

Edward M Gramlich: Maintaining price stability

Remarks by Mr Edward M Gramlich, Member of the Board of Governors of the US Federal Reserve System, before the Economic Club of Toronto, Toronto, Canada, 1 October 2003.

The references for the speech can be found on the Board of Governors of the Federal Reserve System's website.

* * *

One of the remarkable economic developments of the past two decades is the sharp, worldwide reduction in rates of inflation. Before the 1980s, research papers, economic commentary, and textbooks here and abroad were full of discussions of the causes and consequences of high inflation and of the political difficulty of bringing it under control. It looked then like inflation had become a more or less permanent feature of the economic landscape. Concepts associated with deflation such as liquidity traps and the zero bound on nominal interest rates had, for practical purposes, disappeared from economic thought.

But, beginning around 1980, central banks around the world seem to have ganged up on inflation, fought the fight, and won. The International Monetary Fund's consumer price index for industrial countries registered an annual rate of inflation exceeding 12 percent in 1980, but by 2002 the rate had dropped to just 1.4 percent. In the United States, consumer price inflation excluding food and energy fell from 12.4 percent in 1980 to 2.4 percent in 2002. Japan has actually had negative rates of inflation for the past four years. The drop in inflation in many prominent emerging-market countries has been even more breathtaking (Rogoff, 2003).

Perhaps just as remarkable as the abrupt fall in worldwide inflation is the widespread agreement on the reasons for its fall. Virtually all economic analysts would give a primary role to the implementation of plain old economic theory. Macroeconomists had long understood that, in the famous words of Milton Friedman, "inflation is everywhere and always a monetary phenomenon." But somehow this theoretical agreement never was translated into action. Perhaps because of advances in understanding the inflation process, perhaps because of worrisome increases in inflation rates, perhaps because of a discovery that the economic costs of fighting inflation were not as high as previously believed, central bankers in the 1980s suddenly began to use their monetary control to limit the rate of increase in prices. It is not stretching things to call this a revolution in central banking.

Now that something like price stability has arrived, two important policy questions present themselves:

- Is a mere determination to achieve price stability adequate for central banks to stabilize prices, or should their determination be backed up by more formal institutional structures, such as in inflation-targeting regimes?
- Exactly how stable do we want prices to be--rising at a measured rate of zero or at some other rate?

Maintaining Price Stability

Price stability is widely understood to be a feasible objective and to be sustainable with little long-run cost in terms of output and unemployment. Indeed, many believe that greater price stability should also stabilize output, and that the improved allocation of resources coming from price stability should promote productivity and raise living standards in the long run. As a part of this consensus, central banks around the developed world are now firmly determined to achieve and maintain price stability.

But there is a significant debate about whether determination alone is enough. Much of the world--countries in the British Commonwealth (the United Kingdom, Australia, Canada, and New Zealand), in Scandinavia (Finland, Norway, and Sweden), in eastern Europe (the Czech Republic, Hungary, and Poland) and many emerging-market countries (Brazil, Chile, Colombia, Mexico, Peru, the Philippines, South Africa, and Thailand)--have buttressed this determination to maintain price stability by adopting formal inflation-targeting regimes. The regimes vary, but the key element is an explicit commitment to meet a publicly stated numerical target rate of inflation within a particular time frame. The targets and commitments are accompanied by reports giving inflation forecasts and plans for meeting inflation goals. All regimes provide some leeway in meeting their goals--inflation rates do not have to be brought down to the target levels immediately, and special factors can be taken into account. But the

regimes must have a credible plan for meeting inflation targets at least over the medium run. And because these plans are public, they can be embarrassing if not fulfilled. Recently the European Central Bank (ECB) has adopted a system that has some elements of an inflation-targeting regime, though it lacks many of the accountability features.

A few industrial countries--the United States, Japan, and Switzerland--have resisted pressures to adopt formal inflation-targeting regimes. Debates continue in these countries about the desirability of adopting more formal inflation-targeting procedures. The inflation targeters point to the advantages of transparency, commitment, and accountability; non-targeters point to the loss of flexibility and the reduction in the ability to meet alternative goals such as high employment and financial stability. Rudebusch and Walsh (1998) give a good summary of this debate.

To me, the verdict on theoretical grounds alone has always been a close call. I can see the advantages of transparency and commitment and of giving financial markets some indication about the central bank's inflation goals. I can also see the fear that, in practice, such regimes may in fact constrain actions in ways that could be rigid, inflexible, and perhaps unnecessary in bringing down inflation.

Rather than rehash these theoretical issues one more time, let me instead invoke some empirical evidence. There are two ways to evaluate the effect of formal inflation-targeting regimes. One is by time-series analysis; the other by cross-section analysis. In the early days of inflation targeting, a number of researchers took the first approach and compared the inflation performance of targeting countries before and after the inflation-targeting regime was introduced. In these studies, which measured actual inflation rates and inflation premiums on long-term bonds (to measure expected inflation), one could find evidence that inflation targeting seemed to have worked to bring down inflation and its persistence (see, for example, Kuttner and Posen, 1999; and Freeman and Willis, 1995).

But these early studies took place at the time of a worldwide drop in inflation rates. Just as inflation-targeting countries could show elements of improved inflation performance, nontargeting countries could as well. The second approach to studying the effect of targeting, cross-section analysis, compares inflation performance for targeting and nontargeting countries. The advent of the ECB created a huge impediment to such studies because it came into being midway through the relevant data period, and it initiated a centralized regime that had aspects of inflation targeting for many of the large countries that had previously not had such a regime. It is difficult to know whether the ECB countries should be considered inflation targeters in the last half of the time-series sample, and it is especially difficult to interpret the experience of these countries.

For what they are worth, the raw cross-sectional data are shown in table 1. Over the 1994-2002 period, seven prominent inflation-targeting industrial countries had a mean inflation rate of consumer prices of 2.1 percent per year with a standard deviation of 1.7 percent. Both numbers are slightly *higher* than the mean and standard deviation of the eleven prominent nontargeting industrial countries. Eliminating the labeled ECB countries from both groups leads to a targeting mean of 2.0 percent and a targeting standard deviation of 1.7 percent, with both numbers again slightly *higher* than the mean and standard deviation for the three nontargeters. One could consider Japan a special case and eliminate it from the latter group, but the two-country results would still show slightly *lower* and more stable inflation in the nontargeting countries.

Ingenious economists could devise much more elaborate tests. Ball and Sheridan (2003) have done that, also making an adjustment for initial rates of inflation. They again find no significant differences between targeting and nontargeting countries. Their conclusion is similar to an earlier cross-section finding of Cecchetti and Ehrmann (1999) that inflation aversion increased in both targeting and nontargeting countries.

To summarize these international data, one might say that something brought inflation down in the 1980s and 1990s, but success was fairly uniform across both the inflation-targeting and nontargeting countries. When one does "before and after" tests for individual inflation-targeting countries spanning this period, inflation targeting often looks successful in bringing about low rates of inflation and in reducing inflation variability and persistence. But by the late 1990s, inflation had fallen in all industrial countries, and it is hard to find much difference in inflation performance between inflation-targeting and nontargeting countries--if anything, inflation has been lower and more stable in the latter group. One is therefore tempted to conclude that the general understanding of the inflation process and a firm determination to achieve price stability, more than the inflation-targeting regime itself, has been the key element in reducing inflation.

How Stable Should Prices Be?

Collectively, the central banks of the world have done well in reducing inflation, generally moving from an inflationary regime to one of reasonably stable prices. But now the goal question gets harder--how far should these central banks go? All the way to exactly stable prices with a measured inflation rate that averages zero, or somewhat short of that goal? I will argue strongly that central banks should stop short of zero--that a low positive rate of inflation is optimal.

There are two strategic reasons for choosing a low, positive rate, and both have been debated extensively in the economic literature. One involves the so-called zero bound on nominal interest rates; the other involves labor markets. Regarding the zero bound, as inflation drops toward zero, nominal interest rates drop toward real rates, which themselves might be low if the reason inflation is dropping involves negative demand shocks. At these low rates, the central bank is poorly positioned to respond to further negative demand shocks. Because of the zero bound, it is impossible for nominal interest rates to decline further, but it is still possible for inflation to decline to below zero--that is, for deflation to set in. The combination implies that real interest rates may actually start to rise in the face of negative demand shocks. In this range, the central bank loses much of its ability to respond to these negative demand shocks by lowering real rates of interest. As many have pointed out, the central bank could still respond to negative demand shocks by expanding money growth in various "nontraditional" ways. But these nontraditional approaches are untested. Macro-model simulations indicate that if the central bank exercises only traditional policy instruments, the zero-bound problem could significantly increase the severity of recessions. The virulence of the Japanese recession also suggests that zero-bound problems could be serious. Today the standard thinking is that the central bank can and should avoid dealing with these difficulties by guiding the economy toward a low, positive rate of inflation, not a zero rate.

The second argument for avoiding mathematical zero involves labor markets. A good deal of evidence suggests that, in the economy at large, employers are reluctant to cut workers' nominal wages or, at least, downward wage rigidities exceed upward rigidities. The normal market process of reallocation of labor in response to industry or sectoral shifts works by having real wages adjust to keep them in line with workers' marginal products of labor. If nominal wages cannot be cut for workers who, for whatever reason, become less productive, labor market flexibility is enhanced by having low positive overall rates of inflation that effectively cut real wages for these workers.

To these strategic reasons for avoiding mathematical zero can be added a measurement reason. One of the great difficulties in compiling price indexes is in adjusting them for changes in the quality of goods and services. How does one measure the value of a new good or of a service that benefits from a new technology? Casual evidence indicates that failure to make accurate adjustments can create a bias that could be quite large--indeed, King (2002) gives some heuristic illustrations that suggest that the true prices of some consumer durables have not risen at all since as far back as 1915!

The statistical agencies compiling consumer prices in the United States--the Bureau of Labor Statistics for the consumer price index (CPI) and the Department of Commerce for the personal consumption expenditure deflator (PCE)--try valiantly to deal with these quality-change biases. But the weight of professional opinion is that, even after years of trying to make corrections, the indexes are still biased upward. Since no ready methods exist for correcting some of these biases, it seems appropriate simply to adjust for them. It is generally agreed that the biases are about 1 percentage point per year for the CPI and about 1/2 percentage point per year for the PCE deflator (Lebow and Rudd, 2003). Whichever inflation target one uses, there is again a measurement argument for steering away from mathematical zero and toward a measured rate of inflation of from 1/2 point to 1 point per year above the desired economic cushion.

A final argument for steering away from mathematical zero invokes popular preferences, always an elusive topic. When economists have asked people whether inflation is harmful and whether controlling it should be a high priority for government, overwhelming percentages of the population agree. But when people are probed further, their main reason for disliking inflation appears to be the perception that it hurts their standard of living (Shiller, 1997). Popular opinion seems to view personal wage increases as given--that is, not related to inflation--and focuses only on price increases, in which case inflation does look like a destroyer of real income. To economists, on the other hand, inflation should be interpreted to mean a general rise in *all* wages and prices; and the costs of general inflation are the "shoe leather"-type costs of managing cash balances, which the respondents to Shiller's survey viewed as trivial. Hence while inflation is clearly unpopular with the general public, it is not clear that the public is properly identifying the true costs of inflation when making its judgment.

A potentially more appropriate way to ascertain the true degree of inflation was suggested by Richard Ruggles of Yale University. I learned about this test in graduate school many years ago, and Ruggles may have made his suggestion many years before that.

What inflation rates should really measure is the decline in the utility value of a nation's currency. In principle the right test is to offer a sample of the population a constant amount of currency, say \$1,000, along with the opportunity to spend it on a menu of all goods and services available this year or a menu of all goods and services available a while back, say five years ago. If this sample of the population votes in equal numbers for this year's menu and for that available five years ago, one can conclude that prices have been stable over the five-year period. If the majority vote is for the earlier menu of goods and services, one can conclude that prices have risen, or that the utility value of the \$1,000 has decreased. If the majority vote is for the recent menu, one can conclude the reverse--that true prices have actually declined.

I have seen no rigorous polling evidence on this question. But for years in teaching college macroeconomics courses, and recently at the Fed, I have conducted such a poll among my audiences. All audiences have reported that their understanding of what inflation is all about was much improved by this thought experiment. Generally, college students have voted for the current menu even in times when the aggregate rate of price increase averaged 3 percent or more, implying that they felt that true prices had actually declined.

College students may be unusually influenced by fads that do not truly improve goods (narrow or wide ties, etc.), and the implicit bias in measured price indexes may well be overstated by collegiate polls. Since coming to the Fed, I have had the opportunity to talk to and poll many banker groups about inflation, and as one would expect, they are generally more inclined to vote for the earlier menu of goods than were my college students, at any given rate of inflation. But these days, when measured rates of inflation are running at 1.5 percent to 2 percent, even bankers consistently vote for the current menu of goods by fairly wide margins. If even bankers feel that the implicit measurement bias in price indexes exceeds 2 percent, that may be a phenomenon worth noting. The upshot of this highly anecdotal test is that I have long suspected that true price stability might really be achieved in the vicinity of measured inflation rates of 2 percent or even more.

It would be desirable to ground this type of information more firmly in modern-day techniques of data collection. Ideally, consumer survey groups would run polls of currency utility and make periodic reports, the way they already do for indicators of consumer confidence and spending plans. Short of that, there is one econometric calculation that provides some support to my informal voting results. Nordhaus (1998) compared real incomes as measured by the CPI with Survey Research Center data in which people were asked how their financial condition changed over the past year. He fitted a regression to the data, finding that equal shares of the population considered themselves better and worse off when the measured inflation rate was 1.5 percent, a result suggesting a 1.5 percent measurement bias in the CPI. Nordhaus's method requires that the entire income distribution move at the same rate. When Krueger and Siskind (1998) modified the test with more fine-grained distributional calculations, the estimated bias was reduced.

Summary

So where does this leave me on the two policy questions I promised to address? I find myself on the soft side of each of them.

The first question is whether the U.S. central bank should adopt a more formal inflation-targeting regime. I personally would not go that far. My reading of the empirical evidence is that the key ingredient in keeping inflation low and stable is that the central bank be firmly determined to achieve and maintain stable prices in the long run. I believe the Fed is already so determined, with every member of the Federal Open Market Committee (FOMC) since I have been here repeating this mantra often. To me, there does not seem to be huge value in further tying down the committee through a formal inflation-targeting regime, and there could be some costs.

On the other hand, it may be possible to get some of the transparency and accountability advantages of inflation targeting, and to lock in the gains from having reduced inflation, by going to an intermediate approach. The FOMC might simply announce its preferred long-run *range* for inflation. This range should be understood as a preferred range that would not bind the committee or override other important objectives of monetary policy. It should clearly be understood as a long-term objective, not a short-term objective. The FOMC would not have to defend any deviations from the preferred range.

Perhaps such a step would increase transparency without limiting central bank flexibility to any appreciable degree.

If we were to adopt a preferred range, what should it be? In light of the strategic considerations mentioned above, along with quantifiable measurement error, I would personally set the bottom of the range at slightly above 1 percent per year for the core PCE deflator, the Fed's preferred inflation measure. Because of audience polls, and at least until they are replaced by more rigorous information, I would set the top of the range at about 2.5 percent per year. The midpoint of this range is then slightly less than 2 percent per year, which turns out to be about what U.S. core PCE inflation has averaged over the past eight years. But I would stress the range more than the point estimate.

I might close by stressing again the uncertainty involved in answering both questions addressed in the paper. There is theoretical uncertainty about how well formal inflation targeting should work and empirical uncertainty about how well it has worked in those countries that have tried it. If we were to move in the direction of a more systematic approach, I personally would go to a preferred range for inflation rather than a particular target and without all of the other trappings of a formal inflation-targeting regime.

Finally, if the FOMC were to adopt a preferred range, I feel that efforts to quantify measurement bias in price indexes have, on the whole, been too conservative. My own personal preference is that the top of the range could go as high as 2.5 percent per year.

Table 1
Inflation Performance in Selected Industrial Countries, 1994-2002
 (Percent annual average; broad index of prices for all consumer goods)

Country	Mean	Standard deviation
<i>Inflation targeters</i>		
Australia	2.7	2.9
Canada	1.7	1.6
Finland (ECB)	1.6	1.3
New Zealand	2.1	1.8
Spain (ECB)	3.1	1.8
Sweden	1.2	1.5
United Kingdom	2.4	.7
<i>Nontargeters</i>		
Austria (ECB)	1.8	1.1
Belgium (ECB)	1.8	1.4
France (ECB)	1.4	1.0
Germany (ECB)	1.7	1.2
Ireland (ECB)	2.9	1.9
Italy (ECB)	2.9	1.4
Japan	.1	1.3
Netherlands (ECB)	2.6	1.3
Portugal (ECB)	3.3	1.3
Switzerland	1.2	1.5
United States	2.3	1.0
Memo:		
Mean, targeters	2.1	1.7
Mean, nontargeters	2.0	1.3
Mean, non-ECB targeters	2.0	1.7
Mean, non-ECB nontargeters	1.2	1.3

Note: ECB means that as of 1999 the country had joined the European Central Bank and had a common monetary policy.