Forward-looking rules for monetary policy

Nicoletta Batini

and

Andrew G Haldane

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 have benefited greatly from the comments and suggestions of Bill Allen, Andy Blake, Willem Buiter, Paul Fisher, Charles Goodhart, Mervyn King, Tiff Macklem, David Miles, Alessandro Missale, Paul Mizen, Paul Levine, Stephen Millard, Darren Pain, Joe Pearlman, Richard Pierse, John Taylor, Paul Tucker, Ken Wallis, Peter Westaway, John Whitley, Stephen Wright, and especially our discussant Don Kohn and other seminar participants at the NBER Conference on 'Monetary Policy Rules', Islamorada, Florida, January 1998.

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 the Publications Group. (Telephone 0171-601 4030). Working papers are also available from the Bank's Internet site at http:\\www.bankofengland.co.uk.

ISSN 1368-5562

Contents

Abstract

This paper evaluates a class of simple policy rules that feed back from expected values of future inflation — inflation forecast-based rules. These rules are simple, and so are analogous to Taylor rule specifications. Because they are forecast-based, the rules are meant to mimic, albeit imperfectly, monetary policy behaviour among inflation-targeting central banks in practice.

In the paper, inflation forecast-based rules are assessed by evaluating how well they perform when the economy — a small rational expectations macro-model with sticky inflation and forward-looking agents — is buffeted by a combination of shocks, whose distribution is drawn from the Bank of England's forecasting model.

The paper shows that inflation forecast-based rules confer some real benefits: they embody explicitly monetary transmission lags (*lag-encompassing*); they potentially embody all information useful for predicting future inflation (*information-encompassing*); and, suitably designed, they can achieve a high degree of output smoothing (*output-smoothing*). In fact, these rules prove more efficient at minimising inflation and output variability than standard Taylor rule specifications, and almost as efficient as fully optimal rules. These results seem robust across different model specifications.

1 Introduction

It has long been recognised that economic policy in general, and monetary policy in particular, needs a forward-looking dimension. '[I]f we wait until a price movement is actually afoot before applying remedial measures, we may be too late', as Keynes (1923) observes in *A Tract on Monetary Reform*. That same constraint still faces the current generation of monetary policy-makers. Alan Greenspan's Humphrey-Hawkins Testimony in 1994 summarises the monetary policy problem thus:

> 'The challenge of monetary policy is to interpret current data on the economy and financial markets with an eye to anticipating future inflationary forces and to countering them by taking action in advance'.

Or in the words of Donald Kohn (1995) at the Board of Governors of the Federal Reserve System: 'policymakers cannot avoid looking into the future'. Empirically estimated reaction functions suggest that policy-makers' actions match these words. Monetary policy in the G7 countries appears in recent years to have been driven more by anticipated future than by lagged actual outcomes (Clarida and Gertler (1997), Clarida, Gali and Gertler (1997), Orphanides (1997)).

But how best is this forward-looking approach made operational? Friedman's (1959) *Program for Monetary Stability* cast doubt on whether it could be. Likening economic forecasting to weather forecasting, he observes: 'Leaning today against next year's wind is hardly an easy task in the present state of meteorology'. Yet this is just the task that present-day monetary policy-makers have set themselves: in effect, long-range weather-forecasting in a stochastic world of time-varying lags and coefficients. That is a tough problem, even for meteorologists. It is not altogether surprising, then, that attempts to solve the equivalent problem in a monetary policy context have produced several different approaches.

Targeting inflation directly is one of the more innovative of the solutions adopted recently by several countries, including New Zealand, Canada, the

United Kingdom, Sweden, Finland, Australia and Spain (see Haldane (1995) and Leiderman and Svensson (1995)). In the first three of these countries, monetary policy is linked to explicit (and in some cases published) inflation forecasts.[1] These forecasts are the *de facto* intermediate or feedback variable for monetary policy (Svensson (1997a,b), Haldane (1997)). The aim of this paper is to evaluate that particular approach to monetary policy, given the general problem of the need for forward-lookingness.

This is done by evaluating a class of simple policy rules that feed back from expected values of future inflation — inflation forecast-based rules. These rules are simple, and so are analogous to the Taylor rule specifications that have recently been extensively discussed in an academic and policy-making context. Because they are forecast-based, the rules mimic (albeit imperfectly) the monetary policy choices of inflation-targeting central banks.[2] And despite their simplicity, these forecast-based rules have a number of desirable features, which means that they may approximate the optimal feedback rule.

The class of forecast-based rules that we consider takes the following generic form:

$$r_t = \gamma r_{t-1} + (1-\gamma)\, r_t^* + \theta\, [E_t\, \pi_{t+j} - \pi^*] \qquad \textbf{(1)}$$

where r_t denotes the short-term *ex ante* real rate of interest, $r_t \equiv i_t - E_t\pi_{t+1}$, where i_t is the nominal interest rate; r_t^* denotes the equilibrium value of the real interest rate; $E_t(.) = E(\,.|\Phi_t)$, where Φ_t is the information set available at time t and E is the mathematical expectations operator; π_t is inflation

(1) In the other inflation-target countries, inflation forecasts are sometimes less explicit but nevertheless a fundamental part of the monetary policy process.
(2) We discuss below the places in which the forecast-based rules we consider deviate from real-world inflation targeting.

($\pi_t \equiv p^c_t - p^c_{t-1}$ where p^c_t is the log of the consumer price index); and π^* is the inflation target.[3]

According to the rule, the monetary authorities control deterministically the nominal interest rate (i_t) so as to hit a path for the short-term real interest rate (r_t). The short real rate is in turn set relative to some steady-state value, determined by a weighted combination of lagged and equilibrium real interest rates. The novel feature of the rule, however, is the feedback term. Deviations of *expected* inflation — the feedback variable — from the inflation target — the policy goal — elicit remedial policy actions.

The policy choice variables for the authorities are the parameter triplet $\{j,\theta,\gamma\}$. The parameter γ measures the degree of interest rate smoothing (see Williams (1997)). So, for example, with $\gamma = 0$ there is no instrument smoothing. θ is a policy feedback parameter. Higher values of θ imply a more aggressive policy response for a given deviation of the inflation forecast from its target. Finally, j is the targeting horizon of the central bank when forming its policy. For example, in the United Kingdom, the Bank of England feeds back from an inflation forecast around two years ahead (King (1997)).[4] The horizon of the inflation forecast (j) and the size of the feedback coefficient (θ), as well as the degree of instrument smoothing (γ), dictate the speed at which inflation is brought back to target following inflationary disturbances. Because they influence the transition path of inflation, these policy parameters clearly also have a bearing on output dynamics.

As defined in **(1)**, inflation targeting amounts to a well-defined monetary policy rule. That view is consistent with Bernanke and Mishkin's (1997) characterisation of inflation targeting as 'constrained discretion'. There is ample scope for discretionary input into any rule, particularly into **(1)**.

---

(3) The rule could be augmented with other — for example, explicit output — terms. We do so below. This then takes us close to the reaction function specification found by Clarida, Gali and Gertler (1997) to match recent monetary policy behaviour in the G7 countries.

(4) This comparison is not exact, because j defines the *feedback* horizon under the rule, whereas in practice in the United Kingdom, two years refers to the *policy* horizon (the point at which expected inflation and the inflation target are in line).

These discretionary choices include the formation of the inflation expectation itself and the choice of the parameter set $\{j, \theta, \pi^*\}$. So the policy rule **(1)** does not fall foul of the critique of inflation targeting made by Friedman and Kuttner (1996): that it is rigid as a monetary strategy and hence destined to the same failures as, for example, strict monetary targeting.

This is fine as an intuitive description of a forecast-based policy rule such as **(1)**. But what, if any, theoretical justification do these rules have? And, in particular, why might they be preferred to (for example) Taylor rules? Several authors have recently argued that, in certain settings, expected inflation-targeting rules have desirable properties (*inter alia*, King (1997), Svensson (1997a,b) and Haldane (1998)). For example, in Svensson's model (1997a), the optimal rule when the authorities care only about inflation sets the interest rate so as to bring expected inflation into line with the inflation target at some horizon ('strict' inflation forecast-targeting). When the authorities care also about output, the optimal rule is to close less than fully any gap between expected inflation and the inflation target ('flexible' inflation forecast-targeting). (5)

The rules we consider here differ from those in Svensson (*op cit*) in that they are simple feedback rules for the policy instrument, rather than complicated, optimal targeting rules. Simple feedback rules have some clear advantages. First, they are then directly analogous to, and so comparable with, the other policy rule specifications discussed in the papers in this volume, including Taylor rules.(6) Second, simple rules are arguably more robust when there is uncertainty about the true structure of the economy. And third, simple rules may be advantageous on credibility and monitorability grounds (Taylor (1993a)). The last of these considerations is perhaps the most important in a policy context. One way to use these rules is as a cross-check on actual policy in real time. For that to be practical, any rule needs to be simple and monitorable by outside agents.

---

(5) Rudebusch and Svensson (1997) consider empirically rules of this sort.
(6) We refer to the forthcoming NBER Conference volume 'Monetary Policy Rules', edited by Taylor, J B (1998), in which this paper will appear.

At the same time, the simple forecast-based rules we consider do have some clear similarities with Svensson's optimal inflation-forecast-targeting rules. Monetary policy under both rules seeks to offset deviations between expected inflation and the inflation target at some horizon. [(7)] More concretely, even simple forecast-based specifications can be considered **'**encompassing**'** rules, in the following three respects.

(a) Lag-encompassing

The lags between the enactment of monetary policy and its first effects on inflation and output are well known and widely documented. The monetary authorities need to be conscious of these lags when framing policy; they need to be able to calibrate them reasonably accurately; and they then need to embody them in the design of their policy rules. Without this, monetary policy will always be acting too late to head off incipient inflationary pressures. Such myopic policy may itself then become a source of cyclical (in particular, inflation) instability, for the very reasons outlined by Friedman (1959).[(8)]

By judicious choice of j, the lead term on expected inflation in **(1)**, simple forecast-based rules can be designed so as to embody automatically these transmission lags. In particular, the feedback variable in the rule can be chosen so that it is directly under the control of the monetary authorities — inflation j periods hence. The policy-makers' feedback and control variables are then explicitly aligned. Transmission lags are the most obvious (but not the only) reason why monetary policy needs a forward-looking, pre-emptive dimension. Embedding these lags in a formal forecast-based rule is simple recognition of that fact. [(9)] Reflecting this, lag-encompassing was precisely the motivation behind targeting

(7) In particular, since the rules that we consider allow flexibility over both the forecast horizon (j) and the feedback parameter (θ) — both of which affect output stabilisation — their closest analogue is Svensson's flexible inflation-forecast-targeting rule.

(8) Former Vice-Chairman of the Federal Reserve Alan Blinder (1997) observes: **'**Failure to take proper account of lags is, I believe, one of the main sources of central bank error**'.**

(9) Svensson (1997a) shows, in the context of his model, that rules with this lag-encompassing feature secure the minimum variance of inflation, precisely because they guard against monetary policy acting too late.

expected inflation in those countries where this was first adopted: New Zealand, Canada and the United Kingdom.

(b) Information-encompassing

Under inflation forecast-based rules, the inflation expectation in **(1)** can be thought of as the intermediate target variable for monetary policy. It is well suited to this task when judged against the three classical requirements of any intermediate variable: controllability, predictability, and a leading indicator. Expected inflation is, almost by definition, the indicator most closely correlated with the future value of the variable of interest. In particular, expected inflation ought to embody all information contained within the myriad indicators that affect the future path of inflation. Forecast-based rules are, in this sense, information-encompassing. That is not a feature necessarily shared by backward-looking policy rules — for example, those considered in the volume by Bryant, Hooper and Mann (1993).

Of course, any forward-looking rule can be given a backward-looking representation and re-specified in terms of current and previously-dated variables. For example, in the aggregate demand/aggregate supply model of Svensson (1997a), the optimal forward-looking rule can be rewritten as a Taylor rule — albeit with weights on the output gap and inflation that are likely to be very different from one half. But that will not necessarily be the case in more general settings, where shocks do not only come from output and prices. Taylor-type rules will tend then to feed back from a restrictive subset of information variables and so will not in general be optimal.[10] By contrast, inflation forecast-based rules will naturally embody all information contained in the inflation reduced-form of the model: extra lags of existing pre-determined variables and additional pre-determined variables, both of which would typically also enter the optimal feedback rule. For that reason, even simple forecast-based rules are likely to take us close to the optimal state-contingent rule — or at least closer than Taylor-type rule specifications.

(10) Black *et al* (1997) illustrate this in a simulation setting.

(c) Output-encompassing

As specified in **(1)**, inflation forecast-based rules appear to take no explicit account of output objectives. The inflation target, π^*, defines the nominal anchor, and there is no explicit regard for output stabilisation. But π^* is not the only policy choice parameter in **(1)**. The targeting horizon (j) and feedback parameter (θ) — the two remaining policy choice variables — can in principle also help to secure a degree of output-smoothing. These parameters can be chosen to ensure that an inflation forecast-based rule better reflects the authorities' preferences in situations where they care about output as well as inflation variability. To see how these policy parameters affect output stabilisation, consider separately shocks to demand and supply.

In the case of demand shocks, inflation and output stabilisation will in most instances be mutually compatible. Demand shocks shift output and inflation in the same direction relative to their baseline values. So there need not then be any inherent trade-off between output and inflation stabilisation in the setting of monetary policy following these shocks. A rule such as **(1)** will automatically secure a degree of output stabilisation in a world of only demand shocks. Or put differently, because it is useful for predicting future inflation, the output gap already appears implicitly in an inflation forecast-based rule such as **(1)**.

For supply shocks, trade-offs between output and inflation stability are more likely, because they will tend then to be shifted in opposite directions. But inflation-targeting does not imply that the authorities are opting for a corner solution on the output/inflation variability trade-off curve in these situations. For example, different inflation-forecast horizons — different values of j — will imply different points on the output/inflation variability frontier. Longer forecast horizons smooth the transition of inflation back to target following inflation shocks, partly because policy then accommodates (rather than offsets) the first-round effects of any supply shocks.[11] The feedback coefficient (θ) also has a bearing on output dynamics, for much the same reason. So a central bank following an inflation-forecast based rule can, in principle, simply choose its policy parameters $\{j,\theta,\gamma\}$ so as to achieve a preferred point on the output/inflation variability spectrum. Certainly, the simple forecast-based policy rule, **(1)**, ought not to be the sole preserve of monomaniac inflation-fighters.

This paper aims to put some quantitative flesh onto this conceptual skeleton. It evaluates simple forecast-based rules against the three encompassing criteria outlined above.[12] We are then able to address three types of policy questions: What is the optimal degree of policy forward-lookingness, and what does this depend on? Can inflation-only rules secure sufficient output smoothing? And how do simple forecast-based rules compare with the fully optimal rule, and with simple Taylor rules?

To anticipate our conclusions, we find quantitative support for all three of the encompassing propositions. Because inflation forecast-based policy rules embody transmission lags, they generally help improve inflation control (lag-encompassing). These rules can be designed to smooth the path of output as well as inflation, despite not feeding back from the former explicitly (output-encompassing). And inflation forecast-based rules deliver clear welfare improvements over Taylor-type rules, which respond to a restrictive subset of information variables (information-encompassing).

---

(11) This is broadly the practice followed in the United Kingdom. The Bank of England is required to write an open letter to the Chancellor in the event of inflation deviating by more than one percentage point from its target, stating the horizon by which inflation is to be brought back to heel. Longer horizons might be chosen following large and/or persistent supply shocks, so that policy does not disturb output too much *en route* back to the inflation target. That is important because the UK inflation target, while giving primacy to price stability, also requires that the Bank of England take account of output and employment objectives when setting monetary policy. There are other design features of inflation targets that can ensure a sufficient degree of output stabilisation. For example, in New Zealand, there are inflation-target exemptions for 'significant' supply shocks (see Mayes and Chapple (1995)); while in Canada, there is a larger inflation fluctuation margin to help insulate against shocks (see Freedman (1996)).
(12) Previous empirical simulation studies that have considered the performance of forward-looking rules include Black *et al* (1997), Clark *et al* (1995) and de Brouwer and O'Regan (1997).

The rest of the paper is planned as follows. Section 2 outlines our model. Section 3 calibrates this model and conducts some deterministic experiments with it. Section 4 uses stochastic analysis to evaluate the three conceptual properties of forecast-based rules — lag-encompassing, information-encompassing and output-encompassing — outlined above. Section 5 briefly summarises.

2. The model

To evaluate equation **(1)** and variants of it, we use a small, open-economy, log-linear calibrated rational expectations macro-model. It has similarities to the optimising IS/LM framework recently developed by McCallum and Nelson (1997) and Svensson (1997c), and hence indirectly to the stochastic general equilibrium models of Rotemberg and Woodford (1997) and Goodfriend and King (1997). The open-economy dimension is important when characterising the behaviour of inflation-target countries, which tend to be small open economies (see Blake and Westaway (1996) and Svensson (1997c)). The exchange rate also has an important bearing on output/inflation dynamics in our model, in keeping with the results of Ball (1997).

The model is kept deliberately small to ease the computational burden. But a compact model is also useful in helping clarify the transmission mechanism channels at work, and the trade-offs that naturally arise between them. And despite its size, the model embodies the key features of the small forecasting model used by the Bank of England for its inflation projections. Our model is calibrated to match the dynamic path of output and inflation generated by structural and reduced-form models of the UK economy in the face of various shocks.

The model comprises six behavioural relationships, listed below as equations **(2)-(7)**.

$$y_t - y_t^* = a_1 y_{t-1} + a_2 E_t(y_{t+1}) + a_3 [i_t - E_t(p_{t+1})] + a_4 (e_t + p^c_t - p^{cf}_t) + e_{1t} \quad \textbf{(2)}$$

$$m_t - p^c_t = b_1 y_t + b_2 i_t + e_{2t} \quad \textbf{(3)}$$

$$e_t = E_t(e_{t+1}) + i_t - i_t^f + e_{3t} \quad \textbf{(4)}$$

$$p^d{}_t = 1/2\,[w_t + w_{t-1}] \quad \textbf{(5)}$$

$$w_t - p^c{}_t = c_0\,[E_t(w_{t+1}) - E_t(p^c{}_{t+1})] + (1 - c_0)[w_{t-1} - p^c{}_{t-1}] + c_1\,(y_t - y_t^*) + e_{4t} \quad \textbf{(6)}$$

$$p^c{}_t = f\,p^d{}_t + (1 - f)\,e_t \quad \textbf{(7)}$$

All variables, except interest rates, are in logs. Importantly, in the simulations all behavioural relationships are also expressed as deviations from equilibrium. So for example, we set the (log) natural rate of output, y_t*, equal to zero. We also normalise to zero the (log) foreign price level and foreign interest rate, $p^{c\,f}_t = i_t^f = 0$, and the (implicit) mark-up in **(5)** and foreign exchange risk premium in **(4)**.

Equation **(2)** is a standard IS curve, with real output, y_t, depending negatively on the *ex ante* real interest rate and the real exchange rate (where e_t is the foreign currency price of domestic currency), $\{\alpha_3, a_4\} < 0$. The former channel is defined over short rather than long real interest rates. We could have included a long-term interest rate in our model, linking long and short rates through an arbitrage condition, as in Fuhrer and Moore's (1995a) model of the United States. But in the United Kingdom, unlike in the United States, expenditure is more sensitive to short than to long interest rates, owing to the prevalence of floating-rate debt instruments.

Output also depends on lags of itself, reflecting adjustment costs and, more interestingly, a lead term. The latter is motivated by McCallum and Nelson's (1997) work on the form of the reduced-form IS curve that arises from a fully-optimising general equilibrium macro-model. We experiment with this lead term below, though we do not use it in our baseline simulations. e_{1t} is a vector of demand shocks, for example shocks to foreign output and fiscal policy.

Equation **(3)** is an LM curve.[13] Its arguments are conventional: a nominal interest rate, capturing portfolio-balance; and real output, capturing transactions demand.[14] e_{2t} is a vector of velocity shocks. Equation **(4)** is an uncovered interest parity condition. We do not include any explicit foreign exchange risk premium. The shock vector e_{3t} comprises foreign interest rate shocks and other noise in the foreign exchange market, including shocks to the exchange risk premium.

Equations **(5)** and **(6)** define the model's supply side. They take a similar form to that of other staggered contract models.[15] Equation **(5)** is a mark-up equation. Domestic output prices (in logs, p^d_t) are a constant mark-up over weighted average contract wages (in logs, w_t) in the current and preceding period. Equation **(6)** is the wage-contracting equation. Under this specification, wage contracts last two periods.[16] Agents in today's wage cohort bargain over *relative* real consumption wages. Today's real contract wage is some weighted average of the real contract wage of the 'other' cohort of workers: that is, wages already agreed in the previous period and those expected to be agreed in the next period. We do not impose symmetry on the lag and lead terms in the contracting equation, in contrast with the standard Fuhrer and Moore (1995b) model. Instead, we allow a flexible mixed lag/lead specification, which nests more restrictive alternatives as a special case.[17] This flexible mixed specification is found in Fuhrer (1997) to be preferred empirically. It also allows us to experiment with the degree of forward-lookingness in the wage-bargaining process. The lag/lead weights are restricted to sum to unity, however, to preserve price homogeneity in the wage-price system (a vertical long-run Phillips curve). Also in the wage-contracting equation is a conventional output gap term, capturing tightness in the labour market.

---

(13) This is largely redundant in our analysis, since we are focusing on interest rate rules that assume that the demand for money is always fully accommodated at unchanged interest rates.
(14) McCallum and Nelson (1997) show that this form of the LM curve can also be derived as the reduced-form of an optimising stochastic general equilibrium model.
(15) In particular, they are similar to those recently developed by Fuhrer and Moore (1995a) for the United States. For an early formulation of such model, see Buiter and Jewitt (1981).
(16) We could have lengthened the contracting lag — for example to four periods, which in our calibration is one year — to better match real-world behaviour. But two lags appeared to be sufficient to generate the inflation persistence evident in the data, when taken together with the degree of backward-lookingness embodied in the Phillips curve.
(17) See Blake and Westaway (1996).

The shock vector, ε_{4t}, captures disturbances to the natural rate of output and other supply shocks.

This relative wage-price specification has both theoretical and empirical attractions. Its theoretical appeal comes from early work by Duesenberry (1949), who argued that wage relativities were a key consideration in the wage bargain. The empirical appeal of the relative real wage formulation is that it generates inflation persistence. This is absent from a conventional two-period Taylor-contracting (1980) specification (Fuhrer and Moore (1995a), Fuhrer (1997)), which instead produces price-level persistence.(18) Equation **(7)** defines the consumption price index, comprising domestic goods (with weight ϕ) and imported foreign goods (with weight $(1-\phi)$).(19) Note, importantly, that **(7)** implies full and immediate pass-through of import prices (and hence exchange rate changes) into consumption prices —an assumption we discuss further below.

Some manipulation of **(5)-(7)** gives the reduced-form Phillips curve of the model:

$$p_t = c_0 E_t(p_{t+1}) + (1-c_0)\, p_{t-1} + c_1(y_t + y_{t-1}) + m\,[(1-c_0)\, Dc_t - c_0\, E_t\, Dc_{t+1})] + e_{5t} \qquad \textbf{(8)}$$

where c_t º $e_t + p^c_t$ (the real exchange rate); $\mu \equiv 2(1-\phi)$; D is the backward difference operator; and e_{5t} º $e_{4t} + c_0\,[(p^c_t - E_{t-1}p^c_t) - (w_t - E_{t-1}w_t)]$, where the composite error now includes expectational errors by wage-bargainers.

Equation **(8)** is the open-economy analogue of Fuhrer and Moore's (1995a) Phillips curve specification (see Blake and Westaway (1996)). The inflation terms — a weighted backward and forward-looking average — are the same as in the closed-economy case. There is inflation persistence. The specification differs because of additional (real) exchange rate terms,

(18) As Roberts (1995) discusses, Taylor-contracting can deliver inflation-persistence if, for example, expectations are made 'not-quite-rational'. Certainly, there is a variety of mechanisms other than the one adopted here that would have allowed us to introduce inflation persistence in the model.

(19) With the foreign price level normalised to zero in logs.

reflecting the price effects of exchange rate changes on imported goods in the consumption basket. [20]

The transmission of monetary impulses in this model is very different from the closed-economy case, in terms of the size and timing of the effects. We illustrate these effects below. There is a conventional real interest rate channel, working through the output gap and thence onto inflation. But in addition, there is a real exchange rate effect, operating through two distinct channels. First, there is an indirect output gap route running through net exports and thence onto inflation. And second, there are direct price effects via the cost of imported consumption goods, and via wages and hence output prices. The latter channel means that disinflation policies have a speedy effect on consumer prices ($p^c{}_t$), if not on domestically generated prices ($p^d{}_t$) (see Svensson (1997c)). This direct exchange rate channel thus has an important bearing on consumer price inflation and output dynamics, which we illustrate below. Because these direct exchange rate effects derive from the (potentially restrictive) assumption of full and immediate pass-through of exchange rate changes to consumption prices, however, we also experiment below with a model where pass-through is sluggish and/or incomplete. This specification might be more realistic if, for example, we believe that foreign exporters 'price to market', holding the foreign currency price of their exported goods relatively constant in the face of exchange rate changes; or if home-country retail importers absorb the effects of exchange rate changes in their margins.

The model, **(2)-(7)**, is clearly not structural in the sense that we can back out directly from it taste and technology parameters. Nevertheless, as McCallum and Nelson (1997) have recently shown, a system such as **(2)-(7)** can be derived as the linear reduced form of a fully-optimising general equilibrium model, under certain specifications of tastes and technology. That should confer some degree of policy-invariance on model parameters — and hence some immunity from the Lucas critique.

(20) Plus the effects of the composite error term.

3. Deterministic policy analysis

(a) Calibrating the model

To assess the properties of the model described above, we begin with some deterministic simulations. For this, we need to calibrate the behavioural parameters in **(2)-(7)**. As far as possible, we set our baseline calibrated values in line with prior empirical estimates on quarterly data. Where this is not possible — for example, in the wage-contracting equation — we calibrate parameters to ensure a plausible dynamic profile from impulse responses. We also experiment below, however, with some deviations from the baseline parameterisation, in particular the degree of forward-lookingness in the model.

For the IS curve, **(2)**, we set $a_1 = 0.8$, which is empirically plausible on quarterly data. For the moment, we set $\alpha_2 = 0$, ignoring until later any direct forward-lookingness in the IS curve. We set the real interest rate (a_3) and real exchange rate (a_4) elasticities to -0.5 and -0.2 respectively. For the LM curve, we set $b_1 = 1$ and $b_2 = 0.5$, so that money is unit income-elastic and has an interest semi-elasticity of one half. Both of these restrictions are broadly satisfied on UK data (Thomas (1996)).

On the contracting equation, **(6)**, our baseline model sets $c_0 = 0.2$, so that contracting is predominantly backward-looking. This specification matches the pattern of the data much better than an equally-weighted formulation, both in the United States (Fuhrer (1997)) and the United Kingdom (Blake and Westaway (1996)).[(21)] The output sensitivity of real wages is set at 0.2 ($c_1 = 0.2$), in line with previous studies.[(22)] We set ϕ, the share of domestically produced goods in the consumption basket, equal to 0.8, in line with existing shares.

Turning to the policy rule, **(1)**, for consistency with the model, this is also simulated as a deviation from equilibrium. That is, we set π^* (the inflation

---

(21) The lag/lead weights chosen here are very similar to those found empirically in the United States by Fuhrer (1997).
(22) The elasticity of real wages is close to that found by Fuhrer (1997) in the United States of 0.12.

target) and r_t* (the equilibrium real rate) to zero. Because of this, our simulations do not address questions regarding the optimal level of π*. For example, our model does not broach issues such as the stabilisation difficulties caused by the non-negativity of nominal interest rates. We are implicitly assuming that the level of π* has been set such that this constraint binds with only a negligibly small probability. Nor do we address issues such as time-variation in r_t*.

In terms of the parameter triplet $\{j, \theta, \gamma\}$, in our baseline rule we set $\gamma = 0.5$ — a halfway house between the two extreme values of interest rate smoothing we consider; $\theta = 0.5$ — around the middle of the range of feedback parameters used in previous simulation studies (Taylor (1993a), McCallum (1988), Black *et al* (1997)); and $j = 8$ periods. Because the model is calibrated to match quarterly profiles for the endogenous variables, this final assumption is equivalent to targeting the quarterly inflation rate two years ahead. This is around the horizon from which central banks feed back in practice. For example, the Bank of England's 'policy rule' has been characterised as targeting the inflation rate two years or so ahead (King (1996)).[23]

Because the model **(2)-(7)**, and the baseline policy rule **(1)**, are log-linear, we can solve the system using the method of Blanchard and Kahn (1980). Denote the vector of endogenous variables by $\boldsymbol{z_t}$.[24] The model **(1)-(7)** has a convenient state-space representation,

$$\begin{bmatrix} \boldsymbol{q_{t+1}} \\ E_t\, \boldsymbol{x_{t+1}} \end{bmatrix} = \boldsymbol{A} \begin{bmatrix} \boldsymbol{q_t} \\ \boldsymbol{x_t} \end{bmatrix} + \boldsymbol{B e_t} \qquad \textbf{(9)}$$

where $\boldsymbol{q_t}$ is a vector containing $\boldsymbol{z_{t-1}}$ and its lags, $\boldsymbol{x_t}$ is a vector containing $\boldsymbol{z_t}$, $E_t\, \boldsymbol{z_{t+1}}$, $E_t\, \boldsymbol{z_{t+2}}$, etc and where, as usual, E_t is the expectation operator using information up to time t, $\boldsymbol{e_t}$ is a vector of disturbances, and $\boldsymbol{A}$ and $\boldsymbol{B}$ are

(23) Though the United Kingdom's inflation target is defined as an annual percentage change in price levels, which means that this comparison is not exact (see below).
(24) Where the bold font style denotes vector and matrix notation.

matrices of time-invariant coefficients. The solution to **(9)** is obtained by implementing the Blanchard and Kahn (*op cit*) method with a standard computer program that solves linear rational expectations models.[25] This program imposes the condition that there are no explosive solutions, implying a relationship $E_t\,\boldsymbol{x}_{t+1} + \boldsymbol{N}\boldsymbol{q}_{t+1} = 0$, where $[\boldsymbol{N}\ \boldsymbol{I}]$ is the set of eigen-vectors of the stable eigen-values of $\boldsymbol{A}$.

We then evaluate the various rules by conducting stochastic policy simulations and calculating in each case unconditional moments of the endogenous variables. To conduct the simulations, we need a covariance matrix of the shocks for equation disturbances.

There is a variety of ways of generating these shocks. The theoretical model, **(2)-(7)**, does not have enough dynamic structure for its empirically estimated residuals to be legitimate measures of primitive shocks. Alternatively, and at the other end of the spectrum, we could use atheoretic time-series or VAR models to construct structural shocks. But that approach is not without problems either. Identification restrictions are still required to unravel the structural shocks from the reduced-form VAR residuals. Because these restrictions are just-identifying, they are non-testable. Further, in the VAR literature, these restrictions usually include orthogonality of the primitive disturbances, $E_t(\boldsymbol{e}_{it}\,\boldsymbol{e}_{jt}') = 0 \;\forall\; i \neq j$. That is not a restriction we would want necessarily to impose *a priori*.[26]

We steer a middle course between these alternatives, using a covariance matrix of structural shocks derived from the Bank of England's forecasting model.[27] This confers some advantages. First, and importantly, our analytical model can be considered a simplified version of this forecasting model, but without its dynamic structure. This lends some coherence to the deterministic and stochastic parts of the analysis. Second, the structural shocks from the forecasting model permit non-zero covariances.

---

(25) This was conducted within the ACES/PRISM solution software (Gaines, Al'Nowaihi and Levine (1989)).

(26) Though see Leeper, Sims and Zha (1996). Black *et al* (1996) generate identified VAR residuals without imposing these restrictions.

(27) This matrix is available from the authors on request.

For IS, LM and Phillips curve shocks, we simply take the moments of the residuals from the Bank's forecasting model over the sample period 1989 Q1-1997 Q3. Our sample period excludes most of the 1970s and 1980s, during which the variance of shocks for all of the variables was (sometimes considerably) higher. Using a longer sample period would re-scale upwards the variances we report. The exchange rate is trickier. For that, we use quarterly MMS survey data to capture exchange rate expectations over our sample, using the $/£ exchange rate as our benchmark.[28] The exchange rate residuals were then constructed from the arbitrage condition, **(4)**, plugging in the survey expectation and using quarterly data for the other variables. Not surprisingly, the resulting exchange rate shock vector has a large variance, around ten times that of the IS, LM and Phillips curve shocks. Given its size, we conducted some sensitivity checks on the exchange rate variance. Rescaling the variance does not alter the conclusions that we draw about the relative performance of the rules.

(b) A disinflation experiment

To assess the plausibility of the system's properties, we displaced deterministically the intercepts of every equation in the model (the IS equation, the money demand equation, the aggregate supply equation and the exchange rate equation, respectively) by 1%, and traced out in each case the resulting impulse response. Each of these impulse responses gave dynamic profiles that were theoretically plausible. For example, a permanent negative supply shock — for example, a rise in the NAIRU — shifted inflation and output in opposite directions on impact, and lowered output below baseline in steady state; whereas a permanent positive demand shock — such as a rise in overseas demand — shifted output and inflation in the same direction initially, but was output-neutral in steady state.

(28) A preferred exchange rate measure would have been the United Kingdom's trade-weighted effective index. But there are no survey data on exchange rate expectations of this index. We also looked at the behaviour of the DM/£ and yen/£ exchange rates. The variance of the $/£ residuals was somewhere between that of DM/£ and yen/£.

To illustrate the calibrated model's dynamic properties, consider the effects of a shock to the reaction function, **(1)**. Consider in particular a disinflation — a lowering of the inflation target, π^* — of one percentage point. The thick lines in Chart 1 plot the response of output and inflation to this inflation target shock. Impulse response profiles are shown as percentage point deviations from baseline values.

The economy returns to steady state after around 16 quarters (four years). At that point, inflation is one percentage point lower at its new target and output is back to potential. But the transmission process in arriving at this end-point is protracted. Output is below potential for the whole of the period, with a maximum marginal effect of around 0.2 percentage points after around five quarters. Output falls partly as a result of a policy-induced rise in real interest rates (of around 0.14 percentage points) and partly as a result of the accompanying real exchange rate appreciation (of around 0.57 percentage points). The path of output, and its maximum response, are broadly in line with simulation responses from VAR-based studies of the effects of monetary policy shocks in the United Kingdom (Dale and Haldane (1995)).[29] The cumulative loss of output — the sacrifice ratio — is around 1.5%. This sacrifice ratio estimate is not greatly out of line with previous UK estimates (Bakhshi, Haldane and Hatch (1998)), but is if anything on the low side (see below).

Inflation undergoes an initial downward step, owing to the impact effect of the exchange rate appreciation on import prices. Although the effect of the exchange rate shock is initially to alter the price *level*, this effect gets embedded in wage-bargaining behaviour and so has a sustained impact on measured inflation. Thereafter, inflation follows a gradual downward path towards its new target, under the impetus of the negative output gap. The inflation profile, and in particular the immediate step jump in inflation following the shock, is not in line with prior reduced-form empirical evidence on the monetary transmission mechanism.

The simulated inflation path is clearly sensitive to the assumptions that we have made about exchange rate pass-through — namely, that it is

(29) Though the shocks are not exactly the same.

immediate and complete. In particular, it is the full pass-through assumption that lies behind the initial jump in inflation following a monetary disturbance. So one implication of this assumption is that monetary policy in an open economy can affect consumer price inflation with almost no lag (Svensson (1997c)). There may well of course be adverse side-effects from an attempt to control inflation in this way, such as real exchange rate and hence output destabilisation. We illustrate these side-effects below. But more fundamentally, the monetary transmission lag, and hence the implied degree of inflation control, is clearly acutely sensitive to the exchange rate pass-through assumption that we have made.

As a sensitivity check, the dotted lines in Chart 1 show the responses of output and inflation if we assume no direct exchange rate pass-through into consumer prices.[30] Monetary policy impulses are then all channelled through output, either via the real interest rate or the real exchange rate. The resulting output path is little altered. But as we might expect, the downward path of inflation is more sluggish, mimicking the output gap. It is in fact now rather closer to that found from VAR-based studies of the effects of monetary policy in the United Kingdom. Given the clear sensitivity of the inflation profile to the pass-through assumption, we use both pass-through models below when considering the effects of transmission lags on the optimal degree of policy forward-lookingness.

(c) Some limitations of the simulations

The impulse responses suggest that our model is a reasonable dynamic representation of the effects of monetary policy in a small, open economy such as the United Kingdom, Canada or New Zealand — the three longest-serving inflation-targeters. Nevertheless, the simulated model responses are clearly a simplified and stylised characterisation of inflation

(30) Which we reproduce by assuming that the import content of the consumption basket is zero. This would be justified if, for example, all imported goods were intermediate rather than final goods; or more generally, if the effects of exchange rate changes were absorbed in foreign exporters' or domestic retailers' margins rather than in domestic currency consumption prices. See Svensson (1997c) for a comparison of inflation-targeting rules based on consumer and producer prices.

targeting as exercised in practice. Two limitations in particular are worth highlighting.

First, we impose model-consistency on all expectations, including the inflation expectations formed by the central bank, which serve as its policy feedback variable. This is coherent as a simulation strategy, as otherwise we would have to posit some expectational mechanism not consistent with the model in which the policy rule was being embedded. But the assumption of model-consistent expectations has drawbacks too. For example, it underplays the role of model uncertainties. These uncertainties are important, but a consideration of them is beyond the scope of the present paper. Further, the simulations assume that the inflation target is perfectly credible. So the shock to the target shown in Chart 1 is, in effect, believed fully and immediately. This helps explain why the sacrifice ratio implied by Chart 1 is lower than historical estimates; it is the full-credibility case. While the assumption of full credibility is limiting, it is not obvious that it should greatly affect our inferences about the relative performance of various rules, which is the focus of the paper.

Second, and relatedly, under model-consistent expectations, monetary policy is assumed to be driven by the specified policy rule. In particular, the inflation forecast of the central bank — the policy feedback variable — is conditioned on the inflation-targeting policy rule, **(1)**. This differs somewhat from actual central bank practice in some countries. For example, in the United Kingdom, the Bank of England's published inflation forecasts are usually conditioned on an assumption of unchanged interest rates.[31] This means that there is no direct read-across from our forecast-based rules to inflation targeting in practice.

Even among those countries that use it, however, the constant interest rate assumption is seen largely as a short-term expedient. It is not appropriate, for example, when simulating a forward-looking model — as here — because it deprives the system of a nominal anchor and thus leaves the price level indeterminate. So in our simulations, we instead condition

(31) This is also often the case with forecasts produced for the Federal Reserve Board's 'Green Book' (see Reifschneider, Stockton and Wilcox (1996)).

monetary policy (actual and in expectation) on the reaction function, **(1)**. This delivers a determinate price level. Simulations conducted in this way come close to mimicking current monetary policy practice in New Zealand (Reserve Bank of New Zealand (1997)). There, the Reserve Bank of New Zealand's policy projections are based on an explicit policy reaction function, which is very similar to the baseline rule, **(1)**. The Bank of England also recently began publishing inflation projections based on market expectations of future interest rates, rather than constant interest rates. This means that differences between the forecast-based rule, **(1)**, and inflation targeting in practice may not be so great.

4. Stochastic policy analysis

We now consider the performance of the baseline rule, **(1)**, and compare it with alternative rules. This is done by embedding the various rules in the model outlined above and evaluating the resulting (unconditional) moments of output, inflation and the policy instrument — the arguments typically thought to enter the central bank's loss function. Specifically, following Taylor (1993a), we consider where each of the rules places the economy on the output/inflation variability frontier.

(a) Lag-encompassing: the optimal degree of policy forward-lookingness

The most obvious rationale for a forward-looking monetary policy rule is that it can embody explicitly the lags in monetary transmission. But how forward-looking? Is there some optimal forecasting horizon from which to feed back? And if so, what does this optimal targeting horizon depend upon?

Answers to these questions are clearly sensitive to the assumed length of the lag itself. So we experiment below with both our earlier models: one assuming full and immediate import-price pass-through (a shorter transmission lag); the other, no immediate pass-through (a longer transmission lag). Chart 2 plots the locus of output/inflation variability points delivered by the rule **(1)**, as the horizon of the inflation forecast (j) is varied. Two lines are plotted in Chart 2, representing the two pass-through cases. Along these loci, we vary j between zero (current-period inflation-targeting) and sixteen (four-year-ahead inflation-forecast-targeting) periods.[32] Our baseline rule ($j = 8$) lies between these extremes. The two remaining policy choice parameters in **(1)**, $\{\gamma, \theta\}$, are for the moment set at their baseline values of 0.5.[33] Points to the south and west in Chart 2 are clearly welfare-superior, and points to the north and east inferior.

Several points are clear from Chart 2. First, irrespective of the assumed degree of pass-through, the optimal forecast horizon is always positive and lies somewhere between three and six quarters ahead. This forecast horizon secures as good inflation performance as any other, while at the same time delivering lowest output variability. The latter result arises because three to six quarters is around the horizon at which monetary policy has its largest marginal impact. The (integral of the) real interest and exchange rate changes necessary to hit the inflation target is minimised at this horizon. So too, therefore, is the degree of output destabilisation (the integral of output losses). At shorter horizons than this, the adjustment in monetary policy necessary to return inflation to target is that much greater — the upshot of which is a destabilisation of output. Once we allow for the fact that central banks in practice feed back from annual inflation rates, whereas our model-based feedback variable is a quarterly inflation rate, then the optimal forecast horizon implied by our simulations (of three to six quarters) is rather similar to that used by inflation-targeting central banks in practice (of six to eight quarters).[34]

Second, taking either pass-through assumption, feeding back from a forecast horizon much beyond six quarters leads to worse outcomes for *both* inflation and output variability. This is the flip-side of the arguments used above. Just as short-horizon targeting implies ‘too much’ of a policy response to counteract shocks, long-horizon targeting can equally imply

---

(32) Some of the longer-horizon feedback rules were unstable, which we discuss further below. In Chart 2, we show the maximum permissible feedback horizon: 14 periods for the full pass-through case, and twelve periods for the no pass-through case.

(33) We vary them both in turn below.

(34) This comparison is also not exact, because the two definitions of horizon are different: the *feedback* horizon in the rule and the *policy* horizon in practice (the point at which expected inflation is in line with the inflation target) are distinct concepts.

that policy does ‘too little’, thereby setting in train a destabilising expectational feedback. This works as follows.

Beyond a certain forecast horizon, the effects of any inflation shock have been eliminated from the system by the actions of the central bank: expected inflation is back to target. This implies that, beyond that horizon, our forward-looking monetary policy rule says ‘do nothing’: it is entirely ‘hands-off’. In expectation, policy has already done its job. But an entirely ‘hands-off’ policy will be destabilising for inflation expectations — and hence for inflation today — if it is the policy path actually followed in practice. This is because of the circular relationship between forward-looking policy behaviour and forward-looking inflation expectations. The one generates oscillations in the other, which in turn gives rise to further feedback on the first. Beyond a certain threshold horizon — when policy is very forward-looking — this circularity leads to explosiveness. So this is one general instance in which forward-looking rules generate instabilities — namely, when the forecast horizon extends well beyond the transmission lag. [35] The possibility of instabilities and indeterminacies arising in forecast-based rules is discussed in Woodford (1994) and Bernanke and Woodford (1997). The mechanism here is very similar.

Third, the main differences between the two pass-through loci show up at horizons less than four quarters. Beyond these horizons, the full pass-through locus heads due south, while the no pass-through locus heads south-west. With incomplete pass-through, policy forward-lookingness reduces both inflation and output variability. This is because inflation transmission lags are lengthier in this particular case. Embodying these (lengthier) lags explicitly in the policy reaction function thus improves inflation control; it guards against monetary policy acting too late. Pre-emptive policy helps to stabilise inflation in the face of transmission lags. At the same time, it also helps to smooth output, for the reasons outlined above.

(35) We highlight some other cases below.

The same is generally true in the full pass-through case, except that most of the benefits then accrue to output stabilisation. The gains in inflation stabilisation from looking forward are small, because inflation control can now be secured relatively quickly through the exchange rate effect on consumption prices. But the gains in output stabilisation are still considerable, because shorter forecast-horizon targeting induces larger real interest rate and in particular real exchange rate gyrations, with attendant output costs.

All in all, Chart 2 illustrates fairly persuasively the case for policy forward-lookingness. Using a forecast horizon of three to six quarters delivers far superior outcomes for output and inflation stabilisation to, say, current-period inflation-targeting. Largely, this is the result of transmission lags. Forecast-based rules are, in this sense, lag-encompassing. This also provides some empirical justification for the operational practice among inflation-targeting central banks of feeding back from inflation forecasts at horizons beyond one year.

Plainly, the optimal degree of policy forward-lookingness is sensitive to the model (and in particular the lag) specification. In the baseline model, this lag structure hinges critically on the assumed degree of stickiness in wage-setting. This stickiness in turn depends on the nature of wage-price contracting and on the degree of forward-lookingness in wage-bargaining. Given this, one way to interpret the need for forward-lookingness in policy is that it is serving to compensate for the backward-lookingness in wage-bargaining — whether directly through wage-bargaining behaviour or indirectly owing to the effect of contracting. In a sense, forward-looking monetary policy is acting, in a second-best fashion, to counter a backward-looking externality elsewhere in the economy. It is interesting to explore this notion further by considering the trade-off between the degree of backward-lookingness on the part of the private sector in the course of their wage-bargaining, and the degree of forward-lookingness on the part of the central bank in the course of its interest rate setting. [36]

(36) Equivalently, we could have looked at the effects of altering the length of wage contracting.

Chart 3 illustrates this trade-off. Point A in Chart 3 plots the most backward-looking aggregate (wage-setting plus policy-setting) outcome. The central bank feeds back from current inflation when setting policy ($j = 0$), and wage-bargainers assign a weight of only 0.1 to the next period's inflation rate when entering the wage bargain ($c_0 = 0.1$). This results in a very poor macroeconomic outcome, in particular for output variability. In hitting the inflation target, the central bank acts myopically. And the myopia of private sector agents then aggravates the effects of bad policy on the real economy through inflation-stickiness.

The solid line emanating from Point A traces out the locus of output/inflation variabilities as χ_0 rises from 0.1 to 0.9, so that wage-bargaining becomes progressively more forward-looking. Policy, for now, remains myopic ($j = 0$). In general, the upshot is a welfare improvement. With wages becoming a jump(ier) variable, even myopic policy can bootstrap inflation back to target following shocks. Moreover, wage-flexibility means that these inflation adjustments can be brought about at lower output cost. So both inflation and output variability are damped. Fully-flexible wages take us closer to a first-best. There is then little need for policy to have a forward-looking dimension.

The same is not true, of course, when wages embody a high degree of backward-lookingness. The dotted line in Chart 3 plots a j-locus with $c_0 = 0.1$. Though the resulting equilibria are clearly second-best in comparison with the forward-looking private sector equilibria, forward-looking monetary policy does now secure a significant improvement over the bad backward-looking equilibrium at Point A. In this instance, policy forward-lookingness is serving as a welfare-improving surrogate for the backward-looking behaviour of the private sector.

Finally, the two vertical lines in Chart 3, drawn at $j = 6$ and $c_0 = 0.3$, indicate degrees of *economy-wide* forward-lookingness beyond which the economy is unstable. For example, neither of the combinations $\{ j=6, c_0 = 0.4\}$ and $\{j = 7, c_0 = 0.3\}$ yields stable macroeconomic outcomes. This suggests that, just as a very backward-looking behavioural combination yields a bad equilibrium (point A), so too does a very forward-looking combination. It also serves notice of the potential instability problems of forecast-based rules. In general, policy forward-lookingness is only desirable as a second-best counterweight to the lags in monetary transmission. The first-best is for the lags themselves to shrink — for example, because private sector agents become more forward-looking. When this is the case, there is positive merit in the central bank itself not being too forward-looking, because that risks engendering instabilities.

Chart 4 illustrates the above points rather differently. It generalises the baseline model to accommodate forward-lookingness in the IS curve, following McCallum and Nelson (1997). Specifically, we set (somewhat arbitrarily) $a_1 = a_2 = 0.5$, so that the backward and forward-looking output terms in the IS curve are equally weighted. [37] The thick line in Chart 4 plots the j-locus in this modified model, with the dotted line showing the same for the baseline model.

The modified-model j-locus generally lies in a welfare-superior location to that under the baseline model, at least at short targeting horizons. For small j, both inflation and output variability are lower in the modified model. Increasing private sector forward-lookingness takes us nearer to the first-best. Policy forward-lookingness clearly still confers some benefits, since the modified-model j-locus moves initially to the south-west. But these benefits cease much beyond $j = 3$; and beyond $j = 6$ the system is explosive. So again, policy forward-lookingness is only desirable when used as a counterweight to the lags in monetary transmission, here reflected in the backward-looking behaviour of the private sector; it is not, of itself, desirable. The less this intrinsic sluggishness in the economy, the less the need for compensating forward-lookingness through monetary policy.

(b) Output-encompassing: output-stabilisation though inflation-targeting

Although the policy rule, **(1)**, contains no explicit output terms, it is already clear that inflation forecast-based rules are far from output-invariant. Chart 2 suggests that lengthening the targeting horizon up to

(37) McCallum and Nelson's (*op cit*) baseline model has $\{\alpha_1 = 0, \alpha_2 = 1\}$. That formulation is unstable in our model.

and beyond one year ahead can secure clear and significant improvements in output stabilisation. Judicious choice of the forecast horizon should allow the authorities, operating according to **(1)**, to select their preferred degree of output stabilisation.

That is not to say, however, that the output stabilisation embodied in policy rules such as **(1)** cannot be improved upon. For example, might not output stabilisation be further improved by adding explicit output gap terms to **(1)**? Chart 5 shows the effect of this addition. The dotted line simply redraws the full pass-through j-locus from Chart 2. The ray emerging from this line, starting from the base case horizon ($j = 8$) and moving initially to the south, plots outcomes from a rule that adds output gap terms to **(1)** with successively higher weights.(38) These weights, denoted λ, run from 0.1 to 8.(39)

Two main points are evident from Chart 5. First, adding explicit output terms to a forward-looking policy rule does appear to improve output stabilisation, with no costs in terms of inflation control — provided that the weights attached to output are sufficiently small. The ray moves due south for $0 < \lambda < 1$. Second, when $\lambda > 1$, some output/inflation variability trade-off does start to emerge, with improvements in output stabilisation coming at the cost of greater inflation variability. Indeed, for $\lambda > 2$, we begin to move in a north-easterly direction, with both output and inflation variability worsening. At $\lambda = 10$, the system is explosive. In general, though, Chart 5 seems to indicate that the addition of output gap terms to a forward-looking rule does yield clear welfare improvements for small enough λ. Put somewhat differently, it appears to suggest that an inflation forecast-based rule cannot synthetically recreate the degree of output stabilisation possible by targeting the output gap explicitly.

However, this conclusion ignores the fact that the feedback coefficient on expected inflation, θ, can also be altered and that this parameter itself influences output stabilisation. Chart 6 plots a set of j-loci, varying the

(38) The corresponding ray in the no pass-through case is very similar. So we stick here with the full pass-through base case.
(39) Weights much above 8 were found to generate instability — see below.

value of θ between 0.1 and 5.[40] Increasing θ tends to take us in a south-westerly direction — that is, it lowers both output and inflation variability.[41] Aggressive feedback responses are welfare-improving and, in particular, are output-stabilising. This reason is that agents factor this aggressiveness in policy response into their expectations when setting wages. Inflation expectations are thus less disturbed following inflation shocks. Inflation control, via this expectational mechanism, is thereby improved. And with inflation expectations damped following shocks, there is then less need for an offsetting response from monetary policy. As a consequence, output variability is also reduced by the greater aggressiveness in policy responses.[42]

The gains in inflation stabilisation are initially pronounced as θ rises above its 0.5 baseline value. These inflation gains cease — indeed, go into reverse — beyond $\theta \approx 1$. Thereafter, most of the gains from increasing θ show up in improved output stabilisation, usually at the expense of some destabilisation of inflation. The inflation forecast-based rule delivering lowest output variability is $\{j = 5, \theta = 5\}$. This gives a standard deviation of output $\sigma_y = 0.71\%$, and of inflation $\sigma_\pi = 1.32\%$.[43] So can *this* rule be improved upon by the addition of explicit output terms?

The answer, roughly speaking, is no. Adding an explicit output weight to the rule $\{j = 5, \theta = 5\}$ yields unstable outcomes. The ray of trajectories that results from adding output terms to other j-loci with smaller θ is shown in Chart 7. The gain in output stabilisation from adding explicit output terms seems to be very marginal. Moreover, it comes at the expense of a significant destabilisation of inflation. For example, the parameter triplet

(40) At values of $\theta > 5$, the system was again explosive.

(41) This is less clear for high values of θ ($\theta > 1$). The benefits then tend to be greater for output than for inflation stabilisation. Increasing θ also increases instrument variability from 0.27% to 1.35% , as θ moves from 0.1 to 5.

(42) Higher values of θ are not always welfare-enhancing. Larger values of θ also increase the diversity of macroeconomic outcomes at extreme values of j. For example, current-period inflation targeting ($j = 0$) leads to a very high output variance when θ is large. And when j is large, high values of θ increase the chances of explosive outcomes. For example, when $\theta = 5$, simulations are explosive beyond a five-quarter forecasting horizon.

(43) Output variability is then considerably lower than in the $\{j = 8, \theta = 0.5\}$ base case ($\sigma_y = 0.93\%$).

$\{j,\theta,\lambda\}$ delivering the lowest output variability is $\{j = 5, \theta = 4, l = 1\}$. This yields $\sigma_y = 0.69\%$ and $\sigma_\pi = 1.37\%$ — an output gain of only 0.02% points, and an inflation loss of 0.05% points, in comparison with the rule that gives no weight whatsoever to output, $\{j = 5, \theta = 5, \lambda = 0\}$.[44] It is clear that the optimal λ is now smaller even than in the earlier ($\theta = 0.5$) case. Any $\lambda > 1$ now takes us into unambiguously welfare-inferior territory. In forward-looking rules, there would seem to be benefits from placing a *higher* relative weight on expected inflation than on output. Indeed, to a first approximation, a weight of zero on output ($\lambda = 0$) comes close to being optimal.

Chart 7 suggests that there is, in effect, an output variability threshold at around $\sigma_y = 0.70\%$. None of the rules, with or without output gap terms, can squeeze output variability much beyond that threshold. By appropriate choice of $\{j,\theta\}$, inflation forecast-based rules appear capable of taking us very close to that threshold. Almost any amount of output smoothing can be synthetically recreated with an inflation-only rule. Forecast-based rules are, in this sense, output-encompassing. Inflation-nutters and output-junkies may disagree over the parameters in **(1)** — that is a question of policy tastes. But they need not differ over the arguments entering this rule — that is a question of policy technology.

(c) Information-encompassing: a comparison with alternative rules

Another of the supposed merits of an inflation forecast-based rule is that it embodies — and thus implicitly feeds back from — all information that is relevant for predicting the future dynamics of inflation. For this reason, it may approximate the optimal state-contingent rule. Certainly, by this reasoning, forward-looking rules should deliver outcomes at least as good as rules that feed back from a restrictive subset of information variables, such as output and inflation under the Taylor rule. These are empirically testable propositions.

---

(44) It also raises instrument variability from 1.8% to 1.92%.

To assess how close our forecast-based rule takes us to macroeconomic nirvana, we solve for the time-inconsistent *optimal* state-contingent rule in our system. This is the rule that solves the control problem:

$$\underset{r_t}{Min} \ \mathscr{L} = E_t \sum_{i=0}^{\infty} \beta^i \{\omega (\pi_{t+i}-\pi^*)^2 + (1-\omega) (y_{t+i}-y^*)^2 + \xi r_{t+i}{}^2\} \qquad \textbf{(10)}$$

where ω denotes the relative weight assigned to inflation deviations from target *vis-à-vis* output deviations from trend, and ξ is the weight assigned to instrument variability.

Because there are three arguments in the loss function, the easiest way to summarise the performance of the various rules relative to the optimal rule is by evaluating stochastic welfare losses ($\mathscr{L}$), having set common values for the preference parameters {β,ω,ξ}. We (somewhat arbitrarily) set β = 0.998, ω = 0.5 and ξ = 0.1. So inflation and output variability are equally weighted, and both are given higher weight than instrument variability. Table A then compares welfare losses from the optimal rule (OPT) with those from two specifications of the inflation forecast-based (IFB) rule (θ = 0.5 and θ = 5) for various values of j.(45) Table A also shows the standard deviation of output, inflation and (real) interest rates that results from each of these policy rule specifications.

Current-period inflation targeting (j = 0) clearly does badly by comparison with the optimal rule. For example, the rule {j = 0, θ = 0.5} delivers welfare losses that are 85% larger than the first-best. Inflation *forecast*-based rules clearly take us much closer — if not all the way — to that welfare optimum.(46) For example, {j = 6, θ = 0.5} delivers a welfare loss only 30% worse than the optimum. The optimal values of {j,θ} cannot be derived uniquely from Table A, since they clearly depend on the (arbitrary) values we have assigned to the preference parameters,{ω,ξ}, in the

(45) Where the optimal rule, the associated moments of output, inflation and the interest rate, and the value of the stochastic welfare loss are calculated using the 'OPT' routine of the ACES/PRISM solution package. See footnote 24.

(46) As we discuss below, altering the smoothing parameter, γ, takes us nearer still to the first-best.

objective function. But for our chosen preference parameters, the best forecast horizon appears to lie between three and six periods, irrespective of the value of θ.

We can also compare these forward-looking rules with a variety of simple, backward-looking Taylor-type formulations, which feed back from contemporaneous or lagged values of output and inflation. In particular, for comparability with the other studies in this volume, we consider two types of rule:

$$r_t = a\, p_t + b\,(y_t - y_t^*) + c\, r_{t-1} \qquad \textbf{(11)}$$

$$r_t = a\, p_{t-1} + b\,(y_{t-1} - y_{t-1}^*) + c\, r_{t-1} \qquad \textbf{(12)}$$

for a variety of values of $\{a,b,c\}$ listed below.(47) We classify the first T1$\{a,b,c\}$ rules, and the second T2$\{a, b, c\}$ rules. The rule T1$\{a = 0.5, b = 0.5, c = 0\}$ is of course the well-known Taylor rule. A comparison of these rules with the optimal rule (OPT) and forecast-based specifications (IFB) is given in Table B.

We draw several general conclusions from Table B. First, looking only at the performance of the backward-looking rules, it appears that placing a higher relative weight on output than on inflation yields welfare improvements. This differs from our result with forward-looking rules. Second, because they are based on an inferior (time t-1) information set, the T2 rules do worse than the T1 rules. The difference in welfare losses is not, however, very great. This suggests that, at least over the course of one quarter, information lags do not impose much of a welfare cost. Third, both rules yield unstable outcomes in our model when there is a small weight on output ($b < 0.1$) and a large weight on smoothing ($c > 1$). Higher weights on output ($b > 0.5$) or lower weights on smoothing ($c < 1$) are necessary to deliver a stable equilibrium. Fourth, even the best-performing backward-looking rule - interestingly, the Taylor rule —

(47) One difference from the other exercises is that here the policy instrument is the short-term real (rather than nominal) interest rate. This should not affect the relative performance of the rules. But we have subtracted one from the inflation parameter, a, when simulating the backward-looking policy rules, to ensure comparability with the other studies.

delivers a welfare outcome almost 50% worse than the optimum. By comparison, the best forward-looking rule delivers a welfare loss that is around 30% worse than the optimum.

The final conclusion is evidence of the information-encompassing nature of inflation forecast-based rules. A forward-looking rule conditions on all variables that affect future inflation and output dynamics, not just output and inflation themselves. In the context of our simple open-economy model, important additional state variables are (lagged values of) the exchange rate, as well as additional lags of wages and prices. Just as the optimal feedback rule conditions on these state variables, so too will inflation forecast-based rules. That is not a feature shared by Taylor rules. In larger models than the one presented here, these extra conditioning variables would include those other information variables affecting future inflation dynamics, such as (lagged) asset and commodity prices. These variables will be captured in forward-looking rules, but not in Taylor-type specifications. In general, the larger the model, the more diffuse will be the information sets.[48] The welfare differences between forward and backward-looking rules are thus also likely to be larger in these bigger models. So while inflation forecast-based rules cannot take us all the way to the first-best, in general they seem likely to take us further in that direction than Taylor-type specifications, at the same time as retaining the simplicity and transparency of the Taylor rules.

(d) Other policy parameters

Finally, we explore two further design features of inflation forecast-based rules. First, what is the preferred degree of interest rate smoothing, γ, in such a rule? And second, how does a regime of price-*level* targeting compare with the inflation-targeting specifications considered so far? On *interest rate smoothing*, the thick line in Chart 8 replots the j-locus from the baseline rule. The (dotted line) rays emanating from this at $j = \{3,6,9\}$ periods illustrate how output/inflation variabilities are affected as γ varies between zero (no-smoothing) and one. These rays are almost

(48) This is, for example, what Black *et al* (1996) find when simulating the larger-scale Bank of Canada QPM model.

horizontal. Instrument smoothing delivers greater inflation stability, with relatively few countervailing output costs. For example, inflation variability is lowered by 33% when moving from $\gamma = 0$ to $\gamma = 1$, for $\{ j = 6, \theta = 0.5\}$. This arises because rules with higher degrees of smoothing deliver more persistent interest rate responses. These policy responses in turn have a larger impact effect on the exchange rate — and hence on inflation itself.[49] This sharper inflation control comes at some output cost, though our simulations suggest that this cost is fairly small. The benefits of instrument smoothing are smaller (and potentially trivial) at higher values of θ, however, because policy aggressiveness does the same job as instrument persistence in improving inflation control.

If we evaluate welfare losses using the earlier parameterisation of the loss function, then the no-smoothing rule $\{ \gamma = 0, j = 6, \theta = 0.5\}$ delivers a welfare loss that is 14% higher than that from the high-smoothing rule $\{\gamma = 1, j = 6, \theta = 0.5\}$. Indeed, the latter rule now takes us within 25% of the optimal rule. So it seems in general that relatively high degrees of interest rate smoothing are welfare-enhancing, but that the extent of this welfare improvement may be small if policy is already aggressive.

On *price-level targeting*, our baseline rule now takes the modified form:

$$r_t = \gamma r_{t-1} + (1-\gamma) r_t^* + \theta [E_t(p^c_{t+j}) - p^{c*}] \qquad \textbf{(13)}$$

Monetary policy now shoots for a deterministic price-level path, p^{c*}, which we again normalise to zero (in logs). Using the baseline model and the parameter settings $\{ \gamma = 0, \theta = 0.5\}$.[50] Chart 9 plots the j-locus that results from the price-level rule, **(13)**. The baseline inflation forecast-based rule, **(1)**, is also shown for comparison (as the dotted line). For most values of j, the price-level targeting rule delivers welfare-inferior outcomes to the inflation-targeting rule: both output and inflation variability are higher. This is particularly true of short-horizon (for example, current-period) price-level targeting. Other studies have also found this to be the case (Duguay (1994), Fillion and Tetlow (1993), Lebow, Roberts and Stockton

(49) This is even true — though to a lesser extent — in the no pass-through case.
(50) Higher values than this tended to be unstable.

(1992), Haldane and Salmon (1995)). Nevertheless, for large enough *j*, price-level targeting rules still perform little worse (and in some cases perhaps better) than inflation-targeting rules.

Moreover, this comparison may unfairly disadvantage price-level targeting. The baseline model still embodies a relatively high degree of *inflation* persistence. It is questionable whether such persistence would survive the move to a monetary regime that delivered price-level stationarity. In that situation, *price-level* persistence might be a more realistic specification of price dynamics. In the context of our model, wage-contracting might then be better characterised by a conventional Taylor staggered contract wage specification, rather than the Fuhrer/Moore formulation we have used so far.(51) That is, the contracting equation, **(6)**, would be replaced by:

$$w_t = c_0\, E_t(w_{t+1}) + (1 - c_0)\, w_{t-1} + c_1\, (y_t - y_t^*) + e_{4t} \qquad \textbf{(14)}$$

and the Phillips curve equivalent of **(8)** would now be:

$$p_t = E_t\,(p_{t+1}) + c_1(y_t + y_{t-1}) - r\, [E_t\, Dc_{t+1})] + e_{5t} \qquad \textbf{(15)}$$

where $\rho \equiv (1-\phi)/\phi$. Inflation no longer depends on lagged values; it is a jump variable.

The thick line in Chart 10 plots the *j*-locus for the price-level policy rule, **(13)**, with **(15)** now replacing **(8)** in the model. This locus clearly lies to the south of the *j*-locus under inflation-targeting using the baseline model (dotted line). Price-level targeting now does as good (or better) a job in stabilising output as inflation-targeting. This is the result of the increased flexibility in prices. Inflation variability remains higher than under some specifications of inflation-targeting, but never excessively so. In sum, even the *short-term* output/inflation variability costs of price-level targeting appear to be much less pernicious than may have typically been thought likely, under certain parameterisations of the underlying model and policy

---

(51) Though, in principle, the relative wage formulation of Fuhrer and Moore (1995a) is meant to be a structural relationship, and thus immune to the Lucas critique.

rule and assuming perfect credibility of such a regime.[52] For a comprehensive welfare-theoretic comparison, the longer-term benefits of a price-level standard would need to be set against these (potential) short-term costs.

5. Conclusions

It is widely recognised that monetary policy needs a forward-looking dimension. Inflation-targeting countries have explicitly embodied that notion in the design of their forecast-based policy rules. In principle, these rules confer some real benefits: they embody explicitly transmission lags (*lag-encompassing*); they potentially embody all information useful for predicting future inflation (*information-encompassing*); and, suitably designed, they can achieve a degree of output smoothing (*output-encompassing*). This paper has evaluated quantitatively these features of an inflation forecast-based rule using simulation techniques. Our main conclusions are:

(a) On lag-encompassing, an inflation forecast horizon of three to six quarters appears to deliver the best outcomes. Shorter horizons than this risk raising both output and inflation variability — the result of policy lags — while longer horizons risk macroeconomic instability. In general, the greater the degree of forward-lookingness on the part of the private sector, the less the compensating need for forward-lookingness by the central bank. These results support the forecast-based approach to monetary policy-making pursued by inflation targeting central banks in practice.

(b) An inflation forecast-based rule, with an appropriately chosen targeting horizon, naturally embodies a degree of output-stabilisation. Moreover, any degree of output-smoothing can be synthetically recreated by judicious choice of the parameters entering an inflation forecast-based rule. There is no need for any *explicit* output terms to enter this rule. That is evidence of the output-encompassing nature of inflation targeting based around inflation forecasts;

(52) Williams (1997) and Black *et al* (1997) reach similar conclusions in their studies of the United States and Canada respectively. In a theoretical context, Svensson (1996) also argues that price-level targeting need not raise output/inflation variabilities.

(c) While not taking us all the way to the welfare optimum, forecast-based rules do seem capable of securing welfare outcomes superior to backward-looking specifications, which have been the mainstay in the literature to date. That is evidence of the information-encompassing nature of forecast-based policy rules.

We have also evaluated forecast-based price-level rules for monetary policy. Under certain parameterisations, they perform creditably even as a short-run macro-stabiliser. So soon after having secured low inflation, there is an understandable caution about pursuing something seemingly as new as a price-level standard. Inflation-targeting is an embryonic monetary framework, whose performance has yet to be properly evaluated. But price-level targeting, indubitably, is not.

Table A

Comparing optimal (OPT) and inflation forecast-based (IFB{j,q}) rules[(a)] (standard deviation (σ) in %)

	s_y	s_p	s_r	$\mathcal{L}$
OPT rule	0.782	1.103	1.033	41.83
IFB{*j* = 0, *q* = 0.5}	1.52	1.199	0.925	76.37
IFB{*j* = 3, *q* = 0.5}	1.07	1.17	0.61	52.61
IFB{*j* = 6, *q* = 0.5}	0.91	1.34	0.51	54.18
IFB{*j* = 9, *q* = 0.5}	0.94	1.57	0.40	68.04
IFB{*j* = 0, *q* = 5.0}	8.86	1.49	10.33	755.8
IFB{*j* = 5, *q* = 5.0}	0.716	1.32	1.34	53.91

(a) With the value of the smoothing parameter, $\gamma = 0.5$.

Table B

Comparing optimal (OPT), inflation forecast-based (IFB{j,q}) and Taylor (T1/T2{a,b,c}) rules[(a)]

(standard deviation (s) in %)

	s_y	s_p	s_r	$\mathcal{L}$
OPT rule	0.78	1.10	1.03	41.83
IFB{*j* = 6, *q* = 0.5}	0.91	1.34	0.51	54.18
IFB{*j* = 5, *q* = 5.0}	0.72	1.32	1.34	53.91
T1{*a* = 2, *b* = 0.8, *c* = 1}	1.84	0.94	1.79	92.69
T1{*a* = 0.2, *b* = 1, *c* = 1}	0.86	1.56	0.99	68.22
T1{*a* = 0.5, *b* = 0.5, *c* = 0}	1.05	1.38	0.55	61.96
T1{*a* = 0.5, *b* = 1, *c* = 0}	0.92	1.46	0.72	61.97
T1{*a* = 0.2, *b* = 0.06, *c* = 1.3}		Unstable		
T2{a = 2, *b* = 0.8, *c* = 1}	2.24	1.02	2.44	130.9
T2{*a* = 0.2, *b* = 1, *c* = 1}	1.11	1.58	1.40	82.44
T2{*a* = 0.5, *b* = 0.5, c = 0}	1.11	1.38	0.56	64.48
T2{*a* = 0.5, *b* = 1, *c* = 0}	0.99	1.44	0.76	64.21
T2{*a* = 0.3, *b* = 0.08, *c* = 1.3}		Unstable		

(a) With the value of the smoothing parameter, $\gamma = 0.5$.

Chart 1

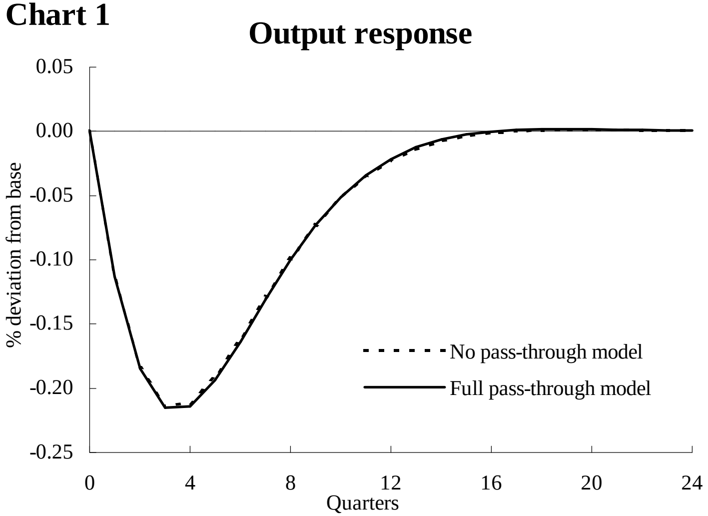

Output response
0.05
0.00
-0.05
-0.10
-0.15
-0.20
-0.25
% deviation from base
0
4
8
12
16
20
24
Quarters
No pass-through model
Full pass-through model

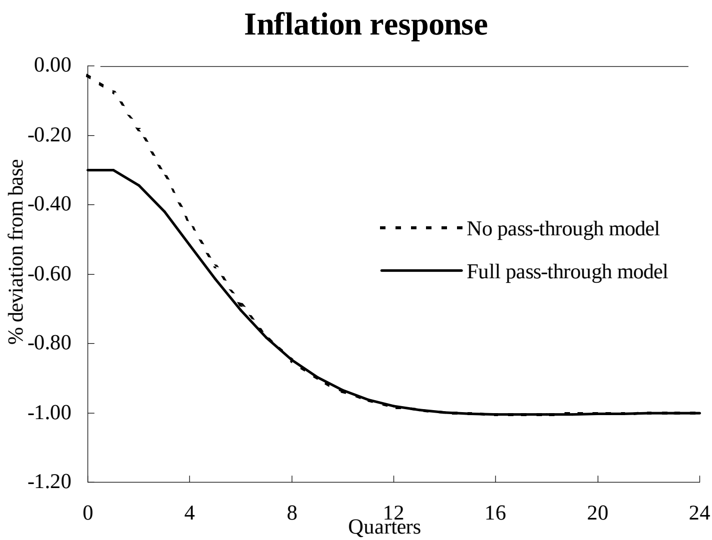

Inflation response
0.00
-0.20
-0.40
-0.60
-0.80
-1.00
-1.20
% deviation from base
0
4
8
12
16
20
24
Quarters
No pass-through model
Full pass-through model

Chart 2

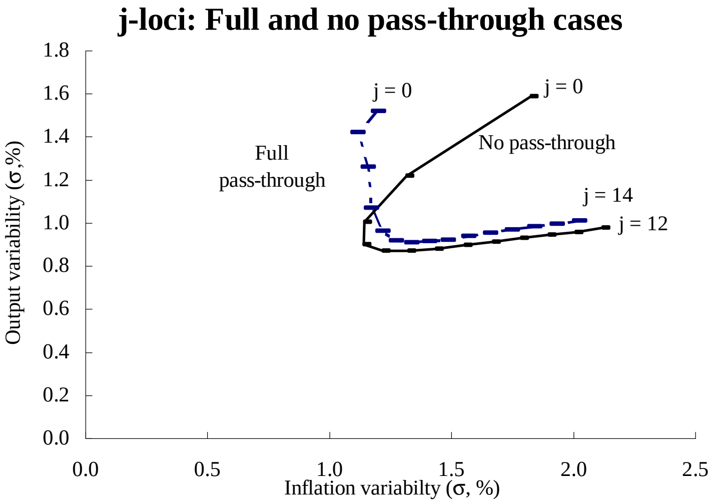

Chart 3

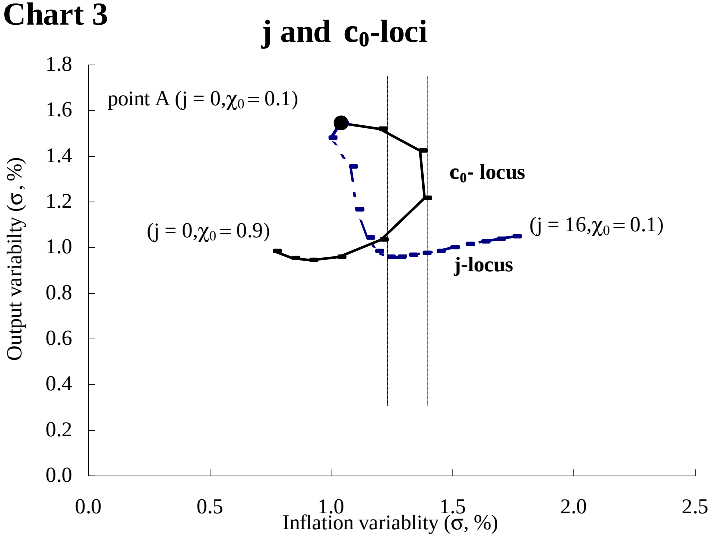

Chart 4

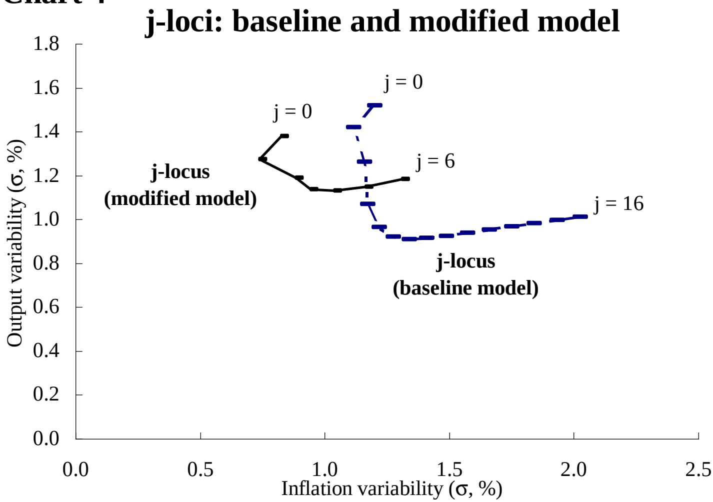

Chart 5

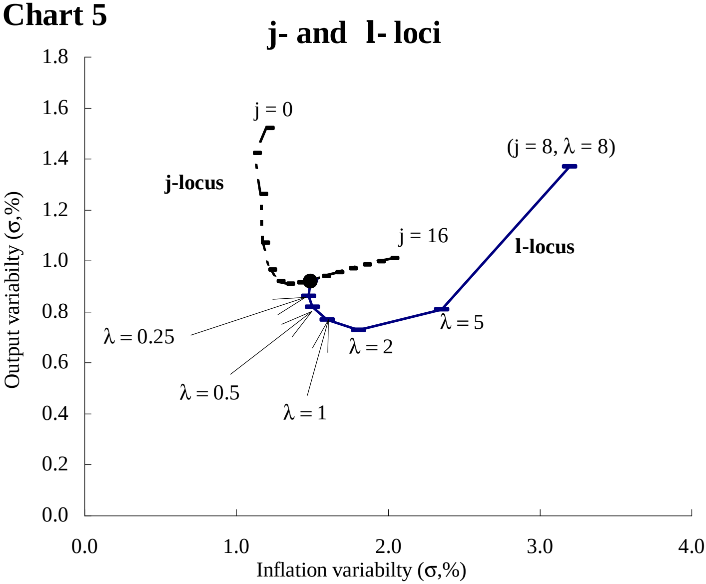

Chart 6

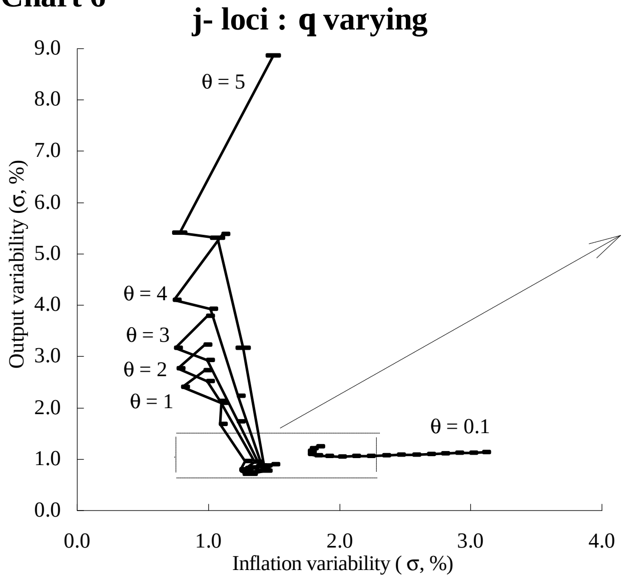
j- loci : q varying
θ = 5
θ = 4
θ = 3
θ = 2
θ = 1
θ = 0.1
9.0
8.0
7.0
6.0
5.0
4.0
3.0
2.0
1.0
0.0
0.0
1.0
2.0
3.0
4.0
Output variability (σ, %)
Inflation variability (σ, %)

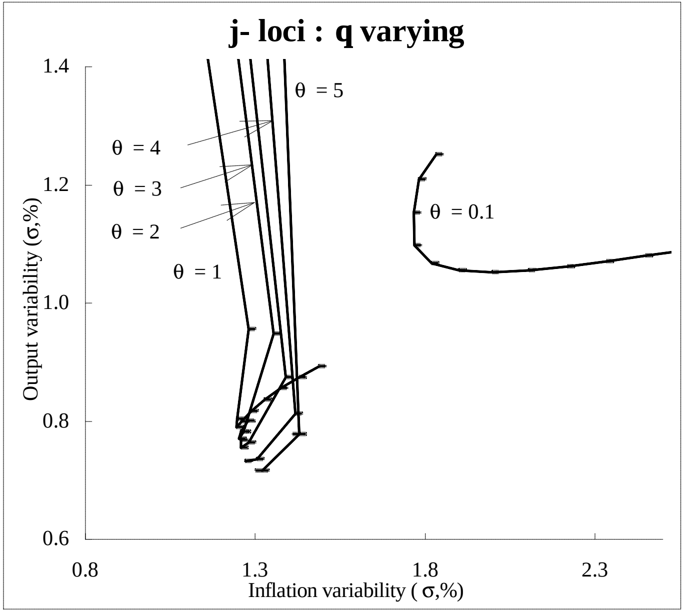
j- loci : q varying
θ = 5
θ = 4
θ = 3
θ = 2
θ = 1
θ = 0.1
1.4
1.2
1.0
0.8
0.6
0.8
1.3
1.8
2.3
Output variability (σ,%)
Inflation variability (σ,%)

Chart 7

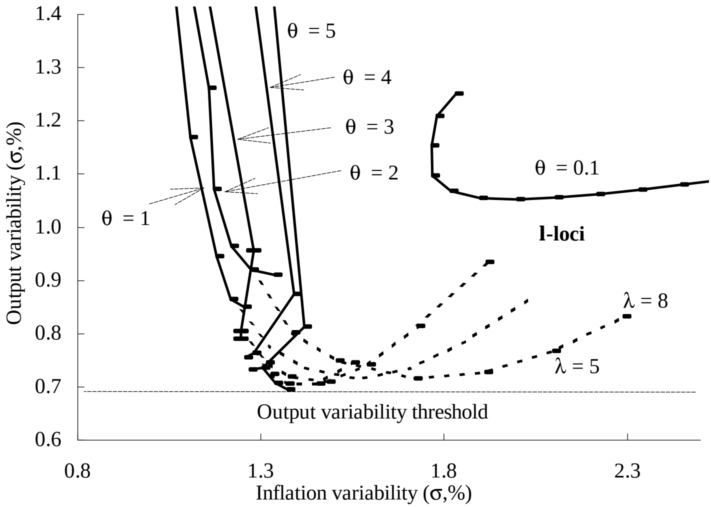

Chart 8

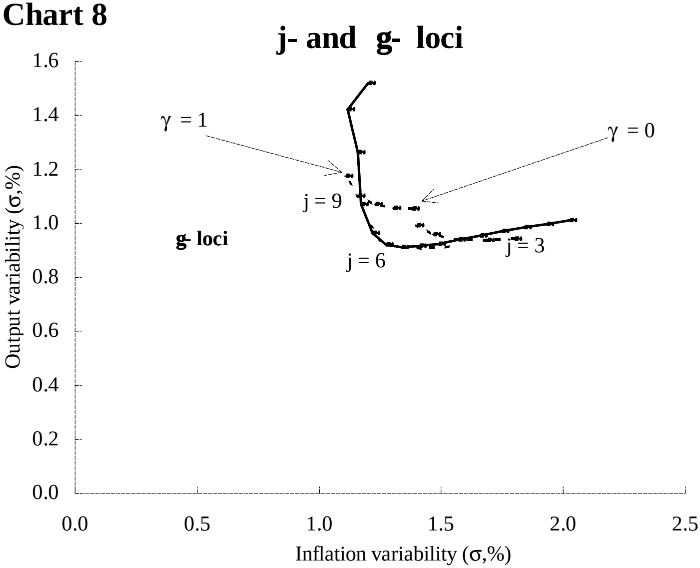

Chart 9

Inflation and price-level targeting

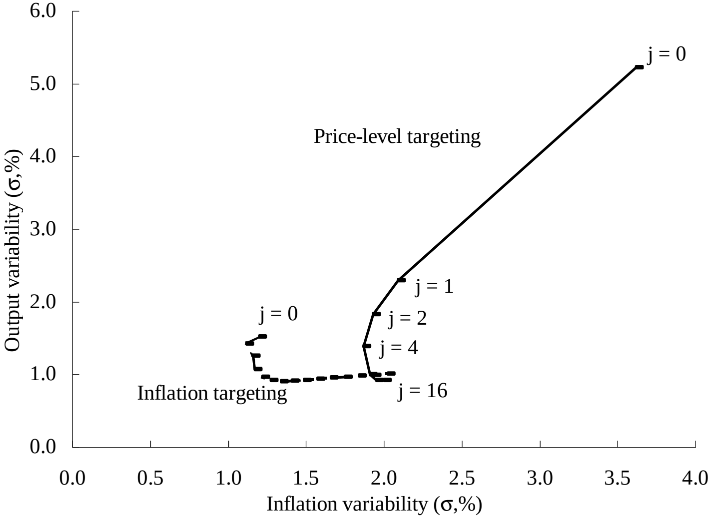

Chart 10

Inflation and price-level targeting

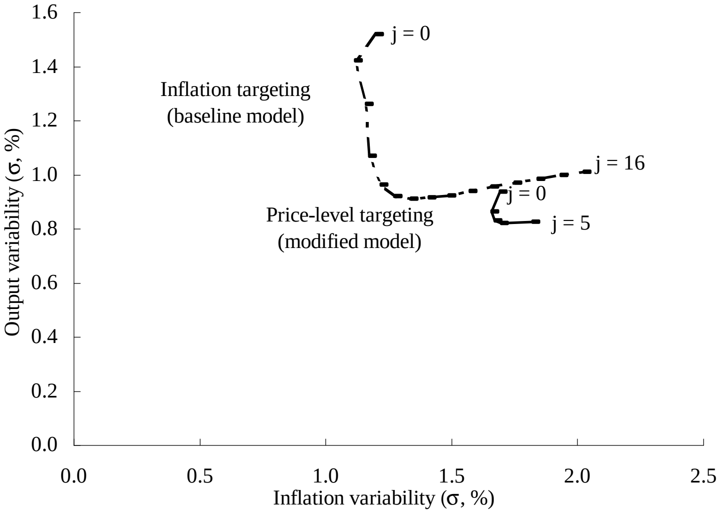

References

Bakhshi, H, Haldane, A G and Hatch, N (1998), 'Some Costs and Benefits of Price-Stability in the United Kingdom', in Feldstein, M (ed), *The Costs and Benefits of Going from Low Inflation to Price Stability*, forthcoming.

Ball, L (1997), 'Policy Rules for Open Economies', in Taylor, J B (ed) (1998), *Monetary Policy Rules*, NBER Conference volume, forthcoming.

Bernanke, B S and Mishkin, F S (1997), 'Inflation Targeting: A New Rule for Monetary Policy?', *Journal of Economic Perspectives*, Vol 11, pages 97-116.

Bernanke, B S and Woodford, M (1997), 'Inflation Forecasts and Monetary Policy', *Journal of Money, Credit and Banking*, Vol 29, pages 653-84.

Black, R, Macklem, T and Rose, D (1997), 'On Policy Rules for Price Stability', in *Price Stability, Inflation Targets and Monetary Policy*, Bank of Canada.

Blake, A (1996), 'Forecast Error Bounds by Stochastic Simulation', *National Institute Economic Review*, May, pages 72-8.

Blake, A P and Westaway, P F (1996), 'Credibility and the Effectiveness of Inflation Targeting Regimes', *The Manchester School*, Vol LXIV, pages 28-50.

Blanchard, O J and Kahn, C M (1980), 'The Solution of Linear Difference Models Under Rational Expectations', *Econometrica*, Vol 48, pages 1305-11.

Blinder, A S (1997), 'What Central Bankers Could Learn from Academics - and Vice Versa', *Journal of Economic Perspectives*, Vol 11, pages 3-19.

de Brouwer, G and O'Regan, J (1997), 'Evaluating Simple Monetary-Policy Rules for Australia', in P Lowe (ed), *Monetary Policy and Inflation Targeting*, Reserve Bank of Australia.

Bryant, R C, Hooper, P and Mann, C (1993), *Evaluating Policy Regimes*, Washington DC, The Brooking Institution.

Buiter, W H and Jewitt, I (1981), 'Staggered Wage Setting with Real Wage Relativities: Variations on a Theme of Taylor', *The Manchester School*, Vol 49, pages 211-28.

Cecchetti, S G (1995), 'Inflation Indicators and Inflation Policy', *NBER Working Paper No 5161*.

Clarida, R and Gertler, M (1997), 'How the Bundesbank Conducts Monetary Policy', in C Romer and D Romer (eds), *Reducing Inflation*, Chicago: University of Chicago Press.

Clarida, R, Gali, J and Gertler, M (1997), 'Monetary Policy Rules in Practice: Some International Evidence', paper prepared for the ISOM Conference.

Clark, P, Laxton, D and Rose, D (1995), 'Capacity Constraints, Inflation and the Transmission Mechanism: Forward-Looking Versus Myopic Policy Rules', *IMF Working Paper No 75.*

Dale, S and Haldane, A G (1995), 'Interest Rates and the Channels of Monetary Transmission: Some Sectoral Estimates', *European Economic Review,* Vol 39, pages 1611-26.

Duesenberry, J S (1949), *Income, Savings, and the Theory of Consumer Behaviour*, Cambridge: Harvard University Press.

Duguay, P (1994), 'Some Thoughts on Price Stability Versus Zero Inflation', *mimeo,* Bank of Canada.

Fillion, J-F and Tetlow, R (1993), 'Zero-Inflation or Price-Level Targeting? Some Answers from Stochastic Simulation on a Small Open-Economy Macro Model', in *Economic Behaviour and Policy Choice Under Price Stability*, Bank of Canada.

Freedman, C (1996), 'What Operating Procedures Should Be Adopted to Maintain Price Stability?', in *Achieving Price Stability*, Federal Reserve Bank of Kansas City.

Friedman, B M (1975), 'Targets, Instruments and Indicators of Monetary Policy', *Journal of Monetary Economics*, Vol 1, pages 443-73.

Friedman, B M and Kuttner, K N (1996), 'A Price Target for US Monetary Policy? Lessons from the Experience with Money Growth Targets', *Brookings Papers on Economic Activity*, Vol 1, pages 77-146.

Friedman, M (1959), *A Program for Monetary Stability*, New York: MacMillan.

Friedman, M and Schwartz, A J (1982), *Monetary Trends in the United States and the United Kingdom: Their Relation to Income, Prices and Interest Rates, 1867-1975*, University of Chicago Press: Chicago.

Fuhrer, J C (1997), 'The (Un)Importance of Forward-Looking Behaviour in Price Specifications', *Journal of Money, Credit and Banking*, Vol 29, No 3, pages 338-50.

Fuhrer, J C and Moore, G R (1995a), 'Forward-Looking Behaviour and the Stability of a Conventional Monetary Policy Rule', *Journal of Money, Credit and Banking*, Vol 27, pages 1,060-70.

Fuhrer, J C and Moore, G R (1995b), 'Inflation Persistence', *Quarterly Journal of Economics*, Vol 110, pages 127-59.

Gaines, J, Al'Nowaihi, A and Levine, P (1989), ACES package for 'The Solution of One-Country Linear Rational Expectations Models', *mimeo*, London Business School.

Goodfriend, M and King, R G (1997), ‘The New Neoclassical Synthesis and the Role of Monetary Policy’, *NBER Macroeconomics Annual 1997*, pages 231-83.

Haldane, A G (1995), *Targeting Inflation*, Bank of England.

Haldane, A G and Salmon, C K (1995), ‘Three Issues on Inflation Targets’, in Haldane, A G (ed), *Targeting Inflation*.

Haldane, A G (1997), ‘Designing Inflation Targets’, in P Lowe (ed.), *Monetary Policy and Inflation Targeting*, Reserve Bank of Australia.

Haldane, A G (1998), ‘On Inflation-Targeting in the United Kingdom’, *Scottish Journal of Political Economy*, Vol 45, No 1, pages 1-32.

Keynes, J M (1923), *A Tract on Monetary Reform*, MacMillan: London.

King, M A (1996), ‘Direct Inflation Targets’, paper presented at the symposium on *Geldpolitische in Europe*, Frankfurt, April 1996.

King, M A (1997), ‘The Inflation Target Five Years On’, *Bank of England Quarterly Bulletin*, Vol 37, pages 434-42.

Kohn, D (1995), ‘Comment’, in B S Bernanke and J J Rotemberg (eds.), *NBER Macroeconomics Annual*, MIT Press, Cambridge, Massachusetts, pages 227-33.

Lebow, D E, Roberts, J M and Stockton, D J (1992), ‘Economic Performance Under Price Stability’, *Board of Governors of the Federal Reserve System Working Paper No 125.*

Leeper, E M, Sims, C A and Zha, T (1996), ‘What Does Monetary Policy Do?’, *Brooking Papers on Economic Activity*, Vol 2, pages 1-63.

Leiderman, L and Svensson, L E O (1995), *Inflation Targets*, CEPR.

Longworth, D and Freedman, C (1995), 'The Role of the Staff Economic Projection in Conducting Canadian Monetary Policy', in *Targeting Inflation*, A G Haldane (ed), Bank of England.

Mayes, D and Chapple, B (1995), 'Defining an Inflation Target', in *Targeting Inflation*, A G Haldane (ed), Bank of England.

McCallum, B T (1988), 'Robustness Properties of a Rule for Monetary Policy', *Carnegie-Rochester Conference Series on Monetary Policy*, Vol 29, pages 172-203.

McCallum, B T and Nelson, E (1997), 'An Optimising IS-LM Specification for Monetary Policy and Business Cycle Analysis', *NBER Working Paper No 5875*.

Orphanides, A (1997), 'Monetary Policy Rules Based on Real-Time Data', *mimeo*, Board of Governors of the Federal Reserve System.

Reifschneider, D L, Stockton, D J and Wilcox, D W (1996), 'Econometric Models and the Monetary Policy Process', *mimeo*, Board of Governors of the Federal Reserve System.

Reserve Bank of New Zealand (1997), *Monetary Policy Statement 1997*, Wellington: New Zealand.

Roberts, J M (1995), 'New Keynesian Economics and the Phillips Curve', *Journal of Money, Credit and Banking*, Vol 27, pages 975-84.

Rotemberg, J J and Woodford, M (1997), 'An Optimisation-Based Econometric Framework for the Evaluation of Monetary Policy', *NBER Macroeconomics Annual 1997, forthcoming*.

Rudebusch, G and Svensson, L E O (1997), 'Policy Rules for Inflation Targeting', in Taylor, J B (ed) (1998), *Monetary Policy Rules*, NBER Conference volume, forthcoming.

Sargent, T J and Wallace, N (1975), ‘‘Rational’ Expectations, the Optimal Monetary Instrument and the Optimal Money Supply Rule’, *Journal of Political Economy*, Vol 83, pages 241-54.

Svensson, L E O (1996), ‘Price Level Targeting Versus Inflation Targeting: A Free Lunch?’, *NBER Working Paper No 5719*.

Svensson, L E O (1997a), ‘Inflation-Forecast-Targeting: Implementing and Monitoring Inflation Targets’, *European Economic Review*, Vol 41, pages 1,111-46.

Svensson, L E O (1997b), ‘Inflation Targeting: Some Extensions’, *NBER Working Paper No 5962.*

Svensson, L E O (1997c), ‘Open-Economy Inflation-Targeting’, *mimeo*, Institute for International Economics.

Taylor, J B (1979), ‘Estimation and Control of a Macroeconomic Model with Rational Expectations’, *Econometrica*, Vol 47, pages 1,267-86.

Taylor, J B (1980), ‘Aggregate Dynamics and Staggered Contracts’, *Journal of Political Economy*, Vol 88, pages 1-24.

Taylor, J B (1993a), ‘Discretion Versus Policy Rules in Practice’, *Carnegie-Rochester Conference Series on Public Policy*, Vol 39, pages 195-214.

Taylor, J B (1993b), *Macroeconomic Policy in a World Economy*, New York: W W Norton.

Taylor, J B (ed) (1998), *Monetary Policy Rules*, NBER Conference volume, forthcoming.

Thomas, R S (1996), ‘Understanding Broad Money’, *Bank of England Quarterly Bulletin*, Vol 36, pages 163-79.

Williams, J (1997) '**Simple Rules for Monetary Policy**', paper presented at the NBER Summer Institute.

Woodford, M (1994), 'Nonstandard Indicators for Monetary Policy: Can Their Usefulness Be Judged from Forecasting Regressions?**'**, in Gregory N. Mankiw (ed), *Monetary Policy*, University of Chicago Press, Chicago.