

DISCUSSION PAPER SERIES

No. 9722

FORWARD GUIDANCE BY INFLATION- TARGETING CENTRAL BANKS

Michael Woodford

INTERNATIONAL MACROECONOMICS



Centre for Economic Policy Research

www.cepr.org

Available online at:

www.cepr.org/pubs/dps/DP9722.php

FORWARD GUIDANCE BY INFLATION-TARGETING CENTRAL BANKS

Michael Woodford, Columbia University and CEPR

Discussion Paper No. 9722
November 2013

Centre for Economic Policy Research
77 Bastwick Street, London EC1V 3PZ, UK
Tel: (44 20) 7183 8801, Fax: (44 20) 7183 8820
Email: cepr@cepr.org, Website: www.cepr.org

This Discussion Paper is issued under the auspices of the Centre's research programme in **INTERNATIONAL MACROECONOMICS**. Any opinions expressed here are those of the author(s) and not those of the Centre for Economic Policy Research. Research disseminated by CEPR may include views on policy, but the Centre itself takes no institutional policy positions.

The Centre for Economic Policy Research was established in 1983 as an educational charity, to promote independent analysis and public discussion of open economies and the relations among them. It is pluralist and non-partisan, bringing economic research to bear on the analysis of medium- and long-run policy questions.

These Discussion Papers often represent preliminary or incomplete work, circulated to encourage discussion and comment. Citation and use of such a paper should take account of its provisional character.

Copyright: Michael Woodford

CEPR Discussion Paper No. 9722

November 2013

ABSTRACT

Forward Guidance by Inflation-Targeting Central Banks*

This paper assesses the value of central-bank communication about likely future policy, with particular reference to the regular publication of projections for the future path of the policy rate, as with the Riksbank's publication of the repo rate path. It first discusses why publication of a projected interest-rate path represents a natural and desirable evolution of inflation-forecast targeting procedures, and the conditions under which the assumptions about future policy underlying such projections will be intertemporally consistent. It then discusses evidence on the extent to which central-bank statements influence private-sector interest-rate expectations. Particular attention is given to the potential use of forward guidance as an additional tool of policy when an executive lower bound for the policy rate is reached, and alternative approaches to forward guidance in this context are compared, including the recent adoption of quantitative "thresholds" for unemployment and inflation expectations by the U.S. Federal Reserve. The potential role of a nominal GDP level target within an inflation-targeting regime is also considered.

Michael Woodford
Department of Economics
Columbia University
420 W. 118th Street
New York, NY 10027
USA

Email: mw2230@columbia.edu

For further Discussion Papers by this author see:
www.cepr.org/pubs/new-dps/dplist.asp?authorid=105523

*Prepared for the conference "Two Decades of Inflation Targeting: Main Lessons and Remaining Challenges," Sveriges Riksbank, June 3, 2013. Portions are based on work previously presented in Woodford (2012a, 2012b). I would like to thank Claes Berg, James Bullard, Charles Evans, Gauti Eggertsson, Narayana Kocherlakota, Argia Sbordone, Lars Svensson, Eric Swanson and John Williams for helpful discussions, while absolving them from responsibility for the arguments presented. I also thank Kyle Jurado and Savitar Sundaresan for research assistance, and the National Science Foundation for supporting my research on this issue under grant number SES-0820438.

Submitted 28 October 2013

One of the notable features of inflation targeting as an approach to the conduct of monetary policy has been the increased degree of transparency on the part of inflation-targeting central banks, not only as to their decisions but also with regard to the goals that policy seeks to achieve and the reasoning behind individual decisions. The degree to which this makes it appropriate, or even necessary, for inflation-targeting central banks to speak in advance about future policy decisions has been a topic of debate,¹ but over time, inflation-targeting central banks such as the Reserve Bank of New Zealand, the Norges Bank, and Sveriges Riksbank have also led the way in increasing the degree of explicit communication about the likely forward path of short-term interest rates on a regular basis.

More recently, many central banks have found immediate cuts in their policy rate an insufficient response to the effects of the global financial crisis, and this has led to increased interest in explicit “forward guidance” about future interest-rate policy as an additional policy tool. This raises questions about the usefulness of this additional dimension of policy in the context of the kind of forecast-targeting procedures already used by many of the leading inflation-targeting central banks. Just in the past year, the Bank of Japan, the Bank of England, and the European Central Bank have all notably increased their use of explicit forward guidance.

Here I first review the general role of discussions of the forward path of the policy rate, and of explicit intermediate targets for policy, as elements of an inflation forecast-targeting approach to monetary policy. I then turn to the special role of forward guidance in the case that a central bank finds itself constrained by a practical lower bound on where it can (or is willing to) set its policy rate. I review recent experience with various approaches to forward guidance in that situation, including the Federal Reserve’s December 2012 introduction of quantitative “thresholds,” and discuss the appropriate role of such intermediate targets in a forecast-targeting framework.

1 The Forward Path of Policy in a Forecast-Targeting Framework

Central banks with explicit inflation targets have emphasized from the start that it is not reasonable to expect a central bank to be able to keep the measured rate of inflation exactly equal to the target rate at all times. They have in particular

¹See, *e.g.*, Goodhart (2005) for a skeptical discussion.

stressed that it is difficult for a shift in monetary policy, even a relatively drastic one, to greatly affect the rate of inflation over the near term (that is, for at least several months following the meeting at which a policy decision is taken); and they have accordingly stressed that the goal of policy should instead be to ensure that inflation can be expected to return to the target rate fairly soon, even when it currently differs from that rate. Hence both policy decisions and communication with the public about those decisions have come to focus on projections for the future path of the economy (and in particular, projections for one or more measures of inflation), and the extent to which these are consistent with the bank's official target.

But while inflation-targeting central banks have in this sense necessarily adopted a *forward-looking* approach to monetary policy, it has not obviously followed that the policy framework requires explicit consideration in advance of an intended forward path for the policy rate, or other policy instruments, still less any communication with the public about the policy committee's thoughts on that matter. Some early discussions of inflation-forecast targeting made it appear that one should be able to determine the appropriate current setting for the policy rate simply by reference to a projection for future inflation conditional on that rate, without having to make any specific assumption about future policy decisions. For example, in the early exposition of inflation-forecast targeting by Svensson (1997), a model is assumed in which the policy rate affects economic activity only with a lag of a year, and activity affects inflation, but only with an additional lag of a year. (Both effects are purely backward-looking; expectations play no role in the determination of either output or inflation.) Hence the model can be reduced to a single structural equation of the form

$$\pi_t = u_t - \gamma i_{t-2}, \tag{1.1}$$

where π_t is the inflation rate, i_t is the policy rate, u_t is a composite of all of the other factors influencing inflation (assumed to evolve independently of the path of the policy rate), periods correspond to years, and $\gamma > 0$ is a constant coefficient.

It is then easily shown that the policy that minimizes the expected squared deviation of the inflation rate from the inflation target π^* is one that sets i_t each period so as to ensure that the inflation forecast satisfies

$$E_t \pi_{t+2} = \pi^*; \tag{1.2}$$

if the forecast is produced using the model (1.1), this will require that

$$i_t = \gamma^{-1} [E_t u_{t+2} - \pi^*]. \quad (1.3)$$

Note that the optimization required in order to determine the setting (1.3) for i_t can be carried out without considering how i_τ will be set for any $\tau > t$. Each meeting of the policy committee can be treated as involving an independent decision, and the inflation target alone suffices to allow a determinate decision on each occasion and to allow the decision to be justified to the public by reference to the target criterion (1.2).

However, these conclusions depend on overly simplistic features of the proposed model. The model (1.1) assumes not merely that interest-rate decisions have delayed effects, but that there are *no* effects until the future horizon (two years later) at which the main effect will suddenly occur. Even if one grants that the largest effects occur with a delay, it is nonetheless more reasonable to suppose that a policy change *begins* to have an effect at some point prior to the date at which the *largest* effect occurs. Yet even this small modification of one's assumptions would have important consequences for the forecast-targeting exercise. In general, one cannot treat the problem of inflation stabilization as a separate problem for each horizon, with an independent instrument choice that can be used to minimize inflation variability at that horizon; instead, one must solve a dynamic optimization problem, in which the optimal interest-rate decision at a given time is not independent of the intended conduct of policy at later dates.

Hence optimal policy, and indeed an internally consistent forecast-targeting exercise, will almost inevitably require a determination at each decision point of what the entire anticipated *forward path* of the policy rate should be, even though this need not mean that a once-and-for-all decision about policy is made at some initial date, and then simply executed thereafter. In practice, the number of future contingencies that may arise will be much too large to make it possible to solve explicitly for a state-contingent policy years in advance, and be content to simply implement it thereafter by deciding which of the contingencies that had been previously foreseen as possible has actually occurred. At the same time, a decision about current policy will require a *forecast* of how policy is expected to be made subsequently, even if it is inevitable that actual future policy will depend on complications that cannot yet be anticipated.

1.1 Medium-Run Forecast Targeting without Choosing a Forward Path

In practice, inflation-targeting central banks have not supposed that their procedures should seek to ensure that forecasted inflation must equal the target rate at the shortest horizon at which inflation can still be influenced, if indeed such a horizon can even be defined. It has generally been recognized that returning inflation to the target rate as quickly as possible would not necessarily be optimal; the focus has instead often been on ensuring that inflation should return to target *over some specified horizon*, where the horizon is chosen to be far enough in the future to ensure not only that inflation can actually be controlled with some accuracy over that horizon, but that always planning to return inflation to the target rate over that horizon should not require excessively sharp adjustments of real variables, while it is still near enough to maintain a reasonably tight bound on the implied variability of the inflation rate around its target value. (Typically, horizons two to three years in the future have been considered suitable.)

However, early discussions of forecast targeting in this vein still often sought to make it possible for a central bank to make a separate interest-rate decision at each decision point without prejudging future policy decisions. For example, the Bank of England's forecast-targeting procedure (Vickers, 1998; Goodhart, 2001) was described as being based on a *constant-interest-rate* forecast, in which forward paths for inflation and other variables were projected under the assumption of a constant value for the policy rate over the forecast horizon. Letting $F_{t,t+8}(i)$ be the forecast of π_{t+8} , the inflation rate eight quarters in the future, under the assumption that the policy rate is kept at an arbitrary level i until then,² then the procedure was described as choosing at each decision point an operating target i_t for the policy rate so as to ensure that

$$F_{t,t+8}(i_t) = \pi^*. \tag{1.4}$$

The policy decision was then justified to the public by presenting, at the beginning of each issue of the Bank's *Inflation Report*, a figure showing the projected path

²Note that this formulation of the exercise is only possible under the assumption that a purely backward-looking model is used to forecast inflation, as was the case at the Bank of England at the time. A similar approach to inflation-forecast targeting was used for some years by Sveriges Riksbank as well (Jansson and Vredin, 2003).

of inflation under the constant-interest-rate assumption, with the interest rate at the level chosen in the most recent meeting of the Monetary Policy Committee. (The projection was presented in the form of a “fan chart,” showing a probability distribution for future inflation outcomes at each horizon, rather than a point forecast.) This figure always included a horizontal line at the target inflation rate, and a dashed vertical line at the horizon eight quarters in the future, so that the eye could easily determine the extent to which the projection was consistent with the target criterion (1.4), by observing whether the modal predicted path of inflation passed through the intersection of the two lines.

This approach had the advantage of allowing an interest-rate decision to be made at each decision point without requiring any explicit consideration of current intentions with regard to future policy. It also had the advantage of allowing definite decisions to be made about the appropriate current level of the policy rate, by making even a quarter-percent change appear quite consequential, insofar it is treated as a *permanent* change of that size in the projection exercise, rather than only a change in the target to be pursued until the next meeting. Nonetheless, there were serious conceptual problems with the approach (Goodhart, 2001; Leitemo, 2003; Honkapohja and Mitra, 2005; Woodford, 2005).

While the assumption of a future policy rate at the same level as the current operating target might seem a natural one, at least in the absence of clear reasons to expect the future to be different from the present, it is actually not at all sensible to suppose that short-term nominal interest rates should remain fixed at some level, regardless of how inflation or other variables may evolve. Indeed, in forward-looking (rational-expectations) models of the kind that are now often used by central banks, the assumption of a constant nominal interest rate typically implies an indeterminate price level, so that it becomes impossible to solve uniquely for an inflation forecast under any such interest-rate assumption.³ In models with backward-looking expectations, the model can be solved, but such policies often imply explosive inflation dynamics. Such difficulties appears 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 1 provides an example. In such a case, it is not obvious why anyone should believe that policy is consistent with the infla-

³See Woodford (2003, chap. 4) for examples of this problem.




Figure 1: The Bank of England’s February 2004 CPI projection under the assumption of a constant 4.0 percent interest rate. Source: Bank of England, *Inflation Report*, August 2004.

tion 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 the internal inconsistency involved in the sequential application of such a procedure. The usefulness of a forecast-targeting procedure as a way of creating confidence that the inflation target should be expected to be satisfied in the medium run — so that it should serve to anchor medium-run expectations — depends on the public’s having reason to suppose that the central bank’s projections do indeed represent reasonable forecasts of the economy’s future evolution. But among the possible grounds for doubt is a tension inherent in the logic of a forecast-targeting procedure itself. Production of projections of the economy’s evolution years into the future requires that the central bank make assumptions about the path of policy variables, such as nominal interest rates, 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).

This raises a question as to whether this decision procedure should be expected to actually 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 farther in the future, private decisionmakers have no reason ever to expect any disinflation at all.

Thus a requirement for credibility of the central bank's projections is that the forecast-targeting procedure be *intertemporally consistent*. This means that 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. But the approach to forecast-targeting represented by requirement (1.4) fails to satisfy this criterion.

The problem is that 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 1, 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.⁴ 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

⁴In the February *Report*, only the projection up to the 8-quarter horizon was shown. The figure that has been extended to a horizon 12 quarters in the future is taken from the August 2004 *Inflation Report*, in which the Bank explained its reasons for abandoning the method of constant-interest-rate projections.

vertical dashed line), and then to continue rising in the following year.

It follows that if the policy rate were to be held at 4 percent for a year, the Bank's expectation in February 2004 should have been that (under the most likely evolution, given what was known then) 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 of the projection. 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. Thus 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.

As these issues have come to be understood, a number of central banks that formerly relied upon constant-interest-rate projections (including the Bank of England, since August 2004) 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. In the case that the projections constructed under this assumption satisfy the target criterion, the correct current interest-rate decision is taken to be the one consistent with market expectations. 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 rates to change soon, while still not expressing any view of its own about the likely future path of interest rates.

But the market expectations approach does not really solve the problem of internal consistency just raised.⁵ One problem is that market expectations can at most supply a single candidate forward path for policy; it is not clear what decision one is supposed to make if that path does *not* lead to projections consistent with the target criterion. Thus the procedure is incompletely specified; and if it is only the projections based on market expectations that are published, even though the central bank has chosen to contradict those expectations, the published projections cannot be expected to shape private decisionmakers' forecasts of the economy's evolution.

Moreover, even if the forward path implied by market expectations does lead to projections that fulfill the target criterion, the exercise is not intertemporally consistent if this path does not in fact correspond to the central bank's *own* forecast

⁵For further discussion of problems with this approach, see Woodford (2005) and Rosenberg (2007).

of the likely future path of interest rates. Why should it count as a justification of a current interest-rate decision that this would be the first step along a path that *would* imply satisfaction of the target criterion, but that the central bank does not actually expect to be followed? And why should anyone who correctly understands the central bank’s procedures base their own forecasts on published projections constructed on such an assumption?

1.2 Sequential Choice of a Forward Path

In fact, there is no possibility of an intertemporally consistent forecast-targeting procedure that does not require the central bank to *model its own likely future conduct* as part of the projection exercise. Approaches like both of those just described — which introduce an artificial assumption about the path of interest rates in order to allow the central bank to avoid expressing any view about policy decisions that need not yet be made — necessarily result in inconsistencies. Instead, a consistent projection exercise must make assumptions that allow the evolution of the central bank’s policy instrument to be projected, along with the projections for inflation and other endogenous variables.

In such a case, it would be possible, but somewhat awkward, for the central bank to remain silent about the implications of its assumptions for the forward path of interest rates; and so it is natural to include an interest-rate projection among the projections that are discussed in the *Monetary Policy Report*. This has been done for the past decade now by the Reserve Bank of New Zealand, and is now done by the Norges Bank (since 2005) and the Riksbank (since 2007) as well. In the case of the latter two central banks, “fan charts” (similar to the one shown in Figure 1) are presented for the policy rate; this (among other things) makes it clear that the path is simply a forecast, rather than a definite intention that has already been formulated, let alone a promise.

But how should future policy be specified in such an exercise? It is sometimes suggested that the monetary policy committee should conceive of its task as the choice of a *path* for interest rates, rather than a single number for the current operating target, in each decision cycle. Discussions of the feasibility of such an approach have often stressed the potential difficulty of committee voting on a decision with so many

dimensions.⁶ And when announcing its intention to begin publishing its own view of the path of the policy rate, the Riksbank (Rosenberg, 2007) indicated that it would publish “forecasts ... based on an interest-rate path chosen by the Executive Board.”⁷

However, the idea that one should simply ask the policy committee to decide which forward path for interest rates they prefer, presumably after asking their staff to produce projections for other variables conditional on each path that is considered, is problematic on several grounds that have nothing to do with the complexity of the decision or the need for a committee to agree among themselves. First of all, the specification of future policy by a simple path for a short-term nominal interest rate, independently of how endogenous variables may develop, is never a sensible choice, and is unlikely to lead to well-behaved results in a sensible model. (The problems mentioned above in connection with the assumption of a constant interest-rate path apply equally to *any* specification of an exogenous path; they do not result from the assumption that the interest rate does not vary with time, but from the assumption that it is independent of outcomes for inflation and other variables.) Moreover, the assumption of a specific path for interest rates, unaffected by future shocks, would seem to require one to publish a specific path for this variable, alongside the fan charts for variables such as inflation; but this would encourage the dangerous misunderstanding that the bank has already committed itself to follow a definite path long in advance.

Even supposing that these technical issues have been finessed,⁸ there remains the

⁶See, for example, Goodhart (2005) for a skeptical view; Svensson (2007) responds by proposing a voting mechanism intended to overcome potential intransitivities in majority preferences over alternative paths.

⁷It is likely, of course, that this was only a loose way of speaking in a statement intended for a non-technical audience, and that the intention was to indicate that the Executive Board would have to endorse the assumptions about future policy involved in generating projections of an endogenous interest-rate path. The change in procedure does seem to have meant that the Executive Board is now required to approve the assumptions made in the projections in a way that was not previously true; this has made it necessary to allow for possible revisions in the projections following the meeting at which the policy decision is made.

⁸For example, one might specify future policy by a policy rule, such a Taylor rule, with some number of free parameters that are optimized, in each decision cycle, so as to result in projections that are acceptable to the monetary policy committee. If only rules that are considered that imply a determinate equilibrium, the first problem is avoided. And since the rule that is chosen would make the interest rate endogenous, an assumption about the distribution of shocks in each future period

more fundamental problem of the intertemporal consistency of the procedure. Here it is important to realize that the mere use of a *consistent criterion* over time to rank alternative projected paths for the endogenous variables — not just a criterion that provides a transitive ordering of outcomes within each decision cycle, but one that ranks different possible paths the same way, regardless of the date at which the decision is being made — is not enough to ensure intertemporal consistency, in the sense defined above. Thus the problems of choosing a forward path for policy are not resolved simply by asking the members of the policy committee to agree on a loss function that they will then use (for an entire sequence of meetings) to rank alternative possible outcomes, as proposed by Svensson (2007).

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 general 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 one will not be lead by the same procedure to continue in subsequent decision cycles, even if there have been no unexpected developments in the meantime.

The reason is the same as in the celebrated argument of Kydland and Prescott (1977) for the “time inconsistency of optimal plans”: the forward path chosen at one time will take account of the benefits at earlier dates of certain expectations about policy at the later dates, but as the later dates approach (and the earlier expectations are now historical facts), there will no longer be a reason to take into account any effect of the policy chosen for those dates on earlier expectations. This problem does not arise solely in connection with the bias in the average rate of inflation chosen by a sequential optimizer, as in the example of Kydland and Prescott (1977). One may solve the problem of “inflationary bias” by assigning the central bank a loss function in which the target level of the output gap is not higher than the level consistent on average with its inflation target, but the optimal dynamic responses to shocks are still not generally the ones that would be chosen under sequential (or discretionary) optimization.

would result in a probability distribution for future interest rates, just as for the future inflation rate.

1.3 Using a Target Criterion to Determine the Forward Path

An alternative approach, that avoids this problem, is to determine the forward path of policy as that path which results in projections that satisfy a *sequence* of quantitative target criteria, one for each of a sequence of future horizons. It is true that a *single* criterion — say, involving the projections for 8 quarters in the future only — can determine only a single dimension of policy, and thus can only determine an entire path if one is constrained to consider only a one-parameter family of possible paths (such as constant-interest-rate paths). But a sequence of similar criteria can independently determine the stance of policy at each of a sequence of dates, and thus can determine the entire forward path of policy. Moreover, if the sequence of target criteria for different horizons are of the *same form* — *i.e.*, if the target criterion is independent of the horizon — then the forecast-targeting procedure will be intertemporally consistent.

As a practical example, consider the targeting procedure used by the Norges Bank in 2005-06. Each issue of the Bank’s *Inflation Report* included a box labeled “Criteria for an appropriate future interest rate path.”⁹ According to the first of the criteria listed, “inflation should be stabilized near the target [*i.e.*, 2.5 percent per year] within a reasonable time horizon, normally 1-3 years,” and moving *toward* that target rate even sooner. This criterion alone would sound similar to the Bank of England target criterion mentioned above, except with greater vagueness about the horizon. But there is then a second criterion: that “the inflation gap [the amount by which actual inflation exceeds the medium-run target rate] and the output gap should be in reasonable proportion to each other until they close,” and in particular that the two gaps “should normally not be positive or negative at the same time.”

The second criterion indicates not only what the projections should look like in some medium run, but also what the *transition path* should look like: there should be an inverse relation between the inflation gap and the output gap, with the two gaps shrinking to zero together. In order to allow visual inspection of the extent to which the projections satisfy this criterion, the Norges Bank presents a figure in

⁹The criteria used starting in 2005, when the Norges Bank first began to announce a forward path for the policy rate as part of its explanation of its recent policy decisions, are discussed in more detail in Qvigstad (2006). Beginning with the 2007/1 issue of the Bank’s *Monetary Policy Report*, the description of the criterion used to select the forward path of policy has been less explicit; see Qvigstad (2008) for a more recent discussion of the criteria.

which the projections for its preferred measures of inflation¹⁰ and of the output gap are superimposed. A criterion of this kind can determine the entire forward path for policy. And with such a criterion, it is not necessary to specify independently the rate at which the inflation rate should be projected to approach the target rate; the appropriate rate is exactly the rate that allows the output gap to remain in the desired proportion to the inflation gap. (Under such a criterion, the inflation gap *will* be projected to close eventually, as long as it is not possible to have a non-zero permanent output gap.)

The criterion just cited applies to each of a sequence of future horizons. It can be represented formally as the requirement that

$$(\pi_{t+h,t} - \pi^*) + \phi x_{t+h,t} = 0 \tag{1.5}$$

for each horizon $h \geq \underline{h}$, for some coefficient $\phi > 0$. Here $y_{t+h,t}$ denotes the projected value at date t of some variable y , at a horizon h periods in the future; $\underline{h} \geq 0$ indicates the shortest horizon at which it is still possible for policy to affect the projections, and I shall assume that a sequence of criteria (1.5) for $h \geq \underline{h}$ suffices to uniquely determine the acceptable projections (including an implied forward path for policy).¹¹

Suppose also that the central bank's forecast of its own forecasts in future decision cycles satisfy the principle that one should expect one's future forecasts to be the same as one's current forecasts (except, of course, as a result of developments that cannot currently be foreseen), so that

$$[y_{t+h_2,t+h_1}]_t = y_{t+h_2,t}$$

for any horizons $h_2 \geq h_1 \geq 0$. Then if at date t a forward path for policy is chosen that leads to projections satisfying (1.5) for each $h \geq \underline{h}$, it should also be projected at that time that at any later date $t + h_1$, the continuation of that same path should lead to projections satisfying a corresponding sequential criterion, since at date t the bank should project that

$$[(\pi_{t+h_2,t+h_1} - \pi^*) + \phi x_{t+h_2,t+h_1}]_t = 0$$

¹⁰The inflation measure emphasized by the Norges Bank in its targeting procedure, CPI-ATE, is a consumer price index that is adjusted for tax changes and energy prices.

¹¹See Svensson and Woodford (2005) for algebraic analysis of a specific example. In the case considered there, prices and spending decisions are each predetermined a period in advance, so that $\underline{h} = 1$.

for all horizons $h_2 \geq h_1 + \underline{h}$. This makes the procedure of choosing a forward path for policy on such a basis intertemporally consistent.

I believe that this kind of targeting procedure provides the most appealing solution to the problem of intertemporal consistency. The way in which the target criterion is used to determine an appropriate forward path for policy is essentially the same as under the procedure used by the Bank of England prior to 2004, as discussed above, except without either the arbitrary emphasis on a single horizon or the arbitrary restriction to forward paths for policy involving a constant interest rate. Since forecast-targeting central banks already publish charts showing their projections for each of a sequence of future horizons, rather than only presenting a set of numerical forecasts for a specific horizon, discussion of a target criterion that should apply at each horizon is fairly straightforward within the existing frameworks for deliberation and communication about policy, as the example of the Norges Bank shows. Moreover, both the Norges Bank and the Riksbank now discuss quite explicitly the fact that their targeting procedures involve the choice of a forward path for policy, and publish “fan charts” for the paths of short-term nominal interest rates implicit in their projections. Hence this aspect of the recommended approach is entirely possible within the context of existing procedures as well.

The main practical obstacle to such an approach, I believe, is that it would require a central bank to adopt a highly structured approach to policy deliberations, and to describe that approach rather explicitly to the public. It would require the bank to be more open about its own view of the likely future evolution of policy than even some forecast-targeting central banks have been willing to be thus far. And it would require the bank to discuss explicitly the nature of the trade-offs that determine an acceptable transition path following a disturbance, and not merely the nature of the “medium-run” targets that one hopes to reach some years in the future.

The latter goal will almost surely require that a bank be explicit about the ways in which projections for variables other than a single measure of inflation are relevant to judgments about the appropriate stance of policy. Even though all inflation-targeting central banks appear to care about projections for real variables as well as inflation, many have been quite cautious about discussing the way in which this may factor into their policy decisions. But this would have to be different if forecast targeting were to be adopted by an institution with a “dual mandate” like the U.S. Federal Reserve (at least, in the absence of a substantial modification of the Federal Reserve Act by

Congress). And even in the case of other central banks, I believe that it would greatly enhance the transparency of policymaking — and ultimately, the credibility of their commitments to inflation control, by making clearer the extent to which temporary failures to return inflation immediately to its medium-run target level are nonetheless consistent with a systematic approach to policy that does indeed guarantee stability of inflation over the medium run.

1.4 Which Form of Target Criterion?

These general considerations do not mean that the specific form of target criterion (1.5) used by the Norges Bank in the period just cited is necessarily the one that should be adopted. In the context of a simple New Keynesian DSGE model, one can show (Woodford, 2003, chap. 7) that an optimal policy commitment involves maintaining proportionality, not between deviations of the inflation rate from its long-run target and the output gap, but between deviations of the inflation rate from target and the *change* in the output gap. That is, rather than requiring that $(\pi_t - \pi^*) + \phi x_t$ be projected to equal zero at all future horizons, one should commit to a forward path of policy under which $(\pi_t - \pi^*) + \phi(x_t - x_{t-1})$ is projected to equal zero at all horizons.¹² Like the Norges Bank criterion, this one implies that both inflation and the output gap should be stabilized, in the absence of “cost-push shocks” that make the two stabilization goals mutually incompatible; and that in the event of such a disturbance, both the inflation gap and the output gap should be allowed to vary, each in order to reduce the amount of adjustment that is required by the other.

The dynamic criterion differs from the Norges Bank criterion, however, in that it implies that if inflation is allowed to increase, and the output gap to decrease, in response to a positive cost-push shock, a *below-target* inflation rate should subsequently be aimed at, as the output gap returns to its normal level (since the output gap is then *increasing*), rather than continuing to aim at an inflation rate *above target* (because the output gap remains negative, albeit to a decreasing extent). If the dynamic response is credible, an expectation of subsequent disinflation should reduce incentives for wage and price increases during the period of the cost-push shock, at

¹²Svensson and Woodford (2005) extend this analysis to an arguably more realistic model in which monetary policy changes can affect inflation and output only with a one-period lag, and show that a target criterion of the same form continues to characterize optimal policy, except that the criterion must be projected to hold only at horizons one period or farther in the future.

any given level of economic activity, and so should shift the short-run Phillips curve tradeoff in a way that tends to offset some of the effects of a cost-push shock. This allows a superior degree of achievement of the stabilization objectives than would be possible under the Norges Bank criterion.

An alternative way of seeing the difference between the two target criteria is to note that the dynamic criterion can alternatively be expressed in a *level* form, as a requirement that the condition

$$p_t + \phi x_t = p_t^*, \quad (1.6)$$

be projected to be satisfied at all future horizons, where p_t is the log of the general price, and p_t^* is a deterministic target path for the log price level, growing at a constant rate π^* each period. Satisfaction of (1.6) each period would imply that

$$\pi_t + \phi(x_t - x_{t-1}) = \pi^* \quad (1.7)$$

each period, and vice versa, assuming that the initial level p_{-1}^* for the target path is chosen so that (1.6) is satisfied by the (historically given) data for the period just before the first period in which either of the target criteria will be enforced.

But (1.6) and (1.7) are only equivalent under the assumption that either target criterion can be precisely satisfied by the realized values of inflation and the output gap each period. Under the more realistic assumption that target misses of some size will constantly occur, even if the target criterion is *projected* at each decision point to be satisfied in all future periods. That is, the requirement that a central bank's projections satisfy

$$[(\pi_{t+h} - \pi^*) + \phi(x_{t+h} - x_{t+h-1})]_{,t} = 0 \quad (1.8)$$

for all horizons $h \geq 0$ at each decision point t is not equivalent to requiring them to satisfy

$$[(p_{t+h} - p_{t+h}^*) + \phi x_{t+h}]_{,t} = 0 \quad (1.9)$$

each period. In the former case, the target p_t^* for the “output-gap adjusted price level” $p_t + \phi x_t$ used in period t is effectively adjusted, relative to the target for the same variable used in the period $t - 1$ projection exercise, by an amount equal to the target miss $p_{t-1} + \phi x_{t-1} - p_{t-1}^*$ in the previous period; in the latter case, instead, the target path $\{p_t^*\}$ remains predetermined. Thus the “level” version of the target criterion incorporates a commitment to subsequent correction of past target misses, while the first-differenced (or “growth-rate”) version does not.

Such a commitment to error-correction increases the robustness of the forecast-targeting procedure to errors of judgment on the part of the central bank.¹³ There is less reason to worry that a sustained departure of the actual inflation rate from the target rate can occur, simply as a result of a persistent bias in the central bank’s inflation forecast, that allows it to project at each decision point that (1.7) will be satisfied, though in fact the output-gap-adjusted inflation rate (*i.e.*, the left-hand side of (1.7)) exceeds π^* each period. Under the level version of the target criterion, a positive overshoot in one period requires the central bank to aim for an output-gap-adjusted inflation rate in subsequent periods that is *less* than π^* , and subsequent overshooting in the same direction (resulting from a systematic bias in the central bank’s projections) will further increase the size of the correction that is called for. Eventually, the central bank will be required to aim at a value of the gap-adjusted inflation rate that is sufficiently far below π^* that the actual outcome will not exceed π^* on average, even given the bias in the central bank’s projections.¹⁴

Hence continuing excess inflation will not result, even if the bias in the central bank’s projections is never recognized and corrected by adjustment of the forecasting model. And even assuming eventual learning on the part of the central bank, the losses that result while the learning takes place are reduced in the case of a forecast-targeting exercise using criterion (1.6) rather than (1.7), as shown in a quantitative example by Aoki and Nikolov (2005).

A level version of the target criterion is also more robust to the occurrence of target misses owing to factors outside of the central bank’s control, as opposed to errors in the central bank’s forecasts. These include the fact that, inevitably, the central bank must choose its instrument setting without full information about the values of the current structural disturbances, so that even if the criterion is (correctly) projected to hold, conditional on the information available to the monetary policy committee at the time of its decision, the actual values of the structural disturbances not exactly known to the committee will almost certainly result in its not holding exactly.¹⁵

¹³See Woodford (2011, 2012a) for further discussion of this issue.

¹⁴Svensson (2012) also discusses the advantages of a level target (a price level target, in his case), but suggests that “a less dramatic change” would be to target a five- to ten-year moving average of inflation, as proposed by Nessen and Vestin (2005). Proposals of this kind have similar virtues as a level target, though I believe that a level target would be simpler both to implement and to explain.

¹⁵For example, Gorodnichenko and Shapiro (2006) discuss the advantages of a level target in minimizing the effects of mis-estimation of the output gap.

Woodford (2011) discusses how to characterize an optimal policy commitment under such an informational constraint, and shows that it involves a commitment to error-correction of the same sign as automatically occurs under a level criterion such as (1.6).¹⁶ The same result applies when the failure to achieve the target criterion results from a constraint on the degree to which the policy instrument can currently be moved, rather than a lack of more precise information about how it should be set. Hence there are substantial advantages to the level version of the target criterion when the central bank is constrained by an effective lower bound on the level of its policy rate, as discussed in section 2.

The numerical value of the coefficient ϕ in the target criterion (1.6) that is best depends on the relative importance assigned to inflation stabilization and output-gap stabilization respectively.¹⁷ In the case that $\phi = 1$, the proposed target criterion has an especially simple interpretation, as it can alternatively be written in the form

$$Y_t = Y_t^*, \tag{1.10}$$

where $Y_t \equiv p_t + y_t$ is the log of nominal GDP (if y_t is the log of real GDP), and the target Y_t^* is given by

$$Y_t^* = y_t^n + \pi^* \cdot t, \tag{1.11}$$

where y_t^n is the log of the natural rate of output (so that the output gap is defined as $x_t \equiv y_t - y_t^n$). In this case, the target criterion can be expressed as a target path for the level of nominal GDP, a concept that is easier to explain than a target path for the output-gap-adjusted price level. Setting the coefficient ϕ equal to 1 might be viewed as representing a “balanced approach” to the dual goals of inflation and output-gap stabilization, and avoids the need to justify using a particular numerical value in the criterion (1.6). Hence this particular form of intermediate target criterion is likely to be an especially practical way of achieving the general objectives discussed above.

¹⁶The optimal commitment actually involves a slightly *stronger* degree of error-correction than the level criterion prescribes; when imperfect information results in a gap-adjusted inflation rate higher than π^* , the subsequent target should be reduced by an amount slightly *greater* than the size of the target overshoot, though the multiplicative factor approaches 1 as the rate of time discounting in the central bank’s stabilization objective approaches zero. Even allowing for discounting, error-correction of the kind prescribed by the level version of the full-information optimal target criterion is clearly desirable relative to the criterion with no such correction at all.

¹⁷Woodford (2003, 2011) shows how the optimal coefficient depends both on the coefficients of the policymaker’s loss function and the slope of the Phillips-curve tradeoff.

2 Forward Guidance at the Interest-Rate Lower Bound

Thus far I have discussed reasons for a central bank to be explicit about its intended future conduct of policy — both in its internal deliberations and in its explanations of its policy decisions to the public — as a routine element of the conduct of monetary policy. But there are special reasons for explicit discussion of future policy in the case that a central bank reaches the effective lower bound for its policy rate, as has occurred for a number of central banks since the fall of 2008.¹⁸ It is no accident that these circumstances have resulted in increased interest in explicit forward guidance as a policy tool.

There are two main advantages for a central bank from talking explicitly about its future policy, rather than simply allowing the public to form its own expectations about policy on the basis of observed behavior. First of all, in the absence of explanations by the central bank itself, misunderstandings of its policy intentions may easily develop, and this should not be left to chance, since uncertainty about how policy will be interpreted implies uncertainty about the effects of the policy. Explicit explanations of policy are most likely to be needed in unusual circumstances, or when a central bank intends to act in ways that could not easily be predicted from its previous behavior.

Hence it is not surprising that explicit forward guidance by central banks has increased precisely in a period when unprecedented policy actions are being taken, so that past rules of thumb are no longer adequate predictors of behavior. At the same time, a situation in which the current policy rate is constrained by floor on the level of short-term rates that the central bank is willing to contemplate¹⁹ is also

¹⁸In some cases, like those of the U.S. Federal Reserve, the Bank of England, and the Bank of Japan, banks have kept their policy rates at levels that they have regarded as lower bounds continuously since late 2008 or early 2009, without achieving a degree of expansion of aggregate sufficient for full utilization of productive capacity, so that the question whether forward guidance can provide further stimulus continues to be relevant. In other cases, like those of the Bank of Canada and Sveriges Riksbank, effective lower bounds were reached in the first half of 2009, but the constraint remained relevant only during 2009-10.

¹⁹The “effective lower bound” to which I refer here is not necessarily a technical constraint on the level of overnight interest rates that could be achieved. None of the central banks that I have described as constrained by their self-imposed lower bounds have actually reduced their targets for

one in which expectational errors should be particularly costly. For one reason, the social cost of an expectational error that makes aggregate demand lower by a given number of percentage points (because of a mistaken expectation that future policy will be tighter than a correct forecast would have indicated) is greater, the greater the extent to which demand already falls short of the efficient level of activity. If a binding interest-rate lower bound results in a larger negative output gap than would be allowed to exist otherwise (since further interest-rate cuts would otherwise occur and reduce the gap), this is reason to be particularly concerned to minimize potential expectational errors at such a time.

In addition, in a situation where the policy rate is expected to remain fixed for a substantial period (because the interest-rate lower bound is expected to continue to bind), but the question is whether people have correct expectations about what will happen *after* that period, New Keynesian models typically imply that changes in expectations about what will happen after the several quarters of constant policy rate will have *larger* consequences for near-term aggregate demand and economic activity than if policy were expected to be conducted over that period in a “standard” way — in accordance with the Taylor rule, or with the reaction function of an inflation-targeting central bank under normal conditions — so that the policy rate would vary with economic activity and with inflation. If increased pessimism about future output or inflation does not lead to anticipated declines in the policy rate, owing to the expectation that the policy rate will already be at its lower bound, the resulting contraction in current demand — and hence the reduction in current output, employment and inflation — will be greater. Furthermore, to the extent that this mechanism is expected to result in lower output and inflation at future dates in the period when the lower bound still binds, such an expectation should produce even lower output and inflation, through a self-amplifying process.

Hence output and inflation in a period when the lower bound is a binding constraint are lower than they would be if their policy rates all the way to zero, the rate of return on currency. (And some would question whether even zero is a genuine lower bound for overnight interest rates, given the existence of at least modest holding costs for currency.) Nonetheless, the floors were *treated* as lower bounds on the targets for the policy rate that would be considered, even if the constraints were prudential rather than technical in nature. The most commonly offered reason for not considering a further immediate cut in the policy rate has been concern for the consequences for private financial intermediaries of a complete elimination of any spread between the return on currency and money-market interest rates.

straint should be particularly sensitive to changes in expectations about macroeconomic conditions once the lower bound no longer prevents the central bank from achieving its normal stabilization objectives.²⁰ This explains the fact that in the numerical example of Eggertsson and Woodford (2003), even a commitment to a modestly expansionary policy after it would become possible to achieve the central bank's normal objectives has a dramatic effect on the severity of the output collapse and deflation that are predicted in the period when the interest-rate lower bound is binding.

A second reason why forward guidance may be needed — that again has particular force when the interest-rate lower bound is reached — is in order to facilitate *commitment* on the part of the central bank. As Krugman (1998) emphasizes using a simple two-period model, and Eggertsson and Woodford (2003) show in the context of a more fully articulated dynamic model, the future policy that one wishes for people to anticipate is one that the central bank will not have a motive to implement later, if it makes its decisions then in a *purely forward-looking* way, on the basis of its usual stabilization objectives. Hence a desirable outcome requires commitment, just as in the analysis of Kydland and Prescott (1977) — even though in this case, the problem is a lack of motive *ex post* to be as expansionary as one wanted people earlier to expect, rather than a lack of motive *ex post* to control inflation as tightly as one wanted them to expect. In practice, the most logical way to make such commitment achievable and credible is by *publicly stating the commitment*, in a way that is sufficiently unambiguous to make it embarrassing for policymakers to simply ignore the existence of the commitment when making decisions at a later time.

These considerations establish a straightforward case for the benefits that should

²⁰The reason for this is closely related to the observation above that New Keynesian models commonly imply that a commitment to a fixed nominal interest rate forever results in indeterminacy of equilibrium. Mathematically, this indeterminacy reflects the fact that when the nominal interest rate is fixed and assumed not to vary with changes in output or inflation, the mapping from expected future macroeconomic conditions into current conditions has an eigenvalue greater than one, so that the deviation from steady-state values that must be expected a period in the future in order to generate a given size deviation from steady-state values now is smaller in magnitude than the current deviation that is produced. If the constant nominal interest rate is extended indefinitely into the future, this makes it possible for bounded departures from the steady-state values to be purely self-fulfilling. But even if the constant nominal interest rate lasts for only a finite time, the same result implies that small changes in expectations about conditions later can generate larger changes in current conditions.

be attainable, at least in principle, from the right kind of advance discussion of future policy intentions. On the other hand, some caution is appropriate as to the conditions under which such an approach should be expected to work. It does not make sense to suppose that *merely expressing* the view of the economy’s future path that the central bank would currently wish for people to believe will automatically make them believe it. If speech were enough, without any demonstrable intention to *act* differently as well, this would be magic indeed — for it would allow the central bank to stimulate greater spending while constrained by the interest-rate lower bound, by telling people that they should expect expansionary policy later, and then *also* fully achieve its subsequent stabilization objectives, by behaving in a way that is appropriate to conditions at the time and paying no attention to past forecasts. But there would be no reason for people believe central-bank speech offered in that spirit.

Hence it is important, under such an approach to policy, that the central bank not merely give thought to the future course of conduct that it would like for people to anticipate, and offer this as a forecast that it would like them to believe. It must also think about how it intends to approach policy decisions in the future, so that the policy that it wants people to anticipate will actually be put into effect, and about how the fact that this history-dependent approach to policy has been institutionalized can be made visible to people outside its own building. These matters are not simple ones, and require considerable attention to the way the central bank communicates about its objectives, procedures and decisions. The problem is all the more difficult when one must communicate about how an unprecedented situation will be dealt with.

2.1 Date-Based Forward Guidance During the Recent Crisis

As mentioned above, the global financial crisis that reached its most intense phase after the fall of 2008 resulted in many central banks slashing their policy rates to their effective lower bounds by early in 2009 (if not even sooner); yet economic activity remained far below potential and unemployment surged. The desire to provide further stimulus to aggregate demand other than through further cuts in the policy rate led to experimentation with a variety of types of “unconventional” policies, including unprecedented uses of explicit forward guidance. In particular, several central banks made statements indicating that they expected to maintain a fixed policy rate for a

specific period of time.

A particularly explicit example of forward guidance was the Bank of Canada's statement on April 21, 2009, which announced the following:

The Bank of Canada today announced that it is lowering its target for the overnight rate by one-quarter of a percentage point to 1/4 per cent, which the Bank judges to be the effective lower bound for that rate.... With monetary policy now operating at the effective lower bound for the overnight policy rate, it is appropriate to provide more explicit guidance than is usual regarding its future path so as to influence rates at longer maturities. Conditional on the outlook for inflation, the target overnight rate can be expected to remain at its current level until the end of the second quarter of 2010 in order to achieve the inflation target.

While the statement included the announcement of a reduction in the current target rate, it *also* offered explicit guidance about where the target should be expected to be, extending more than a year into the future. The release of the statement had an almost instantaneous effect on market expectations about the future path of the policy rate, as indicated by trading in overnight interest-rate swap (OIS) contracts (Figure 2).

The tick-by-tick transactions data plotted in the figure show that market OIS rates fell almost instantaneously at the time that the announcement was made (9:00 AM EST, shown by the vertical line). Moreover, not only did OIS rates for maturities as long as six to twelve months fall, but the longer maturities *fell more*; that is, not only did the OIS yield curve fall in response to the announcement, but it also *flattened*. This implies either that expectations of policy rates for months in early 2010 fell even more than did nearer-term expectations, or that uncertainty about the path of the policy rate over the coming year was substantially reduced (reducing the term premium). Either of these interpretations could be a plausible consequence of the Bank's unprecedented (albeit conditional) commitment to a particular value for the policy rate over the coming year, on the assumption that it was (at least partially) *believed*. Instead, neither would be expected to follow from a simple announcement of a cut in the current policy rate, which would typically steepen the yield curve.²¹

²¹See Chehal and Trehan (2009) and Woodford (2012b) for further discussion of the effects of this




Figure 2: Intraday OIS rates in Canada on April 21, 2009. The dotted vertical line indicates the time of release of the Bank of Canada’s announcement of its “conditional commitment” to maintain its policy rate target at 25 basis points through the end of the second quarter of 2010. Source: Bloomberg.

This seems a fairly clear example of interest-rate expectations being changed by explicit forward guidance from a central bank. It should not surprise one that the clearest such evidence occurs in the case where a central bank most clearly indicated its intention to provide such guidance — both referring to its statement as having made a “conditional *commitment*”²² rather than simply offering a forecast, and stating its intention to “provide more explicit guidance” in order to “influence [longer-term] rates.”

The U.S. Federal Reserve has also made even more extensive use of date-based forward guidance in response to the crisis. In the case of the Fed, statements indicating an expectation that interest rates would remain low “for a considerable period”

announcement.

²²The word “commitment” was used in the title of the press release, as well as in the text.

were first used as a substitute for further interest-rate cuts in 2003.²³ The FOMC began using similar language in its post-meeting statements as soon as its effective lower bound (a band between zero and 25 basis points for the federal funds rate target, with interest paid on reserves at 25 basis points) was reached in December 2008; at that time, it announced not only the target cut, but that it was expected to be maintained “for some time.” In its statement of March 18, 2009, this declaration was strengthened (without any change in the target band), to state that conditions were likely to warrant a low funds rate “for an extended period.” A more aggressive form of forward guidance was first adopted in the statement of August 9, 2011, in which the main news was the line: “The Committee currently anticipates that economic conditions ... are likely to warrant exceptionally low levels of the federal funds rate at least through mid-2013.” The forward guidance was further strengthened in the statement released on January 25, 2012, to say “... at least through late 2014.” On September 13, 2012, the date was moved back to “at least through mid-2015,” in addition to other changes in the forward guidance that are discussed further below.

As discussed in Swanson and Williams (2012), Woodford (2012b), and Raskin (2013), there is considerable evidence that these statements had substantial effects on market expectations regarding the future path of interest rates, particularly the ground-breaking introduction of reference to a specific date, nearly two years in the future, in August 2011. This includes high-frequency data on movements of the OIS yield curve around the announcements, similar to Figure 2; but it also includes survey evidence on professional forecasts of the future path of the federal funds rate, the implied probability distributions for future levels of the funds rate that can be inferred from market pricing of interest-rate options, and changes over time in the sensitivity of interest-rate futures prices to macroeconomic data releases (consistent with an acceptance that the funds rate was likely to remain pinned at the lower bound regardless of these developments).

2.2 Date-Based Forward Guidance in an Inflation-Forecast Targeting Regime

The examples given in the previous section indicate that central-bank statements can influence the interest-rate expectations of market participants, in the case of

²³See Woodford (2005) for a discussion of the earlier episode.

central banks (such as the Bank of Canada and the Federal Reserve) that did not ordinarily offer guidance about the likely future path of their policy rates, outside of the unusual circumstances associated with reaching the interest-rate lower bound. But as discussed in section 1, there is a case for the desirability of *routinely* publishing the central bank’s projections of the forward path of the policy rate, as part of an inflation-forecast targeting procedure. If this is done, is there *also* room for forward guidance of the more special kind that the Bank of Canada and the Fed have sought to use, in the case that a bank finds itself constrained by the effective lower bound for its policy rate? The case of Sveriges Riksbank is of particular interest in this regard. As noted in section 1, the Riksbank has since February 2007 included in each issue of its *Monetary Policy Report* a projected forward path for the repo rate (the Riksbank’s operating target for the overnight rate²⁴), which is the ordinary instrument of policy. While in this sense the Riksbank had begun to routinely use forward guidance as a dimension of policy even prior to the global financial crisis, in the aftermath of the crisis the Riksbank has also announced on more than one occasion that its policy rate would remain fixed for a specified period of time, as a substitute for a larger immediate cut in the policy rate — a form of forward guidance with important similarities to the more *ad hoc* announcements discussed in the previous section.

In a review of Sweden’s experience, Deputy Governor Lars Svensson (2010) argues that, through December 2008, the Riksbank had been relatively successful at “managing expectations” through its policy. Often, he notes, market expectations were already fairly close to the announced forward path for the repo rate prior to the announcement, which he regards as an indication that the bank had succeeded in conducting a predictable policy and in making the systematic character of its policy evident to the public. “When there were some discrepancies,” he writes, “in most cases the market adjusted its expectations towards the [announced] policy-rate path after the announcement” (p. 48).

The effects of the Riksbank’s more recent experiments with announcements of an anticipated duration for the current repo rate have been more mixed.²⁵ On April

²⁴It is called “the repo rate” because at one time the bank’s policy was implemented through lending at that rate under repurchase agreements, though this is not currently the case. It now defines the center of a corridor for the overnight rate, 20 basis points in width, maintained by the Riksbank.

²⁵For the economic context of the experiments with forward guidance discussed here, and information about the Riksbank’s other policy measures during the same period, see Elmér *et al.*




Figure 3: Market expectations of the forward path of the repo rate in Sweden, before and after the Riksbank’s press release on April 21, 2009 that indicated that the repo rate was “expected to remain at a low level until the beginning of 2011.” Source: Sveriges Riksbank.

21, 2009 (a few hours before the Bank of Canada announcement discussed above), the Riksbank announced a cut of the repo rate to 50 basis points, together with a statement that “the repo rate is expected to remain at a low level until the beginning of 2011,” a date nearly two years in the future. The statement was accompanied by the release of a *Monetary Policy Update*, with a projected forward path which showed the repo rate at a constant level of 50 basis points through the end of 2010, as shown in Figure 3.

The figure shows the actual path of the repo rate as a solid black line (a step function); the projected forward path from April onward that was published on April 21; the market expected forward path, as inferred by the Riksbank on the basis of interest-rate forward and swap rates²⁶ the day before the announcement; and

²⁶See Svensson (2010, footnote 7) for more details. The implied forward rates include corrections

the corresponding market expected forward path after the announcement.²⁷ Market participants evidently had expected an even larger cut in the repo rate than occurred, and for the repo rate to remain lower, at least for some months, than was indicated by the projected path. In response to the announcement, the market expected path rose, though still remaining lower than the path projected by the Riksbank, for the first few months after April. By early 2010, market participants had anticipated that the repo rate would already be rising above 50 basis points, whereas the Riksbank projected it to remain at 50 basis points for another year; but in response to the announcement, the market expected path for 2010 rose still further.²⁸

The result is that an announcement that was intended to shift *down* the anticipated forward path of rates, by announcing that a low rate would be maintained until the beginning of 2011, and so to immediately lower longer-term interest rates, had exactly the opposite effect: long rates rose, because the entire anticipated forward path of rates shifted *up*. What went wrong? While many things happened from one day to the next — as noted above, the Bank of Canada introduced its own “conditional commitment” six hours after the Riksbank’s announcement — it seems clear that it was the Riksbank’s announcement that moved market expectations. Figure 4

for credit risk and maturity premia. Of course, this represents only one possible measure of expectations regarding future Riksbank policy. Survey expectations provide another measure (also closely monitored within the bank); but these are not available at high enough frequency for the exercise undertaken in the figure. On the use of both measures by the Riksbank, see for example Elmér *et al.* (2012).

²⁷The figure also shows the Riksbank’s previously announced repo-rate path, from February, so as to show to what extent the new path represented a change from the bank’s own most recent forecast.

²⁸It is important to note that near the zero interest-rate lower bound, uncertainty about the future path of the policy rate is likely to have an asymmetric effect on the forward path that is inferred from forward and swap rates using the method employed in this figure: because it is possible for the rate to be higher than the path viewed as most likely, but not too much lower, the conditional expectation of the forward path is likely to lie above the path actually viewed as most likely by market participants. Hence the paths viewed as most likely by market participants, both before and after the announcement by the Riksbank, may well have been somewhat lower than those shown in the figure. This does not, however, change our conclusion that the market expected path rose in response to the announcement, rather than falling — unless, perhaps, one supposes that there was a substantial increase in uncertainty in response to the announcement. The latter interpretation would also indicate a failure of the Riksbank’s forward guidance to be fully credible, albeit of a different sort than the one discussed in the text.




Figure 4: Intraday Swedish OIS rates on April 21, 2009. The dotted vertical line indicates the time of the Riksbank’s press release (9AM in Sweden, or 3AM EST). Source: Bloomberg.

shows the intraday OIS rates for Sweden on April 21, with the time of the release of the *Monetary Policy Update* shown; the entire term structure of OIS rates moved up within two hours of the release, and well before any news from North America.

What seems to have happened is that market participants took on board *part* of the Riksbank’s forward guidance, and modified their own forecasts to conform more with it: the projection of a path that never fell below 50 points convinced many that (contrary to prior expectations) the Riksbank would not cut the repo rate below that level. This implied an increase in the projected path for the next two quarters. But since the news, as far as market participants were concerned, was that the Riksbank was less inclined toward interest-rate cuts than they had supposed, the *entire* path was also shifted up.

In fact, the Riksbank’s projected forward path contained *two* notable features: it was announced that the repo rate was projected to remain low for nearly two years

into the future, *and*, quite remarkably relative to prior figures, it was projected to remain absolutely constant over that time — the only obvious reason for which would have to have been a decision to treat 50 basis points as the effective lower bound. It is true that the April *Monetary Policy Update* contained no announcement that this was a lower bound; it even referred to “some probability of further cuts in the future.” But as Svensson (2010) notes, it also emphasized that “the repo rate is now close to its lower limit,” and stated that “with a repo rate at this level, the traditional monetary policy has largely reached its lower limit.” Moreover, immediately after admitting the possibility in principle of further cuts, it cautioned: “But when the repo rate is at such low levels, one must consider the fact that this could have negative effects on the functioning of the financial markets.” It is easy enough to see how market participants could have read such remarks as indicating an intention by the Riksbank not to reduce the rate below 50 basis points (at least, under any but exceedingly dire circumstances). Such an announcement would, of course, be precisely the sort that should most affect market expectations: because it was interpreted as revealing something not previously known about the central bank’s *intentions* with regard to policy, rather than the central bank’s judgments about the economic outlook — and so, a matter about which the bank could undoubtedly be regarded as the most knowledgeable authority.²⁹

The Riksbank’s other message — that it expected not to raise the repo rate before 2011 — evidently made less of an impression. One reason might have been an assumption that this reflected the Riksbank’s pessimism about the Swedish economy, and market participants might have been more optimistic, and so expected rate increases to be justified sooner than the bank anticipated. Svensson (2010) argues instead that survey data on traders’ forecasts of inflation and growth indicate that they were no more optimistic than the Riksbank, and hence that market participants simply did not accept the Riksbank’s forecasts about *its own future approach to policy*.

Why might this have been? It is notable that a large (and persistent) discrepancy between the forward paths announced by the Riksbank and those expected by market participants appeared only when the Riksbank began attempting to use projections of a policy rate that would remain fixed for an unusually long time, as a consequence of having reached its (self-imposed) lower bound. One may conjecture that the Riksbank sought, as an alternative to a deeper immediate interest-rate cut, to signal that rates

²⁹Nonetheless, the Riksbank did cut the rate further at its July meeting, as discussed below.




Figure 5: Market expectations of the forward path of the repo rate in Sweden, before and after the Riksbank’s press release on July 2, 2009, announcing an additional cut in the repo rate, and a shortening of the time that the low target was expected to be maintained. Source: Sveriges Riksbank.

would be kept low for a longer time than would ordinarily have been expected; and this supposition about future policy was incorporated into its projections. But this change in the assumption made about future policy was not credible to market participants, perhaps because no adequate explanation was given of how policy decisions would be made in the future. The mere fact that the Riksbank announced that it projected a low path for the repo rate until 2011 was not enough; market participants needed to have a view of how the Riksbank would make decisions in the future that would justify such a path (given their expectations regarding the economy’s evolution), and evidently they were not provided with one.

In July 2009, the Riksbank announced a further cut in the repo rate, to 25 basis points, but now only indicated that the target was expected to remain at its low level “until autumn 2010.” (This might be considered to vindicate skeptics who had not

believed the April projection of a low rate through the beginning of 2011.) As shown in Figure 5, this announcement did shift down market expectations of the forward path,³⁰ but market participants continued to forecast that the repo rate would not remain at that level past the end of 2009, and expected it to be around 100 basis points by autumn 2010. (In fact, it was only raised to 50 basis points in July 2010 and to 75 basis points in September.) This apparent failure to credit the Riksbank's view of the length of time that the target would remain low made policy effectively tighter (in terms of its consequences for longer-term interest rates and hence for spending decisions) during 2009 than the Riksbank's projection assumed it would be.³¹

Why would statements of an apparently similar form by the Bank of Canada and the Federal Reserve have apparently had effects closer to those that were intended? A possible explanation is that forward guidance outside the context of routine predictions about the future path of interest rates is more easily interpreted as revealing central-bank *policy intentions*. Information about policy intentions is likely to affect the expectations of market participants more than information about the central bank's view of the economic outlook, because the way in which the bank intends to conduct policy is a matter about which the bank obviously knows more than do outsiders, no matter how closely they follow economic news. And a statement that is viewed as expressing a *commitment*, that by virtue of its having been stated should at least to some extent constrain future policy decisions, should be most informative of all.

The Bank of Canada's "conditional commitment" in April 2009 seems to have been one of the examples of forward guidance that most clearly changed market expectations, and this is also the case in which a central bank came closest to committing itself to a future course of action. The Bank of Canada did not shy away from using the word "commitment" in its press release, even if this was qualified by the word "conditional," and the nature of the conditionality was not fully spelled out. Other

³⁰It is difficult to be certain how much of the reduction in expectations of the forward path should be attributed purely to the Riksbank's forward guidance regarding the repo rate, and how much might instead be due to signalling effects of the Riksbank's announcement at the same time of its intention to offer fixed-rate loans, with a maturity of a year, at a rate consistent with its declared intentions regarding the repo rate. For a discussion of the Riksbank's loan policies, their possible signalling effects, and an assessment of their quantitative effects, see Elmér *et al.* (2012).

³¹Similar credibility problems persisted once the Riksbank began tightening policy again, as discussed in Svensson (2011) and Woodford (2012b).

central banks, such as the Federal Reserve, have not gone as far; the FOMC's statements have referred only to what the Committee currently anticipates that future conditions will warrant. Yet even in these cases, observers may well have assumed that the unusual announcement made sense only if interpreted as a commitment, and indeed a good deal of commentary interpreted the FOMC's statements this way (and discussed whether the supposed promise was credible). To the extent that reasons are given for a commitment to make sense — as in the case of the Bank of Canada's explicit reference to its desire to “influence rates” through “forward guidance” — the interpretation as a commitment is also more likely.

Releases of central-bank projections of the path of interest rates, in the context of a more general discussion of the central bank's forecast of the economy's evolution over the next few years, are less susceptible to interpretation as a commitment, or even as an expression of a definite intention about future policy that has already been formed. Apart from the fact that the central banks that use this communication strategy take pains to emphasize in the accompanying text that their projections for the policy rate are merely forecasts conditional on current information, the format in which the projections are presented also makes this evident. But to the extent that such projections are viewed simply as following from the bank's forecast of the economy's evolution, including a forecast of the evolution of the policy rate given how it is typically adjusted in response to varying economic conditions, then they provide news that should change other market observers' forecasts of the future path of interest rates *only* to the extent to which they are thought to reflect superior information about the economic outlook that is available to the central bank. Other close observers of the economy may or may not believe this is true; and even when they do believe they can learn something from what the central bank reveals about its information, their own assessment of the best forecast will in general not put a weight of 100 percent on the central bank's forecast.

I have remarked above that the degree to which market participants have regarded the Riksbank's projected repo rate path as informative about the likely future path of the repo rate more than a few months into the future seems to have decreased since April 2009, when the target reached a level that the Riksbank was reluctant to go below, and a statement that the target should remain at that rate for a specific (fairly long) time was offered instead of a sharper immediate reduction. This may well have been interpreted as a departure from the bank's previous practice in the

way it produced its projections — but *not*, evidently, because the bank was now interpreted as making a commitment that it could be counted upon to fulfill.

A possible reason for the reduced credibility of the longer-horizon projections at this point is that this was the first occasion on which the announced path reflected a projection of future policy decisions that were *history-dependent* to any significant extent — that is, an assumption about future policy that differed from what one would expect that policy to be simply on the basis of conditions at the time. The reason why it would be desirable for policy to be expected to be history-dependent, under precisely the circumstances reached by the Riksbank in April 2009, has already been explained above, in section 1: the anticipation *at the time of the binding lower bound* of a lower subsequent repo rate than *would be desirable on purely forward-looking grounds* at the later date could have beneficial (stimulative) effects at the time of the binding constraint, albeit at the cost of less successful stabilization later. This may well be the sort of calculation that led the Riksbank to choose a repo rate path that indicated low rates so far into the future as it did. But in the absence of any intention to actually make policy decisions in a history-dependent way later — or at any rate, in the absence of an explanation of the procedures that would be followed in the future, that made it *credible* that future policy would be made in that way — there would be no reason for market expectations about the future conduct of policy to change.

The Riksbank’s official description of its approach to monetary policy states that “in connection with every monetary policy decision, the Executive Board makes an assessment of the repo-rate path needed for monetary policy to be well-balanced” (Sveriges Riksbank, 2010, p. 14). The document goes on to explain the competing considerations that must be taken into account in such an assessment; there is no suggestion that the exercise is anything but a purely forward-looking consideration, repeated afresh in each decision cycle, of which of the feasible forward paths for the economy from that date onward is most desirable, from the standpoint of a criterion that involves both the rate of inflation (and its distance from the official inflation target of 2.0 percent) and the level of real activity. Indeed, it stresses that the appropriate repo-rate path will be reassessed in each decision cycle, so that “the interest rate path is a forecast, not a promise” (p. 15).

If the model of the economy used in such an assessment of the possible forward paths at a given point in time incorporates forward-looking private-sector behavior —

as the Riksbank’s RAMSES model (Adolfson *et al.*, 2013) certainly does — and if the model is solved under the assumption that the projected forward path of the policy rate is anticipated by those forward-looking decisionmakers, then it might easily be concluded that the most desirable forward path at a given point in time is one which assumes history-dependent policy later. This is particularly likely to be the case when the current policy rate is constrained at its lower bound. But in such a case, repetition of the forward-looking exercise at the later date will not result in a decision to continue the interest-rate path previously projected, *even if there have been no surprise developments in the meantime*; for a forward-looking assessment of “well-balanced policy” at the later date will take no account of the effects of expected policy at that date on decisions expected to be taken in the private-sector earlier, according to the policy projections made at the earlier date.³²

A purely forward-looking forecast-targeting exercise of such a kind would accordingly be intertemporally inconsistent, as discussed in section 1. This means that there would be no reason for market participants to hold the expectations assumed in the projection exercise, even if they perfectly understand the central bank’s decision procedure. The problem might be that they understand it too well — that they have a more accurate forecast of the way that future policy will be made than the one assumed in the projection exercise.

I do not mean to imply that a time-consistent procedure, that assumes that future policy will be determined in a purely forward-looking way, would necessarily be superior. Such a targeting procedure would be intertemporally consistent, but the equilibrium implemented will generally be suboptimal, from the standpoint of the criterion used by the bank itself to rank possible forward paths. In particular, in a situation where the lower bound on the policy rate becomes a binding constraint, an inability to commit to a history-dependent policy would mean acceptance of a low-output trap, and of the fact that interest-rate policy can accomplish nothing more once the lower bound on the current overnight rate is reached. What is needed in order to achieve a better outcome, despite a correct understanding of the determinants of future policy on the part of market participants, is for the central bank to

³²In discussing this pitfall of a forecast-targeting approach to monetary policy, I do not mean to assert that the approach described is necessarily that of Riksbank. It is not clear, however, that current Riksbank policy institutionalizes history-dependence of the kind required by an optimal policy commitment, and still less that market participants have been given a reason to expect this.

adopt procedures under which it will indeed implement a history-dependent policy, and then to make its intentions clear to market participants. In fact, it does need to offer a “promise,” and not merely a “forecast” — though the required form of promise need not be a commitment to a specific pre-announced path for the policy rate.

2.3 The Federal Reserve’s “Thresholds” for Withdrawal of Policy Accommodation

The Federal Reserve’s approach to forward guidance has changed in important respects over the past year. Rather than the simple date-based approach described in section 2.1, the FOMC’s more recent statements have sought to define the future economic conditions that should determine when a withdrawal of current unusually accommodative measures should begin. A first step in this direction occurred with the FOMC’s statement of September 13, 2012. In addition to extending the date until which the “exceptionally low” federal funds rate target was anticipated to be warranted (“at least through mid-2015”), the statement included new language indicating that “if the outlook for the labor market does not improve substantially,” the Committee would continue its program of MBS purchases, “undertake additional asset purchases, and employ its other policy tools as appropriate until such improvement is achieved in a context of price stability.”

While referring primarily to the conditions under which asset purchases would continue or even be increased, this statement indicated for the first time a specific economic goal that would need to be achieved in order for less aggressively expansionary policies to be appropriate. It also indicated for the first time that “the Committee expects that a highly accommodative stance of monetary policy will remain appropriate for a considerable time after the economic recovery strengthens.” This sentence refers more directly to intentions with regard to interest-rate policy, and also mentions a condition relating to the state of the real economy in connection with the timing of the eventual withdrawal of the current highly accommodative policy.

An even more dramatic change in the forward guidance with regard to interest-rate policy came with the statement released on December 12, 2012, in which the reference to a particular date until which the federal funds rate target would remain unchanged was eliminated, in favor of a discussion of the economic conditions under which it would be appropriate to begin raising it. In addition to again indicating the

expectation that accommodation would “remain appropriate for a considerable time after ... the economic recovery strengthens,” the FOMC indicated that it “currently anticipates that [the current] exceptionally low range for the federal funds rate will be appropriate at least as long as the unemployment rate remains above 6-1/2 percent, inflation between one and two years ahead is projected to be no more than a half percentage point above the Committee’s 2 percent longer-run goal, and longer-term inflation expectations continue to be well anchored.” In addition to these necessary conditions for a withdrawal of accommodation, the statement indicated that the timing would also depend on “other information, including additional measures of labor market conditions, indicators of inflation pressures and inflation expectations, and readings on financial developments.” This reference to particular future economic conditions, and above all the specification of precise quantitative “thresholds” for two variables (the unemployment rate and the inflation projection), has attracted considerable comment, and in August 2013 the Bank of England adopted thresholds of a similar form as a basis for its own forward guidance.

The FOMC’s move away from date-based forward guidance has much to recommend it. If viewed as an actual *commitment* not to raise the federal funds rate before the stated date, regardless of what might happen in the meantime, such a policy would be far from best — indeed, in the case of a commitment extending two years or more into the future, it could prove quite reckless. It is important to note that this is *not* the type of policy recommended by theoretical accounts of the desirability of forward guidance. Campbell *et al.* (2012) refer to the “late 2014” statement language introduced in January 2012 as implementing “the policy recommendations of Eggertsson and Woodford (2003),” but Eggertsson and Woodford (2003) do not argue for the desirability of a commitment to keep the policy rate at zero for a fixed period of time. They argue for the desirability of a commitment to conduct policy in a different way than a discretionary central banker would wish to, *ex post*, and show that in their New Keynesian model the optimal commitment involves keeping the policy rate at zero for some time after the point at which a forward-looking inflation-targeting bank (or a bank following a forward-looking Taylor Rule) would begin to raise interest rates. But the date T until which the policy rate should be kept at zero is not a date that can be announced with certainty at the time of the shock that causes the zero lower bound to bind; its optimal value depends on how the economy

develops.³³

The reason for this is simple: the interest rate that will be optimal, simply from the standpoint of its suitability to the conditions that have arisen *ex post*, will generally be state-contingent. An optimal commitment will generally specify a different policy than this, in order to take account of the effects of the anticipation of policy at earlier dates. But, at least to a first approximation, the latter effects depend only on the average level of interest rates that are expected at the later date, averaging over all of the various situations that might arise; so this consideration makes little change in the way in which it is desirable for interest-rate policy to *differ across states* at the later date.

In fact, the FOMC's date-based forward guidance was never expressed as a commitment in any event; the Committee was careful only to offer statements about what it "currently anticipate[d]" — and indeed, not what it anticipated about future policy decisions, but only what it anticipated that future "economic conditions [were] likely to warrant." Thus it offered only its own *predictions* of what was coming, with no indication that this represented a *decision* to behave in a different way.³⁴ It was thus always possible to interpret the FOMC's announcements about future policy as simply reflecting changes in the Committee's view of likely future economic conditions, and hence the path of the funds rate that could be expected under their normal reaction function. For example, when the FOMC announced in January 2012 that "the Committee ... currently anticipates that economic conditions ... are likely to warrant exceptionally low levels for the federal funds rate at least through late 2014," the headline of the *New York Times* online story about the announcement was "Fed Signals That a Full Recovery Is Years Away."

But if an announcement that the date T at which the policy rate will first rise above its lower bound has moved farther into the future is interpreted as meaning that the first date at which a standard (purely forward-looking) Taylor Rule would require a policy rate above the floor has moved farther into the future because of a weakening

³³In their paper, they illustrate numerically how it should depend on the length of time for which the natural rate of interest remains abnormally low; and they give a more general analytical characterization of the optimal policy commitment that implies that T should depend on the evolution of cost-push disturbances as well.

³⁴This was even more obviously true of the forward guidance provided by the FOMC's decision to begin including information about individual committee members' forecasts of the future federal funds rate in the quarterly Survey of Economic Projections.

of the economic outlook — without in any way challenging the expectation that the bank will, as always, follow such a rule — then the announcement (if also believed) should have a *contractionary* effect on aggregate demand, rather than an expansionary one. For rather than implying that, at a certain point in the future, interest rates will be held lower than one would have expected prior to the announcement (so that real incomes at that time will be greater than would previously have been expected, and likely inflation as well), the announcement would instead imply that real incomes at that time will be *lower* than would previously have been expected (and likely inflation as well) — which change in anticipations should reduce current willingness to spend rather than increasing it. Forward guidance of this kind would have a perverse effect, and be worse than not commenting on the outlook for future interest rates at all.

The only way to avoid this pitfall is to accompany any discussion of the forward path of interest rates with an explanation of the considerations behind it — in particular, of the policy commitments that the anticipated forward path reflects. Discussion of the forward path of interest rates implied by a central bank’s policy commitments may well be useful, as discussed in section 1. But this does not mean that presentation of the implied forward path for interest rates suffices as an explanation of the bank’s policy commitments.

The new form of forward guidance used in the FOMC’s statements since September 2012 represents an important step in this direction, by providing information about the *economic conditions* that will need to be observed in order for the removal of policy accommodation to begin. A discussion cast in these terms is more likely to be understood as a *commitment*, and not a mere *forecast* of future conditions, and also represents a more reasonable form of commitment to make. In addition, the explicit statement in September that low rates would “remain appropriate for a considerable time after ... the economic recovery strengthens” sought to counter the interpretation that moving the anticipated “lift-off” date back to “mid-2015” represented merely increased pessimism about the timing of the recovery, and also provided, for the first time, at least an oblique indication of a decision to behave differently than the Committee’s usual reaction function would have dictated. The more specific quantitative criteria included in the December 2012 statement, together with the abandonment of any reference to a particular “lift-off” date, made both the state-contingency of the new guidance, and the extent to which it represented a shift in the reaction function relative to previous policy, even more evident.

Because of the wide attention that this development has received, it is worth commenting on the general desirability of the particular formulation chosen by the FOMC for its new form of forward guidance. The reference to a quantitative “threshold” for the unemployment rate is the feature that has attracted the greatest comment, with some presenting this as a repudiation of an inflation-targeting approach. But I do not believe that reference to an explicit quantitative target that involves the real economy, as opposed to one that refers *only* to the path of some general price index, should be viewed as incompatible with inflation targeting.

As discussed in section 1, the specification of a definite (and non-time-varying) medium-run target for inflation does not in itself suffice to determine how short-run policy decisions should be made, or how they should be expected to be made in various future contingencies; and in my view, a fully-specified inflation forecast-targeting procedure — one that actually makes it possible for the central bank to publicly justify its policy decisions by explaining how they are dictated by its policy targets — requires the medium-run inflation target to be supplemented by an intermediate *target criterion* to determine the short-run policy decision at each decision point. Given an objective for policy that takes into account real stabilization goals alongside the goal of inflation stabilization,³⁵ and the fact that a tradeoff between inflation and real activity does exist in the short run, a desirable intermediate target criterion will involve some measure of real activity or employment, rather than being a function of inflation or a price index alone. The “output-gap-adjusted price level target” proposed by Eggertsson and Woodford (2003) provides an example; the nominal GDP level target path proposed by Woodford (2012b) is another.

A commitment to an intermediate target criterion other than a pure inflation target need not undermine the credibility of a central bank’s claim to conduct policy so as to ensure a definite (and unchanging) medium-run rate of inflation. The intermediate target criterion can (and in my view should) be chosen so as to *imply* a definite long-run rate of inflation, equal to the inflation target, and should furthermore imply that inflation should be expected to return to the vicinity of the target over the span

³⁵Most inflation targeting central banks are clearly expected to take such additional objectives into account, even if they have not been spelled out as precisely as the inflation target has been. In the case of the Federal Reserve, the fact that the goal of “maximum employment” is assigned the same status as “price stability” in the Federal Reserve Account makes it even clearer that concern for the real economy is appropriate under the form of flexible inflation targeting practiced by the Fed.

of a few quarters except under highly unusual circumstances. A target criterion of the form (1.6), for example, where the target path p_t^* grows linearly at the rate π^* (the constant inflation target), has the property that if (1.6) is expected to hold at all times, then the fact that the output gap is not expected to be different from zero on average in the long run, independently of monetary policy, will necessarily imply that in the long run, the log price index p_t should also be expected to grow at the rate π^* . (Even supposing, more realistically, that the central bank will not be able to ensure that (1.6) holds exactly at all times, the conclusion will still follow as long as it is understood that the central bank is committed to prevent discrepancies between the left and right-hand sides of (1.6) from persisting for too long a time.)

As noted above, a nominal GDP level path target, defined by (1.11), is just a special case of this; hence the same argument applies to a target criterion of this latter form. Moreover, even if the term y_t^n in (1.11) is replaced by some other estimate of the potential output trend, as long as the discrepancy between the estimate of potential that is implicit in the NGDP target path will not differ from the actual natural rate of output by an amount that is allowed to *grow cumulatively* over the long run, commitment to the NGDP target path should imply that the long-run inflation rate will necessarily equal π^* .

The kind of thresholds announced by the FOMC are not obviously inconsistent with the Fed's long run inflation target (announced in January 2012, and reaffirmed in January 2013), either. First, the announced thresholds are meant to determine policy only until "liftoff" from the current near-zero level of the federal funds rate occurs; they do not specify how interest-rate decisions will be made thereafter, and so are consistent with an expectation that policy thereafter will be conducted in a way that ensures an average inflation rate of 2 percent per year. Second, even the specification of the economic conditions required to consider raising the funds rate target refers explicitly to "the Committee's 2 percent longer-run goal" for inflation, and indicates that rates could be raised (even with unemployment still above 6-1/2 percent) if inflation is projected to be too far above that rate (or if inflation expectations are too far out of line with the target).

Nonetheless, the kind of thresholds adopted allow more grounds for doubt about the FOMC's long-run policy intentions than was necessary. First, the short-run policy regime that has been announced appears to represent a break from the guidelines used to make decisions about interest-rate policy in the recent past; but the fact that the

reaction function can evidently suddenly change — with no need to justify the new rule as following from the same principles as had underlain past policy, but applied in a different situation than those confronted in the past — might reasonably create doubts about how suddenly and how soon other new policies could be announced in the future.

Second, the new short-run regime is not specified with sufficient completeness for it to be clear how large a cumulative increase in prices might be allowed before the return to a more standard approach to policy. It is true that the FOMC only states that the near-zero federal funds rate will be maintained as long as projected inflation remains below 2.5 percent; but it does not actually commit to *raising* rates in the event that this threshold is breached, it simply does not commit *not* to raise them in that case. Because *two* thresholds are specified — one for the unemployment rate and one for the inflation projection — as determinants of a single decision, it is unclear what should be expected to happen in the event that the two indicators give opposite signals — that is, if the inflation projection were to exceed 2.5 percent while unemployment remains well above 6.5 percent. To the extent that one fears that the FOMC would find it difficult to tighten policy while unemployment remained above the announced threshold, after having offered a precise numerical benchmark, one would have reason to fear a scenario under which inflation could be allowed to run above the long-run target rate for a considerable period or to a considerable extent, as a result of a mis-judgment of the current location of the natural rate of unemployment.

An alternative approach would have had significant advantages on these dimensions. The FOMC might instead have committed themselves to maintain a federal funds rate near zero as long as the level of nominal GDP continues to fall short of a target path, while explaining that they would raise the federal funds rate target when necessary to prevent NGDP from overshooting that path; thereafter, they could explain, the funds rate would be managed so as to keep NGDP close to the path. The target path might be chosen in accordance with (1.11); that is, the target path for the log of NGDP could be chosen to equal the log of the FOMC's estimate of the path of potential real GDP, plus a nominal factor that grows deterministically at a constant rate corresponding to the long-run inflation target. The initial level of the nominal factor could be chosen so that the announced target path would represent a continuation of the path of nominal GDP prior to the crisis — that is, prior to the




Figure 6: US nominal GDP compared with a target path based on the CBO estimate of potential real GDP, as explained in the text. Source: Bureau of Economic Analysis and Congressional Budget Office.

point at which it ceased to be possible for the Fed to keep nominal GDP on its prior trend path using its normal procedures.

Figure 6 illustrates what such a target path might look like, under the current situation of the US economy. The target level shown is equal to the Congressional Budget Office’s 2012 estimate of the path of US real GDP, plus a nominal factor growing at a constant rate of 2 percent per year, with the intercept chosen so that the nominal GDP target exactly equals actual nominal GDP in the first quarter of 2007, the last quarter in which real GDP was (according to the CBO) at potential, and one prior to the sharp drop in nominal GDP relative to its previous trend — hence a quarter in which one might suppose that the FOMC achieved a level of nominal expenditure reasonably close to the one that it desired.³⁶ (The initial level of the target path might be determined in a different way, for example so as to splice the

³⁶This figure has been prepared using data available in April 2013.

target path going forward with an estimated trend for the years immediately prior to the crisis, without materially changing the message of Figure 6.)

One observes that since the onset of the financial crisis, nominal GDP has fallen well below this target path, and continues to run below it — still about 8 percentage points below as of the first quarter of 2013, with little sign that the gap is closing.³⁷ Hence a commitment to maintain the federal funds rate near zero until this gap is closed would imply that the funds rate target would not be increased anytime soon; indeed, a substantial acceleration of the growth rate of nominal GDP would be required in order for the funds rate to be raised before the end of 2015, as currently expected by most members of the FOMC.³⁸

At the same time, it would also achieve the other goal of the FOMC's thresholds, namely, placing a bound on the amount of inflation that the policy might turn out to involve, by strictly limiting the cumulative nominal growth that would be allowed. To the extent that one expects that eventually, real GDP must return to the path of potential estimated by the CBO — which must almost certainly be the case, if the CBO's estimate of potential is correct — then the *cumulative* inflation resulting from the policy, integrating forward from 2007Q1, can be no more than two percent per year. Greater inflation would be possible if the CBO's estimate turned out to be incorrect, and the target path were not adjusted in response to the changed estimate of potential; but even then, the number of percentage points of cumulative growth in the price level that could result would be limited by the number of percentage points by which the CBO has over-estimated potential, and this would be unlikely to be large.

At the same time, this alternative form of intermediate target would avoid the disadvantages of the FOMC's thresholds cited above. Because a threshold for a single variable (albeit one that involves *both* the general level of prices and the real economy) would be offered as the criterion for determining when it is appropriate to tighten policy, the criterion offered would be more complete, and so would allow less ambiguity to remain — both about how much nominal growth might be allowed for the sake of the FOMC's goals for the real economy, and about whether policy might be tightened prematurely owing to an inflation scare. It ought to have a particular

³⁷The cumulative decline in the gap has been less than 0.8 percentage points since 2009Q4, a rate of convergence of less than a quarter of a percentage point per year.

³⁸See the FOMC Survey of Economic Projections, released April 10, 2013.

advantage in bounding uncertainty about future inflation, because it would involve a commitment to a nominal *level* variable, rather than only a growth rate; hence the policy would ensure that *actual* nominal growth would be limited, and not merely the amount of growth that was *forecasted* some years in advance.

And if the NGDP level path were determined in the way proposed above, it would be possible to present the policy as simply extending the principles that have guided FOMC policy in the past to novel circumstances, rather than a break with past policy; the temporary period of unusually accommodative policy would be justified by the fact that nominal GDP had been allowed to fall below its previous trend path to an unusual extent. Even more importantly, the policy announced for the next few years would be completely consistent with the policy that the FOMC would also want the public to anticipate will be followed farther in the future. The near-term commitment would be to increasing nominal GDP fast enough to return to the target path; but since the target path is chosen to be one that, if followed for a period of years, would guarantee an average inflation rate near the declared long-run target rate, an expectation that the FOMC would continue to make interest-rate policy on the basis of that target path would be fully consistent with what the FOMC has said about its longer-term policy intentions.

Thus the particular form of thresholds adopted by the FOMC are not obviously ideal, even as a solution for the special circumstances currently facing the Fed, and under the institutional constraints resulting from the Fed's history, legislative mandate, and declared policy commitments. It is even less apparent that they should be adopted by other inflation-targeting central banks, in the case that they find themselves constrained by an effective lower bound on their policy rate. As I have argued above, a bank that seeks to practice inflation-forecast targeting needs in any event an intermediate target criterion as a basis for the forecast-targeting exercise through which short-term policy decisions are made, and this criterion should be one that is consistent with — indeed, the consistent pursuit of which should imply achievement of — the bank's inflation target in the medium-to-long run. (A fixed target for the unemployment rate, for example, would thus be unsuitable as a proposal of this form.) And if the target criterion that is adopted has the right form — specifically, if it specifies a target path for the level of some nominal variable and not merely the projected rate of growth looking forward from each date — there will be no need to change its form in response to a series of target misses owing to a binding interest-

rate lower bound. A commitment to a target criterion such as a nominal GDP level path would already solve the problem which the FOMC's thresholds are intended to address, so that there would be no advantage to an introduction of temporary, *ad hoc* thresholds as a modification of the standard forecast-targeting procedure.

3 Conclusion

Inflation-targeting central banks have been notable for the amount of information that they provide to the public, not only about their longer-run goals, but about the way in which they expect to conduct policy in the future in order to achieve those goals. For reasons discussed above, I regard this development as a positive one. Greater clarity within the policy committee itself about the way in which policy is expected to be conducted in the future is likely to lead to more coherent policy decisions, and greater clarity on the part of the public as to how policy will be conducted is likely to improve the degree to which the central bank can count on achieving the effects that it intends through its policy. The value of this dimension of policy has become all the more apparent under the conditions recently encountered by many central banks, in which they have found themselves constrained by an effective lower bound on their policy rates.

But while the procedures developed by inflation-targeting central banks over the first two decades of inflation targeting represent important advances in the practice of central banking, and while important progress has been made over the course of that period — especially by methodological innovators like Sveriges Riksbank and Norges Bank — there remain dimensions on which the practice of inflation-forecast targeting could still be improved, relating to the degree to which explicit guidance is given about the way in which future policy decisions will be made. Two of these are of particular importance: the adoption of a more explicit intermediate target criterion to guide short-term policy decisions, that would explain how the projected effects of monetary policy on the real economy are traded off against its projected effects on inflation, and the introduction of a commitment to error-correction into the forecast-targeting procedure, by targeting the cumulative growth in a nominal variable (such as nominal GDP), rather than only its expected growth rate looking forward.

In my view, such changes could be viewed as a completion of the program of

flexible inflation targeting, making the meaning of the central bank's policy commitments more explicit and enhancing the transparency of decisionmaking, rather than a modification of the objectives of policy, or even a fundamental change in the basic approach. Nonetheless, I believe that they would go a considerable way toward answering many of the critics who argue that inflation targeting has failed as an approach, and should be replaced. In particular, they would address two important difficulties exposed by recent developments. The first is the observation that in countries where (implicit or explicit) inflation targeting has achieved considerable stability of the inflation rate over the past two decades, even large variations in the output gap now seem to result in only mild changes in inflation or in inflation expectations; but this raises doubts about whether success in containing inflation and inflation expectations within acceptable bounds should be considered sufficient grounds for regarding monetary policy as successful. Some would draw the conclusion that inflation targets should be abandoned. The adoption of an intermediate target criterion, such as a target path for nominal GDP, that is however chosen to be consistent with (and indeed to deliver) the target inflation rate over the medium run, can address this objection, by providing a basis for short-run policy decisions that clearly would not ignore the level of real activity, while nonetheless retaining the focus on delivering a particular inflation rate over the medium run.

The second difficulty is the possibility that policy can be constrained by an effective lower bound on the policy rate. An expectation that the central bank will remain committed to a strict inflation target limits its ability to create the sort of expectations about future policy that provide the only channel through which interest-rate policy can provide additional macroeconomic stimulus in such a situation (as emphasized by Krugman, 1998); hence some would argue that the recognition that such a situation can arise in practice, and not only in theory, is a ground for abandoning inflation targets. But if inflation targeting is implemented through a commitment to a target path for the level of nominal GDP (or a similar nominal level variable), then a period of persistent target shortfalls owing to the binding lower-bound constraint on policy will require (and should be expected to require) a period of unusually aggressive easing to catch up with the target path again. Thus it should automatically create the kind of expectations regarding future policy that Krugman (1998) and others have called for, but without requiring even a temporary abandonment of standard policy targets. If the practice of inflation-forecast targeting is developed in this way,

there will be less reason to doubt its suitability as an approach to the conduct of monetary policy adequate to the challenges of the twenty-first century.

References

- [1] Adolfson, Malin, Stefan Laséen, Lawrence Christiano, Mathias Trabandt and Karl Walentin, “RAMSES II – Model Description,” Sveriges Riksbank Occasional Paper no. 12, February 2013.
- [2] Aoki, Kosuke, and Kalin Nikolov, “Rule-Based Monetary Policy under Central Bank Learning,” CEPR discussion paper no. 5056, May 2005.
- [3] Campbell, Jeffrey R., Charles L. Evans, Jonas D.M. Fisher, and Alejandro Justiniano, “Macroeconomic Effects of Federal Reserve Forward Guidance,” *Brookings Papers on Economic Activity*, Spring 2012, pp. 1-80.
- [4] Chehal, Puneet, and Bharat Trehan, “Talking About Tomorrow’s Monetary Policy Today,” Federal Reserve Bank of San Francisco *Economic Letter* 2009-35, November 9, 2009.
- [5] Eggertsson, Gauti B., and Michael Woodford, “The Zero Bound on Interest Rates and Optimal Monetary Policy,” *Brookings Papers on Economic Activity* 2003(1): 139-211.
- [6] Elmér, Heidi, Gabriela Guibourg, David Kjellberg and Marianne Nessén, “The Riksbank’s Monetary Policy Measures During the Financial Crisis – Evaluation and Lessons Learnt,” *Sveriges Riksbank Economic Review* 2012(3): 1-24.
- [7] Goodhart, Charles A.E., “Monetary Transmission Lags and the Formulation of the Policy Decision on Interest Rates,” Federal Reserve Bank of St. Louis *Review*, July/August 2001, pp. 165-181.
- [8] Goodhart, Charles A.E., “The Interest Rate Conditioning Assumption,” Financial Markets Group discussion paper no. 547, London School of Economics, October 2005.
- [9] Gorodnichenko, Yuriy, and Matthew D. Shapiro, “Monetary Policy when Potential Output is Uncertain: Understanding the Growth Gamble of the 1990s,” NBER working paper no. 12268, May 2006.

- [10] Honkapohja, Seppo, and Kaushik Mitra, “Performance of Inflation Targeting Based on Constant Interest Rate Projections,” *Journal of Economic Dynamics and Control* 29: 1867-1892 (2005).
- [11] Jansson, Per, and Anders Vredin, “Forecast-Based Monetary Policy: The Case of Sweden,” *International Finance* 6: 349-380 (2003).
- [12] Krugman, Paul R., “It’s Baaack: Japan’s Slump and the Return of the Liquidity Trap,” *Brookings Papers on Economic Activity* 1998(2): 137-206.
- [13] Kydland, Finn E., and Edward C. Prescott, “Rules Rather than Discretion: The Inconsistency of Optimal Plans,” *Journal of Political Economy* 85: 473-491 (1977).
- [14] Leitemo, Kai, “Targeting Inflation by Constant-Interest-Rate Forecasts,” *Journal of Money, Credit and Banking* 35: 609-26 (2003).
- [15] Nessén, Marianne, and David Vestin, “Average Inflation Targeting,” *Journal of Money, Credit and Banking* 37: 837-863 (2005).
- [16] Qvigstad, Jan F., “When Does an Interest Rate Path ‘Look Good’? Criteria for an Appropriate Future Interest Rate Path: A Practitioner’s Approach,” Norges Bank staff memo 2006/5, April 2006.
- [17] Qvigstad, Jan F., “Transparency and Central Bank Communication,” Norges Bank *Economic Bulletin* 2/2008, pp. 71-78.
- [18] Raskin, Matthew D., “The Effects of the Federal Reserve’s Date-Based Forward Guidance,” unpublished, Federal Reserve Board, May 2013.
- [19] Rosenberg, Irma, “Riksbank to Introduce Own Path for the Repo Rate,” a speech at Danske Bank, Stockholm, January 17, 2007.
- [20] Svensson, Lars E.O., “Inflation Forecast Targeting: Implementing and Monitoring Inflation Targets,” *European Economic Review* 41: 1111-1146 (1997).
- [21] Svensson, Lars E.O., “Optimal Inflation Targeting: Further Developments of Inflation Targeting,” in Frederic Mishkin and Klaus Schmidt-Hebbel, eds., *Monetary Policy under Inflation Targeting*, Santiago: Central Bank of Chile, 2007.

- [22] Svensson, Lars E.O., “Policy Expectations and Policy Evaluations: The Role of Transparency and Communication,” *Sveriges Riksbank Economic Review* 2010(1): 43-78.
- [23] Svensson, Lars E.O., “Practical Monetary Policy: Examples from Sweden and the United States,” *Brookings Papers on Economic Activity* 2011(2): 289-332.
- [24] Svensson, Lars E.O., “Differing Views on Monetary Policy,” speech delivered at SNS/SIFR Finanspanel, Stockholm, June 8, 2012.
- [25] Svensson, Lars E.O., and Michael Woodford, “Implementing Optimal Policy through Inflation-Forecast Targeting,” in Ben S. Bernanke and Michael Woodford, eds., *The Inflation Targeting Debate*, Chicago: University of Chicago Press, 2005.
- [26] Sveriges Riksbank, *Monetary Policy in Sweden*, June 3, 2010.
- [27] Swanson, Eric T., and John C. Williams, “Measuring the Effect of the Zero Lower Bound on Medium- and Longer-Term Interest Rates,” Federal Reserve Bank of San Francisco Working Paper no, 2012-02, May 2012.
- [28] Vickers, John, “Inflation Targeting in Practice: The U.K. Experience,” Bank of England *Quarterly Bulletin*, November 1998.
- [29] Woodford, Michael, *Interest and Prices: Foundations of a Theory of Monetary Policy*, Princeton: Princeton University Press, 2003.
- [30] Woodford, Michael, “Central-Bank Communication and Policy Effectiveness,” in *The Greenspan Era: Lessons for the Future*, Kansas City: Federal Reserve Bank of Kansas City, 2005.
- [31] Woodford, Michael, “Optimal Monetary Stabilization Policy,” in B.M. Friedman and M. Woodford, eds., *Handbook of Monetary Economics*, vol. 3B, Amsterdam: Elsevier, 2011.
- [32] Woodford, Michael, “Forecast Targeting as a Monetary Policy Strategy: Policy Rules in Practice,” in E.F. Koenig, R. Leeson, and J.B. Taylor, eds., *The Taylor Rule and the Transformation of Monetary Policy*, Stanford, CA: Hoover Institution Press, 2012a.

- [33] Woodford, Michael, “Methods of Policy Accommodation at the Interest-Rate Lower Bound,” in *The Changing Policy Landscape*, Kansas City: Federal Reserve Bank of Kansas City, 2012b.