stable. I doubt that Taylor's equation would fit over longer periods that include shifts in the target.

The Trade-off between the Variances of Inflation and Output

My fourth and longest set of comments concerns Taylor's central conclusion: Policymakers face a trade-off between the variance of inflation and the variance of output. In Taylor's model, a policymaker's job is to choose a point on this trade-off—to choose how much inflation variance to accept to reduce output variance. In contrast to Taylor, I doubt that this is the right way to think about policy. I am not sure that policymakers really face Taylor's trade-off.

To explain why, I must distinguish between different kinds of macroeconomic shocks. It appears that different trade-offs arise from demand shocks (shocks to the I-S equation or the policy rule) and supply shocks (shocks to the Phillips curve). I will consider these two cases in turn.

In the case of demand shocks, Taylor's model does not support his conclusion: There is no trade-off between the two variances. To see this point, consider the following versions of Taylor's equations (7) and (8). (For simplicity, I set the inflation target π^* to zero and assume that the Fed's target for the real interest rate, r^f , equals the equilibrium rate.)

$$y_t = -c\pi_t + 1/(1 + \beta g)[u_t - \beta v_t]. \quad (7)$$

$$\pi_t = (1 - \alpha c)\pi_{t-1} + \alpha/(1 + \beta g)[u_{t-1} + v_{t-1}] + e_t. \quad (8)$$

These equations imply that policymakers can completely eliminate the effects of demand shocks on *both* inflation and output. They do so by

choosing a very large value of the parameter g : With a large g , the coefficients on the demand shocks u and v approach zero in both (7) and (8). A very large g means that interest rates respond very strongly to output, that is, that policy is very countercyclical. Note that the choice of a very large g does not constrain the parameter c , which determines the effects of supply shocks: Policymakers can always adjust the parameter h to obtain their desired c . Taken literally, the model says that policymakers can costlessly eliminate all effects of demand shocks. Thus, demand shocks do not create any painful trade-off.

In the real world, of course, it is not trivial to eliminate the effects of demand shocks. Problems arise from time lags and uncertainty about the effects of policy, which lead to mistakes. One could add these problems to the model, for example by assuming that certain parameters are unknown. In this case, demand shocks would cause fluctuations in output and inflation. And a huge g would no longer be optimal, because extremely cyclical policy would magnify the effects of mistakes.

Nonetheless, I still doubt that the main problem facing policymakers is a trade-off between output variance and inflation variance. When shocks hit the I-S or interest-rate equations, the main job of policymakers is to minimize the resulting fluctuations in aggregate demand. This task requires that they choose the right degree of countercyclicality: Demand fluctuates excessively if policy is too passive, but also if it is too aggressive and creates large mistakes. A successful policy—one that reduces the variance of demand as much as possible—reduces the variances of *both* inflation and output. Thus, it is less important to weigh the relative costs of the two variances than to develop effective means for reducing both. As Michael Dukakis would put it, the key issue for policymakers is competence, not ideology.

Now consider policymakers' response to supply shocks, which Taylor emphasizes. In this case, the model does imply a policy trade-off: A lower choice of the parameter c reduces the variance of output but increases the variance of inflation. To interpret this result, consider the optimal policy for someone who cares more about output stability than about inflation stability. In Taylor's model, such a policymaker would set c low, accepting large inflation fluctuations to keep output stable. When an adverse supply shock occurs, a low c means that policy is very accommodative and inflation rises far above π^* . Inflation is eventually returned to π^* through tight policy, but this disinflation occurs slowly.

Are Taylor's theoretical results a good guide to practical policy? Can we really stabilize output when supply shocks occur by destabilizing inflation? Two issues make me doubtful. First, in Taylor's model the total output loss from an adverse supply shock is *not* reduced by accommodative policy. Equations (7) and (8) imply that the loss from a unit supply shock, summed over time, is $c\sum_{i=0}^{\infty}(1 - \alpha c)^i$. This expression reduces to $1/\alpha$, and thus is independent of the policy parameter c .

Intuitively, inflation must always return to π^* in the long run, and so the cost of a supply shock is eventually paid in lost output. Non-accommodative policy implies a large recession when a supply shock occurs. Accommodative policy reduces the initial recession, but output is lost when inflation is brought down after its initial rise.

How can we reconcile this result with the result that accommodation reduces the variance of output? The answer is that accommodative policy spreads the output losses from a supply shock over time, whereas non-accommodative policy concentrates the losses when the shock occurs. By spreading the output losses, accommodative policy reduces the sum of *squared* deviations of output even though the sum of absolute deviations is unchanged. Thus, Taylor's result depends on his quadratic loss function. It is crucial that we ascribe greater welfare costs to two points of lost output in one year than to one point in two years. It is not clear whether this assumption is reasonable, and so the benefits of accommodation are unclear.

My second worry about Taylor's result concerns the key parameter α —the slope of the short-run Phillips curve. Taylor's model assumes that α is invariant to policy, but this assumption may fail in important ways. Some empirical evidence suggests that the costs of reducing inflation are smaller if disinflation is quick (Ball 1994). If this is so, then accommodative policy creates large output losses during the long, slow process of bringing π back to π^* . At a deeper level, I think the inflation inertia captured by the Phillips curve arises from the adaptive nature of inflation expectations, which in turn arises because changes in inflation are usually quite persistent. If policy became very non-accommodative, then inflation would revert quickly to its mean after a shock. If people learned they were in this new regime, expectations would become less backward-looking, and this might reduce the costs of stabilizing inflation.

For these reasons, it is not clear to me that accommodative policy really helps to stabilize output. In contrast, the cost of accommodative policy—greater variability of inflation—is clear. So perhaps our presumption should favor non-accommodation. I do not have a firm conclusion, because the benefits of accommodation depend on unresolved issues, such as the shape of the social welfare function and the determinants of the Phillips-curve slope. We need more work on these issues. Again, Taylor's paper is a valuable first step because it asks the right question and presents a tractable model.

Reference

- Ball, Laurence M. 1994. "What Determines the Sacrifice Ratio?" In N. Gregory Mankiw, ed., *Monetary Policy*. Chicago, IL: University of Chicago Press, pp. 155–82.