A Benefit Cost Analysis of Disinflation

by

Christopher J. Neely*

and

Christopher J. Waller**

August 26, 1996

* Economist, Federal Reserve Bank of St. Louis
** Associate Professor of Economics, Indiana University

The authors would like to thank three anonymous referees for their comments. The views in this paper are the authors and not those of the Federal Reserve Bank of St. Louis or the Board of Governors of the Federal Reserve System.

A Benefit Cost Analysis of Disinflation

by

Christopher J. Neely*

and

Christopher J. Waller**

August 26, 1996

This article presents a benefit-cost analysis of disinflation as well as critiques of two common methods for computing those statistics. The costs of disinflation are measured through "sacrifice ratios" (Ball (1994), the quantity of output lost for each percentage-point reduction in the inflation rate. The benefits of disinflation are calculated by Robert Barro (1995), who estimates the growth benefits from lower inflation. Ball's estimates are theoretically objectionable and shown to be very fragile to minor changes in technique, but worthy of study because of the interest they have elicited. Despite the uncertainty associated with Ball's sacrifice ratio estimates and Barro's growth estimates, relating them does permit us to obtain a rough measure of the net cost of disinflation. Our findings indicate that, contrary to popular opinion, disinflation probably produces a net benefit, not a net cost and the output losses associated with a typical U.S. disinflation are likely made up in 10-20 years -- less than one generation.

A BENEFIT-COST ANALYSIS OF DISINFLATION

CHRISTOPHER J. NEELY and CHRISTOPHER J. WALLER*

In recent years, the debate on establishing price stability as the sole objective of monetary policy has intensified. Advocates of price stability argue that monetary policy has only a transitory effect on real variables, such as output or unemployment. Consequently, they argue, the monetary authority should focus on price stability to reap the benefits of a stable and predictable price path. Because the current U.S. inflation rate is positive, the Federal Reserve would have to reduce the rate of inflation -- disinflate, in other words -- to stabilize the price level. Opponents of solely focusing on price stability argue that, because of various rigidities in the economy, such a policy would generate a recession whose costs would exceed the benefits.

One way to measure the costs of disinflation is to estimate "sacrifice ratios," the quantity of output lost for each percentage-point reduction in the inflation rate. A serious problem with sacrifice ratios, however, is that they are not true cost-benefit ratios: They gauge the costs of lost output, but do not assess the benefits from the lower rate of inflation. Consequently, sacrifice ratios overstate the net costs of disinflation. See Ball (1994) and Fuhrer (1994), for example.

Although the concept of sacrifice ratios is not new, the recent pursuit of price stability by central banks around the globe has renewed interest in, and research on, the subject. This article presents a benefit-cost analysis of disinflation as well as critiques of two common methods for computing those statistics. Sacrifice ratios measure the costs of disinflation. The study first replicates Ball's (1994) estimates of these ratios for the United States and other industrialized countries and then extends these results by changing the methodology in several plausible ways, including three alternate measures of trend output and two methods of timing disinflations. Ball's estimates are theoretically objectionable and shown to be very fragile to minor changes in technique, but worthy of study because of the interest they have elicited.

The benefits of disinflation are calculated by Robert Barro (1995), who estimates the growth benefits from lower inflation (or, perhaps more accurately, the costs of inflation). Even though Barro's estimates are better constructed, from an econometric perspective, than Ball's estimates, the authors of this study remain skeptical of imputing a causal interpretation to them.

Despite the fair amount of uncertainty associated with Ball's sacrifice ratio estimates and Barro's growth estimates, relating them does permit us to obtain a rough measure of the net cost of disinflation. Our findings indicate that, contrary to popular opinion, disinflation probably produces a net benefit, not a net cost and the output losses associated with a typical U.S. disinflation are likely made up in 10-20 years -- less than one generation.

II. THE CONCEPT OF A SACRIFICE RATIO

A. *Stylized Time-Series Plot of Inflation and GDP--Ideal Case*

The sacrifice ratio is the cumulative loss of output during a disinflation episode as a percentage of initial output divided by the cumulative reduction in the inflation rate. Thus, a sacrifice ratio of three implies that a one-point reduction in the trend inflation rate is associated with a loss equivalent to 3 percent of initial output. The easiest way to understand sacrifice ratios is to look at a stylized picture of how GDP and inflation evolve over time, according to a simple, textbook model of the economy, in which changes in money growth have real effects.

Figures 1A and 1B illustrate the reaction of the economy over time to a decrease in the growth rate of money. Figure 1A shows output initially growing steadily over time along the trendline, Y. Likewise, in Figure 1B the inflation rate is at π_0 until time t_0, when the monetary authority slows money growth to reduce the inflation rate to π_1. If the slowdown in money growth causes a recession, then output would fall below its trendline. After individuals in the economy adjust to this slower money growth by raising wages and prices less rapidly, output returns to trend while inflation stays at its new steady-state value. The output loss from this disinflation policy is the difference between trend output and actual output -- illustrated by the shaded area in Figure 1A -- divided by the reduction in the inflation rate (π_0 - π_1). If the disinflation policy did not cause a recession, output would stay on the trendline while inflation drops from π_0 to π_1. In this case, the sacrifice ratio would be zero; no output is lost.

B. ***What Does a Sacrifice Ratio Measure?***

An aggregate supply (AS)-aggregate demand (AD) diagram can illustrate the scenario described above for the case of unexpected disinflation. In Figure 2A, the economy initially is at its long-run equilibrium ($AD_0 = SRAS_0 = LRAS_0$). (Although the AS-AD model usually has the price level on the vertical axis, nothing fundamentally changes by respecifying the model with the inflation rate on the vertical axis.) A policy-induced disinflation -- a reduction in the trend growth rate of the money stock -- would shift the aggregate demand curve down to AD_1 as the monetary authority attempts to reduce inflation to π_1. If private agents fail to adjust their wages and prices to reflect this reduction in aggregate demand, nominal wages will grow too quickly, raising real wages. This, in turn, reduces employment, output falls to Y_1 and the economy enters a recession. Over time, wages adjust to the lower level of aggregate demand; the short-run aggregate supply curve shifts to the right (see Figure 2B) until it reaches long-run equilibrium at the point $AD_1 = SRAS_1 = LRAS_0$. During the second period, output returns to its trend value by the end of the period as inflation falls to π_2.

The sacrifice ratio -- $[(Y_n - Y_1)/Y_n]/(\pi_0 - \pi_2)$ -- is often loosely referred to as the (inverse) slope of the aggregate supply curve. (Denoting the slope of the short-run aggregate supply curve by g and the slope of the demand curve by d, then it can be shown that the sacrifice ratio is equal to $(1/Y_n)[1/(g - d)]$.) Countries with large estimated sacrifice ratios are often viewed as having "flat" aggregate supply curves, while countries with small estimated sacrifice ratios are thought to have "steep" aggregate supply curves. According to this interpretation, any variable that influences the slope of the short-run aggregate supply curve will affect the size of the sacrifice ratio. For example, Ball, Mankiw and Romer (1988) argue that high rates of inflation will be associated with a steep aggregate supply curve as people adjust wages and prices more quickly. Thus, the initial value of inflation influences the magnitude of the sacrifice ratio.

In the cases above, an unexpected reduction in the money growth rate was assumed to *cause* a decline in output. It is possible, however, to have an imaginary sacrifice ratio from some exogenous shock, such as an oil price hike, that both reduces output and induces the monetary authority to initially raise money growth to cushion the shock and then slow the rate of money growth to reduce inflation. If inflation declines as output falls, the fall in output will appear to be due to a policy-induced disinflation, despite the fact that the original exogenous shock caused both.

Figure 3 illustrates a simple version of this case. Consider a competitive economy with no frictions or information imperfections, the short-run and long-run aggregate supply curves are the same vertical curve; the economy is always in equilibrium. Now, suppose that a temporary, exogenous, aggregate supply shock shifts the LRAS curve to the left. In the absence of a reduction in aggregate demand, output would fall to Y_1 and the inflation rate would increase (temporarily) as there is suddenly too much money chasing fewer goods. (The authors sympathize with the idea that inflation should be defined as a sustained rise in prices, not any change in the price level. For simplicity, they abuse terminology here.) But if the monetary authority responds to this rise in inflation by contracting aggregate demand to AD_1, then a new equilibrium is reached with the inflation rate at π_1 and the level of output at Y_1.

Although monetary policy reduced inflation, it did not "cause" output to fall. Hence, the true sacrifice ratio is zero. Nevertheless, there appears to be a positive sacrifice ratio associated with this disinflation, given by $S = [(Y_n - Y_1)/Y_n]/(\pi_0 - \pi_1)$. Failing to account for the shift in aggregate supply causes the sacrifice ratio to yield an estimate of the (inverse) slope of an imaginary short-run aggregate supply curve, denoted '$SRAS_0$' in Figure 3. Sacrifice ratio estimates can be misleading if one fails to disentangle supply from demand shocks.

III. CALCULATION ISSUES -- MEASURING SACRIFICE RATIOS

A. *Time Series Plots of U.S. Quarterly Inflation and GDP*

Armed with the concept of a sacrifice ratio and some idea of what it is trying to measure, how straightforward is it to apply this knowledge to the real world? In a nutshell, it is not very easy. Figure 4 plots the behavior of quarterly U.S. inflation rates and Gross Domestic Product (GDP). (The inflation values are the quarterly averages of the monthly Consumer Price Index -

Urban (CPI-U) series.) Obviously, these plots do not resemble the stylized graphs of output and inflation in Figures 1A and 1B. Thus, to estimate sacrifice ratios, researchers must make some assumptions about the world to calculate the trend level of GDP and the trend inflation rate, as well as to identify periods of disinflation.

Ball (1994) systematically investigates cases of significant disinflations since 1960 over a wide range of countries, obtaining a rich set of estimates. This paper examines this method because it does not impose symmetry on the relationship between inflation and output, can be applied simply and mechanistically to a variety of countries, and has produced widely cited estimates of the costs of disinflation. See Posen (1995) or Fischer (1996) for two studies using Ball's estimates.

B. *Estimating Trend Inflation*

An initial step in calculating sacrifice ratios is to identify disinflation episodes, periods during which a *significant* decrease in the rate of inflation occurs. A problem, however, is that the actual quarterly inflation rates fluctuate substantially over short periods of time. Therefore, Ball smoothes the inflation rate series with a nine-quarter moving average to extract the longer-run movements. The top panel of Figure 4 displays the estimated quarterly trend inflation rate and the actual rate. In identifying disinflation episodes attributable to monetary causes, Ball (1994) examined only disinflations that were greater than 2 percentage points.[1] Ball maintains that historical records support his assumption that these episodes represented deliberate disinflationary policies. Even if policy induced, however, these episodes are almost certainly not exogenous and so may confuse real and monetary effects on output, as discussed in section II. Applying this procedure to the entire trend inflation series gives us a set of disinflations. For each disinflation episode, the denominator of the sacrifice ratio is the distance from the trend inflation rate at the peak to the trend inflation rate at the trough.

C. *Estimating Trend GDP*

A measure of trend output is necessary to estimate the output loss from the disinflation. The problems in estimating the trend rate of output are much greater than those involved in estimating the trend rate of inflation. Specifically, it is not clear whether GDP fluctuates around a stable linear trend or a changing trend, or whether its growth is completely unpredictable.

To estimate trend GDP, Ball first assumes that actual output is at its trend value when trend inflation is at its peak. Ball justifies this by arguing that trend output is often defined as that level of output for which inflation is neither rising or falling, and that this condition is satisfied when trend inflation is at a peak. Second, Ball assumes that output is once again back at trend four quarters after inflation reaches its trough. He bases this latter assumption on the argument that the economy continues to adjust even after inflation reaches its new trend level. Trend output is defined by the line connecting the actual level of output at the start of the disinflation with the actual level of output four quarters after inflation reaches its trough. The total output lost, as a percentage of initial output during a disinflation, is the annualized sum of actual GDP each period subtracted from its estimated trend value. This sum is divided by initial output to obtain the numerator of the sacrifice ratio. The second panel of Figure 4 shows output levels along with the measure of trend output (during disinflations) using Ball's approach. The vertical lines indicate the beginning and end of disinflations.

IV. ESTIMATING THE SACRIFICE RATIO

It is important to stress that the assumptions used to estimate the inflation and output trends are not minor; the sacrifice ratios will be very sensitive to changes in these assumptions. The robustness of these estimates is a serious matter: They may overestimate the true sacrifice ratio if they ignore the role of real shocks to the economy, or they may underestimate the true cost for other reasons. Using other estimates of trend output dramatically illustrates that sacrifice ratios are really back-of-the-envelope calculations and are subject to a great deal of uncertainty.[2]

A. *The United States and OECD Countries: 1960-94*

Table 1 presents estimates of sacrifice ratios and disinflation episodes for the United States and eight other OECD countries since 1960, for which quarterly data is available since 1960, calculated by Ball's (1994) methods. (Quarterly data are used here to avoid imprecision with annual data.) In addition to the sacrifice ratios, Table 1 contains the starting and ending dates of each disinflation episode, the level of initial inflation and the total amount of disinflation. The reported sacrifice ratios in Table 1 are typically higher for U.S. disinflations than they are for most OECD countries. The overall average sacrifice ratio was 1.43 while the estimates for the United States, Germany and Switzerland range from about 1.77 to 3.42. Some episodes produce negative sacrifice ratios, implying that output actually *increased* as the trend inflation rate was reduced. Table 1 shows, most other countries had disinflations at about the same time the United States did. All over the world, monetary authorities tried, to one degree or another, to cushion the output effects of the 1970's oil price shocks with expansionary monetary policy. Most, however, were eventually forced to bring inflation down with contractionary monetary policy.

B. *How Robust Are Sacrifice Ratio Estimates?*

Ball's procedure for estimating trend output assumes that output is at trend at the peak of the inflation process and returns to trend 4 quarters after inflation bottoms out. While he claims that this is consistent with the story one tells about the costs associated with disinflations, it is useful to examine the robustness of such an arbitrary assumption. Therefore, as a check, the sacrifice ratios were reestimated under two alternate ways to time the disinflations and three alternate measures of trend output.[3]

To consider how changes in timing affected the calculation of sacrifice ratios, consider first the effect of the (standard business page) assumption that monetary policy works through an output gap channel, a fall in output precedes a fall in inflation. Therefore, the sacrifice ratios were reestimated measuring the output loss starting four quarters *before* the disinflation began and ending at the end of the disinflation. The 8th column of Table 1 shows that most of the sacrifice ratios declined; the mean sacrifice ratio declined from 1.43 with Ball's methods to 0.53 under this alternative. The average sacrifice ratio of the United States also declined, but was still higher than average. Notably, the estimates for the "oil shock" disinflations actually rose for the United States. The second row of Table 2 shows the correlation between this measure and Ball's measure of the sacrifice ratio was only 0.23 across disinflations.

The second change to timing assumptions was that output has returned to trend by the end of the disinflation (instead of 4 quarters later). The 9th column of Table 1 shows that again, the mean sacrifice ratio declined from 1.43 to 0.79 under this assumption. The sacrifice ratios tend to be smaller for almost all the disinflations, including those of Germany, Switzerland and the United States. Table 2 shows that the correlation of this measure with Ball's measure was a little higher than that of the first measure at 0.79.

This paper uses three alternate measures of trend output: a linear trend, the Hodrick-Prescott (HP) filter, and a linear trend that controls for oil price shocks. The linear trend approach allows a structural break in the growth rate in 1973, but attributes all other movements from trend as "cycle." The Hodrick-Prescott filter (see Hodrick and Prescott (1981)) attempts to eliminate low frequency movements from the cyclical part of output by attributing some of the movements in GDP to movements in trend output. The obvious timing of disinflations in the data set with the 1973 and 1979 oil shocks motivated use of the oil price shock adjusted trend. Trend output is the predicted value of output from regressing log output on a time trend and the price of oil in the domestic currency to capture any trend output movements that would have occurred "naturally" in response to the oil shocks. It is necessary to control for oil shocks in order to avoid estimating the imaginary sacrifice ratios discussed earlier.

The 10th column of Table 1 shows that a linear time trend produces sacrifice ratios for the United States' last two disinflations that are half the size of Ball's estimates. The overall average sacrifice ratio falls to 0.83 from 1.43. The most dramatic changes appear to be the estimate for Canada during the 1980's, Germany in the 1970's and Australia in the 1960's. Ball's method estimates that Canada's sacrifice ratio for this disinflation period is near four while using a linear time trend reduces it to zero. The correlation with Ball's estimates is only 0.15.

The Hodrick-Prescott filter results shown in column 11 of Table 1 are similar to those of the linear trend. The clear finding from using this trend procedure is that, once again, the sacrifice ratios are much lower -- mean down to 0.59 -- relative to those of the Ball method.

Most of the sacrifice ratios are less than one and those over one appear to have occurred during the late sixties. Germany, Switzerland and the United States no longer appear to be outliers. Correlation with Ball's estimates is only 0.36.

Finally, the oil price adjusted output trend produces the most dramatic changes in the sacrifice ratios (last column of Table 1). These estimates actually reveal *negative* sacrifice ratios for the U.S. in the sixties and after the first oil shock and a value near zero for the disinflation in the early eighties. This can be interpreted to mean that monetary policy "softened" the fall in output from what it would have been in the absence of policy intervention for the first two disinflations. For the early eighties, policy was actually aimed at lowering inflation slightly further and faster than would have occurred naturally.   Again, the overall average sacrifice ratio (0.55) is about a third of that obtained with Ball's estimate. The correlation with Ball's estimates is actually *negative*, -0.24. The interpretation of data and the policy prescription implied by this model and Ball's model are startlingly different.

In summary, the estimates of sacrifice ratios are extremely sensitive to the timing of the disinflations and the method of estimate trend output deviations. Sacrifice ratio estimates produced by Ball's methods are much higher than, and not well correlated with those produced by other plausible methods. This paper continues to focus on estimates obtained by using Ball's method because they agree with commonly used estimates of the sacrifice ratio, but the authors do not consider any of these methods to be particularly reliable. The simultaneity problem precludes any reasonable way to make these simple methods reliable.

V. COST-BENEFIT RATIOS FOR DISINFLATION

A. *What Does Price Stability Appear to Cost Us?*

Ball's sacrifice ratios estimate that the cost of each percentage-point permanent reduction in U.S. inflation is about 2 to 3 percentage points of output growth lost. This implies the price for a move from 3 percent inflation to zero inflation is around 7.5 percentage points of initial output -- a heavy price in lost output and unemployment.[4] Such estimates convince most economists that the benefits of the move from 3 percent to zero inflation would not be worth the price of getting there. Referring to other (somewhat higher) estimates of sacrifice ratios, Paul Krugman summed up the conventional wisdom:

> "You have to sacrifice four points of output to reduce inflation by one point. That's a high price... It's hard to believe that anyone would be willing to pay the price of bringing the inflation rate down from 10 percent to 4 percent."
>
> - Paul Krugman, *Age of Diminished Expectations* (1994, Chapter 5, p. 64).

Yet, inflation is very unpopular when it exceeds modest levels, and the United States has repeatedly chosen to pay the apparently high price of disinflation. Furthermore, central banks all over the industrialized world are increasingly moving toward policies of price stability. Are voters and central bankers uninformed or irrational? Or do they perceive the benefits from lower inflation that are absent from the sacrifice ratio estimates? The next section presents estimates of the potential growth rate effects from lower inflation that suggest that voters and central bankers may not be irrational after all.

B. *The Effect of Inflation on Growth: The Benefits of Disinflation*

The sacrifice ratio measures the output cost of disinflation. But what are the benefits from disinflating? Economists have long argued that society benefits from reducing the inflation rate in ways that are well known and too numerous to list here (e.g., Fischer (1981)). The central theme is that the elimination of inflation leads to a more efficient allocation of resources and better decision-making by private consumers, workers and investors which can increase total output in two ways: 1) by an increase in the level of output after the disinflation, and/or 2) by an increase in the growth rate of output. (See Marquis and Einarsson (1994) for evidence that money creation may well have a long-term effect on output.) An increase in the level of output may result from one-time adjustments toward more efficient resource allocations arising from lowering the trend inflation rate (such as lengthening of nominal contracts). Output is permanently higher but the trend growth rate is unaffected.

The second possible effect of lower inflation is a higher growth rate. If some of the resources that are saved each period from more efficient allocation are invested in education, training, plants and equipment, then the economy may benefit after the disinflation. In this case, when the disinflation ends, output does not “jump” to a higher level but increases at a faster rate (the slopes are different). While this could go on forever, it may be more plausible to assume that output grows at this higher rate for a long period of time but eventually returns to its long-run growth rate. An increase in the average rate of rightward shift of the AS curve would illustrate it in the AS-AD graphical analysis employed before.

To measure the benefits from disinflation, this study relies on Barro's (1995) estimates of the growth effects of higher average inflation rates. Noting the recent move toward price stability, Barro has attempted to quantify the effects of inflation on growth within a regression framework suggested by a neoclassical growth model using a panel data set of more than 100 countries’ data over 1960-1990. The dependent variable is real per capita GDP growth over 10 year periods, while the independent variables included government consumption, measures of the rule of law, measures of property rights, private decision variables like fertility and the investment ratio, and initial conditions like measures of educational attainment and health at the beginning of the period.

Breaking down the data from each country into three subperiods, Barro examined the correlation between growth and its determinants, including average inflation rates and found that 10 percentage points of higher inflation was associated with 0.2 to 0.3 lower output growth in the data set in a statistically significant way.

Barro tested the robustness of the results in several ways. First, he reestimated the coefficients conditional on low (< 15%), medium (15 - 40%) and high inflation cases (> 40%) and found the statistical significance of these results hinged crucially on the inclusion of high inflation (above 40 percent) countries; with only low inflation countries, the sign and magnitude of the coefficient was similar, but was no longer statistically significant. For example, restricting the analysis to include only inflation rates below 15 percent reduces the point estimate of the cost of inflation to 0.016 percent per year, with a great deal more uncertainty associated with that estimate. Additional tests, however, were unable to reject the hypothesis that the relation is linear, the effect of inflation on growth was independent of the level of inflation. Table 3, excerpted from Barro (1995), shows some of his estimates conditional on inflation ranges and subperiod breakdown.   Second, Barro permitted the coefficient on inflation to vary over subperiods. Again, the hypothesis of equality of coefficients could not be rejected. Barro reports that the results also proved robust to the exclusion of outliers. The association of inflation with growth was found to be stable over inflation rates and time.

As with Ball's results, the fact that inflation and growth are determined simultaneously makes interpretation an important problem. As Barro (1996) implicitly notes, there are at least four ways inflation could be correlated with growth. First, under a classical model in which money is neutral -- see Kocherlakota (1996) -- but output is stochastically determined by exogenous shocks to supply, inflation and growth will vary inversely with one another.   Second, under a standard Keynesian model of demand shocks, inflation and output will vary positively. Third, inflation may be symptomatic of underlying disfunction in the economy, a failure of institutions, that could also influence growth. That is, there may be an omitted third variable such as better enforcement of property rights which contributes to growth and is also likely to put restraints on the monetary authority's ability to inflate excessively. Finally, the channel that Barro presumably considers is a direct effect on efficiency gains through transactions costs.

To show that higher average inflation was directly causing lower growth, Barro uses an instrumental variables approach. This technique attempts to avoid simultaneous equations bias and permit a causal interpretation of the results with carefully chosen instruments. A good list of instrumental variables should be highly correlated with the regressors but uncorrelated with the error term of the structural equation.   Barro followed the common practice of considering lagged values of the regressors as instruments and, in addition, considered adding legal measures of central bank independence and prior colonial status to the list of instruments due to their correlation with inflation. (Prior colonial status is defined by the most recent colonizing country if the nation was not independent in 1776.) A variety of specifications using ordinary least squares (OLS) and instrumental variables (IV) gave similar estimated coefficients on inflation.

There are two good reasons to be suspicious of this interpretation, however. First, the carefully chosen instruments may well still be correlated with the error term in the structural equation and therefore subject to simultaneity bias. For example, Kocherlakota (1996) cites Tamura (1995) to argue that prior colonial status inflation is likely to be correlated with the error term in the structural equation because colonial status determines a host of institutions including a linguistic-cultural network of economic partners that are likely to determine growth. Similarly, lagged inflation may not be a good instrument if inflation is generated by a stationary process. In this case, lagged inflation may proxy for a host of omitted variables that affect growth.

Sims (1996) made a different point, one fundamentally critical of single equation methods. Even if the instruments for inflation are correctly specified, the coefficient on inflation can't be interpreted as the whole result of a policy-induced change in average inflation because inflation is likely to have an effect on output growth through many channels and the regression equation fixes the effect of other variables at their "average" value. That is, a change in inflation may change many other variables such as the black market premium and government consumption, through the budget constraint. Indirect effects through changes in these other variables are very real, yet not accounted for by the coefficient on inflation. Sims recommended mapping the structure of the model into a multiple equation system.

The combination of possible endogeneity of the instruments and the inherent difficulties with single equation methods leaves the causal interpretation that Barro would like to make open to question. If one accepts such an interpretation, then the results imply a 1 percentage-point increase in inflation will reduce real per capita output growth by 0.02 to 0.03 percentage points per year. Alternatively, one could say that a 1 percentage-point decrease in inflation increased per capita output by the same amount.

C. *Timetable of Benefits*

Although most people would consider the size of benefits to growth (0.02-0.03) percentage points a year to be small, these calculations do not indicate that reducing inflation does not benefit society. On the contrary, although the effect on average growth rates appears very small, it will quickly add up over time. Barro calculated that a permanent increase in the average inflation rate of 10 percentage points would tend to lower the level of a country's per capita output by 4 to 7 percent over 30 years. (Barro assumed that the rate of output growth eventually returns to its natural level (in 30 years) after the increase in inflation, but his assumptions are not completely clear. The calculations here discount output at a rate of 0.96 and assume that output growth returns to its long-run value in exactly 30 years.) These estimates can be used to calculate the increase in output each period from a 10 percent decline in the permanent rate of inflation, shown in Table 4.

Another way to consider the output benefits from lower inflation that would be similar to sacrifice ratios is the accumulated discounted gain in output over a period of years from a disinflation. Table 5 shows the *accumulated* gains as a percentage of the original level of output arising from a 10 percentage-point decline in the trend inflation rate. (Table 5 tracks accumulated output increases for 1,000 years -- not, because it is realistic to project that far, but rather to show the limits to output gains.) The values in Table 5 are based on the midpoint of Barro's growth effect estimates (.025), which implies that a 10 percentage-point decline in trend inflation will add 0.25 percentage points to the growth rate of output. Table 5 shows that over a period of 10 years, total discounted output is approximately 12 percent higher than it would be under the higher inflation rate.

D. *Comparing Sacrifice Ratios to Benefit Ratios*

The sacrifice and benefit ratios each tell only half the story. (Here, the term "benefit ratio" denotes the net present value of output gained from the long-term growth benefits of disinflation, rather than Cecchetti's (1994) use of it to denote a short-term output gain from an increase in inflation.) To determine whether disinflation would lead to higher output in the long run, one can compare the costs implied by Ball's sacrifice ratios to the benefits implied by Barro's estimates of the growth costs of inflation. Since sacrifice ratios measure *cumulative* output losses from disinflations, it would be appropriate to compare them to the accumulated gains from higher growth shown in Table 5. If the point estimate of an average sacrifice ratio is 2 percentage points, then the total output cost of a 10 percentage-point disinflation is 20 percentage points of lost output. Table 5 shows that if Barro's growth calculations are correct, the economy will regain this output loss within about 15 years. In other words, if inflation has even small effects on output growth, a policy of disinflation is likely to more than make up for the output losses calculated by sacrifice ratio methods.

Nevertheless, there is substantial uncertainty related to both the size of sacrifice ratios and the benefits of disinflation. The estimates in Table 6 suggest some idea of the importance of this uncertainty. These estimates relate the size of the sacrifice ratio to the number of years it takes to recoup the initial loss of output arising from disinflationary policies using estimates of 0.005 percent to 0.05 percent as the additions to growth from a one point decline in the trend inflation rate. Using the midpoint of Barro's estimates of growth benefits (0.025 percent) and relatively small sacrifice ratios (less than one in magnitude), the authors conclude that any losses from a recession will likely be recouped in less than a decade. Larger sacrifice ratios, such as those estimated by Ball's methods for the United States and Germany (about 2 percent), will probably be recovered in less than a generation.

This timetable of benefits is independent of the methods used to measure the costs and benefits of inflation. That is, whatever the source of the estimated sacrifice ratios -- they could be from a general equilibrium model -- this table compares costs associated with moving to a lower trend inflation rate with the potential benefits (if there are any) of increasing the growth rate of GDP as a result. If disinflation is an "investment" (incurring sunk costs today, for higher output in the future) then the table provides the intertemporal recovery time.

VI. CONCLUSIONS

The objective was to examine and compare two common methods of calculating the costs and benefits of moving an economy toward price stability. The costs were calculated using the concept of a sacrifice ratio and recent techniques employed by Ball (1994). The benefits were calculated using recent work by Barro (1995), which measured the effects of inflation on average GDP growth rates. There is substantial uncertainty associated with each method's estimates, and still more in applying the estimates to current U.S. monetary policy. Ball's (1994) sacrifice ratio methods proved particularly fragile to plausible, small changes in assumptions. Nevertheless, combining the two permits a step toward a balanced benefit-cost analysis of disinflation.

Our best guess is that an economy, such as the U.S. economy, can usually recover the costs associated with disinflations in less than one generation. This suggests that policies aimed at moving toward price stability are potentially beneficial for society. In some sense, the disinflation problem presents us with the same sort of quandary as a fiscal deficit. Current voters bear the costs of responsible policy but some of the benefits will accrue to their children. If this is the case, it argues for putting monetary policy in the hands of forward-looking policymakers who can look past the next 12 months and take the long view.

The profession could advance its understanding of the consequences of disinflationary policies in two ways. The first would be through further research to test the robustness of the simple cost-benefit estimates presented here. The second and better way would be to look at the welfare implications of disinflation in a fully articulated, general-equilibrium model. These lines of research would be useful to policymakers, who must decide whether to make price stability the primary objective of monetary policy.

REFERENCES

Ball, Lawrence, “What Determines the Sacrifice Ratio?,” *Monetary Policy*, University of Chicago Press, Chicago, IL, 1994, 155-182.

Ball, Lawrence, N. Gregory Mankiw and David Romer, “The New Keynesian Economics and the Output-Inflation Tradeoff,” *Brookings Papers on Economic Activity, 1*, 1988, 1-65.

Barro, Robert, “Inflation and Growth,” Federal Reserve Bank of St. Louis *Review*, May/June 1996.

Barro, Robert, “Inflation and Economic Growth,” Bank of England *Quarterly Review*, May 1995, 166-176.

Cecchetti, Stephen, “Comment on “What Determines the Sacrifice Ratio?,” *Monetary Policy*, University of Chicago Press, Chicago, IL, 1994, 188-193.

Faust, Jon and Eric M. Leeper, “When Do Long-Run Identifying Restrictions Give Reliable Results?,” *Federal Reserve Board of Governors, International Finance Discussion Paper, 462*, March 1995.

Fischer, Andreas, "Central Bank Independence and Sacrifice Ratios," *Open Economies Review*, January 1996, 5-18.

Fischer, Stanley, “Towards an Understanding of the Costs of Inflation,” *Carnegie-Rochester Conference on Public Policy, 15*, 1981.

Fuhrer, Jeffery, “Optimal Monetary Policy and the Sacrifice Ratio,” *Goals, Guidelines and Constraints Facing Monetary Policymakers*, Proceedings of the Federal Reserve Bank of Boston Conference, 1994, 43-69.

Hodrick, Robert and Edward Prescott, “Post-War U.S. Business Cycles: An Empirical Investigation,” *Carnegie-Mellon University Working Paper, 451*, 1981.

Kocherlakota, Narayana, “Discussion of Inflation and Growth,” Federal Reserve Bank of St. Louis *Review*, May/June 1996.

Krugman, Paul, *The Age of Diminished Expectations*, MIT University Press, Cambridge, MA, 1994.

Marquis, Milton H. and Tor Einarsson, "Optimal Disinflation Paths When Growth is Endogenous," *Federal Reserve Board of Governors, Finance and Economics Discussion Series, 94-32*, 1994.

Mayes, David and Bryan Chapple, “The Costs and Benefits of Disinflation: A Critique of the Sacrifice Ratio,” Reserve Bank of New Zealand *Bank Bulletin*, March 1995, 9-21.

Quah, Danny, and Shaun P. Vahey, “Measuring Core Inflation,” *The Economic Journal*, September 1995, 1130-1144.

Posen, Adam, “Central Bank Independence and Disinflationary Credibility,” unpublished manuscript, Federal Reserve Bank of New York, 1995.

Sims, Christopher, “Comment on Robert Barro’s Inflation and Growth,” Federal Reserve Bank of St. Louis *Review*, May/June 1996.

Tamura, Robert, “Human Capital and the Endogenous Evolution of Specialization,” *University of Iowa Working Paper*, 1995.

[1] This study has replicated Ball's sacrifice ratio estimates quite closely, but not exactly, because of revisions to the data. Here the authors consider disinflations greater than 1 percent (instead of 2 percent as Ball used) because there were disinflations near the cutoff, that Ball included, that they did not wish to exclude.

[2] Using a vector autoregression (VAR) approach to estimate the sacrifice ratio, Cecchetti (1994) argues that the confidence intervals around these estimates include zero and negative values. See Mayes and Chapple (1995) for another exposition of the problems with sacrifice ratios.

[3] Ideally, one would like to test the robustness of sacrifice ratios to changes in the measure of trend inflation. The authors were unable to find international data on weighted, median or core measures of trend inflation, however, and are quite dubious of measures coming from long-run restrictions on a VAR. See Quah and Vahey (1995) and Faust and Leeper (1995).

[4] An old rule of thumb known as Okun's law suggests that each point of output lost costs about 0.4 percentage points of unemployment. Hence, losing 7.5 percentage points of annual output growth would mean losing almost 3 percentage-point years of unemployment.

TABLE 1
Fragility of Sacrifice Ratios for Quarterly Data

						Measures of the Sacrifice Ratio By Method					
							Timing		Output Measure		
Country	Start Date	End Date	Episode Length	Initial Inflation	Fall In Inflation	Ball's Method	Start 4 Q Before Peak	End At Trough	Trend Stationary	HP Filter	Oil Price Filter
Australia	1965 I	1968 IV	15	3.8	1.3	4.33	1.51	5.37	-0.40	1.44	4.55
Australia	1974 II	1978 I	15	15.5	7.1	0.65	-0.63	-0.45	0.09	0.21	-3.80
Australia	1982 I	1984 I	8	11.0	5.3	0.79	0.62	0.69	1.56	0.82	0.13
Australia	1986 II	1993 I	27	9.0	7.9	-1.44	-1.68	-2.38	-0.66	-0.06	3.72
Canada	1981 II	1987 II	24	11.8	7.7	3.90	1.85	2.63	0.05	0.34	0.32
France	1974 III	1977 I	10	12.4	3.1	1.78	0.99	1.60	-0.22	0.71	-0.20
France	1980 IV	1987 I	25	13.6	11.0	0.71	0.20	-0.11	0.35	0.30	0.72
Germany	1965 III	1968 I	10	3.5	2.0	2.04	-3.97	-1.07	2.66	1.89	8.07
Germany	1973 III	1977 IV	17	7.0	4.0	3.42	1.12	2.67	0.04	0.76	-5.91
Germany	1980 III	1987 I	26	6.0	5.9	1.99	0.90	0.36	2.19	0.54	1.85
Italy	1975 II	1978 II	12	18.5	5.3	-0.39	1.89	-0.58	0.30	0.81	0.29
Italy	1980 II	1987 II	28	20.3	15.6	1.58	0.17	1.08	-0.49	0.12	0.16
Japan	1962 II	1964 III	9	7.6	2.9	-0.56	0.46	0.69	0.54	0.07	19.31
Japan	1965 II	1966 III	5	6.2	2.5	0.63	1.62	0.32	2.11	1.82	6.98
Japan	1970 III	1971 II	3	7.0	1.7	1.48	-1.99	0.24	0.54	-0.35	-17.89
Japan	1974 I	1978 III	18	18.5	14.5	0.53	0.95	0.31	0.34	0.34	-3.66
Japan	1980 II	1987 II	28	6.5	6.6	1.20	-0.02	-0.21	0.42	0.33	0.82
Switzerland	1973 IV	1977 IV	16	9.5	8.5	1.77	0.32	1.30	0.88	0.20	0.84
Switzerland	1981 III	1986 II	19	6.3	5.2	2.34	2.18	2.12	2.59	0.40	0.11
UK	1975 I	1978 II	13	21.3	11.0	0.90	0.41	0.39	-0.10	0.04	-1.02
UK	1980 II	1983 III	13	16.7	12.3	0.52	1.24	0.53	1.63	0.43	0.84
UK	1984 II	1986 III	9	6.4	3.2	0.79	0.61	0.29	1.02	0.67	1.38
US	1969 IV	1971 IV	8	5.6	1.9	2.94	0.23	0.62	3.48	1.58	-2.26
US	1974 I	1977 I	12	10.1	4.1	1.98	2.15	1.56	1.02	1.06	-1.83
US	1980 I	1985 IV	23	12.5	9.9	1.88	2.16	1.91	0.89	0.34	0.19
					Mean Sacrifice Ratios	1.43	0.53	0.79	0.83	0.59	0.55
					Standard Deviation of Sacrifice Ratios	1.34	1.42	1.48	1.09	0.58	6.17

TABLE 2
Correlation Matrix of the Sacrifice Ratio Measures

	Ball's Method	Peak 4 Q Before Start	End at Trough	Trend Stationary	HP Filter	Oil Price Filter
Ball's Method	1.00					
Peak 4 Q Before Start	0.23	1.00				
End at Trough	0.79	0.58	1.00			
Trend Stationary	0.15	-0.04	-0.12	1.00		
Hodrick Prescott Filter	0.36	0.04	0.19	0.54	1.00	
Oil Price Filter	-0.24	0.05	-0.04	0.12	0.29	1.00

TABLE 3
Barro's Estimates Of The Effects Of Inflation On Real Per-Capita Growth

Estimation Procedure	Inflation	Breakdown	Coefficient On	Standard
actual inflation in OLS	none	whole sample	-0.024	0.005
instrument: prior inflation	none	whole sample	-0.020	0.007
instrument: prior colonial status	none	whole sample	-0.031	0.008
actual inflation in OLS	< 15 %	whole sample	-0.016	0.035
actual inflation in OLS	15 to 40 %	whole sample	-0.037	0.017
actual inflation in OLS	> 40 %	whole sample	-0.023	0.005
actual inflation in OLS	none	1965-1975	-0.019	0.015
actual inflation in OLS	none	1975-1985	-0.029	0.010
actual inflation in OLS	none	1985-1990	-0.023	0.005

TABLE 4
Percentage Increase in Annual Output From a 10 Percent Decline in Inflation (Not Accumulated)

Year	Percentage
1995	0.2
2004	3.1
2014	8.0
2024	15.5
2034	19.9

TABLE 5

Increase in Discounted Accumulated Annual Output as a Percentage of Original Output From a 10 Percent Decline in Inflation

Years	Percentage
5	3.4
10	11.8
20	40.8
50	159.4
100	271.8
1000	362.5

TABLE 6

Time (in Years) to Recover for a Given Sacrifice Ratio and Change in Growth

	Change in Growth					
Sacrifice Ratio	.005	.01	.02	.03	.04	.05
0.0	0	0	0	0	0	0
0.5	14	9	6	5	4	3
1.0	22	14	9	7	6	5
1.5	28	18	12	9	8	7
2.0	34	22	14	11	9	8
2.5	40	25	16	13	11	9
3.0	47	28	18	14	12	10
3.5	56	31	20	16	13	11
4.0	65	34	22	17	14	12
4.5	76	37	23	18	15	13
5.0	90	40	25	19	16	14

FIGURE 1
Stylized Measurement of Sacrifice Ratio

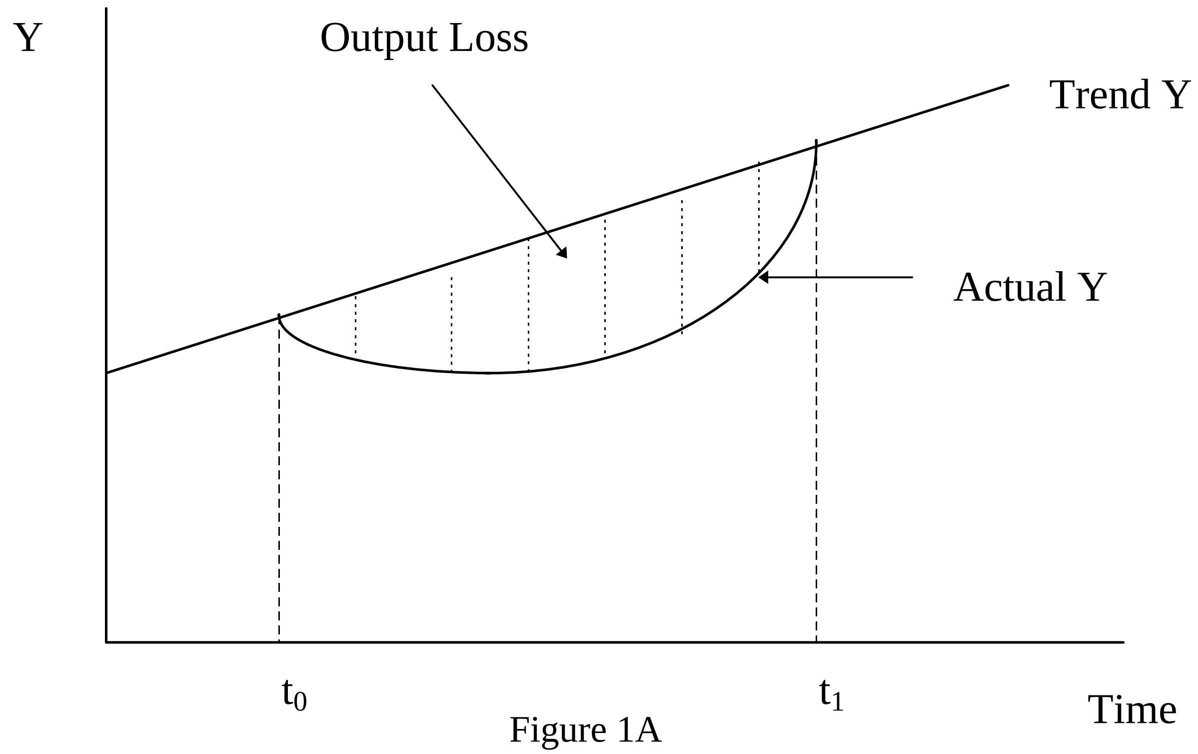

Figure 1A

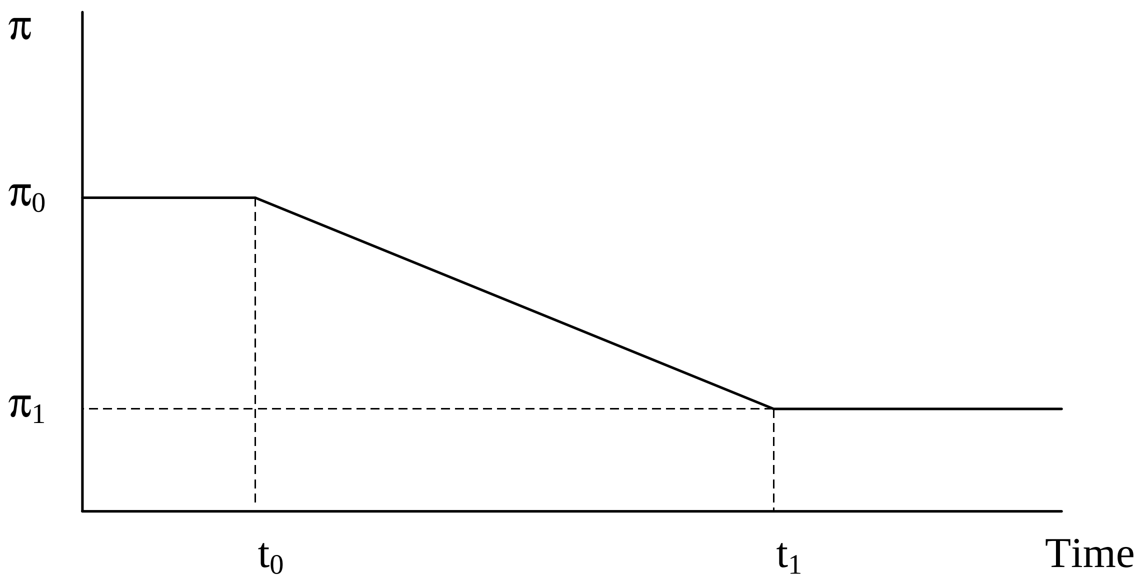

Figure 1B

FIGURE 2

Sacrifice Ratios in an Aggregate Demand-Aggregate Supply Framework

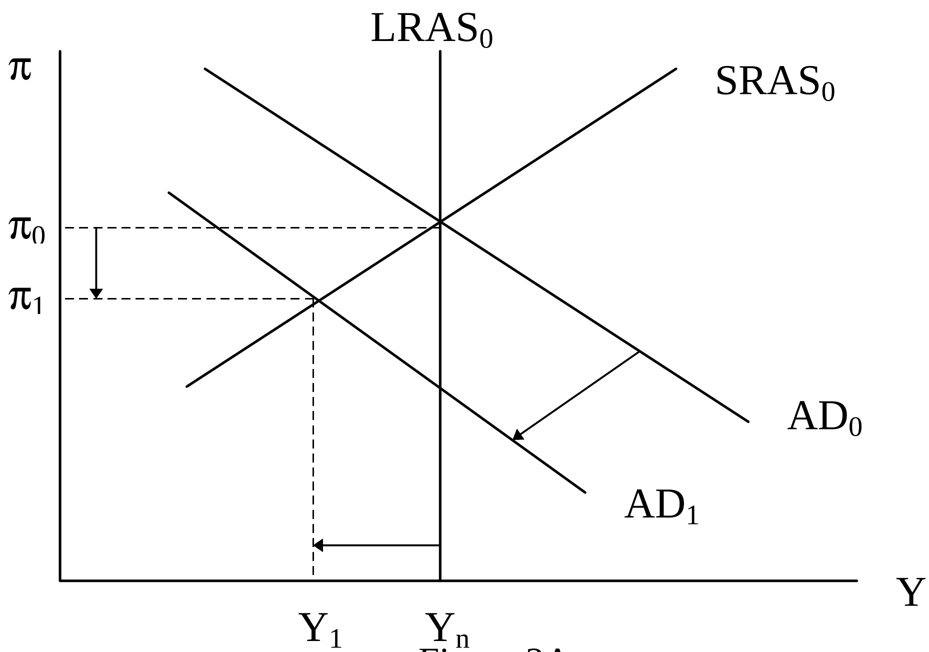

Figure 2A

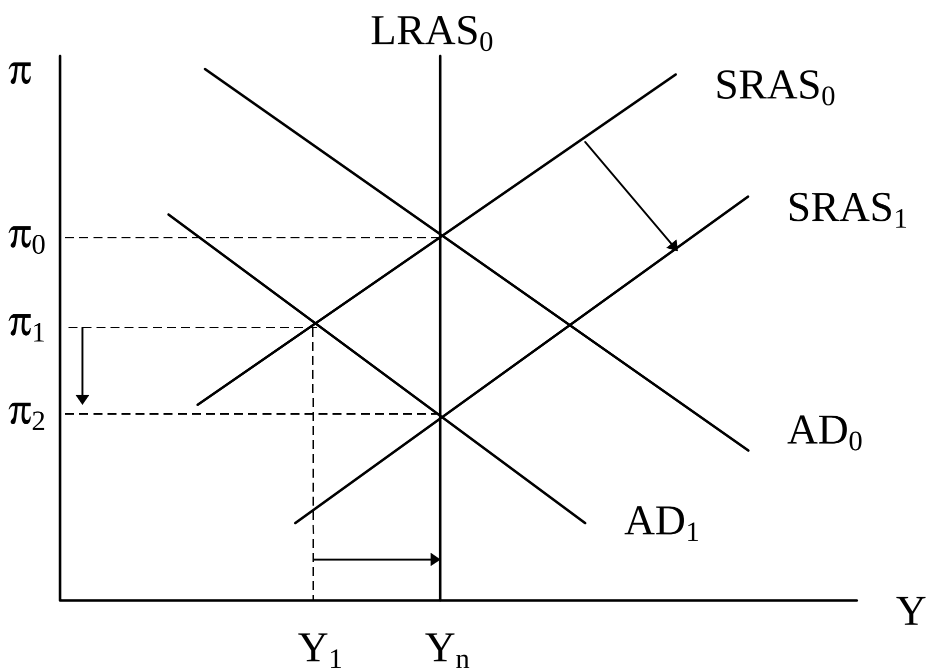

Figure 2B

FIGURE 3

Imaginary Sacrifice Ratio From an Adverse Supply Shock

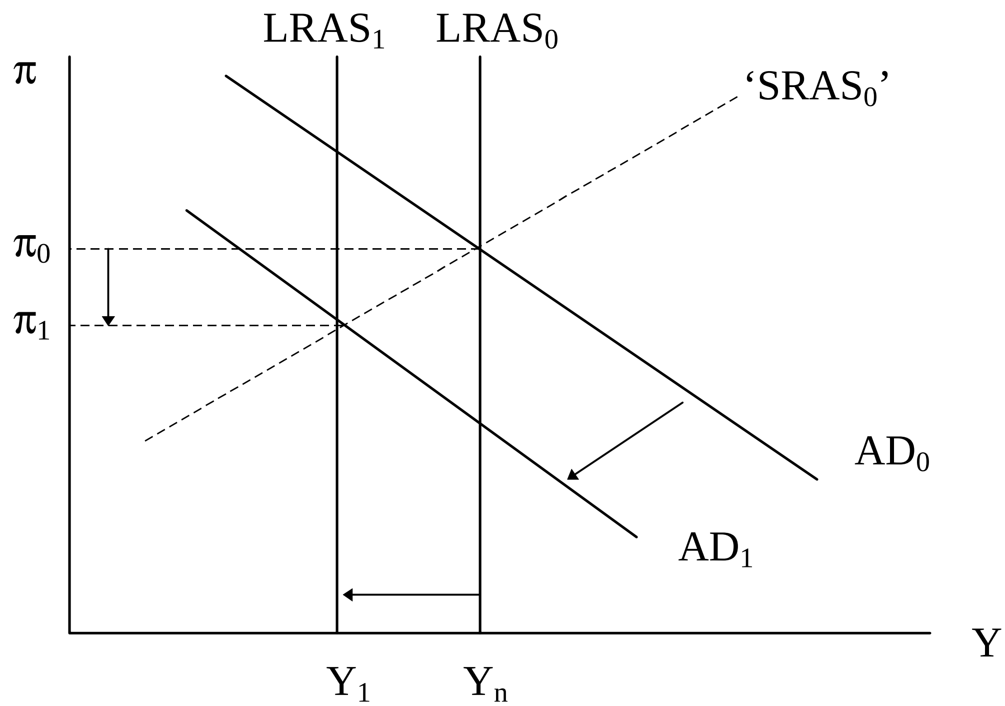

FIGURE 4

U.S. Inflation, GDP and Trend Values