

# The Case for Forecast Targeting as a Monetary Policy Strategy

Michael Woodford

**A**t central banks around the world, forecasts have come to play an increasingly important role both in policy deliberations and in communications with the public. The most striking examples are the Bank of England, Sweden's Riksbank, Norway's Norges Bank, and the Reserve Bank of New Zealand, all of which conduct policy on the basis of a procedure sometimes referred to as "inflation-forecast targeting" (Svensson, 1997, 1999). Under this approach, the central bank constructs quantitative projections of the economy's expected future evolution based on the way in which it intends to control short-term interest rates, and public discussion of those projections is a critical part of the way in which the bank justifies the conduct of policy to the public.

What accounts for the appeal of a forecast-targeting approach, and should it be adopted more widely or more explicitly? I begin by reviewing the long-running debate between the proponents of monetary rules, intended to ensure confidence in the value of money over time, and the proponents of discretionary monetary policy, aimed at stabilizing the real economy. I will argue that inflation-forecast targeting represents a powerful synthesis of the two approaches; in particular, it is an improvement both over simpler rules, such as targeting a monetary aggregate, and over weaker versions of inflation targeting. I shall also argue that a much more extensive communication policy is crucial to escaping from the limitations of the traditional alternatives of rigid rules or rudderless discretion.

I will then explore some common questions that arise about inflation-forecast targeting. Should only the inflation forecast matter, and if not, in what way should forecasts of other economic variables affect policy decisions? What assumptions about the course of future policy should be used in constructing the quantitative projections that are presented to the public? Finally, given that economic forecasts

■ *Michael Woodford is John Bates Clark Professor of Political Economy, Columbia University, New York City, New York. His e-mail address is <michael.woodford@columbia.edu>.*

always have an element of inaccuracy—and may even be subject to persistent bias for a time—is a forecast-based approach really as reliable as adherence to a simpler rule, if one’s primary concern is the avoidance of major errors like the Great Inflation of the 1970s? I conclude with some thoughts about how the U.S. Federal Reserve might move toward an explicit policy of inflation-forecast targeting.

### **Beyond the Antinomy of Rules versus Discretion**

I shall argue that inflation-forecast targeting represents a synthesis of two apparently antithetical conceptions of good monetary policy. The first ideal is the quest for a *monetary standard*, an arrangement that serves to guarantee a stable value for a particular monetary unit over the foreseeable future. The primary aim of a monetary standard is to install confidence regarding the future value of money. Because of the importance attached to instilling and maintaining confidence, rigid rules are often argued to be necessary so that the monetary authority’s conformity with the rules can constantly be made manifest. The second ideal is that of *monetary stabilization policy*, according to which judicious adjustment of monetary policy in response to varying economic conditions can help to stabilize the real economy or, more precisely, to ensure an optimal adjustment to economic disturbances despite lags in the adjustment of nominal wages and prices. Effective monetary stabilization policy requires flexibility in the way that policy may respond to differing circumstances, and as a consequence, the proponents of stabilization policy frequently argue for the importance of allowing central bankers considerable discretion in the decisions that they make at any given time.

These two ideals are often presented as irreconcilable. The choice is often described as being between “rules” and “discretion,” with each term being defined by its denial of the other. Yet a deeper analysis of the two ideals reveals less incompatibility between them. In fact, one of the principal benefits of a monetary standard—which is to say, of a regime that provides a clear anchor for expectations regarding the future value of money—is a greater scope for the use of monetary policy for short-run stabilization of the real economy.

Research in monetary economics over the past few decades—including the work for which Milton Friedman, Edmund Phelps, and Robert Lucas have now all been awarded Nobel Prizes—has shown that the short-run tradeoff between inflation and real activity, which provides the basis for the effects of monetary policy on real activity, depends critically on *inflation expectations*. An increased expected rate of inflation results in a higher rate of current inflation being required for a given level of real activity. It follows that if expectations are not firmly anchored, and thus are easily shifted in response to variations in the observed rate of inflation, then short-run variations in the rate of inflation will *not* produce substantial differences between current inflation and expected inflation, and hence will have only a small effect on real activity. However, if people have reason to believe that inflation will always return fairly quickly to a stable long-run rate, so that an observed departure of the current inflation rate from the average rate has little effect on expected

inflation for the future, the short-run “Phillips-curve” tradeoff between inflation and employment is much flatter, allowing monetary policy a larger short-run effect on real activity.<sup>1</sup> Hence even from the point of view of improved stabilization of the real economy, it is important to find a way of stabilizing inflation expectations.

Some might conclude from this that one goal for monetary policy should have an absolute priority over the other: that one should choose a monetary regime solely with a view to its consequences for long-run price stability, and allow the central bank to disclaim any responsibility for fluctuations in business activity. But there is no necessary conflict between a sensible degree of concern for stabilization of the real economy—one that is consistent with a sound understanding of what monetary policy can actually accomplish—and the pursuit of a policy that ensures stability of the value of money over the medium to long run. Suppositions to the contrary are based on too narrow a view of the possible bases for confidence in the future purchasing power of money.

A traditional view has been that confidence in a monetary standard should derive from a commitment to convertibility of the currency into something else of known value. Typically, this was a precious metal, gold having been the most popular choice, though there were also experiments with silver standards, bimetallic standards, and so on. The classical gold standard did succeed in keeping the cumulative increase in the general price level in the United States near zero from 1776 to 1914. But the relative price-inelasticity of both the supply of gold and of the nonmonetary demand for gold means that the relative price of gold can vary widely for many years at a time, so that stabilizing the dollar price of gold implies little stability for the prices of other goods and services, except over the very long run. Thus, this approach to the maintenance of confidence in the long-run value of the monetary unit has a serious cost: policy is completely subordinated to the achievement of a criterion (a fixed dollar value of gold) that is not very closely related to one’s true stabilization objectives, simply on the ground that conformity to this kind of commitment is easily verifiable.

After the final abandonment of any connection of the world’s currencies to gold in the early 1970s, a widely advocated alternative approach to guaranteeing the value of money was money-growth targeting. Proponents argued that both the likelihood of achieving inflation control and the visibility of a central bank’s commitment to controlling inflation would be increased by the adoption of an “intermediate target” for something that the central bank could more directly control than the rate of inflation itself, but that was nonetheless reliably connected to inflation, at least over long enough periods of time. Once again, however, choice of a target that allows relatively straightforward verification of compliance with the

<sup>1</sup> Bank of England Governor Mervyn King (2005) argues that tighter anchoring of the public’s inflation expectations has made possible greater stability of *both* real activity and inflation since the introduction of inflation targeting in the United Kingdom. Figure 7 of Benati (2006) offers evidence of shifts in the slope of the U.K. Phillips curve across alternative monetary regimes consistent with the mechanism sketched in the text.

commitment conflicts with the choice of one that is closely related to reasonable stabilization objectives, other than at very low frequencies.

A narrow definition of money, like the monetary base, can be closely controlled by the central bank, making verification of compliance simple; but the link between this measure of money and aggregate expenditure can fluctuate considerably. Broader definitions of money, like the M3 measure emphasized by the European Central Bank, are somewhat more reliably connected to the volume of nominal expenditure, but the central bank has less direct influence upon these broader measures of the money supply. Moreover, regardless of the narrowness or breadth of monetary aggregate that one chooses to target, the connection between stable money growth and stability of either inflation or aggregate expenditure has turned out to be shaky, except over fairly long periods of time. Money-demand relations have proven to be particularly unstable in recent decades (Friedman and Kuttner, 1996; Estrella and Mishkin, 1997), as a consequence of rapid innovation in financial arrangements, and as a result, few central banks currently pay much attention to monetary aggregates in judging the appropriateness of their policy stance.<sup>2</sup>

But are mechanical rules of these kinds the only way of creating confidence that a central bank will maintain a predictable rate of inflation over the medium to long run? The problem is surely not a lack of any other effective *means* through which a central bank that wishes to do so can guarantee a desired long-run average inflation rate. This problem is not a difficult one: as long as a central bank possesses a reasonably reliable measure of the price level (with only a modest delay) and some instrument that can influence the rate of inflation (at least eventually) with a known sign, then it is surely possible to offset any past tendency to drift away from the target inflation so as to ensure that no deviation persists unchecked for many years. The only real challenge—at least where long-run price stability is concerned—is one of making *visible* to the public the central bank's commitment to act in this way.

A mechanical rule has a certain advantage in this regard, since it should be possible to observe fairly directly whether it is being followed. But there is an alternative way of addressing the problem, which is to commit to *explaining* publicly the basis for the central bank's decisions. A more sophisticated policy—one that aims to stabilize the real economy to some extent in the short run, but in a way that is consistent with maintaining a relatively constant inflation rate over a period of several years—is possible in principle, and it should be possible to pursue such a policy without sacrificing the possibility of stable medium-run inflation expectations, as long as the central bank can show the public that its actions are consistent with a strategy of this kind and ensure that people understand what the consequences of such a strategy should be.

The key to avoiding the disadvantages of a mechanical rule, without allowing

<sup>2</sup> The European Central Bank is an obvious exception, at least as far as public rhetoric is concerned. But even there, the bank's "reference value" for M3 growth has played at most a minor role in actual policy decisions (Fischer, Lenza, Pill, and Reichlin, 2006).

the drift in inflationary expectations that occurs all too easily under a purely discretionary policy, is thus a commitment to *explain* the policy decisions that are made. While a good policy may well take into account the effects of policy on the real economy, it must do so in a way that does not imply instability in the medium-term inflation outlook. While a good policy may well respond to a wide range of sources of information about the economy's current state, it must do so in accordance with a consistent strategy that private decisionmakers can rely upon in forecasting future conditions. Above all, these aspects of the policy strategy must be made visible to the public. This is where forecast-based monetary policy has a crucial advantage: not simply in helping to improve the accuracy of central bank judgments about how to best achieve the bank's stabilization objectives, but also in explaining the character of that policy to the public.

## **Inflation-Forecast Targeting**

Since the early 1990s, inflation-forecast targeting has rapidly gained popularity as an alternative to monetary policy strategies based either on monetary aggregates or convertibility. When referring to inflation-forecast targeting, I mean not just the public announcement of an inflation target—though that is certainly a crucial element—but also a commitment to a specific structured approach to deliberations about monetary policy actions and a corresponding framework for communication about the justification for those actions. A central bank that practices inflation-forecast targeting is committed to adjusting its instrument or instruments of policy (typically, this means its operating target for an overnight interest rate) in whatever way proves to be necessary in order to ensure that the bank's quantitative *projections* of the economy's future evolution satisfy a specific *target criterion*.


For example, the Bank of England has often stated that its monetary policy is intended to satisfy the requirement that the projection for a particular measure of inflation (currently, one based on a consumer price index) equal 2.0 percent at a horizon eight quarters in the future. Although this description is plainly an oversimplification of the Bank's actions, each issue of the Bank's quarterly *Inflation Report* begins with an overview of the justification of the current stance of policy that contains two charts like those shown in Figure 1. The "fan chart" in Figure 1A indicates a probability distribution of possible future evolutions of GDP over a three-year horizon, while the fan chart in Figure 1B shows a probability distribution of possible future evolutions of inflation, with the modal projection indicated by the most deeply shaded region. Primary emphasis is given to Figure 1B in judging that the evolution of policy assumed in constructing the projections is suitable (more on that below!); the vertical dashed line at a horizon eight quarters in the future and the horizontal line at the inflation target of 2.0 percent help the eye to judge whether the path of deepest shading crosses the intersection.

This forward-looking decision procedure allows the central bank to use *all* available information about the current outlook for the economy, including non-quantitative information ("judgment"), in determining the appropriate level of


Figure 1

**The “Fan Charts” from the *Inflation Report* of the Bank of England**

A: Current GDP Projection based on Market Interest Rate Expectations  
(Percentage increase in output on a year earlier)



B: Current CPI Inflation Projection Based on Market Interest Rate Expectations  
(Percentage increase in prices on a year earlier)



Source: Figures 1A and 1B reproduce Charts 1 and 2 from the introduction of the May 2007 *Inflation Report of the Bank of England*.

Notes: The “fan chart” in Figure 1A indicates a probability distribution of possible future evolutions of GDP over a three-year horizon, while the fan chart in Figure 1B shows a probability distribution of possible future evolutions of inflation, with the modal projection indicated by the most deeply shaded region. The CPI is the Consumer Price Index.

interest rates. There is a specific target criterion, which favors both focus in the decision-making process and predictability of the policy committee's decisions, but the criterion involves the central bank's actual stabilization goals rather than an "intermediate target," like a monetary aggregate that is of little independent interest. Inflation-forecast targeting is not tied to a mechanical formula that makes monetary policy a function of some very small set of present economic variables (like a "Taylor rule" for monetary policy).<sup>3</sup> After all, the relation of current economic variables to the variable that one actually wishes to stabilize may change over time.

Inflation-forecast targeting also involves a commitment to regular publication of the projections on the basis of which policy decisions have been made, typically through reports (like the *Inflation Report* of the Bank of England) published several times per year. Such publications serve the goal of anchoring inflation expectations in several ways. First, they make the central bank's policy commitment verifiable, by allowing the public to see at frequent intervals that policy is still being conducted in a manner consistent with that commitment. In addition, they sharpen expectations about the likely future conduct of policy by allowing people to observe how the central bank processes and responds to developments of various types (the import of which for the bank's projections and decisions are discussed in the report). Finally, publication of the bank's own view of the future outlook for inflation can directly influence inflation expectations. In particular, a chart might show why a current inflation rate that is different from the target rate (and perhaps even moving in the wrong direction, as in the corresponding chart from the Bank of England's November 2006 report) is nonetheless consistent with an expectation that inflation will be close to the target rate within a few years, and this information can help to keep medium-run inflation expectations anchored, despite the high-frequency variations that tend to dominate press coverage. The justification of policy decisions by reference to quantitative projections is a crucial feature of this policy strategy, for these projections are expected to substitute for verification of convertibility (as under a gold standard) or verification of conformity with an "intermediate target" (such as a money-growth target) as a basis for the public's confidence in the future value of money.

Can inflation-forecast targeting really succeed in anchoring inflation expectations in the way that one demands of a monetary standard, while simultaneously allowing the flexibility required for a reasonable degree of short-run stabilization? Some critics suspect that to the extent that inflation-forecast targeting aims to serve as a monetary standard, it will inevitably be too rigid (for example, Friedman and Kuttner, 1996). Others argue that the attempt by inflation-forecast targeting central banks to leave themselves sufficient flexibility to pursue stabilization aims will inevitably undermine the credibility of their commitment to controlling inflation. To address this issue, it is necessary to look more carefully at how inflation-forecast targeting is practiced—and at some of the variations that exist in current practice.

<sup>3</sup> In John Taylor's (1993) celebrated rule of thumb for Fed policy, the federal funds rate should be a linear function of inflation over the previous four quarters and the current output gap.

## Should Only the Inflation Forecast Matter?

Those central banks that have been most explicit about their use of a forecast-targeting procedure (with the exception of the Norges Bank, discussed below) have generally given primary emphasis to the way in which the inflation forecast for a particular future horizon determines the policy decision. This emphasis makes it clear why one speaks of “inflation-forecast targeting”; but is a single-minded focus on inflation an essential feature of the forecast-targeting approach?

As a logical matter, a forecast-targeting approach might involve any of a wide variety of types of target criteria. Nonetheless, there are compelling reasons to choose a criterion that implies a clear commitment to some particular inflation rate (on average) over the medium-to-long term. The stabilization of inflation expectations offers clear benefits, some of which have been sketched above. Moreover, there is no good reason for the public’s inflation expectations to remain anchored other than the existence of a credible commitment on the part of the central bank; for both economic theory and bitter experience teach that the inflation rate *can* vary widely and for long periods of time, depending on the nature of policy. In comparison to fluctuations in inflation, low-frequency movements in the real rate of economic growth are relatively modest and are less obviously dependent on monetary policy.

Some have suggested that if a central bank cares about the stability of the real economy as well as inflation—or if inflation forecast-targeting were to be adopted in countries like the United States, where the Federal Reserve Act assigns the Fed the mandate of promoting “maximum employment” as well as “stable prices”—then there should be an employment target (or growth target) as well as an inflation target. But even if an output or employment stabilization objective is assigned equal weight with inflation stabilization among the *goals* that a monetary policy is designed to serve, it does not follow that equal prominence should be given to *targets* for output or inflation under an ideal decision framework. If by a “target” one means a fixed numerical value of the variable that a central bank should always be aiming to achieve over some medium term, then an inflation target makes a great deal more sense than an employment target. In the case of inflation, monetary policy can achieve pretty much any long-run average rate that is desired. In the case of employment or real activity, monetary policy has short-run effects, but can have little effect on the average levels of such variables over longer periods; hence a fixed long-run target could be futile and in any event unnecessary (insofar as expectations regarding such variables ought not to be based on central bank pronouncements regarding targets).

But the argument that only inflation should have a numerical target does not mean that projections of real variables should not be taken into account in monetary policy decisions. In practice, the medium-run target for inflation does not suffice in itself to determine the appropriate current policy action, owing to the possibility of alternative transition paths by which that medium run might be reached. For example, a central bank facing high inflation might seek to return inflation to the target level more quickly following a disturbance, or more gradu-



ally. In choosing between these alternatives, it is reasonable to take into account the associated alternative paths of real variables that can differ along the transition paths, even if monetary policy cannot affect the long-run levels of such variables.

Norway's Norges Bank has been the most explicit among current practitioners of inflation-forecast targeting about the way in which its target criterion involves real variables as well as inflation. Each issue of its *Monetary Policy Report* contains a box stating the criteria that the Bank looks for in an acceptable set of projections, described as the conditions that identify an appropriate policy intention—"an appropriate future interest rate path" (Norges Bank, *Monetary Policy Report* 2007/2, p. 12). The first criterion is that the inflation projection should show "inflation close to the target [of 2.5 percent per year] in the medium term"; but there is also a second criterion, which is that the projections "should provide a reasonable balance between the path for inflation and the path for capacity utilization." The two criteria are not intended as competing goals that must be balanced with one another; rather, the first indicates the situation that should eventually be reached (in a "medium term" that is not identified with a specific horizon), while the second describes the type of transition path by which it should be reached.

Short-term departures of the inflation rate from the medium-term target are justifiable to the extent that they are associated with a level of capacity utilization that is also temporarily different from its long-run level, in the direction such that faster convergence of inflation to the target rate would be possible only by keeping output away from potential to an even greater extent.<sup>4</sup> The Norges Bank checks whether its projections have this property by superimposing the inflation and output-gap projections in a single figure, like the one in Figure 2. This figure shows that in early 2007, Norway's inflation rate was below the 2.5 percent target level, but that at the same time, output was above potential. The projections show that over the next three years, with output above potential, inflation is projected to rise, eventually approaching the target rate, while at the same time the output gap is projected to be closed at a similar rate.

Other forecast-targeting central banks have been less willing to discuss openly the way in which projections of output growth or other real variables figure in their policy decisions, even though it is fairly obvious that they do, to some extent.<sup>5</sup> These

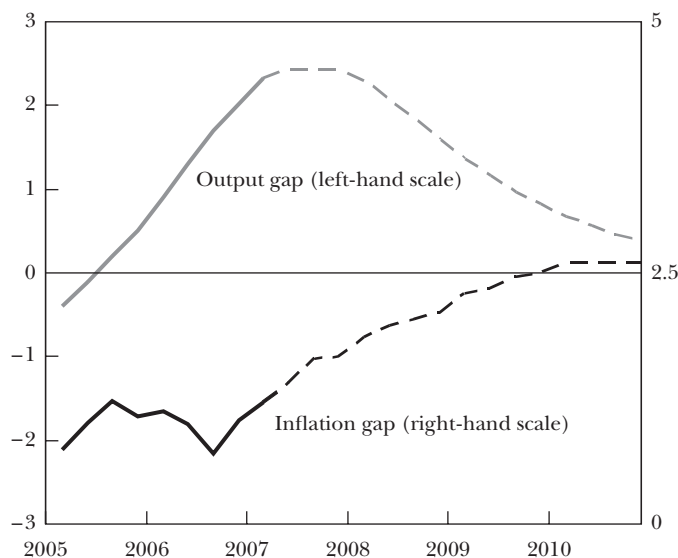
<sup>4</sup> In past issues, the second criterion has stated even more explicitly that "the inflation gap [the departure of projected inflation from the medium-run target] and the output gap [the departure of projected output from potential] should be in reasonable proportion to each other until they close," and "should normally not be positive or negative at the same time further ahead" (Norges Bank, *Inflation Report* 2006/3, p. 10).

<sup>5</sup> For example, in the overview at the beginning of each issue of the Bank of England's *Inflation Report*, the projection for GDP growth (reproduced in Figure 1A) is always presented *first*, and only subsequently the projection for CPI inflation (Figure 1B), though the summary justification then given for the most recent policy decision is much more explicit about how the inflation projection supports the decision. Proponents of inflation targeting frequently stress that inflation-targeting central banks have in practice never been solely concerned with inflation (Bernanke, Laubach, Mishkin, and Posen, 1999; King, 1999; Svensson, 1999). It is perhaps telling that in 2007 both the Norges Bank and Sweden's Riskbank changed the names of their *Inflation Reports*, which are now instead called *Monetary Policy Reports*.

Figure 2

### The Norges Bank Target Criterion: Proportionality between the Projected Inflation Gap and Projected Output Gap

(output gap and inflation gap in percent)



Source: This figure based on Chart 1.15 from the 2007/2 *Monetary Policy Report* of the Norges Bank.

Notes: The dashed part of each line indicates projections for the next three years under the Bank's baseline policy scenario. Each *Monetary Policy Report* also includes a fan chart for each of these variables (not shown here), as well as charts indicating how the projections would be different under specific alternative scenarios. The inflation gap plotted is for a consumer price index (CPI) that is adjusted for tax changes and that excludes energy prices.

banks make it clear that their inflation targets, in their understanding, do not require that inflation be projected to equal the target rate except at horizons two or more years into the future; but the banks are often vague about the criteria used to select among possible paths consistent with that medium-run state.<sup>6</sup> This practice makes it appear as though the inflation target alone determines the policy decision. But other considerations are actually in play, and explicit discussion of the criterion used to select among transition paths would not only increase the transparency of monetary policy, but would also increase the credibility of the central bank's commitment to its medium-run target by making it easier for the public to judge whether departures of the current inflation rate from the target are consistent with a policy that remains committed to that target.

<sup>6</sup> In Svensson's (1997) theoretical argument for inflation targeting, the criterion that expected inflation two years in the future equal the target rate suffices to determine the current interest rate decision completely, because in the simple model of that paper, interest rates have no effect on inflation except with a two-year lag. This is not literally true, however, in the empirical models used for policy simulations in central banks. The use of additional criteria by the Norges Bank, just discussed, makes it clear that there do remain additional degrees of freedom even after convergence of the inflation projection to the target rate is required.

## Is Forecast Targeting Intertemporally Consistent?

The degree to which publication of central bank projections can be expected to shape the expectations of private decisionmakers will depend on how credible these projections are as forecasts of the economy's likely evolution. Among the possible grounds for doubt is a tension inherent in the logic of the forecast-targeting procedure itself. Production of projections of the economy's evolution several years into the future requires that the central bank make assumptions about its conduct of policy not merely in the immediate future, but over the entire forecast horizon (and even beyond, in the case of a forward-looking model). But while the projections must specify policy far into the future each time they are produced, in each decision cycle policy is only *chosen* for a short period of time (say, for the coming month, after which there will be another decision).

Should this decision procedure be expected actually to produce the kind of future policy that is assumed in the projections? One might imagine, for example, a central bank wishing always to choose expansionary policy at the present moment, to keep employment high, while projecting that inflation will be reduced a year or two in the future, so that the expectation of disinflation will make it possible to have high employment with only moderate inflation. But if the procedure is one in which the disinflation is always promised two years in the future, private decisionmakers have no reason ever to expect any disinflation at all.

Thus one requirement for credibility of the central bank's projections is that the forecast-targeting procedure be *intertemporally consistent*; that is, the future policy that is assumed in the projections should coincide with the policy that the procedure itself can be expected to recommend, as long as those aspects of future conditions that are outside the control of the central bank turn out in the way that is currently anticipated. While this requirement may seem obvious, a number of apparently sensible approaches to forecast-targeting fail to satisfy it.

### Constant-Interest-Rate Projections

A popular approach in the early years of inflation-forecast targeting—used, for example, in the *Inflation Reports* of the Bank of England prior to August 2004—was to construct projections conditional upon a *constant interest rate* over the forecast horizon (Vickers, 1998; Jansson and Vredin, 2003). The appropriate current interest-rate decision was then taken to be the interest rate that, if expected to be maintained over the forecast horizon, would lead to projections satisfying the target criterion (for example, 2 percent inflation eight quarters in the future). This procedure had a number of advantages. First, a bank had only to consider variations in policy over a single dimension (alternative constant interest rates), with the consequence that a one-dimensional target criterion would suffice to identify the correct policy. Second, contemplated changes in the current interest-rate decision would be predicted to have nontrivial consequences, given that any change was expected (for purposes of the projection exercise) to be a permanent change. Finally, it was possible to construct projections without the bank's having to tip its hand as to the likely character of future policy.

But constant-interest-rate projections raise a number of conceptual problems (Goodhart, 2001; Honkapohja and Mitra, 2005; Leitemo, 2003; Woodford, 2005). The assumption that the nominal interest rate will remain fixed at some level, regardless of how inflation or other variables may evolve, is not sensible. Moreover, in forward-looking (rational expectations) models of the kind that are now beginning to be used by central banks, the assumption of a constant nominal interest rate often implies an indeterminate price level, so that it becomes impossible to solve uniquely for an inflation forecast under any such interest-rate assumption.<sup>7</sup> In models with backward-looking expectations, the model can be solved, but such policies often imply explosive inflation dynamics.<sup>8</sup> Such difficulties appear to have been a frequent problem with the constant-interest-rate projections of the Bank of England (Goodhart, 2001), which often showed the inflation rate *passing through* the target rate at the eight-quarter horizon, but not *converging* to it. Figure 3A provides an example. In such a case, it is not obvious why anyone should believe this policy is consistent with the inflation target, or expect that inflation expectations should be anchored as a result of a commitment to such a policy.

The most fundamental problem, however, is there will often be no reason to expect interest rates to remain constant over the policy horizon. Indeed, constant-interest-rate projections themselves often imply that the people making the projections should *not* expect the interest rate to be maintained over the forecast horizon. Consider, for example, the inflation projection shown in Figure 3A, a constant-interest rate projection on the basis of which the February 2004 Bank of England *Inflation Report* concluded that a 4 percent policy rate was appropriate at that time.<sup>9</sup> The figure shows that under the assumption of a constant 4 percent policy rate, consumer price inflation was projected (under the most likely evolution, indicated by the darkest area) to pass through the target rate of 2.0 percent at the eight-quarter horizon (indicated by the vertical dashed line), and then to continue rising in the following year. Thus, if the policy rate were to be held at 4 percent for a year, the Bank's expectation in February 2004 should have been (under the most likely evolution, given what was known then) that in February 2005 a similar exercise would forecast consumer price inflation to pass through 2.0 percent at the *one*-year horizon, and to exceed 2.0 percent during the *second* year. Hence, the bank has essentially forecasted that in a year's time, under the most likely evolution, the policy committee would have reason to raise the policy rate. To put it another way, the February 2004 projection itself could have been taken as evidence that the Bank should *not* have expected the policy rate to remain at 4 percent over the following eight quarters.

<sup>7</sup> This was the basis of the critique of interest rate rules by Sargent and Wallace (1975). For discussion of how such policies lead to indeterminacy in modern micro-founded "New Keynesian" models, see Woodford (2003, chap. 4).


<sup>8</sup> This was the basis of the critique of an interest-rate pegging policy by Friedman (1968); the idea goes back to Wicksell's (1898) analysis of the "cumulative process" of inflation under an interest rate policy.

<sup>9</sup> By the "policy rate," I mean the Bank's target for overnight rates in the pound sterling money markets, which it seeks to maintain through its various operations. This is currently called the "official Bank Rate"; it is essentially the equivalent of the federal funds rate operating target in the case of the Fed.


Figure 3

**A Constant-Interest-Rate Projection Compared with a Projection Based on Market Expectations of Interest Rates**

A: February 2004 CPI Projection under Constant (4%) Interest Rates



B: February 2004 CPI Projection under Market Interest Rates



Source: Bank of England *Inflation Report*, August 2004.

As these issues have come to be understood, a number of central banks that formerly relied upon constant-interest-rate projections (such as the Bank of England since August 2004 and the Swedish Riksbank until earlier this year) have switched to an alternative approach; this is the construction of projections based on *market expectations* of the future path of short-term interest rates, as inferred from the term structure of interest rates and/or futures markets. The use of projections based on market expectations allows a central bank to avoid assuming a constant interest rate when there are clear reasons to expect interest rates to change soon, while still not expressing any view of its own about the likely future path of interest rates. In the case just discussed, market expectations were for the policy rate to rise later in the year (as in fact it did), and consequently the Bank of England's February 2004 market-expectations projection shown in Figure 3B implied a less rapid increase in the inflation rate, and consumer price inflation only around 2.0 percent three years in the future.

But the market expectations approach does not ultimately answer the objections just raised to constant-interest-rate projections (Goodhart, 2005; Woodford, 2005). If one simply feeds into a macroeconomic model a specific path for the short-term nominal interest rate that is apparently expected by the markets, it is highly problematic to assume that this interest rate path exists *regardless* of how inflation and other variables evolve. Again, such an assumption will imply indeterminacy in the case of a forward-looking model, and unstable inflation dynamics in the case of a backward-looking model; for these problems arise not from the assumption of a *constant* interest rate, but from the assumption of an *exogenously specified* interest rate, unaffected by the evolution of endogenous variables such as inflation or central bank policy.

An even more fundamental problem is that market expectations can identify only a single candidate assumption about future policy to consider. If the projections based on this particular assumption do not satisfy the target criterion, what alternative policy should one consider instead? One needs a principle for finding a suitable policy assumption other than the mere fact that it is expected by the markets, and hence the issue of intertemporal consistency must be faced.

### **Choosing a Policy Path**

An alternative approach to forecast-targeting uses projections based on the central bank's own forecast of its likely future policy. In the case of both the Norges Bank (since 2005) and the Riksbank (since the beginning of 2007), each issue of the *Monetary Policy Report* now includes a "fan chart" for the evolution of the policy rate, alongside the similar charts for inflation and output. These often do not indicate an expectation that the policy rate will remain constant.

In the approach now used by these two banks, it is necessary in each decision cycle to contemplate alternative *paths* for policy, and not simply alternative choices for the current interest-rate target. The policy path that is chosen is one that is found to result in projections that are desirable from the standpoint of certain criteria. The Norges Bank is fairly explicit about these criteria; this is the significance of the list of "Criteria for an appropriate interest-rate path" in its *Monetary*

*Policy Report*, mentioned earlier. The Riksbank describes the decision as simply reflecting the preferences of the Executive Board over alternative possible projected outcomes (Sveriges Riksbank, *Monetary Policy Report* 2007/1, p. 20). While in each case a choice is effectively made over paths, both central banks emphasize that the published interest-rate paths under their “baseline scenarios” are to be understood as forecasts rather than as *commitments*. The exercise of choosing a path is repeated anew in each decision cycle, and in each case it is only the choice of the level of interest rates in the immediate future that has implications for policy actions.

The question of the intertemporal consistency of such a procedure therefore remains. In fact, the approach of selecting among alternative paths for policy through a vote of the Executive Board is unlikely to result in intertemporally consistent choices, and for reasons unrelated to any incoherence of preferences or failure to optimize correctly by that body. Even in the case of a single decisionmaker who minimizes a well-defined loss function that remains the same over time, using a correct economic model that also remains the same over time, and who never makes any calculation errors, the choice of a new optimal path for policy each period will not lead to intertemporal consistency. For in the case of a forward-looking model of the transmission mechanism, the procedure will lead to the choice of a forward path for policy that this same procedure will *not* lead the decisionmaker to continue in subsequent decision cycles, even if there have been no unexpected developments in the meantime, as explained in Woodford (1999). The reason is the same as in the celebrated critique of discretionary monetary policy by Kydland and Prescott (1977): before making a policy decision, the monetary authority prefers a forward path in which policy is expected to be tight, because of the benefits of low inflation expectations; but after making that decision, a looser policy is preferred, because the inflationary expectations can no longer be affected by the policy that is actually chosen (as opposed to the policy that was anticipated).

Intertemporal consistency can instead be maintained if the path of policy is chosen so that the projections satisfy not a point-in-time target criterion (say, a criterion relating only to projections eight quarters in the future) but rather a *sequence* of target criteria, each taking the same form, but for a sequence of horizons progressively farther in the future. Suppose that, in the spirit of the criterion illustrated in Figure 2, the policy path is chosen so that the projected inflation gap is proportional to the projected output gap (but with an opposite sign), at *every* future date<sup>10</sup>—both in the near term, when both gaps may be substantially different from zero, and in the longer term, when (in order to satisfy the criterion) both will have to approach zero. Then according to such projections, if the economy evolves as anticipated, at a later date (say, a year from now) the same procedure should result in choosing to continue the policy path that is chosen now, as the

<sup>10</sup> More precisely, the criterion should be satisfied at each date after the shortest horizon at which it is still possible for a change in monetary policy to affect the projected evolution of these variables.

continuation path should lead to projections which continue to satisfy the target criterion at all horizons.

The target criterion can be chosen to reflect a quest for balance among competing stabilization objectives. Indeed, one of the advantages of commitment to a target criterion of the kind just discussed is that it allows the central bank to make clear its commitment to a “dual mandate” (the target criterion used to evaluate both the short-term and long-term projections treats the inflation gap and the output gap *symmetrically*) while nonetheless making constantly visible the bank’s expectation that inflation should return to a fixed target value in the medium term. (Acceptable projections would always have the latter property, but as an implication of the bank’s *economic model*, rather than any lexicographic priority of inflation stabilization among its *objectives*.)

In the case of a sufficiently simple model of the monetary transmission mechanism, it is possible to choose a target criterion of this form that implements an optimal state-contingent policy; the sequential target criterion essentially corresponds to a sequence of first-order conditions for dynamic optimization under commitment. Svensson and Woodford (2005) and Giannoni and Woodford (2005) illustrate this for a variety of simple models that incorporate important features of current-generation empirical “new Keynesian” models. In practice, however, a fully optimal target criterion is likely to be too complex for use in explaining policy decisions to the public. It makes more sense to choose a simple criterion that incorporates relatively robust features of desirable policies.

One such example of a robust principle of optimal policy is the rule that higher projected inflation should be accepted when associated with a projected movement to a more negative output gap. Another is the principle that departures of the overall inflation rate from the long-run inflation target should be allowed in connection with adjustments of the structure of relative prices, but should not persist longer than is required for the adjustment of relative prices to occur. Thus it is desirable to allow CPI inflation to increase in response to an increase in the real price of energy, as this allows greater stability of the rate of growth of nonenergy prices; but the increase in CPI inflation should not be greater, or last longer, than can be justified on that ground. While the exact quantitative specification of the target criterion that is optimal depends on details of one’s model, a target criterion that incorporates these qualitative features is likely to approximate optimal policy reasonably well.


## **Are Economic Forecasts Accurate Enough to be Reliable for Inflation-Forecast Targeting?**

An obvious worry about inflation-forecast targeting is that policy decisions may depend on inaccurate forecasts. Indeed, systematic biases in the central bank’s forecasting model might even lead to systematic biases in policy. For example, even if a central bank *aims* consistently at a low rate of inflation, its policy might in fact



Figure 4

### Real-time Estimates of the U.S. Output Gap Compared to “Final” Estimates by the Fed Staff as of 1994



Source: Orphanides (2003).

generate higher inflation year after year, if the bank uses a model that forecasts inflation lower on average than what actually occurs.

One obvious reason why there might well be forecast errors with the same sign for many years arises from mistakes in real-time estimation of productivity trends, which could cause estimates of the output gap to be high or low. Figure 4, taken from Orphanides (2003), illustrates this danger by comparing the real-time estimates of the output gap available to the Fed throughout the 1970s and 1980s to its estimates in 1994 of what the output gap had actually been in each quarter. Of course, the 1970s experienced a productivity slowdown, which was not fully recognized at the time. Thus, in the 1970s, the Federal Reserve believed that the output gap was negative for the entire decade, with the economy more than 5 percent below potential GDP for years at a time, and as much as 15 percent below potential GDP in early 1975. But by 1994, when the productivity slowdown had been recognized, the Federal Reserve came to believe that the output gap during the 1970s had instead been *positive* except during the slump period of 1974–76, and never below 5 percent of potential even then. This kind of mistake could easily lead to monetary policy with an unintended but systematic inflationary bias for years in a row, and Orphanides argues that this pattern explains the “Great Inflation” of the 1970s.

While real-time estimation of productivity trends remains a difficult problem for central banks, I do not feel that this is a sufficient reason to abandon forecast-based policy, or even to insist upon a target criterion that does not involve the output gap. The optimal policy at a given point in time inevitably depends on what is expected to occur later, not only because present actions have delayed effects, but also because the more immediate effects of policy actions depend on what they are

taken to signal about future policy. Hence a policy rule that always prescribes the same response to current conditions, taking no account of what may already be foreseeable about future conditions, will be inferior to one that makes appropriate use of whatever information is available about the future. Forecast-targeting procedures are a natural way of making use of such information so as to achieve a desirable outcome in any of a wide range of possible situations. Svensson and Woodford (2005) discuss in detail the advantages of forecast-targeting procedures as a way of implementing a dynamically optimal state-contingent policy.

Similarly, standard economic models imply that the gap between actual output and the economy's potential is relevant to judging the welfare consequences of alternative monetary policies. Moreover, target criteria that relate inflation measures to an appropriately defined output gap provide characterizations of optimal policy that are especially robust, in the sense of remaining valid in the case of a large number of different types of possible disturbances.<sup>11</sup> If a measure of the output gap is known to be imprecise, the likely degree of imprecision should be taken into account when judging whether there is a basis for accepting a projected near-term inflation rate different from the long-run target. But to a first approximation, an appropriate response is *not* to change the weight placed on one's estimate of the output gap in the target criterion; instead, one need only take care to use an optimal *estimate* of the output gap when checking the target criterion, which adjusts the size of the output estimate for the level of certainty about it (Svensson and Woodford, 2003). For example, if one has little useful evidence about the output gap, the optimal estimate will seldom depart from the unconditional mean of zero, and so the inclusion of an output-gap correction in the target criterion will be of no import; but if one can be fairly sure of the sign of the gap, at least sometimes, it will be appropriate to take this information into account.

Of course, even the best possible forecasts will nonetheless often be inaccurate. The key to avoiding the possibility of an entire decade of inflation well above the target level, due to systematic bias in the underlying model, is to make a commitment to correct past target misses. That is, a year or two of inflation higher than desired should result in policy that deliberately aims at an inflation rate *lower* than the long-run inflation target for the next few years—to correct the overshoot. In this way, even if the central bank uses a model that produces a downward-biased forecast of inflation for many years in a row (due, for example, to a productivity slowdown that leads to a persistent overestimate of potential output), it will not allow excess inflation to occur for very long before policy is tightened.

One way to institutionalize this kind of error-correction would be through commitment to a target path for the *price level*, and not just a target annual inflation rate. Of course, in such a setting, temporary departures from the price-level target path should be allowed in proportion to the output gap. However, target misses—that is, departures of the output-gap-adjusted price level from the target path—

<sup>11</sup> Giannoni and Woodford (2005) derive robustly optimal target criteria for a range of alternative New Keynesian models, and all are “output-gap-adjusted inflation targets” broadly similar to the criterion illustrated in Figure 2.

would then require a policy under which the gap between the adjusted price level and the deterministic trend path would be projected to be eliminated at a certain constant rate. If the required rate of error-correction was quite slow, such a criterion could justify a *price level* persistently above the target path, but not higher by a growing rate (unless the overestimate of potential output were itself to grow steadily as well).

A commitment to error-correction has even greater advantages if one supposes that the private sector should have a better real-time estimate of productivity than the central bank. In this case, even if the central bank overestimates potential output and accordingly pursues too expansionary a policy, the private-sector recognition of this pattern will lead to anticipation of the subsequent disinflationary correction, restraining wage and price increases. Thus, as pointed out by Gorodnichenko and Shapiro (2006), under price-level targeting, mistakes in the central bank's real-time estimates of potential output have less of an effect on inflation outcomes—even in the short run, and not just over longer periods of time. They argue that Fed policy under Alan Greenspan incorporated elements of error-correction and propose that this is why uncertainty about the productivity trend in the late 1990s was so much less of a problem than in the 1970s. Aoki and Nikolov (2005) similarly find that coefficient errors in the central bank's model of the economy result in less deterioration of economic performance in the case of a forecast-targeting procedure with an output-gap-adjusted price-level target than in the case of a gap-adjusted inflation target.

## **Forecast Targeting in the United States: An Idea Whose Time Has Come**

Forecast targeting offers great promise as an approach to the conduct of monetary policy. If properly executed, it can serve to anchor expectations about the future value of a currency—the traditional aim of a monetary standard—while still allowing substantial attention to short-run stabilization concerns. But the realization of this promise requires that the right version of forecast targeting be adopted.

Specifically, forecast-targeting central banks must learn to be more explicit about the near-term target criteria that their projections are expected to satisfy, rather than speaking only about their medium-run targets for inflation. The forecast-targeting exercise is only internally consistent—and therefore able to serve as effective means of communication with sophisticated observers—if the assumptions made about the bank's own future policy are chosen appropriately. Finally, given the inevitable limitations of both the information and the economic models available to central banks, the target criteria should incorporate some degree of commitment to error-correction, rather than being purely forward-looking. Approaches that address these concerns are available in principle, and central banks around the world are making steady progress at the refinement of methods of analysis and communication that can achieve these ambitious goals in practice.

Would a forecast-targeting approach make sense for the U.S. Federal Reserve? Quantitative projections already play an important role in the internal deliberations of the Federal Open Market Committee (FOMC). Ben Bernanke, the current chairman of the Fed, devoted a speech in 2004 (when serving as a member of the Board of Governors) to what he called “forecast-based policy” and asserted that while there was no exclusive commitment to such an approach, “the Federal Reserve relies primarily on the forecast-based approach for making policy.” But full implementation of a forecast-targeting approach of the kind discussed above would require two important further steps: adoption of an explicit target criterion that would provide a structure for policy deliberations, and public discussion of the projections as a central element in the Fed’s communication policy.

Both steps are actively under consideration at the Fed, though no such innovations will be undertaken without thorough discussion. The minutes of the FOMC meeting of March 20–21, 2007, indicate that, as at several other meetings over the past year, the Committee discussed both “the possible advantages and disadvantages of specifying a numerical price objective for monetary policy” and the possibility of “an enhanced role for projections in explaining policy.”

In my view, both steps would have important advantages. While the Fed’s existing procedures have been quite successful at maintaining a low and stable rate of inflation over the past two decades, improving the degree to which medium-run inflation expectations in the United States are anchored remains an important concern for the Fed. The Fed’s concern in late 2002 and early 2003 that deflationary expectations could develop represents one recent instance in which weaknesses of the current approach have been apparent; more recent fears that undue expectations of continuing inflation could be created by relatively transitory increases in commodity-price inflation provide another. The Fed has found it necessary both to give more explicit signals as to the likely forward path of monetary policy and to talk more about its own projections regarding future inflation. Such communication would be both less ambiguous and more credible in the context of an explicit forecast-targeting strategy.

Proposals that the Fed adopt some form of inflation targeting often meet with the objection that this would require legislative authorization. (The minutes of the FOMC discussion just cited are careful to state that “participants emphasized that any such move would need to be consistent with the Committee’s statutory objectives for promoting maximum employment as well as price stability.”) But as I have argued above, a forecast-targeting approach does not require that a central bank close its eyes to the consequences of its policies for employment. The projections that are considered in policy deliberations should include projections for real variables, and indeed a sensible target criterion should involve these projections as much as the projection for inflation.

A forecast-targeting procedure similar to that of the Norges Bank could plausibly be introduced as a framework intended to ensure that policy conforms to the mandates of the Federal Reserve Act and to make this conformity more evident to Congress and to the public. Rather than a further arrogation of power to the Fed to define its own objectives, adoption of forecast targeting would represent a

voluntary decision by the Fed to make itself more accountable. In addition to increasing the effectiveness of the Fed's communication, such a step would help to reconcile the Fed's operational independence with democratic principles.

The conduct of monetary policy under a forecast-targeting framework would probably not have been greatly different than the policy that the Fed has followed in recent years. But in the absence of a clearer commitment to a systematic framework for the conduct of policy, the public has little ground for confidence that the stability achieved over the past decade has not simply been due to luck or the personalities of particular members of the Federal Open Market Committee, so that the situation could change at any time. Adoption of an explicit forecast-targeting framework would instead allow confidence in the value of the dollar to be maintained in the face of changing circumstances and facilitate continued stabilization of the real economy as well. It is time for the Federal Reserve to build on the experience of other central banks and develop a forecast-targeting framework suited to its own policy commitments and institutional setting.

■ *I would like to thank Olivier Blanchard, Rick Mishkin, Julio Rotemberg, Argia Sbordone, Lars Svensson, and the editors of JEP for helpful comments on an earlier draft.*

## References

- Aoki, Kosuke, and Kalin Nikolov.** 2005. "Rule-Based Monetary Policy under Central Bank Learning." Centre for Economic Policy Research Discussion Paper 5056.
- Bank of England.** Various years. *Inflation Report*, various issues.
- Benati, Luca.** 2006. "UK Monetary Regimes and Macroeconomic Stylized Facts." Bank of England Working Paper 290.
- Bernanke, Ben S.** 2004. "The Logic of Monetary Policy." Speech before the National Economists Club, December 2, 2004, in Washington, D.C.
- Bernanke, Ben S., Thomas Laubach, Frederic S. Mishkin, and Adam S. Posen.** 1999. *Inflation Targeting: Lessons from the International Experience*. Princeton: Princeton University Press.
- Estrella, Arturo, and Frederic S. Mishkin.** 1997. "Is There a Role for Monetary Aggregates in the Conduct of Monetary Policy?" *Journal of Monetary Economics*, 40(2): 279–304.
- Fischer, Bjorn, Michele Lenza, Huw Pill, and Lucrezia Reichlin.** 2006. "Money and Monetary Policy: The ECB Experience 1999–2006." <http://www.ecb.int/events/pdf/conferences/cbc4/ReichlinPillLenzaFisher.pdf>.
- Friedman, Benjamin M., and Kenneth N. Kuttner.** 1996. "A Price Target for U.S. Monetary Policy? Lessons from the Experience with Money Growth Targets." *Brookings Papers on Economic Activity*, 1996(1): 77–125, 144–46.
- Friedman, Milton.** 1968. "The Role of Monetary Policy." *American Economic Review* 58(1): 1–17.
- Giannoni, Marc P., and Michael Woodford.** 2005. "Optimal Inflation Targeting Rules." In *The Inflation Targeting Debate*, ed. B.S. Bernanke and M. Woodford, chap. 3. Chicago: University of Chicago Press.
- Goodhart, Charles A. E.** 2001. "Monetary Transmission Lags and the Formulation of the Policy Decision on Interest Rates." *Federal Reserve Bank of St. Louis Review*, July/August 2001, pp. 165–81.
- Goodhart, Charles A. E.** 2005. "The Interest Rate Conditioning Assumption." Financial Markets Group Discussion Paper 547.
- Gorodnichenko, Yuriy, and Matthew D. Sha-**

- piro.** 2006. "Monetary Policy When Potential Output is Uncertain: Understanding the Growth Gamble of the 1990s." National Bureau of Economic Research Working Paper 12268.
- Honkapohja, Seppo, and Kaushik Mitra.** 2005. "Performance of Inflation Targeting Based on Constant Interest Rate Projections." *Journal of Economic Dynamics and Control*, 29(11): 1867–92.
- Jansson, Per, and Anders Vredin.** 2003. "Forecast-Based Monetary Policy: The Case of Sweden." *International Finance*, 6(3): 349–80.
- King, Mervyn.** 1999. "Challenges for Monetary Policy: New and Old." In *New Challenges for Monetary Policy*, 11–58. Kansas City: Federal Reserve Bank of Kansas City.
- King, Mervyn.** 2005. "What Has Inflation Targeting Achieved?" In *The Inflation Targeting Debate*, ed. B. S. Bernanke and M. Woodford, chap. 1. Chicago: University of Chicago Press.
- Kydland, Finn E., and Edward C. Prescott.** 1977. "Rules Rather than Discretion: The Inconsistency of Optimal Plans." *Journal of Political Economy*, 85(3): 473–91.
- Leitemo, Kai.** 2003. "Targeting Inflation by Constant-Interest-Rate Forecasts." *Journal of Money, Credit and Banking*, 35(4): 609–26.
- Norges Bank.** Various years. *Inflation Report and Monetary Policy Report*, various issues. (In 2007 the Norges Bank changed the name of its *Inflation Reports*; they are now called *Monetary Policy Reports*.)
- Orphanides, Athanasios.** 2003. "The Quest for Prosperity without Inflation." *Journal of Monetary Economics*, 50(3): 633–63.
- Sargent, Thomas J., and Neil Wallace.** 1975. "Rational Expectations, the Optimal Monetary Instrument and the Optimal Money Supply Rule." *Journal of Political Economy* 83(2): 241–54.
- Svensson, Lars E. O.** 1997. "Inflation Forecast Targeting: Implementing and Monitoring Inflation Targeting." *European Economic Review*, 41(6): 1111–46.
- Svensson, Lars E. O.** 1999. "Inflation Targeting as a Monetary Policy Rule." *Journal of Monetary Economics*, 43(3): 607–54.
- Svensson, Lars E. O., and Michael Woodford.** 2003. "Indicator Variables for Optimal Policy." *Journal of Monetary Economics*, 50(3): 691–720.
- Svensson, Lars E. O., and Michael Woodford.** 2005. "Implementing Optimal Policy through Inflation-Forecast Targeting." In *The Inflation Targeting Debate*, ed. B.S. Bernanke and M. Woodford, chap. 2. Chicago: University of Chicago Press.
- Sveriges Riksbank.** Various years. *Inflation Policy Report and Monetary Policy Report*, various issues. (In 2007 Sweden's Riskbank changed the name of its *Inflation Reports*; they are now called *Monetary Policy Reports*.)
- Taylor, John B.** 1993. "Discretion versus Policy Rules in Practice." Carnegie-Rochester Series on Public Policy, vol. 39, pp. 195–214.
- Vickers, John.** 1998. "Inflation Targeting in Practice: The U.K. Experience." *Bank of England Quarterly Bulletin*, November.
- Wicksell, Knut.** 1898. *Interest and Prices*. (English translation by R.F. Kahn, London: Macmillan, 1936).
- Woodford, Michael.** 1999. "Commentary: How Should Monetary Policy Be Conducted in an Era of Price Stability?" In *New Challenges for Monetary Policy*, 277–316. Kansas City: Federal Reserve Bank of Kansas City.
- Woodford, Michael.** 2004. "Inflation Targeting and Optimal Monetary Policy." *Federal Reserve Bank of St. Louis Economic Review*, July/August, 2004, pp.15–41.
- Woodford, Michael.** 2003. *Interest and Prices: Foundations of a Theory of Monetary Policy*. Princeton: Princeton University Press.
- Woodford, Michael.** 2005. "Central-Bank Communication and Policy Effectiveness." In *The Greenspan Era: Lessons for the Future*, 399–474. Kansas City: Federal Reserve Bank of Kansas City.

Copyright of *Journal of Economic Perspectives* is the property of American Economic Association and its content may not be copied or emailed to multiple sites or posted to a listserv without the copyright holder's express written permission. However, users may print, download, or email articles for individual use.