Hybrid inflation and price level targeting

*Nicoletta Batini**
and
*Anthony Yates***

* Research Adviser, MPC Unit, Bank of England.
E-mail: nicoletta.batini@bankofengland.co.uk

** Conjunctural Assessment and Projections Division, Monetary Analysis, Bank of England.
E-mail: tony.yates@bankofengland.co.uk

The views expressed are those of the authors and do not necessarily reflect those of the Bank of England. We thank Mervyn King, Spencer Dale and Edward Nelson for the conversations out of which this paper arose, and for their insightful comments on earlier drafts. We also thank Larry Ball, Frank Smets, seminar participants at the Bank of England and at the Federal Reserve Board and two anonymous referees for helpful comments on earlier versions of this paper. Remaining errors and the views expressed herein are those of the authors and not of the Bank of England nor of the Bank of England's Monetary Policy Committee.

Copies of working papers may be obtained from Publications Group, Bank of England, Threadneedle Street, London EC2R 8AH; telephone 020 7601 4030, fax 0207 601 3298, e-mail mapublications@bankofengland.co.uk

Working papers are also available at www.bankofengland.co.uk/workingpapers/index.htm

The Bank of England's working paper series is externally refereed.

ISSN 1368-5562

Contents

Abstract

The previous literature on the benefits of price level versus inflation targeting has, with some qualifications, established that price level targeting entails lower price level variance at the expense of higher inflation and output variance. In this paper we investigate the properties of monetary regimes that combine price level and inflation targeting. We offer two characterisations of these regimes: a set of optimal control rules obtained assuming that policy-makers minimise a loss function that penalises a mixed price level/inflation target; and a set of simple rules feeding back from alternative combinations of (current and future-dated) price level and inflation deviations from target. We derive asymptotic variances of the price level, inflation and output associated with each of these regimes when the economy is modelled as a small-scale open-economy RE model calibrated on UK data. We conclude that: (i) the relative merits of price level and inflation targeting, as well as of mixes of these two, are a function of several modelling and policy assumptions; and (ii) these merits do not change monotonically as we move from one regime to another. It appears also that the probability of nominal interest rates hitting a 'zero bound' under the alternative regimes is model-specific and varies non-monotonically among them.

JEL classification: E52; E37; E58.
Key words: price level targeting, inflation targeting, price stability, optimal policy rules, simple feedback rules.

Summary

The success and spread of inflation targeting (documented by, for example, Julius *et al* (2000)) has stimulated interest in the merits of price level targeting. Under inflation targeting, the expected variance of the price level increases without bound as we look further into the future; under price level targeting, policy acts to reverse shocks to the price level, and the expected variance is constant.

A literature has grown examining the benefits of price level versus inflation targeting, including Lebow *et al* (1992), Fillion and Tetlow (1993), Haldane and Salmon (1995), Black *et al* (1997), Kiley (1998), Svensson (1999a), Smets (2000), Williams (1999), Vestin (1998), and Dittmar *et al* (1999). In the early days of this research effort it was thought that while price level targeting meant lower price level variance, it brought with it the cost of higher variance in inflation (as, for example, below-target misses are inflated back next period) and, in worlds of sticky prices, a greater volatility of output about the natural rate. But more recently, exceptions to this early result have been uncovered.

This paper contributes by first describing and then analysing the consequence of regimes that can be thought to lie 'in between' the extremes of price level and inflation targeting. We describe two ways of characterising the spectrum of regimes. The first is a spectrum of regimes that come from computing optimal rules subject to loss functions that have different relative weights on price and inflation deviations from target. The second spectrum is defined by a set of simple rules where, at one extreme, the real interest rate responds to forecast deviations of prices from target, and at the other from forecast deviations of inflation from target (and a term in the output gap). In between, policy responds to forecast deviations of prices from a moving price level target.

We compute inflation, price and output variability when these rules are followed using a calibrated, rational expectations model of the United Kingdom used in Batini and Haldane (1999).

The paper shows that these 'intermediate' regimes are interesting in that inflation, output and price level variance do not change monotonically as we move from one extreme to another. We also show that the cost benefit analysis of regimes along our spectra depend, not surprisingly, on the degree of forward-lookingness embodied in price-setting, and contrast results obtained using a form of nominal stickiness akin to that in Taylor (1980) on the one hand and Fuhrer and Moore (1995) on the other. We also use our results on how the variance of nominal interest rates changes along the regime spectrum to comment on the probability of hitting a zero band associated with different policies.

1 Introduction

That inflation targeting by central banks is so widespread—55 out of 91 countries in the Julius *et al* (2000) survey follow an explicit inflation target—is testament to the consensus view that price stability, somehow defined, brings with it benefits. As Wicksell put it in his 1935 *Lectures on Political Economy*, '[s]o soon as money becomes a general measure of value, the avoidance of all violent and unexpected fluctuations in its value is of the utmost importance'. There is perhaps less of a consensus as to what, in practice, 'price stability' means.

One particular focus of debate is whether monetary authorities should choose paths for the price level, or for the inflation rate. Irving Fisher's (1922) original proposal for 'The regulation of the value of money' was to stabilise the price level. The first known example of an explicit target for price stability was also in terms of the price level, in Sweden in the 1930s (Jonung (1979), Berg and Jonung (1998)). Milton Friedman's celebrated prescription for optimal monetary policy, however, was in terms of inflation (actually deflation), and modern-day definitions of price stability bear more resemblance to this proposal (see Haldane (ed) (1995), and Svensson (1995), and Bernanke *et al* (1999)).

The debate has perhaps been given new life by the remarkable success in eliminating high inflation across the industrialised world by central banks (see, among others, Haldane (*op cit*) and Bernanke *et al* (*op cit*).[(1)] Chart 1, showing an average of the annualised quarterly CPI inflation rates across a core group of inflation-targeting countries for the period 1985 Q1-1999 Q1, quantifies this success.[(2)] Average inflation performance has improved dramatically in these countries since the early 1990s, with inflation rates falling, in some individual cases, to below 1%.

Chart 1: Inflation performance 1985 Q1-1999 Q1

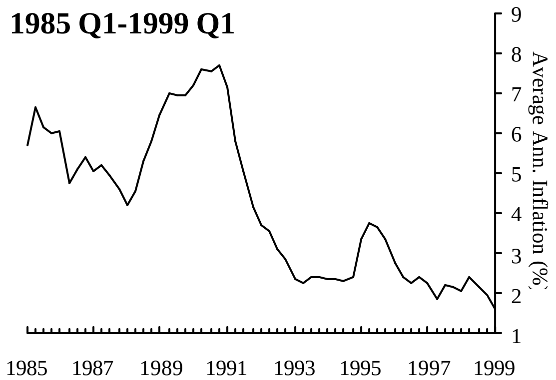

In practice, the difference between the price level and inflation-targeting regimes is that under 'price level targeting', the expected future *price level* and the variance of the price level do not increase over time. Under inflation targeting, the expected level of future *inflation* and the variance of inflation do not increase over time, but the mean and variance of future prices do. Which regime is most advantageous depends on policy-makers' views of the costs and the benefits of one as against the other. This subject has received much academic attention in recent years. Contributions by Lebow *et al* (1992), Fillion and Tetlow (1993),

---

[(1)] Although it is not clear how much of this success is attributable to the ability of policy-makers in these countries to conduct monetary policy versus the incidence of favourable shocks in the international economy.
[(2)] Including United Kingdom, Canada, Sweden, Finland, Spain, Australia and New Zealand. In the chart, inflation rates across countries are averaged weighting countries by their relative 1998 Q4 GDP levels.

Haldane and Salmon (1995), Black *et al* (1997), Kiley (1998), Svensson (1999a), Smets (2000), Williams (1999), Vestin (1999) and Dittmar *et al* (1999) all sought to evaluate the consequences of pursuing price level stability on the one hand or inflation stability on the other. With some qualifications, a recurrent finding of some of these earlier works is that the cost of inflation targeting as against price level targeting is that it results in increased, and ever-increasing, variability in the price level. The benefit is that inflation targeting leads to lower inflation and output volatility than price level targeting.(3) As Svensson (1999a) remarks, the reason for this is intuitive. By construction, inflation targeting allows drift in the price level. Thus when a shock pushes inflation above its average, inflation needs to be returned only to its average level. By contrast, price level targeting penalises base level drift. Hence, periods of 'above-average' inflation must be followed by periods of 'below-average' inflation. If there are nominal rigidities in the economy, the higher inflation variability of price level targets will naturally translate into higher output variability.

In line with earlier work, in this paper we also examine what happens to price level, inflation and output volatility as we move from price level to inflation targeting. But in addition, we investigate the implications for monetary policy when the target lies *between* these two extremes. In other words, we also look at policies that accommodate some portion of—but not all—the shocks to the price level, or, equivalently, at combinations of policies that accommodate every one-off shock to the price level (pure inflation targeting), and policies that reverse every one-off shock to the price level (pure price level targeting).(4) Because they allow temporary drift in the price level, these policies may be relevant, for instance, in the event of temporary supply-side shocks (eg an oil shock) or temporary changes in indirect taxes that enter the calculation of the inflation target.

We offer two characterisations of this spectrum of policies—under 'optimal' rules and under a family of 'simple' rules—and study how the volatilities of inflation, the price level and the output gap change as we move from one extreme to the other, under the assumption that the central bank can credibly commit to its policy strategy.

Our results reveal that studying these intermediate regimes is important. This is because—at least for our analytical set-up—the variabilities of the price level, output and inflation do not change monotonically along the spectra between the two extremes (price and inflation targeting). This implies that we cannot simply extrapolate from the results in the previous literature on price and inflation targeting to infer the welfare implications of intermediate regimes. We find that these implications tend to change as we vary the assumptions about the model, just as the previous literature found that the welfare comparisons of price and

(3) This result is partially overturned by Svensson (1999a), who contrasts time-consistent and time-inconsistent monetary policy rules for a simple model of the inflationary process. He finds that time-consistent rules can generate the result that price level targeting does not increase the variability of inflation as compared with inflation targeting. Vestin (*op cit*) shows that Svensson's (1999a) findings hold in a forward-looking model with Calvo-Taylor style price stickiness. On the other hand, Kiley (1998) found that in the kind of sticky-price models analysed in Vestin (*op cit*), discretionary price level targeting generates more output variability than inflation targeting—the conventional result once again. Here we abstract from time-consistency issues and assume that the central bank can credibly precommit. Under commitment, Williams (*op cit*) also partially overturned the result that inflation targeting leads to lower inflation and output volatility than price level targeting by using simple three-parameter rules with optimised coefficients and constrained interest rate smoothing within the FRB/US model. Finally, Smets (2000) finds that for a given horizon a price level target always gives lower inflation variability, but higher output and inflation variability. However, if one endogenises the horizon, a price level target can also be superior to an inflation target in terms of output and interest rate variability.
(4) Hybrid policies of a similar kind were analysed for nominal income targeting by McCallum (1993) and Hall and Mankiw (1994).

inflation targeting are model-specific. In addition, we find that the welfare implications of intermediate regimes change when: (i) with optimal rules, we alter the weight on the variation of the output gap in the loss function; (ii) with simple rules, we make the rule more or less forward-looking.

The remainder of the paper proceeds as follows. In Section 2 we offer two alternative characterisations of the policy spectrum between pure price and pure inflation targeting. The model used to evaluate policies along the spectra is described in Section 3. In Section 4 we present the results, which include an analysis of the problem of zero bounds under hybrid regimes. Section 5 concludes.

2 From price level to inflation targeting: a spectrum of monetary policies

2.1 The optimal control spectrum

One way of thinking about the inflation target-price level target spectrum is as a continuum of loss functions, ranging between the two target extremes, that policy-makers minimise. At one extreme, the loss function penalises the deviation of inflation from an inflation target. At the other, it penalises the deviation of the price level from a price level target.

More formally, we can express this as a family of loss functions in **(1)** below, where targets are normalised to zero for convenience:

$$L_t = E_t \sum_{j=0}^{\infty} b^j [l_p (p_{t+j} - h\, p_{t+j-1})^2 + l_y {y_{t+j}}^2] \qquad \textbf{(1)}$$

where p_t is (the log of) the consumer price index; y_t is (the log of) the output gap; E_t is the expectational operator defined over information available at time t; and lp and ly are the weights on real and nominal deviations from their respective targets. We set these to 0.5 in our baseline loss function specification, but we also experiment with different weights below.

Finally, h is the parameter that defines the spectrum of targets between price level and inflation targeting. It varies between zero and one. For h = 0, policy-makers target the price level. When h = 1, policy-makers target the level of the inflation rate. For h taking any value between zero and one, policy-makers target a 'hybrid' regime, under which the price level target is allowed to drift temporarily by a portion (h) of the change in the price level observed between one period and the next. As h tends to 1, the policy-maker's loss function penalises fewer shocks to the price level—so that temporary deviations of the price level from its target get larger—and *vice versa*.

This can perhaps best be seen by re-writing the quasi-difference that represents our 'price stability term' in equation **(1)** as follows:

$$\left[h(p_t - p_{t-1}) + (1-h)\,p_t\right]^2 \qquad \textbf{(1a)}$$

which shows that the family of loss functions in equation **(1)** is in practice a family of weighted combinations of an inflation and a price level target (with h again representing the amount of price level drift endorsed by the policy-makers).[(5)] In this sense, a more formal way of thinking of the case when $1 > h > 0$ in our 'generalised price stability' target in **(1a)** (and **(1)**), is as a price level target (with no drift in the price level in the long run)[(6)] with dynamic adjustment.[(7)] Note that because this spectrum of regimes ranges between a stable price level and a zero inflation target, the extreme targets in our set of regimes are comparable with those in Svensson (1997a, 1997b, 1999b).[(8)]

Interpretation of the optimal control spectrum

How should we interpret the parameter h in the optimal control spectrum? Technically, h specifies the relative price of the unconditional variance of the price level (deviations from the price target) with respect to the variance of the price level conditional on last period's price (deviations from the zero inflation target). Restricting h to between zero and one amounts to saying that the conditional variance of the price level is not infinitely costly to society relative to the unconditional variance (or *vice versa*).

In practice, the optimal value of h will depend on many factors, such as: (i) the size of the inflation tax on money balances (which may bias h towards giving an inflation target); (ii) the cost of indexation (menu costs, which may make it costly to update the price level in line with the drift associated with an inflation target); (iii) the length of nominal contracts (the longer the average contract, the greater the cost of a given amount of drift in the price level); and (iv) the information set of the private sector. (Strictly speaking, in full information rational expectations models, price level drift may not be costly. However, in models where consumers find it hard to distinguish between general and relative price changes, it might.) Deriving analytically a value for h is beyond the scope of this paper. Woodford (1999) investigates whether h should be zero or one for a variety of sticky price models—but does not dwell on intermediate cases like

---

(5) As both Larry Ball and Frank Smets pointed out to us, the expression in **(1a)** implies that a covariance term between ($p_t - p_{t-1}$) and (p_t) enters the loss function. This would not be the case if we had expressed the central bank's concern for inflation and price level variability as a weighted average of separate targets in the loss function, ie like $[h(p_t - p_{t-1})^2 + (1-h)(p_t)^2]$. As we explain later, we abstract throughout from attempting to map between the posited central bank's loss function and a (deeper) representative agent utility function, so we do not examine which one of these alternative specifications is to be preferred in practice.

(6) In fact the price level will be anchored not only under optimal (and simple) pure price level targeting rules, but also under rules that penalise a mixture of price level and inflation deviations from target, as long as h is strictly smaller than one. Technically, this is because, whenever $h < 1$, the price level still enters the state vector: it is one of the variables that enters the optimal decision rule for interest rates. Unless the state vector contains another nominal variable on which the instruments also feeds back with the same (but oppositely signed) coefficient, the price level will not drift.

(7) Alternatively, one could define the hybrid regime as an inflation target plus an error correction term that gradually eliminates the period drift of the price level as in Dittmar, Gavin and Kydland (1999) and King (1999).

(8) This exercise is analogous in some respects to that in Dittmar *et al* (*op cit*). They look at the case in which the policy committee in charge of monetary policy is composed of two types of individual: the first willing to target inflation; and the second willing to target the price level. The loss function that they come up with entails an inflation target with an error-correction term to capture a long-term price objective of the form. In this sense, their analysis differs from ours because their loss function does not contain a covariance term relating variations in the inflation target to variations in the long term price level target.

those examined in this paper. He shows that for a variety of sticky price models it is optimal to target the inflation rate (at zero) at all times. However, his analysis abstracts from the factors (i)-(iv) above. These and other factors—like the specificity of the utility functions with which he endows the private sector—suggest that other policies apart from inflation targeting might be optimal, and that his results may not be general.

Even if the optimal value of h were unique, as Woodford (*op cit*) concludes, there are at least two reasons why it is still useful to compare ranges of alternative h s. Assume that the socially optimal value of h is $\tilde{h}$, say, and that this value is known with certainty. Then by comparing the performance of the socially optimal regime with regimes based on $h \neq \tilde{h}$, we can assess the welfare consequences of departing from the social optimum. This is in line with the work of Svensson (1999a), Vestin (*op cit*) and Smets (*op cit*). They assume that equation **(1)** represents a contract delegated to the central bank according to some underlying welfare function, (whose h and l are unknown), and address questions like 'what would be the welfare consequences of delegating a price level target to the central bank if society inherently favours inflation targeting?'. Alternatively, assume that the socially optimal value of h is unknown (because of uncertainty about the true model of the economy). Then exploring whether some h s unequivocally dominate others may be informative about the socially (unknown) optimal value of h.

2.2 The simple rules spectrum

The second way of thinking about the inflation-price level target spectrum is within simple rules. Simple rules, where policy is conditioned on some subset of the variables that the central bank observes in the real world, have the virtue that they may be easily computed and monitored—and therefore are usually associated with credibility building. Although the performance of simple rules can still be model-specific, typically they are also more robust to uncertainties about the underlying model of the economy than optimal, complex rules that tend to be highly model-specific. The use of simple rules to guide or inform policy has been advocated by many, notably Friedman (1959), Currie and Levine (1993), McCallum (1988, 1990, 1994), Taylor (1993a, 1996), Henderson and McKibbin (1993), Isard *et al* (1999), Christiano and Gust (1999) and Williams (*op cit*).

In this spirit, imagine that policy follows a simple, generalised Taylor rule:

$$R_t = E_t p_{t+1} + f_p[Ep_{t+k} - hEp_{t+k-1}] + f_y[y_t - y^T] \qquad \textbf{(2)}$$

where R_t denotes the short-term nominal interest rate, and p_t , π_t, y_t and E_t are defined in the same way as in **(1)**. Rule **(2)** assumes that policy-makers set the nominal interest rate so that the *ex ante* short-term real interest rate (R_t - $E_t p_{t+1}$) responds to deviations of a price/inflation term from target and of output from potential. We assume that the policy-makers have no tendency to smooth rates, so we do not include any lagged interest rate term on the RHS of equation **(2)**.

Once again, as for the optimal control spectrum, $h \in [0, 1]$ is the parameter that defines the spectrum of rules from inflation to price level targeting. As before, when $h = 0$, the feedback variable is the price level; when $h = 1$, the feedback variable is inflation; and finally, when h is between zero and one, the feedback variable is a moving price level target, where the target moves (just like the penalty in the loss function above) by some portion h of the change in the price level from one period to the next. k is the feedback horizon of the central bank[9] which we set to equal either 0 or 8. When $k = 0$ the central bank feeds back from current-dated variables only.[10] When $k = 8$ the central bank feeds back instead from deviations of two year ahead forecasts of variables from target. By combining different values of k and h we can illustrate the performance of a continuum of backward and forward-looking policy rules from price level (current and forecast-based) rules to inflation (current and forecast-based) rules.[11] The other parameters in the rule are set such that $r_R = 0.5$, $f_P = 0.5$, and $f_Y = 0.5$. So, in all cases, policy-makers respond also to output deviations from potential (as in Taylor (1993a)).

3 The model

To evaluate the performance of the above spectra of monetary policy regimes/rules we need a model of the economy. We take the open-economy adaptation of Fuhrer and Moore (1995)[12] that Batini and Haldane (*op cit*) have used to analyse the performance of alternative rules for the policy instrument.[13]

Once we normalise to zero capacity output and foreign variables and remove the constants in each equation, so that variables are expressed in terms of deviations from equilibrium, the model—excluding the policy rule—can be written as follows:

$$y_t = a_1 y_{t-1} + a_2 E_t y_{t+1} + a_3 [R_{t-1} - E_{t-1} p_t] + a_4 q_{t-1} + e_{IS,t-1} \quad \textbf{(4)}$$

$$p_t = c_1 E_t p_{t+1} + (1 - c_1) p_{t-1} + c_2 [y_t + y_{t-1}] + c_3 [(1 - c_1)\Delta q_t - c_1 E_t \Delta q_{t+1}] + e_{pt} \quad \textbf{(5)}$$

$$E_t \Delta q_{t+1} + E_t p_{t+1} = R_t + e_{UIP,t} \quad \textbf{(6)}$$

where again, y_t is (the log of the) output gap, p_t is (log) inflation, R_t is the nominal interest rate, and q_t and Δ are the (log) real exchange rate and the first-difference operator respectively. e_{ISt}, ep_t, and e_{UIPt} are disturbance terms whose properties are described below. The aggregate demand and aggregate supply shocks are assumed to be white noise, while the UIP shock is modelled a first-order autoregressive process, as in Batini and Nelson (*op cit*), to match actual exchange rate data.[14] Equation **(4)** is the model's IS equation, with real output depending positively on leads and lags of itself and the real exchange rate, and negatively on the *ex ante* real interest rate. Importantly, we set $a_1 > 0$, so that there is persistence in the output gap (a necessary condition for the results in Svensson (1999a)). Equation **(5)** is an open-economy aggregate supply curve; and equation **(6)** is an uncovered interest parity relationship linking interest rate differentials (domestic relative to abroad) to expected changes in the exchange rate and to a time-varying, stochastic risk premium.

We choose two variants of this basic model for our analysis. In the first variant, we set $a_1 = 0.8$,

---

(9) See Batini and Nelson (2000) for alternative definitions of the 'horizon'.
(10) When $k = 0$, this rule is similar to the one in King (1999).
(11) 'Inflation forecast-based rules' where $k > 0$, are familiar from the work of, de Brower and O'Regan (1997), Black, Macklem and Rose (1998), Batini and Haldane (*op cit*) and McCallum and Nelson (1999, Appendix A).
(12) Buiter and Jewitt (1981) in fact first posited the overlapping real-wage contracting structure in Fuhrer and Moore (1995).
(13) See also Blake and Westaway (1997).
(14) Assuming that the UIP shocks follow an AR(1) process is common in macroeconomic modelling (see Taylor (1993b)).

$a_2 = 0$ and $c_1 = 0.2$. With this parameterisation, the model has a backward-looking IS function and a partially forward-looking AS equation (as in Fuhrer (1997)). As Batini and Nelson (*op cit*) point out, this may be the most appropriate setting of the parameters if expectations play a more important role in wage/price-setting decisions than in household spending decisions. Batini and Haldane (1999) show that with these parameters the model's transmission mechanism is rather sluggish and is broadly in line with simulation responses from VAR-based studies of the effect of monetary shocks in the United Kingdom.[15] In the second variant, we set a_1 and a_2 as before, but set $c_1 = 1$, so that the aggregate supply equation becomes an open-economy version of Taylor's (1980) staggered contracts model, with inflation depending only on leads of itself, rather than on both lags and leads of itself. As Batini and Haldane (*op cit*) emphasise, in this case the transmission of policy impulses is swifter and, hence, policy needs to be less forward-looking than otherwise. In general, the extent of forward-looking behaviour appears to be important in determining the relative costs of inflation and price level stabilisation (see, for example, Batini and Haldane (*op cit*), and Levin *et al* (1999)), and, therefore, may matter when it comes to establishing the relative merits of policies at intermediate points along the spectrum. So looking at both variants enables us to establish the robustness of the results about the spectra.

The model is calibrated on UK data as in Batini and Haldane (*op cit*).[16] And the shock processes are modelled as in Batini and Nelson (*op cit*), who generated residuals from the model equations **(4)**-**(6)** using UK data over the period 1981 Q1-1998 Q1.

4 Results

The two variants of the model were simulated under both the optimal control and the simple rule spectrum with a grid of values for h ranging between zero and one. For each value of h we solved the model and derived asymptotic variances of inflation, output and the price level using the Hansen-Sargent (1999) doubling algorithm. Within the optimal control experiment, we computed results for various h—the spectrum—also using different weights on output variability in the loss function (equation **(1)**), to test how changes in the output preferences by the policy-makers affected our results. Following Taylor (1993a), we consider where each of the rules places the economy on the output-inflation volatility frontier. Henceforth we refer to this frontier as to the 'Taylor curve'.[17] Table A below summarises the experiments that we have conducted. In what follows we comment on these in some detail.

---

(15) For instance, they show that it takes around five quarters for output to touch its lowest level after a permanent negative 1% shock to the inflation target.

(16) In particular, the remaining parameter values not quantified in the text are: $a_1 = 0.75$, $a_2 = 0$, $a_3 = -0.5$, $a_4 = 0.16$, $c_2 = 0.2$ and $c_2 = 0.5$.

(17) We prefer this to presenting results using a simplistic ranking system like a loss function because this would put rules that penalise targets different from those in the ranking loss artificially at a disadvantage.

Table A: Experiments

Rule	Policy regime or spectrum[1]	Chart
Fuhrer and Moore contracting model (slower transmission mechanism)		
Rule	*Policy regime or spectrum*[(1)]	*Chart*
Optimal time inconsistent	h = 0 (price level target); l_y = 0.5 h = 1 (inflation target); l_y = 0.5 $0 \geq h \geq 1$ (hybrid regimes); l_y = 0.5	Charts 1a & 1b
Simple (current-dated variables)	Same as above	Charts 2a & 2b
Simple (future-dated variables, k = 8)	Same as above	Charts 3a & 3b
Simple (future-dated variables, $k < 4$; $k = 4$; $k > 4$)	Same as above	Chart 4
Optimal time inconsistent with different output gap preferences	Same as above where, in turn: l_y = 0, 0.5, 4	Charts 8a & 8b
Taylor contracting model (swifter transmission mechanism)		
Rule	*Policy regime or spectrum*[(1)]	*Chart*
Optimal time inconsistent	h = 0 (price level target); l_y = 0.5 h = 1 (inflation target); l_y = 0.5 $0 \geq h \geq 1$ (hybrid regimes); l_y = 0.5	Charts 4a & 4b
Simple (current-dated variables)	Same as above	Charts 5a & 5b
Simple (future-dated variables, k = 8)	Same as above	Charts 6a & 6b
Optimal time inconsistent with different output gap preferences	Same as above where, in turn: l_y = 0, 0.5, 4	Charts 9a & 9b

(1) Here we use the word 'target' loosely to indicate either a target in the loss function (in the optimal rules case), or a target from which we compute deviations when specifying the feedback term of a simple rule.

4.1 Fuhrer and Moore contracting

Charts 1a–3b plot results derived employing the first model variant, assuming that wage/price-setting decisions are moderately forward-looking but that household spending decisions are entirely backward-looking. Specifically, Chart 1a plots the locus of output/inflation variability points delivered by a fully optimal rule minimising **(1)** as h is varied (the 'h-locus' hereafter). Points that lie south and west in the chart are welfare superior, and points to the north and east inferior.

In this first chart, the shape of the h-locus indicates that the variance of output increases monotonically as policy tends towards price level targeting. This is in line with the notion—and contrary to results under discretion[18]—that pure price level targeting brings with it the cost of higher output variability. The chart also shows that the variance of inflation is higher under price level targeting than it is under inflation targeting, in line with that notion. Interestingly, however, this seems not to increase monotonically throughout the spectrum: rather, it falls for $h \leq 0.5$, and rises thereafter. This suggests that even if a policy based on inflation targets moves almost halfway towards a pure price level target this may have no adverse consequences for inflation stabilisation, rather, it could reduce the volatility of inflation.

Again in line with results in the earlier literature, we also find that price level variability increases when policy moves closer to the inflation-targeting end of the spectrum. But although eventually the variance of prices 'explodes' when $h = 1$, its increase is not dramatic initially when moving from pure price level targeting to a hybrid inflation/price level target regime (ie for $h < 0.5$) (Chart 1b). This implies that inflation targeting could realise most of the reduction in price level variability achievable with this model by moving only some way in the direction of pure price level targeting.

For simple Taylor-like rules, such as **(2)**, that feed back on *current*-dated values of prices or inflation, results are somewhat different. As Chart 2a illustrates, in this case both inflation variability and output variability *fall* monotonically—apart from when h is exactly 1—as policy-makers become more like 'price level targeters'. This contradicts the previous literature on the benefits of price level targets. Likewise, even if in line with that literature, departures from pure price level targeting initially *worsen* price level control, a pure inflation target appears to give lower variance of the price level than a pure price level target with these rules (Chart 2b).

For simple rules that feed back on *future*-dated values of prices or inflation, the results are very similar to the optimal case, and more generally agree with what we have labelled the 'conventional' result in the literature. For instance, if we constrain policy to feed back on forecasts of endogenous variables four or more periods ahead, then we find that as policy moves towards inflation targeting, both inflation and output variability improve (Chart 3a). So in this case the Taylor curve slopes up as in Chart 1a. Chart 4 shows what happens when $k < 4$, $k = 4$ and $k > 4$. With policy feeding back on variables dated between the current date and four periods ahead, the Taylor curve rotates by 180 degrees. Thereafter, at least as far as the curve that has feedback on variables eight periods ahead, it stays upward-sloping, but with inflation targeting now being closer to the origin, ie welfare superior in comparison with price level targeting, rather than the reverse. From Chart 3b it emerges that price level variability increases at an increasing rate as h increases, rather than—as in the case of feeding back from current-dated variables—deteriorating initially but then improving as we progress towards a pure inflation-targeting rule.

In summary, using our first model specification, a comparison of alternative targeting regimes (pure price level versus pure inflation) under optimal and simple future data based rules (with $k > 4$) appears to confirm the wisdom that price level targets generate better price level control at the expense of higher inflation and output variability. Hybrid regimes, especially those where $h \geq 0.6$, may dominate pure inflation targets in terms of inflation control, delivering at the same time most of the price level stabilisation achievable via a pure price level target. These results however are reversed when we solve the model using current data based simple rules: price level targets give lower inflation and output variability, but higher price level variability than inflation targets. In this case, hybrid regimes give worse outcomes than pure regimes.

(18) See Svensson (*op cit*) and Vestin (*op cit*).

4.2 Taylor contracting

When we alter the specification of aggregate supply, keeping other things equal, and assume that wages are set *à la* Taylor (1980) rather than as in Fuhrer and Moore (*op cit*)—so that inflation is a function of only leads of itself rather than leads and lags of itself—our simulation results change somewhat. The Taylor contracting results are shown in Charts 5a-7b. Below we summarise the key differences between these and the Fuhrer and Moore results.

The first thing to notice is that for optimal rules, inflation, output and price level variabilities are lower when we assume the Taylor specification. For instance, output variability associated with a pure inflation targeting under Taylor contracting is around a third of that under Fuhrer-Moore contracting. This is a well-known result. Shocks have effects for a long time to come and are hard to iron out with monetary policy. This is not true however of simple (either current or forecast-based) rules, which now show a much higher output variability than under Fuhrer-Moore, perhaps reflecting the fact that our original choice of feedback parameters f_p, f_y and k is particularly inadequate for this parameterisation of the model.

For the optimal control rules, with Taylor contracting we again find that price level variability also falls monotonically when we move towards a price level target (Chart 5b). And like before, here the largest gains in controlling the price level accrue for $h \leq 0.5$ (the Taylor curve is quite vertical for h between 0 and 0.5), suggesting that hybrid regimes close to a pure price level regime deliver most of the reductions in price level variability achievable.

However, under optimal rules and with Taylor contracts we find that a pure inflation target (ie $h = 1$) no longer unequivocally dominates a price level target in terms of inflation and output variability: inflation variability is, in fact, worse under inflation targeting than under price level targeting (Chart 5a). What can explain the difference in results with respect to the Fuhrer-Moore contracting case?

One explanation, suggested by Williams (*op cit*), is that, as models become more forward-looking—for example moving from the Fuhrer-Moore contracting model to the Taylor contracting model in our case—the channels via which price level targets operate to stabilise inflation change. In his words, a price level target ensures that the sum of all future inflation rates is equal to the negative of today's gap between the level of prices and the target. Typically, this gap is proportional to current inflation. Hence, when inflation is high, a target on the price level will induce expectations of undershooting inflation and *vice versa* when inflation is low. In both circumstances this will help stabilise inflation today. The beneficial expectational effects of price level targets on inflation are more powerful with forward-looking models, where—in contrast to backward-looking models—expectations play a major role. That is why we find that price level targets outperform inflation targets in stabilising inflation only when we experiment with the (jumpier) Taylor contracting model variant.

Note that the same argument does not apply to price level or output stabilisation. More stable inflation does not entail more stable prices. Imagine again a model as the one in variant 2, with forward-looking inflation but sticky prices. With prices still diverging from target after a shock, the rule will dictate further changes in rates, which in turn, will have adverse implications for output variability.

In short, when agents' nominal contracts become more forward-looking, inflation stabilisation can be achieved more easily via the expectational channel that we described—the lower inflation variability result associated with a price level target ($h = 0$) in Chart 5a relative to the corresponding target in Chart 1a. But price level control still requires interest rate gyrations, which can be destabilising for output. This explains why, under price level targets, output variability is higher than under inflation targets in both Chart 5a and Chart 1a.

With respect to simple rules, some results using *forecast-based* rules are similar under the two contracting specifications. Moving from price level targeting to inflation targeting we still obtain an improvement in output variability and a deterioration in the variability of inflation (diagrammatically, the Taylor curve in Chart 7a slopes downwards as that in Chart 3a). A forecast-based rule has the property of making the Fuhrer-Moore contracting model variant altogether more forward-looking. So, in this case, this model too will benefit from an expectational channel similar to the one at work in the Taylor contracting variant (as well as be subject to the ensuing output variability costs).

However, differences emerge when we compare the price level stabilisation properties of forecast-based rules under different contracts. We find that rules feeding back on inflation give a better price level control than rules responding to the price level itself (Chart 7b). This may again reflect the inadequate choice of feedback horizon in the rule, which lets the rule respond to movements in prices at time $t + 8$. If, as with Taylor contracts, the assumed contracting specification implies a shorter than eight quarters adjustment dynamics for prices, our forecast-based rule may destabilise, rather than stabilise the price level.

Results for *current-dated* rules also differ under the two contracting specifications. In particular, with forward-looking aggregate supply, we now observe explosive outcomes for output variability—and thus inflation variability—once the rule becomes one that corrects some, but not all, the drift in the price level after a shock (see Charts 6a and 6b); this is particularly true for h equal to 0.6 and 0.7. Indeed, for these rules it appears optimal to feed back from a pure price level ($h = 0$).

An explanation for the divergence in results under the two model specifications is as follows. Simple rules that respond to current-dated variables tend to be highly destabilising because they attempt to affect variables that are pre-determined. Typically, this leads to more aggressive policy responses than those implied by rules that act pre-emptively, due to the lower efficacy of the interest rate at a close range. In turn, such aggressive policies are both costly in terms of output and can lead to cycles in inflation—which increases inflation variability in the short run. When the private sector is strongly forward-looking, as in the Taylor model specification, the instability in inflation can become explosive as higher inflation today leads agents to expect even higher inflation tomorrow.

To summarise, as we expected, making the length of the transmission mechanism shorter by assuming a greater degree of forward-looking behaviour in the economy changes our results on pure and hybrid regimes quite dramatically. Optimal and simple future data based rules no longer provide analogous outcomes. Contrary to before, results on optimal rules now suggests that price level targets may give *lower* inflation volatility than inflation targets, even if they are still more output-costly to operate. And results on simple future data based rules now indicate that price level targets—at least for our choice of feedback horizon—may destabilise rather than stabilise the price level. Finally, although they now appear to mitigate output volatility under optimal rules, here hybrid regimes are never optimal. Shortening the transmission lag seems to emphasise the 'kink' in the h-locus when we close the model with simple current data based rules, implying that these regimes may be highly destabilising when agents are strongly forward-looking.

4.3 Implications of alternative output preferences

As a final test of the robustness of our results comparing inflation and price level targeting, we checked how sensitive these were to changes in the weight on output gap variability in the loss function. Because this test only involves changing the weights in equation (**1**), its results relate uniquely to the optimal control spectrum. Charts 8a-9b plot h -loci when the weight on output variability l_y varies from 0 to 4. For comparison, they also include the case when $l_y = 0.5$—our baseline loss function specification in Sections 4.1 and 4.2 above. Charts 8a and 8b present results for the Fuhrer-Moore contracting specification; Charts 9a and 9b do the same for the Taylor contracting case.

A key result from this exercise is that, regardless of the assumption on the contracting specification, moving from a zero to a sizeable weight on output volatility in the loss function (ie from $l_y = 0$ to $l_y = 4$), the variance of output drops considerably for any level of inflation and price level variance. With a zero or low weight on output variance relative to price level/inflation variance in the loss function, the optimal policy rule prioritises price level and inflation control at the expense of output control; and *vice versa* when the weight on output stability is large relative to that on price stability.

Yet varying the weight on output stabilisation also affects our results in other ways.

Let us consider first the Fuhrer-Moore contracting case, where the inflation process is governed by a mixture of backward and forward-looking elements. In Chart 8a, the dashed h -locus for $l_y = 0.5$ is the same as in Chart 1a, our baseline output preference case; there price level targeting entails higher inflation and output volatility. The solid h -locus associated with $l_y = 0$ shows instead what happens when policy-makers display no interest in stabilising output. In this case, inflation and output standard deviations also decline monotonically when we move from a price level towards an inflation target. In the end, as for the baseline preferences case, with no concern for output volatility, pure inflation targets still dominate pure price level targets in the inflation-output volatility space—the conventional result.

On the contrary, when policy-makers place a high weight on output stabilisation, the conventional literature result is partially overturned. The tilted grey h -locus associated with ($l_y = 4$) in Chart 8a indicates that a pure inflation target warrants lower output variance relative to a pure price level target, but this comes at the cost of higher inflation variability. In effect, with strong output stabilisation preferences a trade-off emerges between output and inflation stabilisation as policy-makers forgive more and more of the amount of drift in the price level.[19]

As one would expect, changing the weight on the output gap in loss function **(1)** has implications also for price level control. Chart 8b plots the previous h -loci in price level-output volatility space. When output variance is highly rated by the policy-makers ($l_y = 4$), a small price level variance can only be achieved by using a target that penalises at least a half of the drift in the price level ($h \leq 0.5$). Moving towards regimes that allow more than this drift in the price level secures very low levels of output variability, but makes it

[19] Williams (*op cit*) conducts a similar experiment. Comparing just price level and inflation targets, he finds that price level dominates inflation targeting under both inflation and output variability criteria for weights on output deviations from potential in the loss function above 0.1.

much harder to control the price level. On the other hand, when we assume that policy-makers have no concern for output stabilisation ($l_y = 0$), we find that price level variability does not increase significantly until policy moves very close to the inflation targeting end of the spectrum. In other words, the h-locus is very steep. This suggests that a policy that prioritises inflation stabilisation and that is based on inflation targets would have to move only marginally in the direction of a pure price level target in order to obtain the reduction in price level variability achievable with this model.

Finally, Charts 9a and 9b plot h-loci derived conducting the same experiments above but when the model variant involves a fully forward-looking aggregate supply curve (as emerges when we assume Taylor contracting). As for the previous experiments that employed this model variant, here the scale of the variances is much smaller (around a half for output and a fifth for inflation and the price level) than that in the corresponding charts for the model variant with a mixed backward/forward-looking Phillips curve (Charts 8a-8b). A 'jumpier' inflation facilitates inflation, output and price level control, a point originally made by Ball (1994).

In this setting, an exclusive concern for price stability almost completely eliminates inflation variability (Chart 9a, h-locus for $l_y = 0$), as well as granting perfect control of the price level—the h-locus lies on the *y*-axis for $l_y = 0$. Stronger desire to stabilise output enhances the ability to minimise output variability, reducing it from a half to a fifth of its original value under strict price level/inflation targeting. However, this is accompanied by higher inflation and price level variability, with more of each materialising at the inflation-targeting end of the spectrum ($h > 0.5$).

In summary, we found that different output preferences change our baseline case results in three ways. First, a zero weight on output deviations from potential in the loss function reduces the possibility to lessen the variance of output, regardless of the nature of the target for the stabilisation of prices. *Vice versa*, when the concern for output is high, it seems as if the amount of output variation can be reduced uniformly for any kind of price stability target.

Second, when the preference for output stabilisation changes relative to the one for price stability, a pure inflation target no longer necessarily entails lower inflation variability than a pure price level target. When policy-makers strongly dislike output variability, for instance, policy becomes automatically restricted in the ability to operate via the output gap channel, the upshot of which is worsened inflation control.

And third, the result that moving from a price level to an inflation target may not involve big costs in terms of worsened price level control until a lot of drift in the price level is allowed for appears to be progressively overturned as policy-makers become more concerned about output fluctuations. With strong weights on output variability, moving towards a pure inflation target worsens almost immediately (ie even for low values of h) price level control.

4.4 Zero bounds and the h-loci

In this final section we focus on one potential additional property of price level targeting regimes, ie the ability to reduce the probability of nominal interest rates hitting a zero level for a given level of steady state inflation. Woodford (1999) has investigated this subject and has found that price level targets could reduce the variability of nominal rates, thereby minimising the chance that these hit zero. The probability that interest rates touch zero is a relevant issue in assessing the performance of alternative monetary policy rules/regimes, because zero is generally considered a lower bound for nominal rates. When policy reaches this bound, real rates cannot be reduced further and hence policy can no longer be expansionary—the 'zero bound problem' (see Buiter and Panigirtzoglou (1999), Goodfriend (2000), Krugman (1999), McCallum (2000), and Svensson (2000)).

To investigate how the probability of hitting this zero bound varies along the h-loci in our analysis, we have calculated probabilities for the implied variance of nominal interest rates under each rule contemplated in the previous sections.[20] We found that this probability is a function of: (1) the type of nominal rigidity embedded in the price-setting mechanism; (2) whether policy feeds back on a current or future dated target; (3) in line with previous findings in the paper, it turns out that the consequences of intermediate regimes for this probability cannot be extrapolated between the two extremes. In particular, for the Fuhrer and Moore model specification—which we believe offers a more realistic portrait of the functioning of price setting in the United Kingdom—it seems as if the probability of hitting a zero bound, and hence of policy becoming constrained, is maximum under price level targeting. This is true both for optimal and simple rules. So under this criterion, inflation targets and hybrid regimes close to these targets on the spectra may be superior to price level targets and regimes close to those targets on the spectra.

[20] For simplicity, we preferred this to re-simulating the models under non-linearity constraints. More specifically, we assumed that the steady-state growth rate is 2.5% and that this is unaffected by the choice of the monetary policy rule. We also assumed that whichever rule is followed generates an inflation rate of 2.5%. This gives steady-state nominal rates of approximately 5%. We take the distribution of nominal rates to be normal. This is appropriate since we have already assumed normal distributions for the shocks and the models we employ are linear. We can then use the formulae for the normal distribution to calculate the probability mass of nominal rates below zero, which we consider as an approximation to the risk of hitting the zero bound under those rules. These results are available on request.

Conclusions

The striking success of inflation-targeting countries in reducing inflation has encouraged debate about the relative merits of price level and inflation targeting.

The literature to date suggests that inflation targeting implies lower output volatility at the expense of higher price level volatility. Recently, Svensson (1999a) showed that this conventional wisdom can be overturned when policy is assumed to be discretionary, rather than committed to a time-inconsistent rule. By assuming that the central bank can credibly pre-commit, our paper sheds light on other factors that inform the price level/inflation targeting comparison.

In particular, the paper looks at intermediate regimes between inflation and price level targeting, and defines two spectra of intermediate regimes. We define the first spectrum as a set of optimal control rules for the interest rate obtained assuming that policy-makers minimise a loss function that penalised departures from a linear combination of a price level and an inflation target ('optimal control spectrum'). We define the second spectrum as a set of simple, Taylor-like rules for the interest rate, feeding back from the deviations of actual prices and inflation from a linear combination of a (current, or forecast) price level and an inflation target. Where policy feeds back from current-dated variables, the rule was related to that in King (1999).

We show that the relative merits of regimes located along these spectra depend on at least three things.

First, they depend on the degree of forward-looking behaviour of agents in the economy. As Batini and Haldane (1999) have illustrated, both the forward-lookingness of the private sector (ie the parameter c_1) and that of the central bank (ie the feedback horizon parameter k) matter in this respect. The way in which forward-looking behaviour affects the relative performance of various regimes depends, in turn, on the way in which this affects the expectational channel of monetary transmission.

Second, they depend on the way in which policy is actually operated (as in King (*op cit*)). Under optimal control both inflation targeting, price level targeting and hybrid regimes give solutions that converge to stable and unique equilibria. This holds for simple rules responding to forecasts of future variables, inasmuch as these rules respond indirectly to most states in the state vector as it is done by optimal rules. This is not true, however, of simple current data based rules *à la* Taylor, where restrictions on the set of feedback variables seem to lead to instabilities particularly for hybrid regimes. It is possible, though, that these conclusions are affected by the fact that we have chosen the simple rules' parameters arbitrarily, rather than optimally.

Third, the relative merits of regimes that we analyse hinge on what are the other preferences of the central bank in addition to price stability. In the paper, we examine the implications of different degrees of central bank's aversion for variability in the output gap when policy follows an optimal control, time-inconsistent rule. In general, when the preference for output stabilisation increases relative to that for price stabilisation, inflation control deteriorates, as policy becomes restricted in the ability to operate via the output gap channel. *Ceteris paribus*, this may impair the inflation stabilisation properties of inflation targets relative to price level targets.

Fourth, we find that intermediate regimes between extremes of the spectra can be important in informing the comparison of price level and inflation targeting, since it appears that the benefits associated with each regime do not change monotonically as we move from one extreme to another. For example, when simulating the spectrum of optimal rules with Fuhrer-Moore contracting, we found that most of the reduction in price level variability occurs in the first few increments along the spectrum from inflation targeting towards price level targeting.

Finally, the probability of hitting a 'zero bound' appears also to be a function of the degree of forward-looking behaviour in the economy. In general, with Fuhrer and Moore contracting, price level targets and hybrid regimes close to those targets lead to higher volatility in the nominal interest rate and so maximise that probability. In this respect, inflation targets and hybrid regimes close to those targets prove superior.

Charts

Chart 1a: h**-locus under optimal rules (FM)**

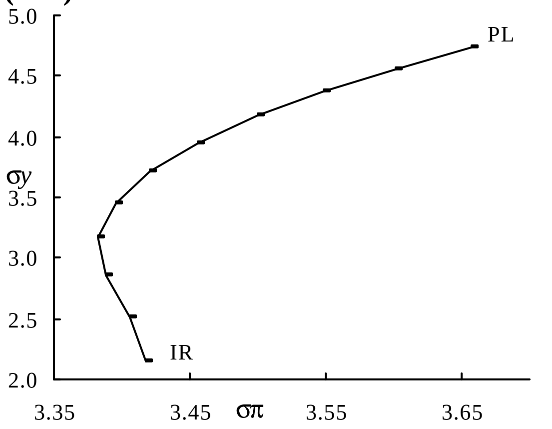

Chart 1b: h**-locus under optimal rules (FM)**

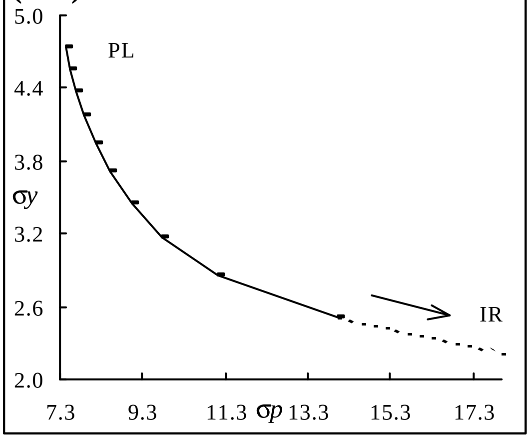

Chart 2a: h**-locus under simple current data based rules (FM)**

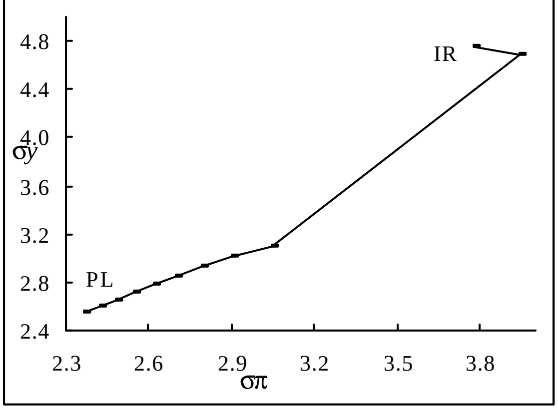

Chart 2b: h**-locus under simple current data based rules (FM)**

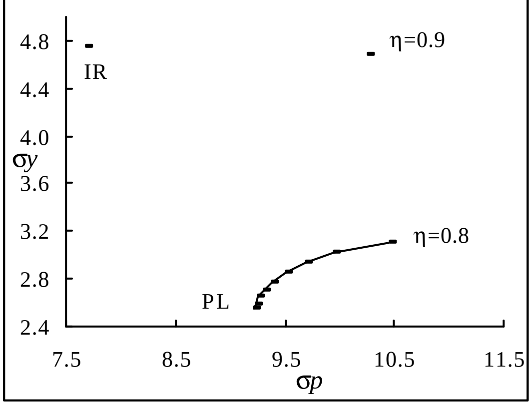

Chart 3a: h**-locus under simple future data based rules (FM)**

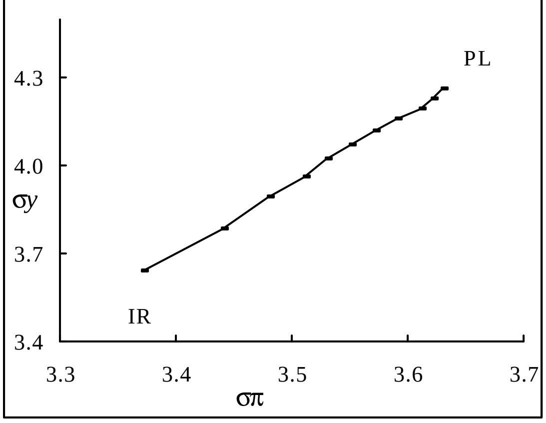

Chart 3b: h**-locus under simple future data based rules (FM)**

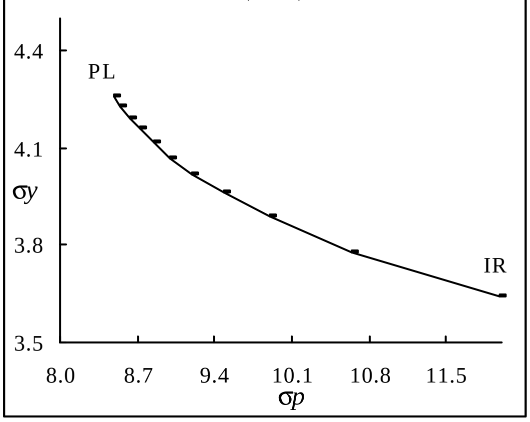

Chart 4: h **- loci current vs. future data based simple rules**

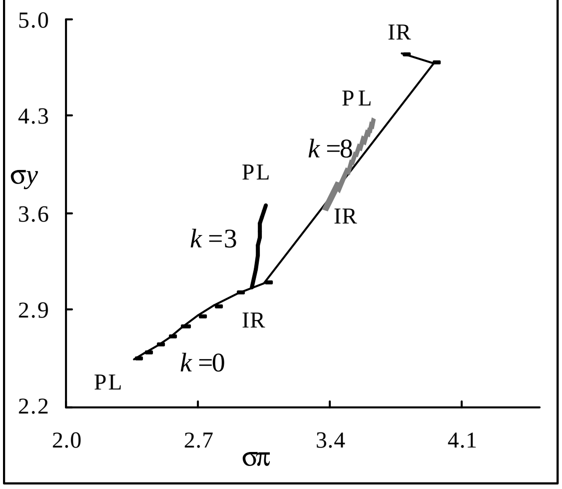

Chart 5a: h**-locus under optimal rules (T)**

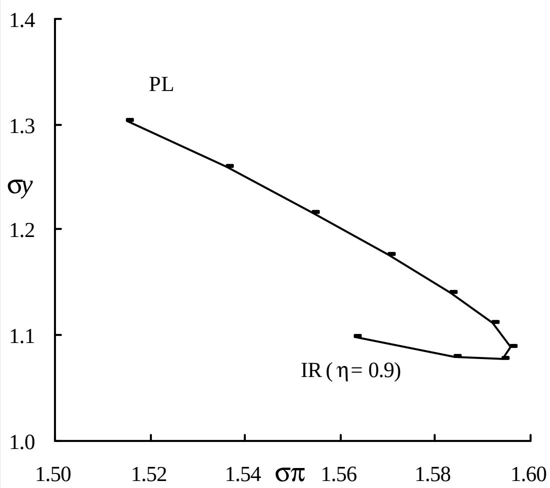

Chart 5b: h**-locus under optimal rules (T)**

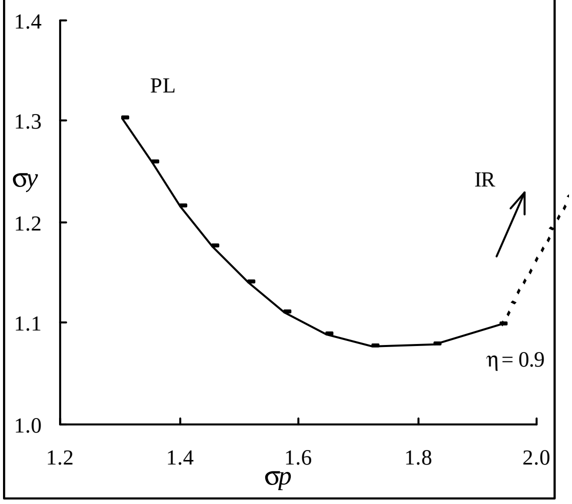

Chart 6a: h**-locus under simple Current data based rules (T)**

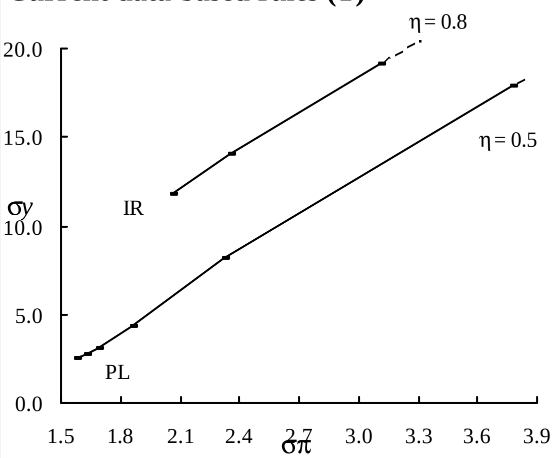

Chart 6b: h**-locus under simple Current data based rules (T)**

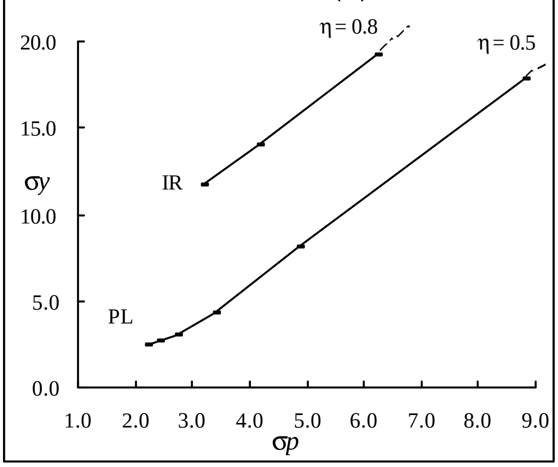

Chart 7a: h-locus under simple future data based rules (T)

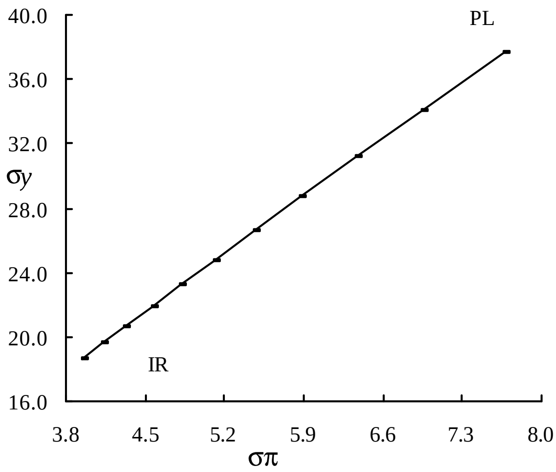

Chart 7b: h-locus under simple future data based rules (T)

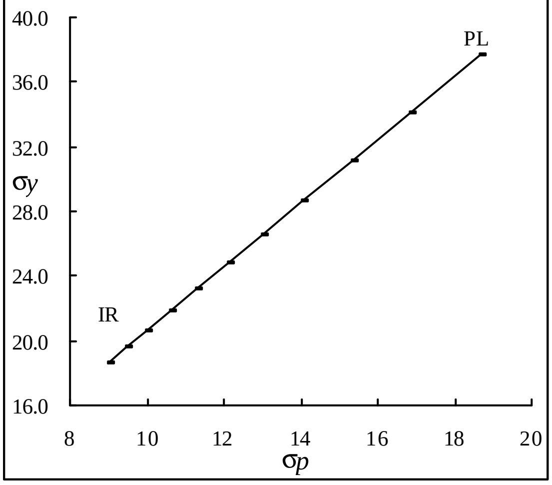

Chart 8a: h-loci for various ly (FM)

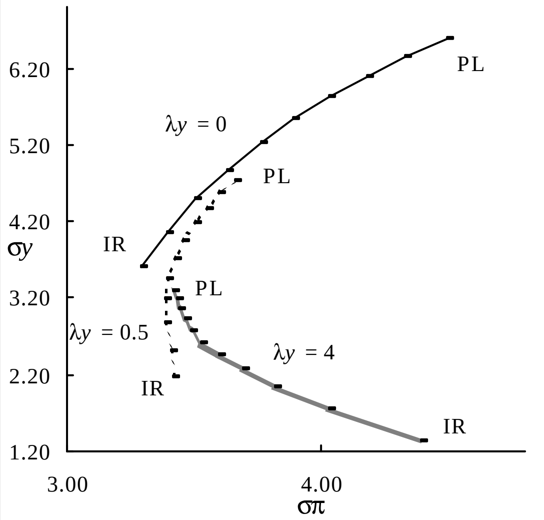

Chart 8b: h-loci for various ly (FM)

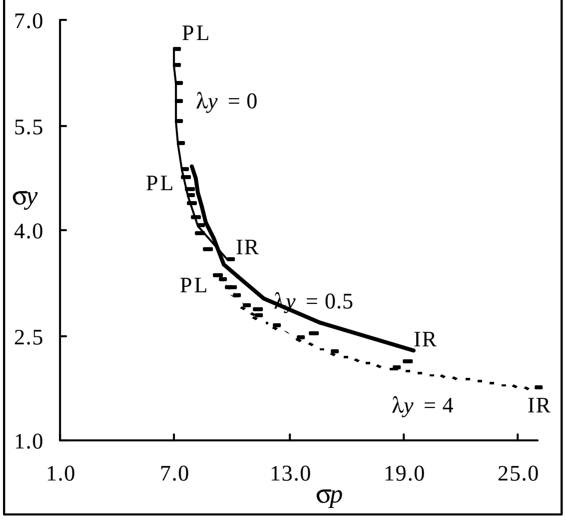

Chart 9a: h-loci for various ly (T)

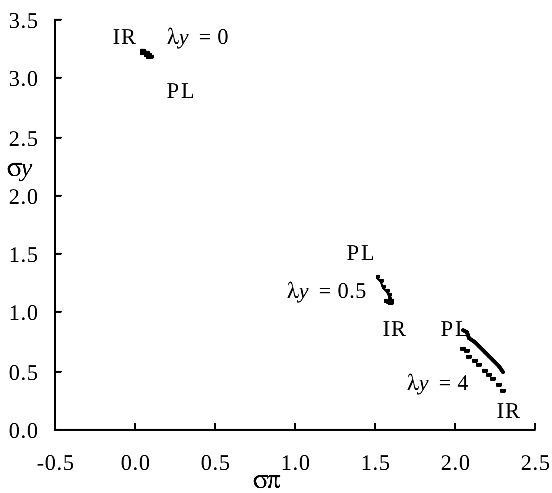

Chart 9b: h-loci for various ly (T)

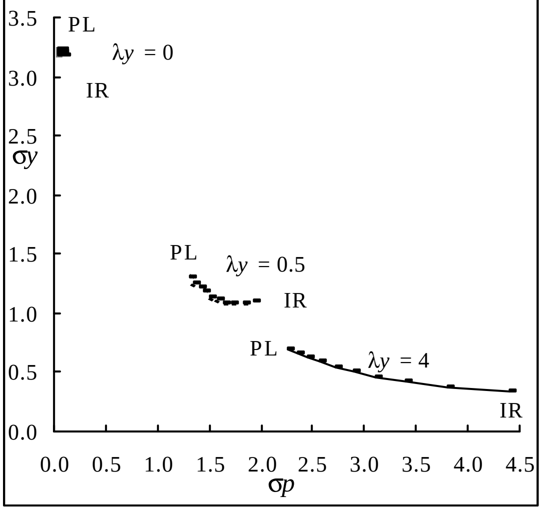

References

Ball, L (1994), 'Credible disinflation with staggered price-setting', *American Economic Review*, Vol. 84, No. 1, pages 282-89.

Batini, N and Haldane, A G (1999), 'Forward-looking rules for monetary policy', in Taylor, J B (ed), *Monetary Policy Rules*, NBER, University of Chicago Press, Chicago.

Batini, N and Nelson, E (2000), 'Optimal horizons for inflation targeting', *Bank of England Working Paper*, no. 119.

Berg, C and Jonung, L (1998), 'Pioneering price level targeting: the Swedish experience 1931-1937', *Sveriges Riksbank Working Paper 63*, June.

Bernanke, B, Mishkin, R, Laubach, T and Posen, A (1999), *Inflation targeting: lessons from the international experience*, Princeton University Press, Princeton.

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 and Westaway, P (1997), 'Credibility and the effectiveness of inflation targeting regimes', *Manchester School*, Vol. 64, pages 28-50.

de Brower, J and O'Regan, J (1997), 'Evaluating simple monetary policy rules for Australia', in Lowe, P (ed), *Monetary policy and inflation targeting*, Reserve Bank of Australia.

Buiter, W and Jewitt, I (1981), 'Staggered wage setting with real wage relativities: variations on a theme by Taylor', *The Manchester Schoo1*, Vol. 49, pages 211-28.

Buiter, W and Panigirtzoglou, N (1999), 'Liquidity traps! How to avoid them and how to escape them?, *NBER Working Paper*, No.7245, July.

Christiano, L J and Gust, C J (1999), 'Taylor rules in a simple limited participation model', forthcoming in a conference volume.

Currie, D and Levine, P (1993), 'The design of feedback rules in linear stochastic rational expectations models', in *Rules, reputation and macroeconomic policy coordination*, Cambridge; New York and Melbourne: Cambridge University Press, pages 95-121.

Dittmar, R, Gavin, W T and Kydland, F E (1999), 'Price level uncertainty and inflation targeting', *Review of the Federal Reserve Bank of St. Louis*, Vol. 81(4), pages 23-34.

Fillion, J F and Tetlow, R (1993), 'Zero inflation or price-level targeting? Some answers from stochastic simulations on a small open-economy macro model', in *Economic behaviour and policy choice under price stability*, pages 129-66, Bank of Canada.

Fisher, I (1922), *The purchasing power of money*, MacMillan.

Friedman, M (1959), *A program for monetary stability*, Fordham University Press, New York.

Fuhrer, J C and Moore, G R (1995), 'Inflation persistence', *Quarterly Journal of Economics*, Vol. 111, No. 1, February, pages 127-59.

Fuhrer, J C (1997), 'The (un)importance of forward-looking behaviour in price specifications', *Journal of Money, Credit and Banking*, Vol. 29, pages 338-50.

Goodfriend, M (2000), 'Overcoming the zero bound on interest rate policy', *Federal Reserve Bank of Richmond Working Paper*, No. 00-3, August.

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

Haldane, A G and Salmon, C K (1995), 'Three issues on inflation targets', in Haldane A G (ed), *op cit*.

Hall, R E and Mankiw, N G (1994), 'Nominal income targeting', in Mankiw, N G (ed), *Monetary policy, studies in business cycles*, Vol. 29, Chicago and London: University of Chicago Press, pages 71-93.

Hansen, L P and Sargent, T J (1999), *Recursive linear models of dynamic economies*, Chicago: University of Chicago Press.

Henderson, D and McKibbin, W (1993), 'An assessment of some basic monetary policy regime pairs', in Bryant, R *et al* (eds), *op cit*, *Evaluating policy regimes: new research in macroeconomics*, Washington DC, Brookings Institution.

Isard, P, Laxton, D and Eliasson, A C (1999), 'Simple monetary policy rules under model uncertainty', forthcoming in, Isard, P, Razin, A and Rose, A (eds), *International finance and financial crisis: essays in honor of Robert P Flood, Jr,* (Boston, Kluwer and Washington, International Monetary Fund).

Jonung, L (1979), 'Knut Wicksell's norm of price stabilisation and Swedish monetary policy in the 1930s', *Journal of Monetary Economics*, Vol. 5(4), pages 459-96.

Julius, D, Fry, M, Mahadeva, L, Roger, S and Sterne, G (2000), *Monetary policy frameworks in a global context*, Bank of England, October.

Kiley, M T (1998), 'Monetary policy under neoclassical and new-Keynesian Phillips Curves with an application to price level and inflation targeting', Board of Governors of the Federal Reserve System, *mimeo*.

King, M A (1999), 'Challenges for monetary policy: new and old', paper for the symposium on 'new challenges for monetary policy', sponsored by the Federal Reserve Bank of Kansas City at Jackson Hole, Wyoming, 27 August 1999.

Krugman, P (1999), 'Itsbaaack: Japan's slump and the return of the liquidity trap', *Brookings Papers on Economic Activity*, Vol. 2, pages 137-87.

Lebow, D, Roberts, J and Stockton, D (1992), 'Economic performance under price stability', Working Paper No 135, Board of Governors of the Federal Reserve System, Washington DC.

Leiderman, L and Svensson, L E O (1995), Inflation targets, CEPR, London.

Levin, A, Wieland, V and Williams, J C (1999), 'Robustness of simple monetary rules under model uncertainty', in Taylor, J B (ed), *op cit*.

McCallum, B T (1988), 'Robustness properties of a rule for monetary policy', *Carnegie-Rochester Conference Series on Public Policy*, Vol. 29, pages 53-84.

McCallum, B T (1990), 'Could a monetary base rule have prevented the Great Depression?', *Journal of Monetary Economics*, Vol. 26(1), pages 3-26.

McCallum, B T (1993), 'Specification and analysis of a monetary policy rule for Japan', *Bank of Japan Monetary and Economics Studies*, Vol. 11, pages 1-45.

McCallum, B T (1994), 'Monetary policy rules and financial stability', NBER Working Paper, No. 4692, page 33.

McCallum, B T (2000), 'Theoretical analysis regarding a zero lower band on nominal interest rates', *NBER Working Paper*, No. 7677, April.

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

Nelson, E (2000), 'UK monetary policy 1972-97: A guide using Taylor rules', *Bank of England Working Paper*, No 120.

Smets, F (2000), 'What horizon for price stability?', *European Central Bank Working Papers*, No. 24.

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', *mimeo*, IIES, Stockholm University.

Svensson, L E O (1999a), 'Price level targeting versus inflation targeting: a free lunch?', *Journal of Monetary Economics*, forthcoming.

Svensson, L E O (1999b), 'Inflation targeting as a monetary policy rule', *Journal of Monetary Economics*, forthcoming.

Svensson, L E O (2000), 'The zero bound in an open economy. A foolproof way of escaping from a liquidity trap', *NBER Working Paper*, No. 7957, September.

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 (1996), 'Policy rules as a means to a more effective monetary policy', *Bank of Japan Monetary and Economic Studies*, Vol. 14(1), pages 28-39.

Vestin, D (1999), 'Price level targeting versus inflation targeting in a forward-looking model', IIES, Stockholm University, *mimeo*, Stockholm.

Wicksell, K (1935), *Lectures on Political Economy*, Routledge, London.

Williams, J (1999), 'Simple rules for monetary policy', Board of Governors of the Federal Reserve System, finance and economics discussion papers series.

Wolman, A L (1997), 'Zero inflation and the Friedman rule: a welfare comparison', *Federal Reserve Bank of Richmond Economic Quarterly*, 83/4, Fall.

Woodford, M (1999), 'Inflation stabilisation and welfare', *mimeo*, Princeton University, June.