

DISCUSSION PAPER SERIES

No. 3365

THE INFLATION FORECAST AND THE LOSS FUNCTION

Lars E O Svensson

INTERNATIONAL MACROECONOMICS



Centre for **E**conomic **P**olicy **R**esearch

www.cepr.org

Available online at:

www.cepr.org/pubs/dps/DP3365.asp

THE INFLATION FORECAST AND THE LOSS FUNCTION

Lars E O Svensson, Princeton University and CEPR

Discussion Paper No. 3365
May 2002

Centre for Economic Policy Research
90–98 Goswell Rd, London EC1V 7RR, UK
Tel: (44 20) 7878 2900, Fax: (44 20) 7878 2999
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 a private 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. Institutional (core) finance for the Centre has been provided through major grants from the Economic and Social Research Council, under which an ESRC Resource Centre operates within CEPR; the Esmée Fairbairn Charitable Trust; and the Bank of England. These organizations do not give prior review to the Centre's publications, nor do they necessarily endorse the views expressed therein.

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: Lars E O Svensson

May 2002

ABSTRACT

The Inflation Forecast and the Loss Function*

This Paper argues that inflation-targeting central banks should announce explicit loss function with numerical relative weights on output gap stabilization and use and announce optimal time-varying instrument rate paths and corresponding inflation and output gap forecasts. Simple voting procedures for forming the Monetary Policy Committee's aggregate loss function and time-varying instrument rate paths are suggested.

Announcing an explicit loss function improves the transparency of inflation targeting and eliminates some misunderstandings of the meaning of 'flexible' inflation targeting. Using time-varying instrument rate paths avoids a number of inconsistencies and other problems inherently associated with constant interest rate forecasts.

JEL Classification: E42, E52 and E58

Keywords: inflation targeting and interest rate paths

Lars E O Svensson
Princeton University
Department of Economics
Fisher Hall
Princeton, NJ 08544-1021
USA
Tel: (1 609) 258 0329
Fax: (1 609) 258 6419
Email: svensson@princeton.edu

For further Discussion Papers by this author see:
www.cepr.org/pubs/new-dps/dplist.asp?authorid=101049

* This Paper was presented at the Goodhart Festschrift, Bank of England, November 15–16, 2001. I thank Charles Bean, Charles Goodhart, Laurence Meyer, Edward Nelson, Marianne Nessén, Torsten Persson, and participants in the Festschrift and in a seminar at Sveriges Riksbank for comments, and Kathleen DeGennaro for editorial and secretarial assistance.

Submitted 15 April 2002

1. Introduction

In Goodhart (2001), Charles defends the current use by Bank of England’s Monetary Policy Committee (MPC) of inflation forecasts conditional on a constant interest rate. He also expresses misgivings about the appropriateness and feasibility of the MPC specifying an explicit loss function for monetary policy. This paper criticizes Charles’s views and argues the opposite: that the MPC can and should specify an explicit loss function, and that it should abandon the constant-interest-rate forecasts for those conditional on time-varying interest-rate paths.

Announcing an explicit loss function improves the transparency of inflation targeting and eliminates some common misunderstandings of the meaning of “flexible” inflation targeting. Using time-varying instrument-rate paths avoids a number of inconsistencies and other problems inherently associated with constant-interest-rate forecasts.

Since I end up criticizing some of Charles’s views, I am afraid that this paper may to some readers not correctly convey how deeply I admire Charles as a scholar, policy-maker and person, and how much I appreciate the opportunity (over many years in the past and hopefully many years in the future) to learn from his deep knowledge and vast experience, both by reading his work and by many discussions with him. Although I do not always agree with Charles, I always learn from him.

Section 2 deals with specifying the loss function. Section 3 deals with forecasts conditional on constant or time-varying interest rate paths. Section 4 presents some conclusions on how I believe practical inflation targeting should be further developed.

2. The loss function

There is by now widespread agreement among central bankers and academics that inflation targeting in practice is “flexible” inflation targeting, as is apparent in, for instance, several contributions in Federal Reserve Bank of Kansas City (1996, 1999): The objective is to stabilize inflation around the inflation target, but also to put some weight on stabilizing the real economy, for instance, as represented by the output gap, the difference between actual output and the “natural” output level, potential output (the level of output that would result with flexible prices).

However, without further specification, the precise monetary-policy objectives under inflation targeting are still open to interpretation and suffer from a lack of transparency. For instance,

how much weight is put on stabilizing the real economy relative to stabilizing inflation around the inflation target? Indeed, the objectives can be misunderstood. For instance, Meyer (2001b), although arguing strongly in favor of a numerical inflation target, interprets the inflation-targeting regimes in New Zealand, Canada and the United Kingdom as having a “hierarchical” mandate for price stability and contrasts this with a “dual” mandate (which he favors) in Australia and the United States. Although, as explained below, I believe this distinction between a hierarchical and dual mandate is a misunderstanding of the nature of flexible inflation targeting (and I argue in Svensson (2001a) that New Zealand is currently a prime example of flexible inflation targeting), as long as inflation targeting central banks do not announce a precise loss function for monetary policy, misunderstandings of the precise objectives are invited.

However, the objectives corresponding to flexible inflation targeting can be described precisely by a quadratic intertemporal loss function in period t ,

$$\mathcal{L}_t = (1 - \delta)E_t \sum_{\tau=0}^{\infty} \delta^\tau L_{t+\tau}, \quad (2.1)$$

where δ ($0 < \delta < 1$) is a discount factor, E_t denotes expectations conditional on information available in period t , and L_t denotes the “period” loss function. The period loss function is quadratic and given by

$$L_t = \frac{1}{2}[(\pi_t - \pi^*)^2 + \lambda x_t^2], \quad (2.2)$$

where π_t and x_t denote inflation and the output gap in period t , respectively, π^* is the inflation target, and $\lambda > 0$ is the relative weight on output-gap stabilization. Thus, inflation and the output gap are the “target variables,” that is, the variables that enter the loss function. The corresponding “target levels” are π^* and zero.

The zero target level for the output gap corresponds to an output target equal to potential output. There is general agreement that inflation-targeting central banks do normally not have overambitious output targets, that is, exceeding potential output. Thus, discretionary optimization does not result in average inflation bias, counter to the case in the standard Kydland-Prescott-Barro-Gordon setup. Since the inflation target is subject to choice but not the output target, there is an asymmetry between the inflation and output target, consistent with the inflation target being the “primary objective.” In this sense, flexible inflation targeting can be interpreted as a “hierarchical” mandate. On the other hand, given the inflation target, the objective is to minimize an expected weighted sum of squared inflation deviations from the inflation target *and* squared output deviations from potential output. In this sense, flexible

inflation targeting can be interpreted as a “dual” mandate. Thus, flexible inflation targeting can be interpreted as having *both* a hierarchical and a dual mandate, and no conflict need arise between them.

Regarding the two parameters, the discount factor and the relative weight, δ and λ , the discount factor is for all practical purposes likely to be very close to one, especially when the period is a quarter. Interestingly, when the discount factor approaches one (and the intertemporal loss function is scaled by $1 - \delta$ as in (2.1)), the intertemporal loss function approaches the weighted sum of the unconditional variances of inflation and the output gap,

$$\lim_{\delta \rightarrow 1} \mathcal{L}_t = \frac{1}{2}(\text{Var}[\pi_t] + \lambda \text{Var}[x_t]) \quad (2.3)$$

(when the unconditional means of inflation and the output gap equal the inflation target and zero, respectively: $E[\pi_t] = \pi^*$ and $E[x_t] = 0$).¹ As mentioned, flexible inflation targeting corresponds to $\lambda > 0$. “Strict” inflation targeting would be the unrealistic case of $\lambda = 0$.

Charles (2001, p. 173) states that there is a variety of problems with both establishing and minimizing a loss function. “First, formally establishing such a loss function, unless it was agreed by the Chancellor, might be seen as the MPC abrogating the right to select its own (short-term) goals; it could be thought of to involve a ‘democratic deficit.’” On the other hand, Charles implies that the Chancellor’s letter “... does provide some tightly limited room for discretion, recognizing that ‘the actual inflation rate will on occasions depart from its target as a results of shocks and disturbances. Attempts to keep inflation at the inflation target in these circumstances may cause undesirable volatility in output.’” My own view is that the Chancellor’s words are completely consistent with a loss function of the form (2.1) and (2.2), and that the MPC definitely has the right to give an operational interpretation to the Chancellor’s letter in the form of an explicit loss function. Therefore, I think that the MPC should go ahead and make the loss function explicit. Then, the Chancellor also gets the chance to voice any objection to the interpretation.²

“Second, it might be difficult for a committee to agree on any formal functional representation. The coefficients in the function would be somewhat arbitrary (and what would be done about the standard central bank practice of interest rate smoothing?). Moreover, membership

¹ However, a fine point to remember is that, since (2.3) does not allow derivatives with respect to inflation and output gap in a particular (future) period, when such derivatives are needed, they must be computed before the limit is calculated.

² Indeed, my guess is that the Chancellor’s eminent advisors had exactly such a quadratic loss function in mind, when they proposed the precise wording in the Chancellor’s letter.

of the committee is time-varying, and existing members may find that their views about the (short-run) loss function shift as the context changes.”

I find the quadratic loss function above very attractive and intuitive. In particular, the marginal loss of deviating from the target, $\partial L_t / \partial |\pi_t - \pi^*| = |\pi_t - \pi^*|$ and $\partial L_t / \partial |x_t| = \lambda |x_t|$, is rising in the distance from the target and close to zero when the distance is close to zero. Thus, it is more important to return to the target the further away from it, and inflation and the output gap close to their targets is almost as good as them being on target. The quadratic form makes the marginal loss linear, the simplest functional form for an increasing function. Therefore, I believe that the MPC members would easily see its advantages and agree to this form.³

It remains for the MPC to decide on the parameters δ and λ , as well as the estimate of potential output that is used in constructing the output gap. Regarding the inflation target, π^* , the Chancellor has already specified that inflation should be measured as 12-month increases in the RPIX and that the inflation target is 2.5 percent.⁴ As for the discount factor δ , as noted above, it can uncontroversially be approximated by one. The main sticking point is the relative weight, λ . For the MPC, voting is a natural mechanism for aggregating decisions and preferences. Then the MPC’s λ would be the *median* of the MPC members’ individual λ s.⁵ Of course, the MPC members may need some introspection and assistance in deciding what their individually favored λ is, for instance, by ranking a few potential outcomes. As a revealed-preference exercise, this is relatively trivial. Regarding the estimate of potential output, voting on it would result in the MPC’s estimate of potential output being the median of the MPC members’ individual estimates.

Indeed, this aggregation procedure into an MPC loss function is quite simple, compared to many other real-world committee decisions. After the MPC’s loss function has been specified,

³ Charles (2001, p. 179) states that his loss function is not quadratic but corresponds to the absolute deviation from the target level. Consequently, he did not focus on the mean forecast (which becomes the relevant focus with a quadratic loss function and a linear model) but the median forecast (which becomes the focus with an absolute-deviation loss function and a linear model). This would here correspond to $L_t = |\pi_t - \pi^*| + \lambda |x_t|$. But this loss function has the counter-intuitive property that it is as important to reduce inflation 0.1 percentage point when it is 0.1 percentage point away from the target as when it is 2 percentage points away from the target (since $\partial L_t / \partial |\pi_t - \pi^*| = 1$ and $\partial L_t / \partial |x_t| = \lambda$). Put differently, the derivative of the loss function with respect to one of its arguments is discontinuous at the target level, $\partial L_t / \partial \pi_t = 1$ for $\pi_t > \pi^*$ and $= -1$ for $\pi_t < \pi^*$. I find this awkward and counter-intuitive and am convinced the majority of any MPC would rather settle for a quadratic loss function.

⁴ In the inflation-targeting countries like Sweden, where the government does not determine the inflation target and the objective for monetary policy is more generally specified as “price stability,” the central bank also has to specify both the price index and the level of the inflation target that it deems consistent with price stability.

⁵ Let $\text{median}(v)$ denote the median of the elements of the vector v . For an MPC with J members, let λ_j denote the individually preferred relative weight of member j , $j = 1, \dots, J$. Then the MPC’s aggregate relative weight, $\bar{\lambda}$, will simply be given by $\bar{\lambda} = \text{median}(\lambda_1, \lambda_2, \dots, \lambda_J)$.

the members should agree to apply that loss function rather than their individual ones. The loss function can be interpreted as the reflecting the interpretation of the majority of the MPC of the Chancellor’s letter. Any (minority) MPC member who does not want to comply with the majority loss function can of course explicitly dissent and note his/her own loss function as a better interpretation of the Chancellor’s letter.

Regarding additional objectives, like interest-rate stabilization and/or smoothing (corresponding to additional terms $\lambda_i(i_t - i^*)^2 + \lambda_{\Delta i}(i_t - i_{t-1})^2$), my own view is that there is no good reason why they should enter a loss function corresponding to inflation targeting (this is discussed in some detail in Svensson (2001d, section 5.6)). I believe the observed serial correlation in actual instrument-rate settings can be explained by other circumstances (gradual updating of unobservable state of the economy, implicit history-dependence corresponding to a commitment to “continuity and predictability” or optimal policy in a time-less perspective, etc.). However, if MPC members interpret the Chancellor’s letter as implying such additional objectives, they should vote on how to specify them and make the corresponding terms in the loss function explicit. Furthermore, I believe that additional objectives like financial stability and a functioning payment system are best handled as constraints that do not bind under normal circumstances. In exceptional circumstances, when they do bind, this should be explicitly announced and entered into the motivation for policy decisions. Under normal circumstances, the constraints can be disregarded and do not affect policy.

The period loss function (2.2) is a symmetric loss function (the value for $\pi_t - \pi^*$ is the same as the value for $-(\pi_t - \pi^*)$, etc.). Some researchers have argued that asymmetric preferences are relevant in monetary policy and also examined their implications. This would require a more complex loss function. Put differently, a second order approximation is not enough, and higher order terms are needed. I find a symmetric loss function for monetary policy very intuitive, especially since these days not only too high inflation but also too low inflation is considered undesirable, due to the risk of falling into liquidity traps and deflationary spirals. Furthermore, more complex loss functions and more complicated tradeoffs may be too sophisticated to be both operational and sufficiently verifiable for reasonable accountability.⁶

⁶ Nobay and Peel (1998), ali-Nowaihi and Stracca (2001) and Ruge-Mercia (1999) examine alternative asymmetric monetary-policy loss functions. Asymmetric loss functions are frequently motivated from a descriptive rather than prescriptive point of view, for instance, corresponding to observed deviations from rational behavior. I believe an informed and competent MPC deciding on the appropriate loss function should approach the issue from a prescriptive point of view and select the most appropriate and rational loss function.

3. Forecasts

3.1. Forecast targeting

Monetary policy affects the economy with considerable lags. Current inflation and output are, to a large extent, determined by previous decisions of firms and households. Normally, current monetary-policy actions can only affect the future levels of inflation and the output gap, in practice with substantial lags and with the total effects spread out over several quarters. This makes forecasts of the target variables crucial in practical monetary policy. By “forecast targeting,” I mean using forecasts of the target variables effectively as intermediate target variables, as in King’s (1994) early characterization of inflation targeting. This means minimizing a loss function where forecasts enter as arguments.

Let us assume that the transmission mechanism is approximately linear, in the sense that the future target variables depend linearly on the current state of the economy and the instrument rate. Furthermore, assume that any uncertainty and any deviation from the simple models enter additively. Finally, let the intertemporal loss function be quadratic, as above in (2.2). It is then a standard result in optimal-control theory that so-called certainty-equivalence applies, and that optimal policy need only focus on conditional *mean* forecasts of the future target variables, that is, mean forecasts conditional on the central bank’s current information and a particular future path for the instrument rate.⁷ Since this means treating the forecasts as (intermediate) target variables (that is, putting forecasts of the target variables in the loss function), the procedure can be called “forecast targeting.”⁸

Let me be more specific. Let $i^t = \{i_{t+\tau,t}\}_{\tau=0}^{\infty}$ denote an instrument-rate plan in period t . Conditional on the central bank’s information in period t , I_t (including its view of the transmission mechanism, etc.), and its “judgment,” z^t (to be discussed further below) and conditional on alternative instrument-rate plans i^t , consider alternative (mean) forecasts for inflation, $\pi^t = \{\pi_{t+\tau,t}\}_{\tau=0}^{\infty}$, and the output gap, $x^t = \{x_{t+\tau,t}\}_{\tau=0}^{\infty}$ (consisting of the difference between y^t , the (mean) output forecast, and y^{*t} , the (mean) potential-output forecast). That is,

⁷ Wallis (1999) and Vickers (1998) have examined symmetric alternatives to the quadratic loss function, relating the form of the loss function to the “central tendency” of the forecast. Thus, (under the assumption of a linear model of the transmission mechanism) a quadratic loss function corresponds to a mean forecast, an “absolute deviation” loss function (Charles’s favorite, cf. footnote 3) corresponds to a median forecast, an “all or nothing” loss function (a so-called Dirac delta function) corresponds to a mode forecast, and a “zone of indifference” loss function corresponds to a condition of equality of the probability densities of the forecast at the edges of the zone. Again, I believe a majority of an informed and competent MPC would quickly see the advantages of a quadratic loss function and the corresponding mean forecast.

⁸ In cases when the assumptions of a linear model and quadratic loss function are not fulfilled, as discussed in Svensson (2001b), one can still apply “distribution forecast targeting,” where the forecasts are explicit probability distributions and the intertemporal loss function is the explicit or implicit integral over those distributions.

$\pi_{t+\tau,t} = \mathbb{E}[\pi_{t+\tau} | i^t, I_t, z^t]$, etc.⁹ Furthermore, consider the intertemporal loss function in period t applied to the *forecasts* of the target variables, that is, when the forecasts are substituted into the intertemporal loss function (2.1) with (2.2),

$$\mathcal{L}_t = (1 - \delta) \sum_{\tau=0}^{\infty} \delta^\tau \frac{1}{2} \left[(\pi_{t+\tau,t} - \pi^*)^2 + \lambda x_{t+\tau,t}^2 \right]. \quad (3.1)$$

Each period t , conditional on the central bank’s forecasting model, information I_t and judgment z^t , the bank should then find the combination of forecasts π^t and x^t and instrument-rate plan i^t that minimizes (3.1), and then makes the current instrument-rate decision according to the current optimal instrument-rate plan. The process will result in an endogenous reaction function for the current instrument-rate decision, a function $F(I_t, z^t)$ of the central bank’s information and judgment in period t . This reaction function need not be specified explicitly, however, and it need not be followed mechanically.¹⁰

Forecast targeting requires that the central bank has a view of what the policy multipliers are, that is, how instrument-rate adjustments affect the conditional inflation and output-gap forecasts. But it does not imply that forecasts must be exclusively model-based. Instead, it allows for extra-model information and judgmental adjustments, as well as very partial information about the current state of the economy. It basically allows for any information that is relevant for the inflation and output-gap forecasts.

How would the central bank find the optimal forecasts and instrument-rate plan? One possibility is that, conditional on the information I_t and the judgment z^t , the central-bank staff generates a set of alternative forecasts for a set of alternative instrument-rate plans. This way, the staff constructs the “feasible set” of forecasts and instrument-rate plans. The MPC would then select the combination of forecasts that “looks best,” in the sense of achieving the best compromise between stabilizing the inflation gap and stabilizing the output gap, that is, implicitly minimizes (3.1). This can be done informally with visual inspection of the forecasts. It can also be done more formally with an explicit loss function, since then the loss for each

⁹ Constructing conditional forecasts in a model without forward-looking variables is straightforward. Constructing such forecasts in a model with forward-looking variables raises some specific difficulties, which are explained and resolved in the appendix of the working-paper version of Svensson (1999). The conditional forecasts for an arbitrary interest-rate path derived there assume that the interest-rate paths are “credible”, that is, anticipated and allowed to influence the forward-looking variables. Leeper and Zha (1999) present an alternative way of constructing forecasts for arbitrary interest-rate paths, by assuming that these interest-rate paths result from unanticipated deviations from a normal reaction function.

¹⁰ For simplicity I here abstract from a time-consistency problem that arises with models with forward-looking variables. Even in the absence of an average inflation bias, this time-consistency problem results in “stabilization bias” (non-optimal coefficients in the implicit reaction function) and a lack of history-dependence. The magnitude of the problem may be small in realistic models with relatively strong backward-looking elements. The nature of the problem and possible solutions, including “a commitment to continuity and predictability” and a commitment to an optimal specific targeting rule are discussed in Svensson and Woodford (1999) and Svensson (2001d).

combination of inflation and output-gap forecasts can easily be calculated numerically. (To do the latter is one of my suggestions to the Reserve Bank of New Zealand in Svensson (2001a).)

Another possibility is that the MPC determines a “specific targeting rule,” a condition that the forecasts of the target variables must fulfill. Conditional on the information and the judgment, the staff then has to generate a combination of forecasts and instrument-rate plan that fulfills the specific targeting rule. The Bank of England and the Riksbank have formulated a simple specific targeting rule to guide policy, which can be expressed approximately as “set the instrument-rate so a constant-interest-rate inflation forecast about two years ahead is on target” (Goodhart 2001 and Heikensten 1999). As Charles puts it (2001, p. 177): “When I was a member of the MPC I thought that I was trying, at each forecast round, to set the level of interest rates so that , without the need for future rate changes, prospective (forecast) inflation would on average equal the target at the policy horizon. This was, I thought, what the exercises was supposed to be.” With the period being a quarter, this targeting rule can be written

$$\pi_{t+8,t} = \pi^*,$$

with the understanding that the inflation forecast is constructed under the assumption of a constant interest rate. Although this specific targeting rule is both simple and operational, it is not likely to be optimal.¹¹ As is discussed in Svensson (2001d), an *optimal* specific targeting rule instead expresses the *equality of the marginal rates of transformation and substitution* between the target variables in an operational way.

As an example (the details are explained in Svensson (2001d)), consider a variant of the popular New Keynesian model, where inflation and the output gap are predetermined one period (a small concession to realism relative to the standard variant when both inflation and the output gap are treated as forward-looking variables, a.k.a. jump variables) and, in particular, “judgment” matters. The aggregate-supply/Phillips curve is

$$\pi_{t+1} = \pi_{t+2|t} + \alpha_x x_{t+1|t} + \alpha_z z_{t+1}, \tag{3.2}$$

where $\pi_{t+2|t}$ denotes expectations in period t of inflation in period $t + 2$, etc., α_x is a positive constant, α_z is a row vector, and z_{t+1} is a column vector, the “deviation”, to be explained below. Thus, inflation in period $t + 1$ is determined by expectations in period t of inflation in

¹¹ Bank of England (2000, p. 67) actually states that “[h]owever, there is no mechanical link between the projected level of inflation in two years time based on constant interest rates and the appropriate current setting of monetary policy.”

period $t + 2$ and of the output gap in period $t + 1$ and by the deviation in period $t + 1$. The aggregate-demand/IS curve is

$$x_{t+1} = x_{t+2|t} - \beta_x(i_{t+1|t} - \pi_{t+2|t}) + \beta_z z_{t+1}, \quad (3.3)$$

where β_x is a positive constant, $i_{t+1|t}$ is the expectation in period t of the nominal interest rate in period $t + 1$ and β_z is a row vector. Thus, the output gap in period $t + 1$ is determined by expectations in period t of the output gap in period $t + 2$ of the output gap in period $t + 2$ and of the real interest rate in period $t + 1$ and by the deviation in period $t + 1$.

The deviation represents the difference between the true model and this simplified New Keynesian model and includes all other determinants of inflation and the output gap. For simplicity it is treated as an exogenous variable. The central bank's "judgment," $z^t \equiv \{z_{t+\tau,t}\}_{\tau=0}^{\infty}$, is the central bank's best forecast of the deviation. This is a way to represent the importance and inevitability of judgment in monetary policy. Conditional on the central bank's judgment, the bank's forecasting model in period t is then given by

$$\pi_{t+\tau,t} = \pi_{t+\tau+1,t} + \alpha_x x_{t+\tau,t} + \alpha_z z_{t+\tau,t}, \quad (3.4)$$

$$x_{t+\tau,t} = x_{t+\tau+1,t} - \beta_x(i_{t+\tau,t} - \pi_{t+\tau+1,t}) + \beta_z z_{t+\tau,t}, \quad (3.5)$$

for forecast horizons $\tau \geq 1$ (where $\pi_{t+\tau,t}$ refers to the central bank's τ -period-ahead forecast of inflation in period t , etc.).

The optimal specific targeting rule for the loss function (3.1) and the model (3.4) and (3.5) can then be found by finding the marginal rate of transformation (MRT) and substitution (MRS) between (the forecasts of) the target variables (inflation and the output gap), and setting these equal.¹² A marginal increase in inflation two periods ahead only, $d\pi_{t+2,t} > 0$, $d\pi_{t+j,t} = 0$, $j \neq 2$, by the aggregate-supply relation (3.4) requires a fall in the output gap one period ahead, $dx_{t+1,t} = -d\pi_{t+2,t}/\alpha_x < 0$, and an equal increase in the output gap two periods ahead, $dx_{t+2,t} = -dx_{t+1,t} > 0$. We can then define the marginal rate of transformation of the linear combination $\tilde{x}_{t+1,t} \equiv (x_{t+1,t}, x_{t+2,t}) \equiv (1, -1)x_{t+1,t}$ into $\pi_{t+2,t}$, $\text{MRT}(\pi_{t+2,t}, \tilde{x}_{t+1,t})$, which will equal

$$\text{MRT}(\pi_{t+2,t}, \tilde{x}_{t+1,t}) \equiv \left. \frac{d\pi_{t+2,t}}{dx_{t+1,t}} \right|_{dx_{t+2,t} = -dx_{t+1,t}} = -\alpha_x.$$

¹² For simplicity, inflation π_t in (3.1) is then taken to refer to quarterly inflation, as in (3.4) and (3.4), and not 4-quarter inflation. Nessén (2001) and Nessén and Vestin (2000) examine the consequences of explicit multi-period averages of inflation in the loss function.

From the loss function (3.1) follows that the marginal rate of substitution of $\pi_{t+2,t}$ for $x_{t+j,t}$ is given by $\text{MRS}(\pi_{t+2,t}, x_{t+j,t}) \equiv d\pi_{t+2,t}/dx_{t+j,t}|_{d\mathcal{L}_t=0} = -\lambda x_{t+j,t}/(\pi_{t+2,t} - \pi^*)$ (in the limit when $\delta \rightarrow 1$, for simplicity). From this it is easy to show that the marginal rate of substitution of $\pi_{t+2,t}$ for the above linear combination $\tilde{x}_{t+1,t}$, $\text{MRS}(\pi_{t+2,t}, \tilde{x}_{t+1,t})$, will be given by

$$\text{MRS}(\pi_{t+2,t}, \tilde{x}_{t+1,t}) \equiv \frac{d\pi_{t+2,t}}{d\tilde{x}_{t+1,t}} \Big|_{d\mathcal{L}_t=0, dx_{t+2,t}=-dx_{t+1,t}} = \frac{\lambda(x_{t+2,t} - x_{t+1,t})}{\pi_{t+2,t} - \pi^*}.$$

Redoing this for $\pi_{t+\tau,t}$ for all $\tau \geq 1$ and setting the marginal rates of transformation equal to the marginal rates of substitution leads to the optimal specific targeting rule,

$$\pi_{t+\tau,t} - \pi^* = -\frac{\lambda}{\alpha_x}(x_{t+\tau,t} - x_{t+\tau-1,t}), \quad (3.6)$$

where $x_{t,t}$ for $\tau = 1$ is understood to be $x_{t,t-1}$, the one-period-ahead forecast of the output gap in period $t - 1$. Thus, the optimal targeting rule in this example can be expressed as “find an instrument-rate path so the inflation-gap forecast is $-\lambda/\alpha_x$ times the change in the output-gap forecast.”¹³

As discussed more thoroughly in Svensson (2001d), the optimal specific targeting rule has the attractive properties that it only depends on the marginal tradeoffs between the target variables. Therefore, it only depends on the loss function (via the relative weight λ) and the form of the aggregate supply/Phillips curve (via the slope of the short-run Phillips curve, α_x). In particular, judgment does not enter explicitly in the optimal targeting rule. Still, judgment will be incorporated in the construction of the forecasts. Furthermore, the targeting rule solves the time-consistency problem, so that it corresponds to the full commitment equilibrium “in a time-less perspective” (Woodford 1999a and Svensson and Woodford 1999).

In this example, inflation-forecast targeting can then be described as: (1) Conditional on the judgment $z^t \equiv \{z_{t+\tau,t}\}_{\tau=0}^{\infty}$ and the aggregate-supply relation (3.4), find inflation and output gap forecasts that fulfill the specific targeting rule (3.6). (2) Substitute these forecasts into the aggregate-demand relation (3.5) so as to find the corresponding instrument-rate plan. (3) Announce these forecasts and the instrument-rate plan, and set the current instrument rate accordingly. This results in the optimal/appropriate instrument-rate setting, conditional on the judgment, z^t , without having to specify the optimal reaction function.¹⁴

¹³ As is explained in Svensson (2001d), (3.6) also applies for $\tau = 1$, when $x_{t,t}$ is interpreted to be $x_{t,t-1}$. Formulating the targeting rule this way leads to “optimality in a time-less perspective,” corresponding to a situation of commitment to optimal policy far in the past, as discussed in Woodford (1999a) and Svensson and Woodford (1999).

¹⁴ As is shown in Svensson (2001d), even for this relatively simple model, the optimal reaction function is overwhelmingly complex, especially since it must specify how to respond optimally to judgment, making verifiability and commitment directly to the optimal reaction function completely unrealistic.

We note that the optimal specific targeting rule (3.6) refers to the whole future path of the inflation and output-gap forecasts. It does not refer to a specific horizon, like the two-year horizon emphasized by Bank of England and the Riksbank at which the inflation forecast shall be on target. Indeed, the focus on a specific horizon is not supported by this approach.

Furthermore, as discussed in Svensson (1999) and (2001c), inflation-forecast targeting, either in the general form of minimizing a loss function over forecasts or in the specific form of fulfilling a specific targeting rule is generally *not* the same thing as implementing a “forecast-based” instrument rule, as

$$i_t = \gamma(\pi_{t+T,t} - \pi^*), \quad (3.7)$$

where the instrument rate responds to a T -period-ahead inflation forecast, or the variants thereof that originated in Bank of Canada’s Quarterly Projection Model and are examined by, for instance, Batini and Haldane (1999), McCallum and Nelson (1999) and Batini and Nelson (2001).¹⁵

3.2. Optimal time-varying instrument-rate paths instead of constant ones

The above decision-making process centers on finding the optimal combination of inflation and output-gap forecasts and instrument-rate plan, the (π^t, x^t, i^t) that minimizes the intertemporal loss function or fulfill the specific targeting rule, conditional on current information, including the view of the transmission mechanism, and current judgement. There is no reference to the forecasts conditional on unchanged instrument rates used by Bank of England and the Riksbank (and by the Eurosystem as well). Thus, the process involves the MPC agreeing on forecasts and instrument-rate plans that normally are time-varying.

Using forecasts conditional on constant-interest-rate forecasts leads to a variety of problems and inconsistencies. Charles (2001, p. 171) lists most of them: (1) Such instrument-rate paths are normally not optimal. (2) In backward-looking models they lead to Wicksellian instability, in that inflation eventually veers off its target level, rather than approaches this asymptotically as optimal paths tend to do. In models with forward-looking models, they imply indeterminacy of the forecasts. A frequent remedy is to apply a stabilizing interest-rate reaction function from a period beyond the 8-quarter forecast horizon, at which point the interest rate frequently jumps substantially. This is “spatchcocked,” as Charles put it. (3) Market expectations of interest rates

¹⁵ Batini and Nelson (2001) discuss these two very different definitions of the policy horizon in monetary policy, calling them the “optimal policy horizon” and the “optimal feedback horizon.” The former refers to the horizon at which inflation reaches the target after a shock away from the target; the latter refers to the optimal horizon T for a forecast in a forecast-based instrument rule. In general, there is no specific relation between the leads of inflation that appear in the optimal specific targeting rule and the leads that correspond to these optimal-horizon definitions. Put differently, there is no specific “optimal horizon.”

normally correspond to a time-varying path. One alternative is then to use market expectations and exchange rates and other asset price as inputs in the forecasts that are inconsistent and hence systematically falsified by the interest-rate path used and inflation and forecast paths constructed. The other alternative (which I favor, see footnote 9) is to use hypothetical market expectations and asset prices consistent with the constant interest-rate path assumed, as if the private sector actually believed that the interest-rate path would be constant. This allows internal consistency (which I favor) but implies using inputs in the forecasts that are normally quite different from observed market prices. (4) Since the constant interest rate is not the best forecast for the actual interest rate, the corresponding inflation and output-gap forecasts are not the best forecasts of actual outcomes. This makes it more difficult and less relevant to compare actual outcomes to central bank forecasts.¹⁶

In spite of all these problems, and to the surprise of Charles's discussant of that paper (Meyer 2001a) (and to my own surprise), Charles ends up defending constant-interest-rate forecasts, on the ground that there are no feasible better alternatives. Charles (2001, p. 172-173) argues that "it is hard to see how a committee could ever reach a majority for any particular time path. A great advantage of restricting the choice to what to do now, this month, is that it makes the decision relatively simple, even stark. Given the difficulties involved already in achieving majority agreement in the MPC on this simple decision, the idea of trying to choose a complete time path by discretionary choice seems entirely fanciful and counter-productive."

I am afraid that I don't buy this argument. The MPC is already agreeing on time-varying inflation and output forecasts, so agreeing on a time-path does not seem to be impossible at all. It is true that there are general problems aggregating preferences in a MPC and that it is easiest to vote about a one-dimensional issue, like an instrument-rate level (or the parameter λ as discussed above). In particular, majority voting will lead to the median-voter outcome, in which the median of the MPC members' individually favored levels of the instrument-rate will be chosen.

Along these lines, I have a simple proposal for agreeing on an instrument-rate plan: Let each MPC member plot his/her preferred instrument-rate plan in the same graph with the future periods (quarters) on the horizontal axis and the instrument rate on the vertical axis (the resulting set of curves might cross each other at several future dates). Form the MPC's aggregate instrument-rate plan by taking the median of the instrument rates for each future

¹⁶ Leitemo (2001) provides additional analysis and criticism of constant-interest-rate forecasts.

quarter.¹⁷ This median instrument-rate plan can be seen as the result of a majority vote in a particular voting procedure.¹⁸ Conditional on this instrument-rate plan, agree on the inflation and output-gap forecasts. If necessary, let each MPC plot his/her conditional inflation and output-gap forecast, and pick the median outcome of these. If these do not look good, let each MPC member consider new individual instrument-rate plans, and then take the median of these. I would be surprised if this procedure does not converge very quickly.

The observant reader will realize that a median instrument-rate plan and inflation and output-gap forecasts picked this way need not be entirely consistent, in the sense that the median inflation forecast may include segments that correspond to instrument-rate plans differing from the median instrument-rate plan. Still, I believe any such inconsistency must be a minor problem, and a final round of adjustments in the MPC's decision may explicitly aim to reduce or eliminate any such inconsistency.

The resulting instrument-rate plan and inflation and output-gap forecast should then be seen as reflecting the majority view of the MPC. Dissenters then have the option to explicitly dissent in the minutes of the meeting. The general setup with the MPC's decision reflecting the majority view and the possibility of dissent is already used by the Riksbank's Executive Board. I think it is more logical and easy to understand than the idea of the "the best collective judgment" used by Bank of England's MPC.

As an alternative to agreeing on a time-varying interest-rate path, Charles considers "optimal control procedures." After dismissing the idea of specifying a loss function, which I have dealt with above, he states:

"Third, it is not clear that optimal control procedures could be applied in practice to larger, messy forecasting models incorporating a wide variety of subjective assumptions, residual adjustments, and such other discretionary adjustments as the MPC adjusts to its own forecast."

"Fourth, if such optimal control [OC] procedures were applied to the forecast, the resulting outcome of time paths for interest rates, inflation, etc., would become a hideously complex interaction of forecast and OC procedure. ... [Introducing such OC procedures] might lead

¹⁷ Let each member j , $j = 1, \dots, J$, of the MPC individually prefer the instrument-rate plan $i^{tj} \equiv \{i_{t+\tau,t}^j\}_{\tau=0}^{\infty}$ in period t . Then the MPC's aggregate instrument-rate plan, $\bar{i}^t \equiv \{\bar{i}_{t+\tau,t}\}_{\tau=0}^{\infty}$, fulfills $\bar{i}_{t+\tau,t} = \text{median}(i_{t+\tau,t}^1, i_{t+\tau,t}^2, \dots, i_{t+\tau,t}^J)$ for all $\tau \geq 0$.

¹⁸ The proposal can be seen as a mechanism for aggregating preferences that avoids the so-called Condorcet paradox, that with multiple policy alternatives there may not be a policy that commands a majority vote against all alternatives (see, for instance, Person and Tabellini (2000)). The proposal means that the MPC members vote *simultaneously* on the instrument rate for all future periods, by each member first writing down his/her preferred instrument rate for each period. The aggregate instrument rate for each period, the median rate for that period, can then be seen as the result of voting on the instrument rate in that period, *independently* of the outcome of the voting for other periods.

MPC members to regard the whole exercise as a mysterious ‘black box’ whose entrails are only comprehensible to a tiny number of staff academic specialists.”

“Fifth, if the MPC should find it more difficult to understand how the resultant outcome for the relevant variables had been determined, how would it be possible to explain it to the public, or to justify the decisions that would hang in part from it? To say that we have done what our model told us was best to do is not very convincing, especially given the track record of fancy economic models.”

I cannot help interpreting this as a misunderstanding of “optimal control procedures,” or at least of the procedures that I have in mind and have discussed above. Indeed, the only thing required is that the staff and the MPC generate a few alternative combinations of inflation and output-gap forecasts and instrument-rate plans, conditional on the information available and the MPC’s judgment, and that the MPC then picks the combination that minimizes the loss function, or finds a combination that fulfills the “specific targeting rule,” as discussed above.

The advantages in using the optimal inflation and output gap forecasts and instrument plans seem overwhelming to me: (1) This combination of forecasts and instrument plan are the best forecasts of inflation, the output gap and the instrument rate conditional on the information available and the central bank’s judgment. This means that it makes sense to compare these to private-sector forecasts, and to actual outcomes. (2) When policy is credible, there would be little difference between the central-bank forecasts and market expectations for inflation, the output gap and interest rates. This means that market values of exchange rates and asset prices can without inconsistency be used as inputs in the forecasts.

Charles (2001, p. 175) presents an additional argument in favor of unchanged-interest-rate forecasts, namely that a time-variable path would imply some degree of undesirable commitment to future policy actions, and that such commitments would be burdensome and unhappy. I am afraid that I don’t buy that argument either. Observers of inflation-targeting central banks are already used to seeing published graphs of time-varying inflation and output-gap forecasts, and they have already learned that new information may warrant revisions of previously announced forecasts. There is no difference between revising a forecast of optimal time-varying interest rates due to new information and revising other forecasts. Furthermore, the Reserve Bank of New Zealand already publishes time-varying interest-rate forecasts that are revised when new information arrives, and during my review (Svensson 2001) of its policy I did not notice that this created any problems or misunderstanding by observers of the bank.

In addition, to the extent that published instrument-rate paths would be understood as some degree of commitment, this may actually be a good thing. It is a well-known result that optimal policy with forward-looking variables require a degree of history-dependence and inertia (Sachs and Oudiz 1985, Backus and Driffil 1986, Currie and Levine 1983, Woodford 1999a and c, Svensson and Woodford 1999).

4. Conclusions

Finally, let me summarize the conclusions I draw for practical inflation-targeting monetary policy from the above discussion:

- Inflation-targeting central banks should specify explicit loss functions. Without putting a specific relative weight on output-gap stabilization relative to inflation stabilization, the objectives under inflation targeting remain somewhat nontransparent and invite misunderstanding. The MPCs should simply vote on the form of the loss function and the value of its parameters, and then make the result public.
- If the MPC chooses a quadratic loss function, the advantage of which I have explained above, consistency requires that it should let the “central tendency” of the forecasts be the mean rather than the mode (that is, the most likely) forecasts.
- The use of the problematic constant-interest-rate paths should be discontinued. Inflation and output-gap forecasts should be constructed conditional on time-varying instrument-rate paths. The MPC should each decision period decide on its optimal instrument-rate path and the corresponding inflation and output-gap forecasts and make those public. Such decisions are feasible with the aggregation and voting procedures I suggest.
- The use of the simple specific targeting rule (a constant-interest-rate inflation forecast two-years ahead on target) applied by Bank of England and the Riksbank should be discontinued. The two-year horizon should be deemphasized, and the emphasis should be on finding inflation and output-gap forecast paths that minimize the loss function.
- More generally, and as argued in more detail in Svensson (2001d), I believe that inflation-targeting, neither from a descriptive nor prescriptive perspective, should be described or interpreted as a commitment to a simple instrument rule, like a Taylor rule or a forecast-based instrument rule. Instead, inflation-targeting central banks should formulate opera-

tional and approximately optimal specific targeting rules, which can be derived from their loss functions and their estimated aggregate supply relations/Phillips curves.¹⁹

References

- al-Nowaihi, Ali, and Livio Stracca (2001), “Non-Standard Central Bank Loss Functions, Skewed Risks, and the Certainty-Equivalence Principle,” working paper.
- Backus, David, and John Driffill (1986), “The Consistency of Optimal Policy in Stochastic Rational Expectations Models,” CEPR Discussion Paper No. 124.
- Bank of England (2000), *Inflation Report, November 2000*, Bank of England.
- Batini, Nicoletta, and Andrew G. Haldane (1999), “Forward-Looking Rules for Monetary Policy,” in John B. Taylor, ed., *Monetary Policy Rules*, Chicago University Press, 157–192.
- Batini, Nicoletta, and Edward Nelson (2001), “Optimal Horizons for Inflation Targeting,” *Journal of Economic Dynamics and Control* 25, 891–910.
- Currie, David, and Paul Levine (1993), *Rules, Reputation and Macroeconomic Policy Coordination*, Cambridge University Press, Cambridge
- Federal Reserve Bank of Kansas City (1996), *Achieving Price Stability*, Federal Reserve Bank of Kansas City, Kansas City.
- Federal Reserve Bank of Kansas City (1999), *New Challenges for Monetary Policy*, Federal Reserve Bank of Kansas City, Kansas City.
- Goodhart, Charles A.E. (2001), “Monetary Transmission Lags and the Formulation of the Policy Decision on Interest Rates,” *Federal Reserve Bank of St. Louis Review* July/August 2001, 165–181.

¹⁹ Optimal specific targeting rules are the Euler conditions corresponding to optimizing monetary policy. I believe it is much better to describe and prescribe inflation targeting as goal-directed, optimizing policy than as following a mechanical instrument rule. Monetary policy by the world’s more advanced central banks these days is at least as optimizing and forward-looking as the behavior of the most rational private agents. I find it strange that a large part of the literature on monetary policy still prefers to represent central-bank behavior with the help of mechanical instrument rules, like a Taylor rule or a forecast-based instrument rule. The literature long ago ceased representing optimizing households and firms as following mechanical consumption and investment functions, and instead represents their behavior by Euler conditions, optimal first-order conditions. The concept of general and specific targeting rules is designed to provide a discussion of monetary policy rules that is fully consistent with the optimizing and forward-looking nature of modern monetary policy. From this point of view, general targeting rules essentially specify operational objectives for monetary policy, and specific targeting rules essentially specify operational Euler conditions for monetary policy. In particular, an optimal targeting rule expresses the equality of the marginal rates of transformation and the marginal rates of substitution between the target variables in an operational way. I hope there will be more research along these lines in the future.

- Heikensten, Lars (1999), “The Riksbank’s Inflation Target—Clarification and Evaluation,” *Sveriges Riksbank Quarterly Review* 1/1999, 5–17.
- King, Mervyn A. (1994), “Monetary Policy in the UK,” *Fiscal Studies* 15, No. 3, 109–128.
- Leeper, Eric M., and Tao Zha (1999), “Identification and Forecasting: Joint Inputs to Policy Analysis,” working paper .
- Leitemo, Kai (2001), “Targeting Inflation by Constant-Interest-Rate Forecasts,” *Journal of Money, Credit, and Banking*, forthcoming.
- McCallum, Bennett T., and Edward Nelson (1999), “Performance of Operational Policy Rules in an Estimated Semi-Classical Structural Model,” in John B. Taylor., ed., *Monetary Policy Rules*, Chicago University Press, 1–45.
- Meyer, Laurence H. (2001a), “Comment,” *Federal Reserve Bank of St. Louis Review* July/August 2001, 183–186.
- Meyer, Laurence H. (2001b), “Inflation Targets and Inflation Targeting,” Speech at University of California at San Diego Economics Roundtable, San Diego, July 17, 2001.
- Nessén, Marianne (2001), “Targeting Inflation over the Short, Medium and Long Term,” *Journal of Macroeconomics*, forthcoming.
- Nessén, Marianne, and David Vestin (2000), “Average Inflation Targeting,” working paper.
- Nobay, Robert P., and David A. Peel (1998), “Optimal Monetary Policy in a Model of Asymmetric Central Bank Preference,” FMG Discussion Paper No. 306.
- Oudiz, Gilles, and Jeffrey Sachs (1985), “International Policy Coordination in Dynamic Macroeconomic Models,” in William H. Buiter and Richard C. Marston, eds., *International Economic Policy Coordination*, Cambridge University Press, Cambridge.
- Persson, Torsten, and Guido Tabellini (2000), *Political Economics: Explaining Economic Policy*, MIT Press, Cambridge, MA.
- Ruge-Murcia, Francisco J. (1999), “Inflation Targeting under Asymmetric Preferences,” working paper.

- Svensson, Lars E.O. (1999), "Inflation Targeting as a Monetary Policy Rule," *Journal of Monetary Economics* 43, 607–654. Working-paper version with unabridged appendix available at www.princeton.edu/~svensson/.
- Svensson, Lars E.O. (2001a), "Independent Review of the Operation of Monetary Policy in New Zealand: Report to the Minister of Finance," www.princeton.edu/~svensson/.
- Svensson, Lars E.O. (2001b), "Price Stability as a Target for Monetary Policy: Defining and Maintaining Price Stability," in Deutsche Bundesbank, ed., *The Monetary Transmission Process: Recent Developments and Lessons for Europe*, Palgrave, New York, 60-102. (Also available as CEPR Discussion Paper No. 2196.)
- Svensson, Lars E.O. (2001c), "Requiem for Forecast-Based Instrument Rules," working paper.
- Svensson, Lars E.O. (2001d), "What is Wrong with Taylor Rules? Using Judgment in Monetary Policy through Targeting Rules," working paper.
- Svensson, Lars E.O., and Michael Woodford (1999), "Implementing Optimal Policy through Inflation-Forecast Targeting," working paper.
- Vickers, John (1998), "Inflation Targeting in Practice: the UK Experience," *Bank of England Quarterly Bulletin* 38, 368–375.
- Wallis, Kenneth F. (1999), "Asymmetric Density Forecasts of Inflation and the Bank of England's Fan Chart," *National Institute Economic Review* No.167 (January), 106-112.
- Woodford, Michael (1999a), "Commentary: How Should Monetary Policy Be Conducted in an Era of Price Stability?" in Federal Reserve Bank of Kansas City (1999), 277–316.
- Woodford, Michael (1999b), "Inflation Stabilization and Welfare," working paper.
- Woodford, Michael (1999c), "Optimal Monetary Policy Inertia," working paper.