

Optimal horizons for inflation targeting

Nicoletta Batini

and

Edward Nelson

Bank of England, Threadneedle Street, London EC2R 8AH.

The views expressed in this paper are those of the authors and do not necessarily reflect those of the Bank of England. We thank Claes Berg, Spencer Dale, Andrew Haldane, Peter Ireland, DeAnne Julius, Stephen Millard, Glenn Rudebusch, Frank Smets, Lars Svensson, Paul Tucker, Anders Vredin, Peter Westaway, Tony Yates, and two anonymous referees for useful comments on earlier versions of this paper. We also thank participants at Bank of England seminars, the Fifth International Conference of the Society for Computational Economics, and the Sveriges Riksbank Workshop on Inflation Targeting and Exchange Rate Fluctuations.

Issued by the Bank of England, London, EC2R 8AH, to which requests for individual copies should be addressed; envelopes should be marked for the attention of Publications Group (telephone 020-7601 4030). *Working Papers* are also available from the Bank's Internet site at <http://www.bankofengland.co.uk/wplist.htm>

Bank of England 2000
ISSN 1368-5562

Contents

Abstract	5
1. Introduction	7
2. Optimal horizons for inflation targeting: two definitions	8
3. Objective function and models	13
4. Optimal policy horizons (OPHs)	19
5. Optimal feedback horizons (OFHs)	24
6. Parameter uncertainty	28
7. Conclusions	30
References	31
Data appendix	34
Technical appendix — solution methods	37

Abstract

In this paper we investigate the problem of selecting an optimal horizon for inflation targeting in the United Kingdom. There are two key ways of thinking about an optimal horizon, so we look at optimal horizons for both of these interpretations. In addition, to see whether our results are robust in the face of model uncertainty, we compute optimal horizons for two different models with divergent structural and dynamic characteristics.

1 Introduction

Many central banks, including the European Central Bank, and those of Australia, Canada, Israel, New Zealand, Sweden and the United Kingdom, pursue an explicit inflation target. In practice, this usually involves ‘targeting’ the conditional forecast of inflation — the inflation rate expected to prevail in the future given presently available information — rather than current inflation.

A crucial issue is how to choose the horizon — the appropriate value of k when the operational target is expected inflation k periods ahead.

There are two key ways of thinking about an optimal horizon for inflation targeting, depending on the way that inflation targeting is modelled. If policy is represented, for instance, by a simple feedback rule on expected future inflation, one way is to think of it as the best horizon for which the authorities should form a forecast for inflation to use in the rule. If, instead, policy is represented by an optimal rule for the instrument, the optimal horizon can be thought of as the time at which inflation should be on target in the future when the authorities aim at minimising their loss function, and a shock occurs today. In what follows, we refer to the first kind of horizon as the ‘optimal feedback horizon’ and to the second kind as the ‘optimal policy horizon’.

This paper calculates optimal horizons for inflation targeting, using both definitions described above. Since the results may well hinge on the features of the model used for the calculations, the paper derives parallel results for two models: a vector autoregression (VAR) estimated on quarterly UK data; and a small-scale structural open-economy model based on Ball (1999), Batini and Haldane (1999), and McCallum and Nelson (1999a). A key difference between the two models is the importance that the second model assigns to forward-looking behaviour in spending and pricing decisions.

The paper is organised as follows. In Section 2 we discuss alternative definitions of horizons for inflation targeting. In Section 3 we describe the policy-makers’ objective function and the macroeconomic models that we employ. In Section 4 we compute optimal policy horizons for each model,

and discuss the results. In Section 5 we consider optimal feedback horizons, and Section 6 provides a brief exploration of sensitivity of the results to parameter uncertainty. Concluding remarks follow in Section 7.

2 Optimal horizons for inflation targeting: two definitions

In countries that pursue inflation targets, the formal wording of the mandate for the central bank usually provides that the inflation target be achieved each year.⁽¹⁾ For instance, the mandate for the European Central Bank, the most recently established central bank, states that: ‘[P]rice stability shall be defined as a year-on-year increase in the Harmonised Index of Consumer Prices for the euro area of below 2%...’ (ECB (1999), page 46). But central banks that have an inflation target need an operating strategy for achieving it. In practice, many do so by focusing on expected future inflation. The main reason for this is the existence of lags in the transmission of monetary policy to inflation.

The recognition of the existence of lags from monetary policy changes to inflation — and attempts to quantify these lags — have a long history. Jevons (1863), using UK data, concluded that: ‘An expansion of the currency occurs one or two years previous to a rise in prices...’. Milton Friedman (1972) similarly found, using postwar US M2 data, ‘a lag varying between eleven and thirty one months’ from monetary growth to inflation. These authors also recognised that prices reacted more slowly than output.

More recent empirical work, primarily using VAR analysis, has employed interest rate based monetary policy measures. Using US data, Christiano, Eichenbaum, and Evans (1996), for example, estimate that a monetary policy shock affects real GDP with a two-quarter lag, and the GDP deflator with a four-quarter lag.

It is possible that in more open economies, such as those currently pursuing an inflation target, the transmission lag is shorter because policy

⁽¹⁾ Bernanke *et al* (1999) provide a recent overview of several countries’ inflation targeting arrangements.

operates not only via the conventional ‘output gap channel’, but also via a potentially swift ‘exchange rate channel’. But even in economies that are open, the exchange rate channel may not be as effective, or as quantitatively important, as the output gap channel. Neither the pass-through of exchange rate changes to import prices, nor the propagation of import price changes into aggregate price level changes, can be taken for granted.⁽²⁾ In the United Kingdom, for example, the substantial exchange rate depreciations in 1982 and 1992 failed to produce appreciable increases in the inflation rate (measured by the RPIX index). In other words, even in economies that are more open than the United States, the time it takes for monetary policy to have its main impact on nominal demand may still be lengthy.

In the presence of transmission lags, returning inflation to target immediately after a shock may involve considerable costs. This is because immediately offsetting the inflationary consequences of a shock may require large movements in the policy instrument, resulting in unduly large output losses. One obvious way to avoid this is to try to anticipate inflationary events and react to them pre-emptively in a more gradual fashion. Acting before it is too late permits central banks to minimise these losses by reducing the extent to which the instrument has to be moved in the short run in order to control inflation.

In the literature on inflation targets, this forward-looking approach to policy has been represented in two ways.

Rudebusch and Svensson (1999) argue that real-world inflation targeting can be approximated by viewing the central bank as carrying out an optimisation exercise, where the welfare function penalises inflation departures from a target, and policy is thus set according to the ensuing optimal rule.⁽³⁾ In their words (1999, page 204): ‘In examining policy rules that are consistent with inflation targeting, we consider two broad class of rules: instrument rules and targeting rules... A targeting rule may be closer to the actual decision framework under inflation targeting. It is

⁽²⁾ The estimates of our VAR for the United Kingdom (see Section 3 below) support caution about the empirical importance of the exchange rate channel for output and inflation.

⁽³⁾ See also Svensson (1997, 1999).

represented by the assignment of a loss function over deviations of a goal variable from a target level...’.

Rudebusch and Svensson’s definition has strengths and weaknesses. One strength is that it reflects a comprehensive approach to policy — a ‘look-at-everything’ approach. It also implies that information is processed in the most efficient way — which is the way central banks often visualise their policy. However, this definition implies that policy is set in a complex manner that the public may find difficult to understand, particularly when there is no consensus on either the model or the objective function.

A second definition of inflation targeting is used in work by Batini and Haldane (1999), McCallum and Nelson (1999b), and others. This views targeting expected future inflation simply as setting the policy instrument in response to deviations of future inflation from target. In other words, it defines targeting as the use of a policy rule such as:

$$R_t = \text{constant} + \psi_p [E_t \pi_{t+k} - \pi^*] \quad (1)$$

where R_t is the short-term nominal interest rate, $E_t \pi_{t+k}$ is the period t forecast of inflation in $t+k$, π^* is the inflation target, and $\psi_p > 1$.

Strengths of this approach are that it does not necessarily require agreement on every aspect of model specification other than in the construction of the forecast (for which there is no need for a full model). This is because (1) has the property that it will rein in deviations of inflation from target in a variety of models, subject to weak conditions which will be met by most standard model specifications.⁽⁴⁾ Consequently, (1) is a rule that policy-makers could agree would control inflation, even if they had no consensus on the appropriate objective function or model of the economy. Furthermore, (1) is consistent, for $k > 0$, with the inflation forecast at some specific horizon being a key input into policy-makers’ decisions. By contrast, in most cases a standard optimisation exercise would not give special weight to the inflation

⁽⁴⁾ The condition is that the long-run response of the inflation rate to monetary policy tightenings is negative.

forecast at a particular horizon, especially as that forecast may not even appear in the first-order conditions for optimality.

The principal weakness of this approach is that it is based on a simple rule that uses information less efficiently than an optimal rule.

The above considerations suggest that, in either case, inflation-targeting central banks need to decide how forward-looking they should be in order to bypass the transmission lags. That is, when setting policy, they must choose an ‘optimal horizon’ over which to pursue the goal of price stability. Since the concept of horizon means different things under the above two interpretations of inflation targeting, we provide two operational definitions of ‘optimal horizon’ to investigate this issue: the ‘optimal policy horizon’ and the ‘optimal feedback horizon’.

We define the *optimal policy horizon* (‘OPH’ hereafter) as the number of periods it takes for inflation to settle on target after a shock, under the optimal rule for the instrument. This is in line with Rudebusch and Svensson’s interpretation of inflation targeting — an optimisation exercise that penalises deviations of inflation from target and of output from potential. (More specifically, this corresponds to what Svensson (1997, 1999) calls ‘flexible inflation targeting’.) We call this horizon the *optimal policy horizon* because of its intimate connection with the optimal policy rule; but an equally valid, and perhaps more descriptive, label would be the *optimal stabilisation horizon*.

There are two things to note about this definition. First, we treat the underlying optimisation exercise undertaken by the central bank as one subject to *commitment*. Svensson (eg 1999) instead generally views it as optimisation subject to *discretion*, but Woodford (1999) demonstrates that discretion has drawbacks when (as in one of our models below) the model’s structural equations contain forward-looking components.⁽⁵⁾ More recently, Svensson and Woodford (1999) have proposed several candidate modifications to standard discretionary optimisation, which — by connecting current and past policy actions — reduce the problems associated with discretionary behaviour. Future work could compare our

⁽⁵⁾ These problems arise from the fact that discretionary optimisation generally leads to predictable future deviations by policy-makers from their currently announced plans.

results using commitment with results based on Svensson and Woodford's modified discretionary framework.⁽⁶⁾ Second, the optimal policy horizon in our definition is an *output* of the optimisation exercise, rather than a literal constraint on the optimisation problem. See Smets (1999) for an alternative approach, in which the requirement that $E_t \pi_{t+k} = 0$ (for a specified $k \geq 0$) is a constraint on the policy-makers' optimisation problem (ie an *input* in the optimisation process).

Turning to our second concept of horizon, we define the *optimal feedback horizon* (OFH hereafter) as the best period in the future for which the authorities should form the inflation forecast that enters their policy rule. The OFH is an optimal *response* horizon — the best horizon to focus on when designing a simple, inflation forecast based rule. The OFH is therefore the k associated with the minimisation of the policy-makers' loss function, when policy follows a simple rule such as **(1)**.

In the next section, we define the policy-makers' preferences and describe the models that we use to derive optimal horizons under each of the above definitions.

⁽⁶⁾ As discussed further in the technical appendix, the specific commitment policy that we assume is one in which the central bank optimises, once and for all, subject to the structural model of the economy. Solutions for the endogenous variables are then implied by the structural equations plus the policy-makers' time-invariant first-order conditions for optimality. This procedure follows King and Wolman (1999) and corresponds closely to what Woodford (1999) calls a commitment policy with a 'timeless perspective'.

3 Objective function and models

Objective function

We assume that policy-makers wish to minimise deviations of inflation from target and departures of output from potential. We also assume that policy-makers dislike instrument volatility. For computational convenience, these preferences are represented by a quadratic loss function. In the optimisation exercises used to derive optimal policy horizons, this is the function that is being minimised. And when we derive optimal feedback horizons by comparing the performance of rules like (1) for various ks , this loss function is used to compute welfare losses in all experiments. Formally, the loss function is given by:

$$L_t = E_t \sum_{j=0}^{\infty} \beta^j [\lambda_{\pi} (4^* \pi_{t+j} - 4^* \pi_{t+j}^T)^2 + \lambda_y (y_{t+j} - y_{t+j}^T)^2 + \lambda_{\Delta R} (4^* \Delta R_{t+j})^2] \quad (2)$$

where β is the discount factor, $4^* \pi_t$ is annualised quarterly inflation, π_t^T is the inflation target, y_t is log output, y_t^T is log capacity output, and where λ_{π} , λ_y and $\lambda_{\Delta R}$ denote the weights assigned to inflation deviations from target, output deviations from potential, and volatility in the first difference of the nominal interest rate, respectively.

We set $\beta = 0.99$, $\lambda_{\pi} = 1$, $\lambda_y = 1$ and $\lambda_{\Delta R} = 0.5$, so that inflation and gap variability are penalised equally. The interest rate volatility term, which rules out extremely large movements of the instrument in response to shocks, receives a penalty half that of the other terms. These weights are similar to those used in Rudebusch and Svensson (1999).

Models

To explore the optimal horizon issue, we look at two models: a vector autoregression (VAR) estimated on quarterly UK data; and a calibrated, forward-looking small structural model. These models are described below.

A VAR model

Our first model is a one-lag VAR with a linear trend, estimated over the period 1981 Q1–1998 Q1.⁽⁷⁾ There are four endogenous variables in the VAR: log output (y_t); the deviation of annual RPIX inflation from the inflation target (π_t^{DEV}); the log-change in the nominal effective exchange rate (Δe_t); and the nominal interest rate (interbank lending rate), measured as an annualised fraction ($4 * R_t$).⁽⁸⁾

As Rudebusch (1998) observes, the interest rate equation in a VAR has a structural interpretation as a monetary policy reaction function. As he also notes, however, there is a danger that policy regime shifts may produce non-constant parameter estimates. In our sample period, there have been two major breaks in the United Kingdom's monetary policy regime: the membership of the Exchange Rate Mechanism (ERM) from 1990 to 1992; and, following its exit from the ERM, the United Kingdom's adoption of an inflation targeting regime from 1992 Q4. If we allow both the slopes and the intercepts of the interest rate equation in our VAR to vary across these regimes, the restriction of no structural change is rejected ($F(12,51) = 2.51$ [p value = 0.01]). However, the convenient restriction that the parameter non-constancy is isolated to the equation's intercepts is not rejected ($F(10,51) = 1.69$ [p value = 0.11]). Hence we proceed under that assumption, augmenting each equation of the VAR with two regime-shift intercept dummies, $DERM_t$ and $D924_t$.⁽⁹⁾

⁽⁷⁾ The linear trend is included to detrend the output variable.

⁽⁸⁾ The data appendix provides a detailed description of these series and their time series properties. In our work in Sections 4 and 5 with this VAR, the estimated y_t equation will serve as the output gap equation, and the π_t^{DEV} equation, converted to quarterly units, will be used as the inflation equation.

⁽⁹⁾ These variables take the value 1.0 for 1990 Q4–1992 Q3 and 1992 Q4 onwards, respectively.

Estimates of this system are reported in Table A. The lag length of one quarter is not rejected by a χ^2 test against the alternative of a two-lag VAR (p value = 0.10).⁽¹⁰⁾

Several features of the dynamic properties of the system emerge. First, consider the output equation. Although the estimated coefficient on 4^*R_{t-1} in the y_t regression is quite small (-0.0908), suggesting a minor initial impact of monetary policy on real demand, the estimated coefficient on lagged y_t in the equation is large, implying that the *long-run* response to 4^*R_t is much greater (-1.3373%). The exchange rate term in the y_t equation, on the other hand, is small and insignificant.

Second, in line with economic intuition, output has a significant positive coefficient in the inflation equation; and so interest rates have a negative effect on inflation, apparently via a conventional output gap channel.

Finally, the estimated equation for 4^*R_t looks like a Taylor (1993)-type policy rule, with the interest rate responding positively both to lagged output (long-run coefficient = 0.10) and to inflation (long-run coefficient = 0.91, not significantly below unity).⁽¹¹⁾ In addition, the coefficient on the lagged dependent variable (0.58) suggests a strong tendency by policy-makers to smooth interest rates; and the coefficient on the exchange rate change (equal to -0.042) suggests that the interest rate responds positively to exchange rate depreciations, as one would expect.⁽¹²⁾

⁽¹⁰⁾ Due to the large lagged dependent variable coefficients, the one-lag specification is consistent with a long lag between monetary policy changes and their peak effect on output or inflation.

⁽¹¹⁾ Estimating UK policy rules over sub-samples of the 1972-1997 period, Nelson (2000) finds point estimates of the long-run response to inflation below unity, with the exception of the post-1992 period. Nelson also finds output gap responses of the same magnitude as that estimated in Table A.

⁽¹²⁾ Note that no variables enter the exchange rate equation itself significantly, indicating that a random walk model for e_t is not rejected.

Table A: VAR estimates				
Sample period: 1981 Q1 – 1998 Q1				
	Y_t	π_t^{DEV}	Δe_t	4^*R_t
y_{t-1}	0.9321 (0.0464)	0.1250 (0.0575)	-0.1583 (0.3273)	0.0444 (0.0782)
π_{t-1}^{DEV}	-0.0926 (0.0731)	0.8559 (0.0905)	0.3544 (0.5150)	0.3867 (0.1231)
Δe_{t-1}	-0.0029 (0.0175)	-0.0023 (0.0217)	0.1526 (0.1234)	-0.0416 (0.0295)
4^*R_{t-1}	-0.0908 (0.0529)	0.0336 (0.0655)	-0.0438 (0.3731)	0.5759 (0.0892)
Constant	0.7821 (0.5177)	-1.4072 (0.6409)	1.7167 (3.6487)	-0.4786 (0.8722)
Time trend	0.00046 (0.00040)	-0.00091 (0.00049)	0.002477 (0.00280)	0.000287 (0.00067)
$DERM_t$	-0.00891 (0.00563)	0.0091 (0.0070)	-0.0509 (0.0397)	-0.0300 (0.0095)
$D924_t$	-0.0091 (0.0073)	0.0108 (0.0090)	-0.0573 (0.0515)	-0.0422 (0.0123)
R^2	0.9984	0.8880	0.1017	0.9259
S.E.E.	0.0052	0.0064	0.0364	0.0087
S.D. dep. var.	0.1212	0.0181	0.0364	0.0303

Note: Standard errors in parentheses.

In this paper, we subject the VAR to hypothetical policy rules different from the estimated one. This requires us to identify the VAR model's responses to shocks. We do this by means of a Cholesky decomposition, where the equation disturbances are assumed to follow the causal ordering (output innovation → inflation innovation → exchange rate innovation → interest rate innovation). Under this identification scheme, no variable

other than the interest rate responds contemporaneously to the monetary policy shock.

A forward-looking structural model

The second model that we consider is a small-scale, forward-looking open-economy model that incorporates elements of Ball (1999), Batini and Haldane (1999), and McCallum and Nelson (1999a). The structural equations of the model are:

$$y_t = E_t y_{t+1} - \sigma(R_t - E_t \pi_{t+1}) + \delta \tilde{q}_{t-1} + e_{y_t} \quad (3)$$

$$\pi_t = \alpha \pi_{t-1} + (1 - \alpha) E_t \pi_{t+1} + \phi_y y_{t-1} + \phi_q \Delta \tilde{q}_{t-1} + e_{\pi_t} \quad (4)$$

$$E_t q_{t+1} = q_t + R_t - E_t \pi_{t+1} + \kappa_t \quad (5)$$

where y_t is log output, R_t is the nominal interest rate (again, a quarterly fraction), π_t is quarterly inflation, q_t is the log real exchange rate (measured so that a rise is a depreciation), and $\tilde{q}_t = 1/4 \sum_{j=0}^3 q_{t-j}$ is a four-quarter moving average of q_t . These variables are all expressed relative to steady-state values. e_{y_t} , e_{π_t} , and κ_t are exogenous IS, Phillips curve, and uncovered interest parity (UIP) shocks, respectively.

Equation (3) is the model's IS relationship, giving y_t as a function of its expected future value, the real interest rate, and lags of the real exchange rate. Apart from the term in \tilde{q}_t , this equation corresponds to the optimisation-based IS function in McCallum and Nelson (1999a), and we choose parameter values based on their estimates: $\sigma = 0.2$ and an AR(1) process for e_{y_t} with coefficient 0.3 and 1% innovation standard deviation. Our choice of $\delta = 0.05$ then produces the same ratio of interest rate to exchange rate coefficients in the IS curve as used in Batini and Haldane (1999).

Equation (4) is a quarterly version of Ball's (1999) open-economy Phillips curve, modified to allow for some forward-looking behaviour. While Ball has lagged inflation appearing on the right-hand side of (4) with coefficient 1.0, we replace this with the mixed backward-forward looking

term $\alpha\pi_{t-1} + (1 - \alpha)E_t\pi_{t+1}$, and calibrate α to 0.8, close to estimates in Fuhrer (1997) and Rudebusch (1999). We calibrate the coefficient ϕ_y to 0.1, the quarterly counterpart of Ball’s choice. We choose $\phi_q = 0.025$; this is considerably lower than Ball’s 0.10, but a relatively conservative value of ϕ_q seems prudent in light of the failure of the VAR to pick up *any* effect of depreciation on inflation. We assume that e_{π_t} is white noise with standard deviation 1%.

The exchange rate enters both the IS and Phillips curve relationships in a backward-looking manner, as a lagged four-period average. A more forward-looking specification of the model’s open-economy elements would put Δq_t and $E_t\Delta q_{t+1}$ in (4).⁽¹³⁾ We found, however, that this scheme produced an implausibly tight and mechanical relationship between exchange rate changes and inflation.⁽¹⁴⁾ Thus we have followed Ball (1999) by allowing q_t to enter only with lags; this might be rationalised by ‘gradual pass-through’ of exchange rate changes to export and import prices, which might be realistic for the United Kingdom (Bank of England (1999)).⁽¹⁵⁾ While q_t enters (3) and (4) only in a backward-looking manner, this is compensated by the fact that the exchange rate itself is a highly forward-looking variable, as equation (5) indicates. The shock term κ_t that produces deviations from strict UIP in (5) is assumed to be AR(1) with coefficient 0.753 and innovation standard deviation 0.92%; these choices are based on our estimates of this process using quarterly UK data. The shocks in (3)-(5) are assumed to be mutually uncorrelated.

Note that in this model — hereafter referred to as the FLSM (forward-looking structural model) — monetary policy has some effect on contemporaneous inflation due to the fact that π_t responds to $E_t\pi_{t+1}$ in (4). This effect may be small in relation to the long-run effect of policy on inflation, but the presence of at least *some* effect means that the model

⁽¹³⁾ It might also put q_t and E_tq_{t+1} into equation (3).

⁽¹⁴⁾ In Bank of England (1999), it is shown that the inclusion of Δq_t and $E_t\Delta q_{t+1}$ in Batini and Haldane’s (1999) Phillips curve makes exchange rate movements the dominant determinant of inflation, and virtually removes any inflation persistence from the model.

⁽¹⁵⁾ In his annual model, Ball specifies Δq_t as entering the inflation equation with a one-period lag; in our quarterly model, we approximate this by having the prior year’s average of Δq_t entering equation (4).

does not have the same property as that in Svensson (1997), where there is no scope for current policy changes to affect inflation today.

4 Optimal policy horizons (OPHs)

In line with our discussion in Section 2, we define the *optimal policy horizon* (OPH) as the time at which it is least costly, for a given loss function, to bring inflation back to target after a shock. More intuitively, the OPH is the horizon-analogue of the optimal speed of disinflation — the optimal time required for the dissipation of a shock. Operationally, the OPH is given by the number of periods after a shock when inflation is back on target under an optimal rule.

Below, we derive OPHs for both models. In this respect, an important question is how to interpret the idea of being ‘on target’. Since, in these models, inflation tends to fluctuate around target before settling definitely on a particular number in the wake of a shock, a point target (eg 2½%) is not very meaningful. Instead, we consider OPHs as referring to target *ranges* — so the OPH is the time when inflation returns to a specified band around the target. This is not an argument for target ranges rather than point targets, but a device to make model experiments useful.

We use two operational definitions of an OPH: an absolute and a relative horizon concept. The first definition interprets an OPH as the number of periods ahead, k , at which inflation has returned permanently to within a target range of ± 0.1 percentage points, following a shock today. The second definition is based on what fraction of a shock’s effect policy has succeeded in eliminating. Specifically, it interprets the optimal policy horizon as the number of periods ahead, k , at which 90% of the peak effect of the shock on inflation has been extinguished. We denote OPHs under the absolute criterion by ‘ k_A^* ’, and OPHs under the relative criterion by ‘ k_R^* ’. Since, typically, both k_A^* and k_R^* will vary according to the nature of the economic shock, we compute OPHs under the two criteria for different kinds of shocks. In our discussion, we focus primarily on k_R^* . This is because, in these two linear models, the k_A^* value depends on the size of the initial shock, but k_R^* — defined in terms of a percentage of the effect of the shock — does not.

After solving each model of Section 3 under its optimal rule, we computed impulse responses for inflation.⁽¹⁶⁾ Charts 1(a) to 1(c) show the optimal paths of inflation associated with each model in the face of 1% shocks to aggregate demand, aggregate supply and the exchange rate, respectively.

Chart 1(a)

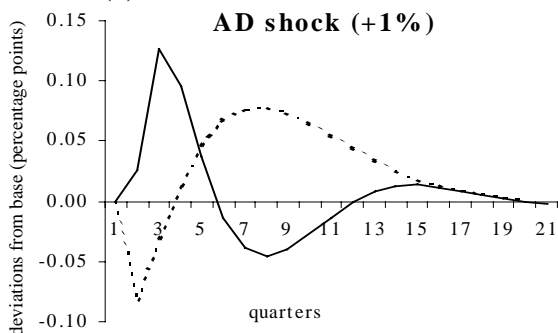
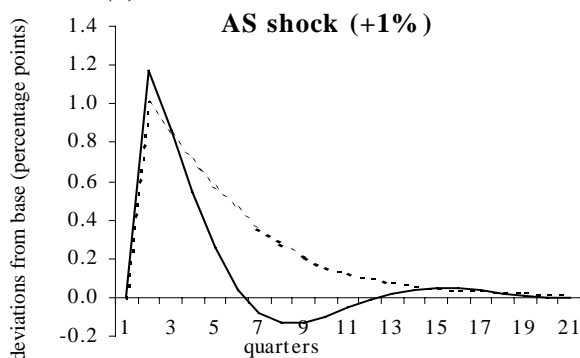


Chart 1(b)



⁽¹⁶⁾ We follow King and Wolman (1999) by augmenting the model's structural equations with the policy-makers' first-order conditions for optimality, and solving the resulting system of expectational difference equations. The Lagrange multipliers for the policy-makers' problem form part of the state vector in this commitment solution. The technical appendix provides details.

Chart 1(c)

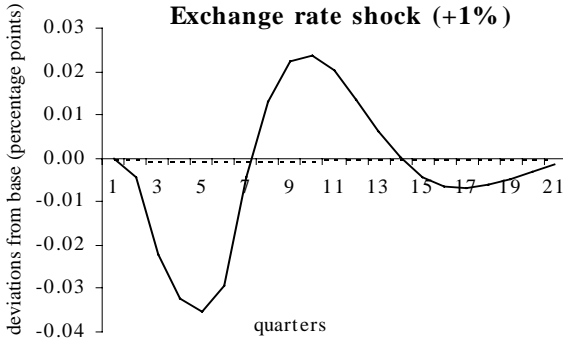


Table B gives OPHs for each model under the two definitions (k_A^* and k_R^*). We discuss these results in turn as we comment on the impulse response function of inflation for each shock.⁽¹⁷⁾

⁽¹⁷⁾ Since the OPHs are obtained from impulse response functions, the OPHs for models in previous papers, Rudebusch and Svensson (1999) for example, can be deduced provided the papers include plots of impulse responses under optimal policy (eg Rudebusch and Svensson's Figs. 5.3 and 5.4).

Table B: Optimal policy horizons (OPHs)				
	k_A^*		k_R^*	
<i>Shock</i>	VAR	FLSM	VAR	FLSM
AD	14	9	17	14
AS	18	16	10	8
Exchange rate	0	5	16	19

Aggregate demand shock

Chart 1(a) shows the impulse response functions of inflation for the two models described in Section 3, in the wake of a temporary positive 1% shock to aggregate demand.

In the VAR (the dashed line in Charts 1(a)-1(c)), an AD shock has no effect on inflation until period $t+1$; and by that time the policy-maker (whose actions in t affect inflation in $t+1$) is in a position to provide offsetting pressure on inflation through interest rates. Consequently, the initial net response of inflation to the positive demand shock is negative in Chart 1(a). Within a couple of periods, inflation has returned to zero, but inflation then overshoots for several periods — leading to a long OPH (defined by k_R^*) of 17 quarters. This is partly due to the presence in the loss function of a penalty for volatility in the policy instrument. If this were absent, the interest rate would be raised much more sharply in response to the shock, restraining the response of inflation, and leading to a lower OPH ($k_R = 9$).

In the FLSM (the solid line), the presence of forward-looking elements, together with a contemporaneous effect of the real interest rate on output, reduces the length of the policy transmission lag and increases the capacity of policy to offset shocks. The OPH under the relative criterion is less than that for the VAR, and suggests that it is optimal to carry out disinflation within fourteen quarters of the shock (Table B).⁽¹⁸⁾

Aggregate supply shock

Chart 1(b) shows the impulse responses of inflation for the two models, in the wake of a temporary positive 1% shock to aggregate supply. For this disturbance, both models display quite smooth inflation dynamics. In the VAR case, monetary policy cannot affect inflation in the period of the shock, so the 1% supply shock raises inflation by a full 1% in the first period. Strong inflation persistence inhibits policy-makers' ability to remove the effect of the shock without unduly large output costs; it takes two and a half years for this shock to be reversed under the optimal rule ($k_R^* = 10$).

For the FLSM, the response of inflation to the supply shock is visually close to that in the VAR. This reflects some similarities between the Phillips curve in this model and the VAR's inflation equation, notably the fact that lagged inflation enters with a sizable coefficient. As with the AD shock, the OPH is shorter in the FLSM than in the VAR ($k_R^* = 8$ instead of 10) — partly due to forward-looking dynamics in the IS curve (**I**), which make a more rapid disinflation optimal in this case.

Exchange rate shock

Chart 1(c) shows the impulse responses of inflation, in the face of a temporary positive 1% shock to the exchange rate equation of each model — constructed such that it would lead to a 1% appreciation in both cases, *ceteris paribus*.

⁽¹⁸⁾ If the IS shock were white noise in this model, the OPH would be five quarters shorter ($k_R = 9$).

The estimated VAR contains an exchange rate equation that is essentially detached from the rest of the model, in the sense that there is virtually no feedback to other variables from the exchange rate. This explains the very flat response of inflation to such a shock in Chart 1(c). Under the relative criterion, the OPH is long ($k_R^* = 16$) because it demands that 90% of an already negligible inflation response be eroded. In this case, the OPH as measured by the absolute criterion (k_A^*) provides useful auxiliary information: it is zero quarters, and would be zero even if the initial shock were 10% or larger, rather than 1%.

The appearance of the exchange rate in both the IS and the AS functions implies that the exchange rate has an important role in the FLSM. In response to the UIP shock, the exchange rate appreciates on impact and (partly because of the shock's persistence) this appreciation is not completely reversed for more than a year. The response of inflation is also protracted, as inflation depends on long lags of the exchange rate (both via the output gap channel and via the 'direct' exchange rate channel). The optimal policy response is to cut interest rates, which reduces the extent of the appreciation. This also stimulates demand, offsetting some of the contractionary effects of the appreciation, at the cost of creating a positive output gap and a rise in inflation above target for a few quarters. Overall, the combined effect on inflation of the UIP shock and the policy reaction is quite small — inflation is never more than 0.05 percentage points from target. This turns the optimal horizon measured by k_R^* almost into a point target, so the OPH (at 19 quarters) exceeds the OPHs for the other two shocks.

5 Optimal feedback horizons (OFHs)

The previous section obtained optimal horizons assuming that the policy-makers followed a complex optimal rule — a function of the entire state vector. Suppose instead that the policy-makers operate via a simple rule that involves changing the policy instrument in response to deviations of expected inflation from its target value, as in equation (1).

By suitable choice of the feedback horizon, this rule can be designed so as to incorporate monetary transmission lags. In particular, in the case where lags hinder control of current inflation, the date of the inflation forecast in the rule can be chosen so that inflation at that date is indeed affected by monetary policy.

When inflation targeting is implemented through rules like **(1)**, the best k period ahead forecast of inflation will be the one that minimises the costs of inflation control according to loss function **(2)**. As explained in Section 2 above, we define this horizon as the *optimal feedback horizon* (OFH). Two things are worth noting here. First, we consider the OFH to be the k that minimises **(2)** when the feedback coefficient in **(1)** is itself optimally chosen. That is, the choice of the optimal k is conditioned on ψ_p (and possibly an interest rate smoothing coefficient) being optimal. Second, in contrast with the previous section, the optimal horizon is not a concept that can be bracketed by a range. Rather, it can only be a specific period (ie the best k at which to form the forecast of inflation that enters the rule). The OPHs and OFHs are thus very distinct concepts: the first is a metric associated with an optimal rule, the second is an optimised parameter of a simple rule.

In this section we derive OFHs for our two models. We generalise rule **(1)** by including an interest rate smoothing term (a coefficient on the lag of the nominal interest rate). This gives equation **(6)** below, where the degree of interest rate smoothing is governed by the parameter $\rho_R \in [0,1]$.

$$R_t = \rho_R R_{t-1} + \psi_p (E_t \pi_{t+k} - \pi^*) + \text{constant} \quad \mathbf{(6)}$$

Table C summarises the results on OFHs. To obtain them, we closed the models with rule **(6)**, where the parameters (ψ_p , ρ_R , and k) were chosen optimally by minimising loss function **(2)**, evaluated using analytical formulae for the model moments. We contemplated values of k of 0, 1, ..., 15 — horizons up to four years ahead.

Table C: Optimal feedback horizons (OFHs)

<i>Model</i>	ψ_p^*	ρ_R^*	k^* (OFH)
VAR	35	0.85	15
FLSM	1.24	0.98	2

Table C indicates that both the VAR and the FLSM favour a positive feedback horizon. This is in line with Batini and Haldane (1999), who find that responding to expected future rather than current inflation is beneficial when there are lags in the effect of monetary policy. Comparing results on OFH for our two models, it appears as if forward-looking behaviour ‘brings forward’ the optimal feedback horizon; forward-looking agents take into account current and prospective interest rate decisions in their spending and pricing decisions, which reduces the transmission lag (but does not eliminate it, since inflation still has an inert component). Specifically, the OFH is $k = 2$ for the FLSM, compared with 15 for the VAR.

The optimal value of ψ_p is very large in the case of the VAR, whose backward-looking, data-based nature might instead have led us to expect instrument instability when policy becomes too aggressive. Such a large feedback parameter is consistent with relatively low interest rate variability because, in equilibrium, the standard deviation of the variable in the rule ($E_t \pi_{t+15}$) is very low (0.02% annualised).⁽¹⁹⁾ The short-run feedback parameter is much smaller for the FLSM. In a forward-looking model, agents’ actions take the expected *long-run* policy response into account. An integral component of policy in the FLSM is therefore high interest

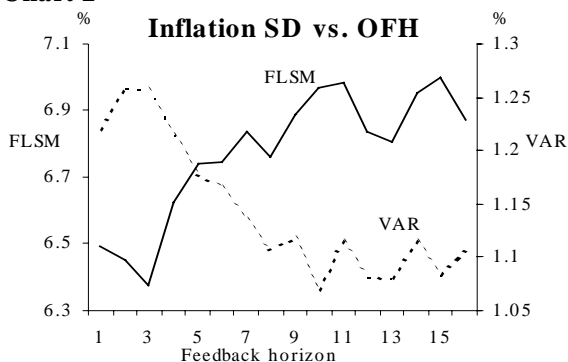
⁽¹⁹⁾ Evaluating the VAR model’s properties where the feedback parameter takes such large values may take the model into ranges where it ceases to be a useful approximation for policy analysis. This is a caution when interpreting our results.

rate smoothing, which, for reasons discussed in Rotemberg and Woodford (1999), exerts a restraining effect on inflation via the forward-looking IS function.

To get an impression of the inflation cost of responding to the ‘wrong’ horizon, Chart 2 plots (for both models) the standard deviation of inflation against the horizon included in rule (6).⁽²⁰⁾ For the VAR, the inflation outcome seems equally good for all reasonably long horizons (six quarters or longer); $k = 15$ is optimal in Table C largely because of the lower interest rate and gap volatility associated with that long horizon. By contrast, for the FLSM, there is a much sharper increase in inflation volatility from using horizons other than the OFH (both shorter and longer). The common message from both models is that inflation control is sacrificed if the chosen horizon is too short.

The optimised simple rules in Table C go a long way to approaching the minimum of the loss function achieved by the optimal rules. For the VAR, the optimal-rule loss function value is 0.0220 compared with a loss function value of 0.0224 for the OFH rule, a difference of only 1.8%. Similarly, for the FLSM the loss from optimal policy is 0.5945, versus 0.6080 for the optimised simple rule, a difference of 2.2%.

Chart 2



⁽²⁰⁾ In each case, the coefficients in (6) were re-optimised for the fixed k .

For this reason, we found little welfare gain from amending rule (6) to respond to *multiple* horizons. We undertook two experiments in this regard. First, we examined optimised simple rules that responded to annual inflation one, two, three, or four years ahead (eg responding to a moving average of π_t , $E_t\pi_{t+1}$, $E_t\pi_{t+2}$, and $E_t\pi_{t+3}$, in the case of the one-year rule). But we found no welfare improvement for either model from any of these annual horizons, relative to the optimised simple rules in Table C.

Second, we looked at simple rules that responded to two distinct horizons. Specifically, for the FLSM with its low OFH, we examined rules that responded to horizon 3, 4, 5, or 6 in addition to horizon 2; and for the VAR, we examined rules that responded to horizon 11, 12, 13 or 14 in addition to horizon 15. Response coefficients were re-optimised, but again we found no significant gain. For the FLSM, responding to two and three quarter ahead inflation gave a welfare improvement relative to the rule that responded to two quarters alone, but the gain was less than 0.05%. For the VAR, we found that a single horizon dominated all the multiple-horizon alternatives.

6 Parameter uncertainty

As our comparison of results from the VAR and FLSM shows, the specification of model structure is an important factor determining the OFH. Another issue is how sensitive the OFH is to the choice of parameter values for a *given* model structure. In this section we briefly explore this issue for the FLSM.

We focus on the most contentious parameter in the FLSM — namely, σ in equation (2).⁽²¹⁾ In a closed-economy model, σ can be interpreted as the intertemporal elasticity of substitution for consumption. Values of σ in policy rule studies vary drastically. For example, the studies of Estrella and Fuhrer (1998) and McCallum and Nelson (1999a) (the source for our

⁽²¹⁾ Of course, the parameter α in (4) varies across studies but (unlike σ) it is at least bounded in $[0,1]$.

calibration) suggest quite low values — $\sigma = 0.2$ in the latter — while at the opposite extreme, Rotemberg and Woodford (1999) find $\sigma = 6.0$.

How then does the OFH change if we adopt a higher value of σ ? Table D examines this issue by presenting OFHs for a version of the FLSM that uses a value of $\sigma = 1.0$ — popular in the literature because it is associated with logarithmic preferences.

Table D: OFHs under different assumptions about σ			
	ψ_p^*	ρ_R^*	k^* (OFH)
FLSM, $\sigma = 0.2$	1.24	0.98	2
FLSM, $\sigma = 1.0$	5.02	0.00	8

Intuitively, a higher value of σ could shorten the OFH by making spending more sensitive to current monetary policy changes. In practice, raising σ to 1.0 lengthens the OFH from 2 to 8. The reason is that a longer horizon means responding to a lower variance variable — a far-ahead forecast of inflation. The resulting rule features low interest rate volatility and consequent low output gap volatility, producing welfare gains.⁽²²⁾ When σ is 0.2 (our baseline parameterisation), a given amount of interest rate volatility has less of an effect on gap volatility than when $\sigma = 1.0$. Responding to inflation eight periods ahead reduces interest rate volatility, but the gains in terms of reduced gap volatility are not as substantial, and policy sacrifices too much control over inflation. A shorter horizon then becomes optimal.

⁽²²⁾ The standard deviations of $4^*\Delta R_t$ and the output gap in the $\sigma = 1.0$ model are 0.8% and 3.2% when the optimised rule in Table D is used. If, on the other hand, the coefficients in that rule were used but a horizon of 2 were put in the rule instead of 8, the standard deviations of $4^*\Delta R_t$ and the gap would rise to 8.8% and 3.9% respectively.

Thus, a policy-maker that used the baseline version of our model, but was uncertain about the true value of σ (aware that it could be too low in the baseline version), might want to respond to a longer inflation horizon than it would if it knew that $\sigma = 0.2$ with certainty.

7 Conclusions

In this paper we investigated the problem of selecting an optimal horizon for inflation targeting. For this purpose, we provided two operational definitions of ‘optimal horizon’, corresponding to two different interpretations of how inflation targeting works in practice. Results were obtained from two small-scale models with divergent dynamic properties: a VAR and a forward-looking structural model.

For the *optimal policy horizon*, a definition based on the assumption that inflation targeting involves an optimal policy (obtained in a ‘flexible inflation targeting’ framework that penalises inflation, interest rate, and output gap volatility), we found that it is optimal to remove the effects of the various shocks considered here over a period of 8 to 19 quarters. For the *optimal feedback horizon*, a definition based on the view that inflation targeting is well approximated by a simple forward-looking policy rule, we found that the best horizon to focus on depends crucially upon the degree of forward-looking behaviour in the economy. With no forward-looking behaviour (the VAR), long feedback horizons — responding to forecasts of far-ahead inflation — are desirable. With at least some forward-looking behaviour (the FLSM), the appropriate feedback horizon is much shorter. Even in this case, however, it appears sub-optimal to feed back on current or next-quarter inflation.

To summarise, our analysis supports the view that inflation targeting in practice should be designed so that the target is achieved over the medium term. In other words, central banks wishing to act optimally should not attempt to neutralise inflationary shocks immediately, but instead should respond gradually to those shocks. This becomes particularly important when the economy adjusts sluggishly to economic shocks. Further research on optimal horizons could investigate the robustness of our results to different models.

References

- Ball, L (1999)**, ‘Policy rules for open economies’, in Taylor, J B (ed), *Monetary Policy Rules*, University of Chicago Press, pages 127–44.
- Bank of England (1999)**, *Economic models at the Bank of England*, Bank of England.
- Batini, N and Haldane, A G (1999)**, ‘Forward-looking rules for monetary policy’, in Taylor, J B (ed), *Monetary Policy Rules*, University of Chicago Press, pages 157–92.
- Bernanke, B, Laubach, T, Mishkin, F S and Posen, A S (1999)**, *Inflation Targeting: Lessons from the International Experience*, Princeton University Press.
- Christiano, L J , Eichenbaum, M and Evans, C L (1996)**, ‘The effects of monetary policy shocks: some evidence from the flow of funds’, *Review of Economics and Statistics*, Vol 78(1), pages 16–34.
- Estrella, A and Fuhrer, J C (1998)**, ‘Dynamic inconsistencies: counterfactual implications of a class of rational expectations models’, *mimeo*, Federal Reserve Bank of Boston.
- European Central Bank (1999)**, *Monthly Bulletin*, January.
- Friedman, M (1972)**, ‘Have monetary policies failed?’, *American Economic Review*, Vol 62(2), pages 11–18.
- Fuhrer, J C (1997)**, ‘The (un)importance of forward-looking behaviour in price specifications’, *Journal of Money, Credit, and Banking*, Vol 29(3), pages 338–50.
- HM Treasury (1980)**, ‘The stability of the income velocity of circulation of money supply’, in *Third Report from the Treasury and Civil Service Committee, Session 1980-81: Monetary Policy, Volume II: Minutes of Evidence*. HM Stationery Office, pages 126–29.

Jevons, W S (1863), ‘Serious fall in the value of gold ascertained and its social effects set forth’ (pamphlet). Reprinted in Jevons, W S, *Investigations in currency and finance* (London, 1884).

King, R G and Wolman, A L (1999), ‘What should the monetary authority do when prices are sticky?’, in Taylor, J B (ed), *Monetary Policy Rules*, University of Chicago Press, pages 349–98.

Levin, A T, Wieland, V W and Williams, J C (1999), ‘The performance of forecast-based monetary policy rules under model uncertainty’, *mimeo*, Federal Reserve Board, Washington DC.

McCallum, B T and Nelson, E (1999a), ‘Performance of operational policy rules in an estimated semi-classical structural model’, in Taylor, J B (ed), *Monetary Policy Rules*, University of Chicago Press, pages 15–45.

McCallum, B T and Nelson, E (1999b), ‘Nominal income targeting in an open-economy optimising model’, *Journal of Monetary Economics*, Vol 43(3), pages 553–78.

Nelson, E (2000), ‘UK monetary policy 1972-1997: a guide using Taylor rules’, *mimeo*, Bank of England.

Rotemberg, J J and Woodford, M (1999), ‘Interest rate rules in an estimated sticky-price model’, in Taylor, J B (ed), *Monetary Policy Rules*, University of Chicago Press, pages 57–119.

Rudebusch, G D (1998), ‘Do measures of monetary policy in a VAR make sense?’, *International Economic Review*, Vol 39(4), pages 907–31.

Rudebusch, G D (1999), ‘Assessing nominal income rules for monetary policy with model and data uncertainty’, *mimeo*, Federal Reserve Bank of San Francisco.

Rudebusch, G D and Svensson, L E O (1999), ‘Policy rules for inflation targeting’, in Taylor, J B (ed), *Monetary Policy Rules*, University of Chicago Press, pages 203–53.

Smets, F (1999), ‘What horizon for price stability?’, *mimeo*, European Central Bank.

Svensson, L E O (1997), ‘Inflation targeting: some extensions’, *mimeo*, IIES, Stockholm University.

Svensson, L E O (1999), ‘Inflation targeting as a monetary policy rule’, *Journal of Monetary Economics*, Vol 43(3), pages 607–54.

Svensson, L E O and Woodford, M (1999), ‘Implementing optimal policy through inflation-forecast targeting’, *mimeo*, Princeton University.

Taylor, J B (1993), ‘Discretion versus policy rules in practice’, *Carnegie-Rochester Series on Public Policy*, Vol 39(0), pages 195–214.

Woodford, M (1999), ‘Commentary: how should monetary policy be conducted in an era of price stability?’, *mimeo*, Princeton University.

Data appendix

Definitions of the variables used in the VAR

The variables we use in the VAR are:

y_t : log of real GDP (quarterly, seasonally adjusted).

π_t^{DEV} : $\pi_t - \pi_t^{TAR}$, where:

π_t : $\log(RPIX_t / RPIX_{t-4})$, where $RPIX_t$ is the RPIX deflator in quarter t .

Our use of the four-quarter inflation rate rather than the quarterly change is motivated by the fact that, historically, targets for UK inflation or other nominal aggregates have been expressed in terms of annual changes rather than quarter-to-quarter movements. It is also conceivable that a four-quarter inflation rate may be a better empirical measure of underlying quarterly inflation than actual quarterly inflation.

π_t^{TAR} : Target (annualised) inflation rate, calculated as follows:

- 1976 Q3 – 1985 Q1: The implicit inflation target was the one implied by the monetary aggregate (£M3) targets. Using the assumption about trend velocity growth of +1.25% per annum stated by the Treasury in 1980,⁽²³⁾ we backed out the nominal income growth targets implicit in the 1976–1980 £M3 targets and in the targets for 1980–1984 announced in the Medium Term Financial Strategy (MTFS) in 1980. To obtain the implied inflation target, we subtracted 2.5% (to allow for output growth) from each nominal income growth target.

A new set of £M3 targets for 1982–1985 was announced in March 1982; for the years also covered by the 1980 MTFS, the £M3 growth targets now had a 3 percentage point higher mid-point, which we take as a change in the annual trend velocity growth assumption from +1.25% to –1.75% rather than changed targets for nominal income growth. Subsequent revisions to the 1982 MTFS are again assumed to represent changed

⁽²³⁾ The Treasury, in its submission to the Treasury and Civil Service Committee (1980), estimated trend £M3 velocity growth at 1.0–1.5% per annum.

assumptions about velocity rather than changed targets for nominal income growth.

From 1985 to 1990, we use the announced nominal income growth target, then subtract 2.5% for output growth.

- 1990 Q4 – 1992 Q3: given the ERM agreement, the implicit inflation target coincided with the German 2% inflation target;
- 1992 Q4 onwards: 2.5% (explicit inflation target).

Δe_t : Nominal exchange rate change, $\log(ERI_t / ERI_{t-1})$, where ERI is the exchange rate index. This variable is measured such that an observation of -0.10 indicates a depreciation of sterling of 10%.

R_t : Quarterly average of the annualised nominal interbank lending rate, measured as a fraction.

Time series properties of the data

We now show that all of the variables in our VAR are adequately described as stationary (I(0)) or trend-stationary processes, and hence, a Johansen-style cointegration approach to our VAR is not appropriate.

We model log output (y_t) as trend-stationary, with π_t^{DEV} and Δe_t treated as I(0) series, and R_t as an I(0) series after controlling for key shifts in monetary policy regime. As evidence, in Table E we present augmented Dickey Fuller (ADF) statistics that test the null of a unit root for the variables in our model. Since our contention is that output can satisfactorily be modelled as trend-stationary, Table E actually gives a test for a unit root in *detrended* output.⁽²⁴⁾ Two ADF test statistics are

⁽²⁴⁾ Detrended output is the residual from a prior regression of y_t on a constant and linear trend over 1980 Q3 – 1998 Q1. No constant is included in the regression used in calculating the ADF test for the stationarity of detrended output because the dependent variable is mean zero by construction.

calculated for the nominal interest rate, R_t : the first includes only a constant in the ADF regression, whereas the second includes a constant, $DERM_t$, and $D924_t$. Excluding these dummy variables may bias the test toward suggesting a unit root in R_t .

Table E: ADF Tests for VAR

<i>Variable</i>	ADF statistic
Detrended y_t	-2.427
π_t^{DEV}	-3.297*
Δe_t	-3.817**
R_t (no shifts)	-1.984
R_t (shifts included)	-3.675*

Note: A lag length of four is used in the ADF regressions for each variable except R_t (one lag). A * denotes significance at 0.05 level according to the Dickey-Fuller distribution's critical values; a ** significance at the 0.01 level according to these values.

The tests generally reject the null of a unit root in favour of the alternative of stationarity (or, in the case of R_t , an $I(0)$ series with structural breaks). So we believe it is satisfactory to treat the elements of our VAR as all $I(0)$, and therefore we do not apply cointegration analysis.⁽²⁵⁾

⁽²⁵⁾ There is also an economic reason for not favouring a cointegration-based approach to this VAR. The only real variable in the VAR is y_t and we do not want to explain the trend in y_t with nominal variables.

Technical appendix – solution methods

To demonstrate how we solve for the optimal rule, consider the policy-makers' problem for determining optimal policy when the economy's structure is described by the forward-looking structural model (FLSM). The optimal policy exercise for this model is not a standard optimal control problem because the model's equations, and therefore the policy-makers' constraints, contain forward-looking elements. A similar problem is faced by King and Wolman (1999), who assume certainty equivalence and then follow the standard Lagrangean optimisation approach. We follow King and Wolman's procedure here.

For convenience, we restate the equations of the FLSM here.

$$y_t = E_t y_{t+1} - \sigma(R_t - E_t \pi_{t+1}) + \delta \tilde{q}_{t-1} + e_{yt} \quad (\text{A1})$$

$$\pi_t = \alpha \pi_{t-1} + (1 - \alpha) E_t \pi_{t+1} + \phi_y y_{t-1} + \phi_q \Delta \tilde{q}_{t-1} + e_{\pi} \quad (\text{A2})$$

$$E_t q_{t+1} = q_t + R_t - E_t \pi_{t+1} + \kappa_t \quad (\text{A3})$$

If we let Λ_t , ϕ_t , and Ω_t denote the Lagrange multipliers on (A1)-(A3) respectively, the policy-makers' first-order conditions (with respect to y_t , π_t , q_t , and R_t) for minimising the loss function (2) are:

$$\beta \phi_y E_t \Gamma_{t+1} = 2\lambda_y y_t + \lambda_t - (1/\beta) \lambda_{t-1} \quad (\text{A4})$$

$$\alpha \beta E_t \Gamma_{t+1} = 2\lambda_{\pi} \pi_t + \Gamma_t - ((1-\alpha)/\beta) \Gamma_{t-1} - (\sigma/\beta) \lambda_{t-1} + (1/\beta) \Omega_{t-1} \quad (\text{A5})$$

$$\beta^4(\delta/4)E_t \lambda_{t+4} + \beta^3(\delta/4)E_t \lambda_{t+3} + \beta^2(\delta/4)E_t \lambda_{t+2} + \beta(\delta/4)E_t \lambda_{t+1} \quad (\text{A6})$$

$$-\beta^5(\phi_q/4) E_t \Gamma_{t+5} + \beta(\phi_q/4)E_t \Gamma_{t+1} = -\Omega_t + (1/\beta)\Omega_{t-1}$$

$$2\beta\lambda_{\Delta R} E_t R_{t+1} = \sigma\lambda_t - \Omega_t + (2\lambda_{\Delta R} + 2\beta\lambda_{\Delta R})R_t - 2\lambda_{\Delta R} R_{t-1} \quad (\text{A7})$$

In deriving these first-order conditions, we have suppressed the shocks to equations (A1)-(A3), in effect exploiting certainty equivalence. The FLSM equations (A1)-(A3), combined with the policy-makers' first-order conditions (A4)-(A7) for optimal policy, can be cast in the vector form:

$$A E_t \mathbf{y}_{t+1} = B \mathbf{y}_t + C \mathbf{z}_t \quad (\text{A8})$$

where \mathbf{y}_t is the vector of endogenous variables and \mathbf{z}_t is the vector of exogenous shocks. Once this system of expectational difference equations has been solved, we restore the shocks to (A1)-(A3), and can calculate impulse responses and analytical variances under the optimal rule. In the model's solution, the state vector consists of the exogenous shocks in (A1)-(A3); the lagged endogenous variables R_{t-1} , π_{t-1} , q_{t-1} , q_{t-2} , q_{t-3} , q_{t-4} , and q_{t-5} that appear in equations (2), (A1), and (A2); and the lagged Lagrange multipliers Γ_{t-1} , λ_{t-1} and Ω_{t-1} that appear in the policy-makers' first-order conditions (A4)-(A6).

For the optimal *feedback* rules, we simply supplemented equations (A1)-(A3) with rule (6), and optimised over the parameters (ψ_p , ρ_R , and k), with the criterion being minimisation of loss function (2). In conducting this search over parameter values, whenever we found multiple solutions to the model, we used the minimal state variable criterion to pick the bubble-free solution. An alternative approach, used in Levin, Wieland, and Williams (1999), is to exclude from consideration those regions of the parameter space that are associated with multiple solutions.