

Rule-based monetary policy under central bank learning*

Abstract

The paper evaluates the performance of three popular monetary policy rules when the central bank is learning about the parameter values of a simple New Keynesian model. The three policies are: (1) the optimal non-inertial rule; (2) the optimal history-dependent rule; (3) the optimal price-level targeting rule. Under rational expectations rules (2) and (3) both implement the fully optimal equilibrium by improving the output-inflation trade off. When imperfect information about the model parameters is introduced, it is found that the central bank makes monetary policy mistakes, which affect welfare to a different degree under the three rules. The optimal history-dependent rule is worst affected and delivers the lowest welfare. Price level targeting performs best under learning and maintains the advantages of conducting policy under commitment. These findings are related to the literatures on feedback control and robustness. The paper argues that adopting integral representations of rules designed under full information is desirable because they deliver the beneficial output-inflation trade-off of commitment policy while being robust to implementation errors.

keywords: monetary policy rules; learning
*JEL:*E31; E5

Kosuke Aoki
Department of Economics, LSE and CEPR.
E-mail: k.aoki@lse.ac.uk

Kalin Nikolov
Conjunctural Assessment and Projections Division
Bank of England
E-mail: kalin.nikolov@bankofengland.co.uk

*We would like to thank Peter Andrews, Luca Benati, Andy Blake, Fabio Canova, V. V. Chari, Tim Cogley, Martin Ellison, George Evans, Jordi Galí, Jan Groen, Albert Marcet, Kaushik Mitra, Katharine Neiss, Thomas Sargent, Lars Svensson, David Vines, Jan Vlieghe, Peter Westaway, Mirko Wiederholt, and Michael Woodford for helpful suggestions and comments. We would also like to thank the participants at the Conference “Phillips Curve Revisited” in Berlin and ISOM Conference in Barcelona for helpful comments and discussions. Aoki thanks to the MCYT and FEDER (project No. SEC2002-03816) for financial support. This paper represents the views and analysis of the authors and should not be thought to represent those of the Bank of England or Monetary Policy Committee members. This is a revised version of our paper under the same title in the Bank of England Working Paper Series (No. 235).

1 Introduction

In this paper, we evaluate the performance of a number of monetary policy rules when the central bank's knowledge about model parameters is imperfect but improves through learning. The recent literature on monetary policy has emphasised the role and importance of the systematic component of monetary policy.¹

The literature has shown that, by committing to conduct monetary policy in a systematic way, the central bank can stabilise inflation and the output gap more efficiently than would otherwise be the case. The advantage of rule-based policy stems from the fact that by committing the central bank can internalise the effect of its predictable policy on private sector expectations.

However, the usefulness of rule-based policies is sometimes criticised from a practical point of view, because optimal rules in general depend on the structure of the economy, such as the slope of the Phillips curve, which in practice is not known with certainty. In the face of this uncertainty, a central bank is continuously learning about the structure of the economy. As a consequence, its best estimates of key model parameters may change. In light of these issues, the question of how to think about policy rules under central bank learning is a key area of research.

In this paper, we use three popular monetary policy rules which perform well in rational expectations models and evaluate their performance in an environment in which the central bank is learning about the structural parameters of the economy. We argue that rules that perform well under learning should be adopted.

We examine the policy problem of a boundedly rational central bank which learns the structural parameters of a simple New Keynesian model. Following Sargent (1999) we assume 'anticipated utility' behaviour of the central bank. More specifically, we assume that both the central bank and private agents learn about the slopes of the IS and Phillips curve by recursive least squares² and adjust expectations and policy in the light of the most recent parameter estimates. The three policy rules we evaluate are (1) the optimal non-inertial rule, (2) the optimal history-dependent rule, and (3) the optimal price level targeting

¹ See, for example, Woodford (1999a) and Taylor (1999).

² Thus we assume symmetric information throughout the paper. See Section 3 for more discussion.

rule.³ In our basic model, under rational expectations and perfect knowledge about model parameters, the last two rules are optimal from a ‘timeless’ perspective, in the sense defined by Woodford (1999a) and Svensson and Woodford (2004). These rules involve history dependence because they optimally internalise the effects of the predictable part of policy on private expectations formation. Both of these rules deliver the optimal equilibrium, although they introduce history dependence in different ways.

However, when we test the performance of these policies under learning we find that imperfect knowledge about the structure of the economy adversely affects the performance of certain policy rules that deliver good economic outcomes in a world of perfect knowledge. The paper shows that imprecise parameter estimates lead to policy mistakes that affect the performance of some history-dependent rules much more than others. In particular, the performance of the optimal history-dependent rule deteriorates substantially. In contrast, the optimal price level targeting rule maintains the benefit of history dependence without generating undesirable feedback from past policy mistakes that worsens the performance of the optimal history-dependent rule.

We also draw some insights from the literature on feedback control. We argue that the optimal price level targeting rule performs best because it has elements of integral control, in the sense that the policy reacts to the integral of past deviations of inflation from its target value. Integral control terms can reduce the propagation of policy mistakes in two ways. First, when mistakes are persistent (due to, for example, persistent mismeasurement of the natural interest rate), the integral terms can reduce the impact effects of the policy errors. Second, history dependence in optimal policy rules generate endogenous persistence in equilibrium dynamics that is independent of persistence in exogenous disturbances, as emphasised in Woodford (1999b). This endogenous persistence may work as a propagation mechanism of policy mistakes for certain policy rules. Policy rules that have elements of integral control terms can reduce this endogenous propagation.

The paper is organised as follows. The next section reviews related literature. Section 3 presents the model of the economy and the problem of optimal monetary policy, and discusses least squares learning

³ These rules are only optimal under rational expectations. They do not capture fully optimal policy under parameter uncertainty.

and candidate policy rules. Section 4 presents some numerical examples to evaluate the policy rules, and give intuition for our results. Section 5 extends the basic model in two ways; it includes an interest rate variability term in the central bank's loss function; and then considers a Phillips curve with inflation inertia. Section 6 concludes.

2 Related literature

We draw insights from three kinds of related literature: optimal monetary policy rules, adaptive learning, and feedback control. In the literature on optimal monetary policy, Giannoni and Woodford (2002a,b) develop a method to derive optimal monetary policy rules that implement the optimal equilibrium in the New Keynesian framework. This approach takes two steps. First, one characterises an optimal allocation. Then one looks for a policy rule that implements the optimal allocation. In general, there are several optimal rules that implement the optimal equilibrium under rational expectations. In other words, those rules perform equally well. However, we show that this equivalence breaks down when the bank's model is not perfect.

Following the literature on adaptive learning, we borrow an idea that stability or robustness under learning can be used as a selection criterion when there are multiple equilibria or multiple ways of implementing policy. The literature on adaptive learning grew rapidly in the 1990s, following the seminal paper by Marcet and Sargent (1989). The literature, which is thoroughly covered by Evans and Honkapohja (2001), explores the consequences of expectations formation by boundedly rational agents. Rational expectations equilibria which are asymptotically reached by boundedly rational agents are known to be 'stable under learning' or 'E-stable'. The criterion of stability under learning therefore narrows the class of plausible outcomes. Bullard and Mitra (2002) and Evans and Honkapohja (2003) use E-stability as a selection criterion for good monetary policy rules in the New Keynesian framework. Evans and Honkapohja (2003) considers several optimal policy rules, and show that a central bank that directly responds to private expectations can achieve a determinate and E-stable equilibrium. By contrast, an instrument rule defined in terms of the underlying economic shocks cannot achieve a learnable equilibrium.

In this paper, we focus on the welfare evaluation of alternative policy rules under learning, rather than long-run E-stability of rational expectations equilibria. We show that the rules which are equivalent under rational expectations equilibrium generate very different outcomes under learning, and we order the rules according to how well they perform under learning. A rule that works well only under rational expectations is less interesting and useful than one that also works well under learning. Our paper can be interpreted as providing a selection criterion for narrowing down the optimal class of policy rules by insisting on a certain kind of robustness to implementation errors under learning.

Finally, in order to explain why some rules are more robust than others, we relate our result to the literature on feedback control. One can usually express a rule in various forms: proportional, derivative, and integral forms. Proportional form depends on the level of goal variables, derivative form on the difference in goal variables, and integral form on sums of the goal variables. In the context of monetary policy, when inflation is one of central bank's goal variables, a policy rule responding to inflation deviation from its target represents proportional feedback, while responding to price level represents integral feedback. As is shown in Franklin et al. (2002), integral control elements improve the performance of feedback rules when, for example, there are errors in estimating the steady state of the system. Phillips (1954) offers an early economic application using a simple dynamic model. In our paper, a rule involving integral term performs better because it reverses past policy mistakes.

3 Model

3.1 Structural equations

The model is a simple variant of the dynamic sticky price models which have often been used in the recent research on monetary policy. The structure of the economy is described by a log-linearised Phillips curve and an expectational IS curve (Kerr and King (1996), Woodford (1996), Clarida et al. (1999), McCallum and Nelson (1999)).

The expectational IS equation is given by

$$y_t = E_t y_{t+1} - \sigma [i_t - E_t \pi_{t+1}] + g_t, \quad \sigma > 0 \tag{1}$$

where y_t , π_t , i_t are, respectively, time t output, inflation, and the nominal interest rate.⁴ The parameter σ can be interpreted as the inverse of the intertemporal elasticity of substitution of expenditure, and the exogenous disturbance g_t represents a demand shock.

The aggregate supply equation is represented by an expectational Phillips curve of the form

$$\pi_t = \kappa(y_t - y_t^n) + \beta E_t \pi_{t+1} + u_t, \quad \kappa > 0, \quad 0 < \beta < 1 \quad (2)$$

where y_t^n is an exogenous supply shock at time t . It represents the natural level of output, which would be the equilibrium level of output if prices were fully flexible. The exogenous disturbance u_t is a cost-push shock.⁵ This Phillips curve can be derived from a log-linear approximation to the first-order condition for the optimal price-setting decision of a firm in Calvo (1983)'s staggered price-setting model. The parameter κ is a function of the speed of price adjustment, and β can be interpreted as the discount factor of price-setters. The expectation operator E_t is discussed in detail in Section 3.4.

3.2 Least squares learning about structural parameters

Next we discuss learning. We assume that the central bank estimates the parameters of its econometric model. Following the literature on adaptive learning, we assume that the central bank updates those parameters by recursive least squares estimation. In our framework the central bank updates the parameters of the model using recursive least squares and take them as the true values. It believes that the parameter estimates will remain unchanged in future and does not take into account the fact that it is likely to revise them subsequently. Sargent (1999) shows that this behaviour is interpreted as what Kreps (1998) calls an anticipated utility model. This is a deviation from full rationality, because the bank does not take account of the effects of their current decisions on future learning, and they ignore period-by-period model misspecification.

We also assume symmetric information between the private agents and the central bank. This implies that the private agents use the same econometric model to make projections about future inflation, output,

⁴ All variables are measured in percentage deviations from their equilibrium values in a steady state with zero inflation.

⁵ See Clarida et al. (1999) and Giannoni (2000).

and the policy rate. This implies that the projections about the future evolution of the economy are identical between the private agents and the central bank. Combined with the anticipated utility behaviour, this implies that the private agents expect that the bank's commitment to a policy rule based on the current knowledge of the economy is credible and time-invariant over time. The assumption of symmetric information is not essential to our results reported below, but it simplifies our analysis. If the private agents were fully rational and had perfect knowledge of the economy, they should be able to predict that the bank will eventually adjust its policy when it learns more. This would imply that the private agents no longer expect the bank would follow its current policy rule for indefinite future. Our assumption rules out this possibility.⁶

We assume that the central bank's econometric model is given by

$$y_t = E_t y_{t+1} - \sigma(t) [i_t - E_t \pi_{t+1}] + g_t^o + \varepsilon_{y,t} \quad (3)$$

$$\pi_t = \kappa(t) (y_t - y_t^{no}) + \beta E_t \pi_{t+1} + u_t^o + \varepsilon_{\pi,t} \quad (4)$$

where each variable is explained below. Here we assume that the bank and private agents do not know the true value κ and σ in (1) and (2), and need to estimate them.⁷ Those estimates are denoted by $\kappa(t)$ and $\sigma(t)$, respectively. For simplicity, we focus on κ and σ by assuming that β is known. We choose this assumption because, compared with κ and σ , it seems that there is less disagreement in the literature about the estimate of β .⁸ We also assume that the aggregate structural shocks involve white-noise unobservable components. In equations (3) and (4), variables g_t^o , y_t^{no} , u_t^o are observable components of the structural shocks. The unobservable components are represented by $\varepsilon_{y,t}$ and $\varepsilon_{\pi,t}$.⁹ We impose this assumption in

⁶ Although in a different context, Hansen and Sargent (2003) also assume that private agents and policymakers have the common model of the economy when they make robust decisions. Our assumption is similar to theirs. As an alternative, one could assume that the private sector projection is based on the true model (1) and (2), but that monetary policy is credible and believed to be time-invariant in each point in time. We have checked that our main results presented below remain robust under this alternative assumption.

⁷ One might think that the private agents should know σ and κ because those represent the parameters of their behavioural equations. The underlying assumption here is that, although each agent knows its own σ and κ , it does not know the values of σ and κ of the other agents, and their average values in the economy. The agents therefore need to estimate them in order to make projections about future aggregate output, inflation, and the interest rate. We also assume agents are identical, but agents do not know that they are all identical.

⁸ A typical estimate of β is around 0.99 quarterly.

⁹ Specifically, we assume that

$$g_t = g_t^o + \varepsilon_{g,t}$$

order to make estimation non-trivial. Without those unobservable components, the bank would be able to identify the true parameter values immediately.

We assume that the bank estimates $\kappa(t)$ and $\sigma(t)$ recursively. Following the literature on adaptive learning, we assume that, at time t , the central bank updates its parameter estimates $(\sigma(t), \kappa(t))$ based on time $t - 1$ data.¹⁰ The estimation equation is therefore given by recursive OLS estimators:

$$\sigma(t) = \sigma(t-1) + t^{-1} R_{\sigma,t}^{-1} X_{\sigma,t-1} (Y_{\sigma,t} - \sigma(t-1) X_{\sigma,t-1}) \quad (5)$$

$$\kappa(t) = \kappa(t-1) + t^{-1} R_{\kappa,t}^{-1} X_{\kappa,t-1} (Y_{\kappa,t} - \kappa(t-1) X_{\kappa,t-1}) \quad (6)$$

where

$$Y_{\sigma,t-1} \equiv y_{t-1} - E_{t-1} y_t - g_{t-1}^o$$

$$X_{\sigma,t-1} \equiv i_{t-1} - E_{t-1} \pi_t$$

$$R_{\sigma,t} \equiv R_{\sigma,t-1} + t^{-1} (X_{\sigma,t-1}^2 - R_{\sigma,t-1})$$

$$Y_{\kappa,t-1} \equiv \pi_{t-1} - \beta E_{t-1} \pi_t - u_{t-1}^o$$

$$X_{\kappa,t-1} \equiv y_{t-1} - y_{t-1}^{no}$$

$$R_{\kappa,t} \equiv R_{\kappa,t-1} + t^{-1} (X_{\kappa,t-1}^2 - R_{\kappa,t-1})$$

In order to focus our analysis on the implications of learning about structural parameters, we assume for simplicity that expectations are also observable.

3.3 Policy problem under learning

We turn to the analysis of monetary policy. We consider policy implementation by committing to policy rules which are optimal from a ‘timeless perspective’ defined in Woodford (1999a) and Svensson and Woodford (2004). ‘Timeless optimality’ provides us with a convenient and practically realistic way of characterising optimal policy under central bank learning. In our setting, a timelessly optimal policy rule

where $\varepsilon_{g,t}$ is white-noise unobservable components. Then we define $\varepsilon_{y,t} \equiv \varepsilon_{g,t}$. Similarly, we define $\varepsilon_{y,t} \equiv \kappa \varepsilon_{yn,t} + \varepsilon_{u,t}$, where $\varepsilon_{yn,t}$ and $\varepsilon_{u,t}$ are white-noise unobservable components of cost-push shock and natural output, respectively.

¹⁰ This assumption is commonly used in the literature on adaptive learning to avoid circularity between learning and equilibrium. See, for example, Evans and Honkapohja (2003).

is a time-invariant policy rule, conditional on the bank's current model of the economy, that implements the optimal state-contingent plan, subject to the constraint on the economy's initial condition that prevents the bank from exploiting existing private sector expectations at the time the policy is chosen. The central bank commits to following a certain fixed rule indefinitely, and is expected to follow this rule unless its best knowledge of the economy improves in subsequent periods.

The welfare loss of the central bank at time t is the expected discounted sum of period loss functions:

$$W_t \equiv E_t \left[\sum_{s=t}^{\infty} \beta^s L_s \right] \quad (7)$$

The period loss function is the weighted sum of the squared output and inflation gaps, given by¹¹

$$L_s = \frac{1}{2} \left[\pi_s^2 + \lambda_y(t) (y_s - y_s^n)^2 \right], \quad \lambda_y(t) \equiv 16\kappa(t)/\theta \quad (8)$$

Rotemberg and Woodford (1997) show that the loss measure (8) is a quadratic approximation to the expected utility of the representative household in the Calvo model, where $\theta > 0$ is the price elasticity of demand for differentiated goods. Since $\lambda_y(t)$ depends on the slope of the Phillips curve, κ , it changes as the bank's estimate of κ changes over time.

Consider a central bank in period t whose most recent estimates of the two structural parameters are $\kappa(t)$ and $\sigma(t)$. Based on those parameters, the central bank at time t seeks to minimise (7) subject to the bank's model of the economy:

$$y_s = E_s y_{s+1} - \sigma(t) [R_s - E_s \pi_{s+1}] + g_s^o \quad (9)$$

$$\pi_s = \kappa(t) (y_s - y_s^{no}) + \beta E_s \pi_{s+1} + u_s^o \quad (10)$$

Note that there are two kinds of misspecification in (9) and (10). First, $\sigma(t)$ and $\kappa(t)$ are not necessarily equal to their true values, σ and κ . Second, only the observable components of the structural shocks are in (9) and (10).

It is worth making clear that policy rules derived in this section are not designed to be optimal under the imperfect knowledge that we assume. The design of fully optimal policies is not the main objective of the

¹¹ Following Woodford (2003), the weight is given by $16\kappa(t)/\theta$ rather than $\kappa(t)/\theta$. Thus the weight is appropriate when inflation is measured as annualised percent changes.

paper. Our objective is to use learning dynamics as a criterion for distinguishing between policy rules from policy rules that are equally good under perfect knowledge. There are two kinds of deviations from full optimality (or rationality). First, the bank takes the estimated parameters as if they are the true parameter values and time-invariant. In other words, the bank is not Bayesian. This is a standard assumption in the literature on adaptive learning, and we are concerned with the question of how robust each rule is under adaptive learning.¹² Second, when choosing policy, the bank regards the white-noise unobservable components of the structural shocks equal to zero. When $\sigma(t) = \sigma$, $\kappa(t) = \kappa$, and when the white-noise components of the shocks are observable, the policy derived below becomes fully optimal.

The problem can be solved by the Lagrange method (Woodford (1999b)). The first-order conditions are given by

$$\pi_s - \beta^{-1}\sigma(t)\phi_{1s-1} + \phi_{2,s} - \phi_{2s-1} = 0 \quad (11)$$

$$\lambda_y(t)(y_s - y_s^{no}) + \phi_{1s} - \beta^{-1}\phi_{1s-1} - \kappa(t)\phi_{2s} = 0 \quad (12)$$

$$\sigma(t)\phi_{1s} = 0 \quad (13)$$

for each $s \geq t$. Here ϕ_{1s} and ϕ_{2s} respectively represent the Lagrange multipliers associated with (9) and (10).

It follows from the first-order conditions that, in the case of the optimal commitment plan that has been implemented since $s - 1$, we can infer ϕ_{1s-1} and ϕ_{2s-1} from the past observable variables π_{s-1} and $y_{s-1} - y_{s-1}^{no}$. Those are given by $\phi_{1s-1} = 0$ and $\phi_{2s-1} = \lambda_y(t)/\kappa(t)(y_{s-1} - y_{s-1}^{no})$. In this case, the first-order conditions imply¹³

$$\pi_s = -\frac{\lambda_y(t)}{\kappa(t)} [(y_s - y_s^{no}) - (y_{s-1} - y_{s-1}^{no})], \quad s \geq t + 1 \quad (14)$$

Next we turn to the initial conditions on ϕ_{1t-1} and ϕ_{2t-1} . A once-and-for-all commitment from period t onwards requires the initial conditions $\phi_{1t-1} = 0$ and $\phi_{2t-1} = 0$, implying that the bank exploits

¹² A fully rational policy would be an optimal policy problem under parameter uncertainty and Bayesian learning. For example, Wieland (2000) studies optimal Bayesian learning in a simple dynamic model.

¹³ In this case, the weight on the output gap term, $16\lambda_y(t)/\kappa(t) = 16/\theta$ and becomes constant. As shown in Section 5, this is not a general result. In general, policy parameters change as the estimated parameters change.

the existing expectations at the time the policy is chosen. However, as Woodford (1999a) and Svensson and Woodford (2004) argue, the once-and-for-all commitment implicit in the optimal equilibrium is not a practically useful concept, when the bank's knowledge about the economy improves over time and, hence, it wishes to revise its policy rule. If the bank exploited the existing expectations whenever they revise the optimal plan, the bank's commitment would not be credible, leading to a suboptimal equilibrium. We instead restrict our policy regime to policies that are optimal from a timeless perspective, following Woodford (1999a) and Svensson and Woodford (2004). In our case, this would imply that we impose the initial conditions on ϕ_{1t-1} and ϕ_{2t-1} as $\phi_{1t-1} = 0$ and $\phi_{2t-1} = \lambda_y(t)/\kappa(t)(y_{t-1} - y_{t-1}^{no})$. Then we obtain a targeting criteria

$$\pi_s = -\frac{\lambda_y(t)}{\kappa(t)} [(y_s - y_s^{no}) - (y_{s-1} - y_{s-1}^{no})], \quad s \geq t \quad (15)$$

Notice that the rule involves the lagged output gap $(y_{t-1} - y_{t-1}^{no})$. As stressed in Woodford (1999b), optimal monetary policy involves history dependence. This results from the fact that the central bank internalises the effects of its predictable policy on private sector expectations, which in turn affect current inflation and the output gap.

Since (15) involves the targeting variables in the central bank's loss function, it is a targeting rule in the sense of Svensson (1997, 2002, 2003). Targeting rules are consolidated first-order conditions, and specified in terms of a criterion that certain target variables should satisfy. The central bank chooses the interest rate so as to satisfy this criterion, and it must use its estimated model of the economy in setting the interest rate. We will discuss the details of implementation in Section 3.4. Finally, it is easy to show that the rational expectations equilibrium is determinate under this rule.¹⁴

Another form of targeting rule we consider which is optimal from a timeless perspective is given by

$$p_s = -\frac{\lambda_y(t)}{\kappa(t)} (y_s - y_s^{no}), \quad s \geq t \quad (16)$$

Rational expectations equilibrium is determinate under this rule, too. Notice that, since $\pi_s \equiv (1 - L)p_s$,

¹⁴ See, for example, Giannoni and Woodford (2002b).

where L represents lag operator,

$$p_s = \sum_{i=0}^{\infty} \pi_{s-i}$$

In our case where the target inflation rate is set equal to zero, p_s represents integral deviations of past inflation from its target. Thus we can express an optimal rule in terms of an integral feedback rule, as is shown in Currie and Levine (1987).¹⁵ Since the price level depends on past history of inflation, this rule also introduces history dependence to monetary policy, by forcing the central bank to compensate any shock that affected inflation in the past.¹⁶ Compared with (15), (16) introduces history dependence without explicitly committing to respond to the lagged output gap. We call it the optimal price level targeting rule. Policy rules (19) and (18) are closely related. Under perfect information, commitment to rule (18) and commitment to rule (19) result in the same rational expectations equilibrium, as is shown in Clarida et al. (1999).

It is possible to find other optimal policy rules than (15) and (16). However, we focus our analysis on those two rules because they are among the simplest rules and extensively analysed in the literature.¹⁷ Evaluation of other optimal rules is left for future research. Instead, as a comparison, we consider a rule that does not involve history dependence. This is called the optimal non-inertial rule:

$$\pi_{t|t} = -\frac{\lambda_y(t)}{\kappa(t)} (1 - \beta\delta_u) (y_{t|t} - y_t^{no}) \quad (17)$$

where δ_u is the persistence parameter of the cost-push shock (defined more precisely later in Section 3.4.) Again, rational expectations equilibrium is uniquely determined. Conceptually, this is similar to a simple Taylor rule that involves only contemporaneous feedback from inflation and the output gap. The coefficient $\frac{\lambda_y(t)}{\kappa(t)} (1 - \beta\delta_u)$ is chosen to implement the optimal non-inertial plan defined in Woodford (2003).¹⁸ Since our structural equations are purely forward-looking and the rule does not involve lagged endogenous vari-

¹⁵ Currie and Levine (1987) show that an optimal commitment policy can be expressed in terms of integral of past predetermined variables. Here we interpret (16) as a targeting rule which involves an integral term of one of the goal variables, namely, π .

¹⁶ See, also Giannoni (2000) and Vestin (1999) for the desirability of price level targeting in forward-looking models.

¹⁷ See, for example, Clarida et al. (1999), Vestin (1999), Svensson (2003), Svensson and Woodford (2004), Woodford (2003).

¹⁸ More specifically, (17) implements the optimal non-inertial plan discussed in Section 3.1 of Woodford (2003). His model is identical to the model employed here. See Woodford (1999b) and Giannoni and Woodford (2002b) for further discussion of the optimal non-inertial plan.

ables, in the resulting equilibrium, the endogenous variables will depend only on the current and expected future shocks and the current model estimation errors.

3.4 Sequence of events

Here we describe the sequence of events at time t . At the beginning of time t , the central bank and the private sector use recursive OLS to update the estimates of the slopes of the IS and Phillips curves, $\sigma(t)$ and $\kappa(t)$. The estimation equations are given by (5) and (6). After they update their parameter estimates, they observe time- t shocks. Their economic model is given by (9) and (10). Based on their model, they form expectations and set policy. Expectations are consistent with the estimated model (9) and (10).¹⁹

When implementing a policy rule, (15) for example, the central bank needs to make its decisions on the projections of the current endogenous variables, such as $\pi_{t|t}$ and $y_{t|t}$ using its model. Here a variable $x_{t|t}$ denotes the projection of x_t based on the central bank's model at time t . Therefore, the bank chooses the nominal interest rate in order to satisfy

$$\pi_{s|t} = -\frac{\lambda_y(t)}{\kappa(t)} \left[\left(y_{s|t} - y_{s|t}^{no} \right) - \left(y_{s-1|t} - y_{s-1|t}^{no} \right) \right], \quad s \geq t \quad (18)$$

When the bank has perfect information, $\pi_{t|t} = \pi_t$, $y_{t|t} = y_t$ (ie the bank's projections and the realisations coincide). For the implementation of the price level targeting rule (16), the bank chooses the nominal interest rate to satisfy

$$p_{s|t} = -\frac{\lambda_y(t)}{\kappa(t)} \left(y_{s|t} - y_{s|t}^{no} \right), \quad s \geq t \quad (19)$$

Implementation of the optimal non-inertial policy (17) is also similar.

After expectations and policy rates are set based on (9) and (10), those are put into the true model (1) and (2), and inflation and the output gap are realised. In other words, the expectation terms in (1) and (2) are generated by (9) and (10).²⁰ In the next section, we provide some numerical examples to evaluate the

¹⁹ Model-consistent expectations by private agents are assumed to keep symmetry. The existing literature on adaptive learning, such as Bullard and Mitra (2002) and Evans and Honkapohja (2003), assumes that agents' expectations are based on their 'perceived law of motion'. Here we do not mean to argue that our approach is more appropriate. However, our primary objective is to investigate performance of policy rules under central bank learning, and it is plausible to assume that the bank's projection is consistent with its model.

²⁰ Recently, Preston (2005) argues that it may not be appropriate to use as a representation of aggregate dynamics the log-linearised IS and Phillips curve equations with expectations terms replaced by arbitrary subjective expectations. He argues that this 'Euler

performance of the three rules under learning.

4 Numerical examples

4.1 Simulation exercises

In this section we use stochastic simulations using our model to evaluate our three candidate ‘targeting rules’. We obtain expected values for the loss function as well as representative paths for our parameter estimates as averages across 1,000 simulated paths. Each simulation length is 500 periods ($\beta^{500} = 0.0066$).

4.1.1 Parameter values

Because our exercise involves a set of wrong initial beliefs about parameter values, we need to calibrate these as well as the true parameter values in our economy. We set the initial parameter values of our model in line with Rotemberg and Woodford (1997) and Woodford (1999b). The intertemporal elasticity of substitution of consumption, σ , is set to 6.365. The slope of the Phillips curve κ is set to 0.0238. β , the rate of time preference is equal to 0.99, giving an annual steady-state real interest rate of 4%.²¹ The corresponding weight on the output gap in the loss function, $\lambda_y = 16\kappa/\theta$, is set to 0.062.

For concreteness, we assume the following stochastic processes for the exogenous shocks:

$$y_t^n = \delta_y y_{t-1}^n + e_{yt}, \quad 0 < \delta_y < 1 \quad (20)$$

$$g_t = \delta_g g_{t-1} + e_{gt}, \quad 0 < \delta_g < 1 \quad (21)$$

$$u_t = \delta_u u_{t-1} + e_{ut}, \quad 0 < \delta_u < 1 \quad (22)$$

where e_{yt} , e_{gt} and e_{ut} are independent and serially uncorrelated at all leads and lags. The autocorrelation

equation approach’ is consistent with optimising behaviour of agents only when expectations are rational (ie model-consistent). He argues that, even after expectations converges to rational expectations, output (consumption) chosen by the Euler equation approach may remain to be suboptimal. See Evans et al. (2003) for further discussion. In our model, expectations are consistent with the estimated model. This means that, as the estimated model converges to the true model, it is possible to show that the level of output chosen by (1) and (2), with the expectation terms generated by (9) and (10), converge to the level consistent with agents’ optimisation. As a robustness check, we also conducted the same simulation exercise for an economy in which the private agents are endowed with the true economic model and the private expectations are based on the true parameter values. We confirmed that the results given in the subsequent sections remain the same. This implies that our results mainly stem from central bank’s learning and are robust to changes in assumption on private sector behaviour.

²¹ We set the initial values of the variances of the estimates of κ and σ equal to, respectively, $2 * 10^{-3}$ and $1.5 * 10^{-4}$. These are very close to the values, to which the variances converge after approximately 100 periods of recursive estimation. So, we are effectively assuming a training sample of 25 years.

coefficients of the shocks are set as follows: δ_g is 0.92, δ_{y_n} is 0.15, δ_u is 0.35.²² We assume that those persistence parameters are known to the central bank, in order to focus on the estimation of σ and κ . The ‘true’ value of σ and κ in the model is equal to 4.490 and 0.031, respectively. These are cited from Giannoni (2001) and are two standard errors away from Woodford (1999b)’s central estimates.²³ In order to check the robustness of the subsequent results, we also examined the other three cases: (σ^{-1} : two standard errors above, κ : two standard errors below); (σ^{-1} : two standard errors below, κ : two standard errors above); (σ^{-1} : two standard errors below, κ : two standard errors below). The welfare ranking of the policy rules reported below is not affected by the assumed starting values of σ and κ , therefore we do not report those simulation results.²⁴ The standard deviations of the innovations in the demand and supply disturbances in the Rotemberg-Woodford model are calculated as 1.022% and 1.906%, respectively.²⁵ The standard deviation of the innovation in the cost-push shock is set 0.37%, cited from Giannoni (2000). Finally, as a benchmark case, we set the variances of unobservable shocks, $\varepsilon_{y,t}$ and $\varepsilon_{\pi,t}$ equal to zero.²⁶ Thus we focus on the implications of imprecise parameter estimates. Allowing non-zero $\varepsilon_{y,t}$ and $\varepsilon_{\pi,t}$ makes learning slower and makes fluctuations in inflation and the output gap larger, but we find that the welfare ranking of the policy rules is not affected.²⁷

4.1.2 Welfare results

Table A presents the expected values of the welfare function under the optimal non-inertial rule (ONP) and the optimal history-dependent rule (OHDP). The first three rows of the table (under the Learning heading)

²² Our definition of the demand disturbances follows Bernanke and Woodford (1997). That is,

$$g_t = E_t G_{t+1} - G_t$$

where G_t is Rotemberg and Woodford (1997)’s definition of the demand disturbance. If we assume that G_t follows an AR(1) process with autoregressive root 0.92, it is easily shown that g_t also follows an AR(1) process with the same autoregressive root.

²³ The estimate in Giannoni (2001) is σ^{-1} instead of σ . We inverted the upper bound of his σ^{-1} .

²⁴ The simulation results not reported in the text are available from the authors upon request.

²⁵ Specifically, these are the standard deviations of the innovations in the processes $E_{t-2} [G_{t+1} - G_t]$ and $E_{t-2} Y_t^S$ where G_t and Y_t^S are Rotemberg-Woodford’s demand and supply disturbances. This is because their structural equations coincide with our simpler model only when conditioned on information available two periods earlier.

²⁶ However, this does not imply that the bank can identify the structural parameters immediately in our simulation, as recursive estimation is given by (5) and (6).

²⁷ Again, those results are available from the authors upon request.

Table A: Welfare loss (The results under the benchmark case)^(a)

Learning (sample average of 1,000 simulations)					
	ONP	OHDP	PLT	ONP/OHDP	ONP/PLT
$W = \sum_{s=0}^{500} \beta^s L_s$	0.0037	0.0114	0.0032	-67.7%	17.0%
$\sum_{s=0}^{500} \beta^s \pi_s^2$	0.0034	0.0063	0.0026		
$\sum_{s=0}^{500} \beta^s \lambda_y (y_s - y_s^n)^2$	0.0003	0.0051	0.0005		
Full information (sample average of 1,000 simulations)					
	ONP	OHDP	PLT	ONP/OHDP	ONP/PLT
$W = \sum_{s=0}^{500} \beta^s L_s$	0.0035	0.0029	0.0029	21.0%	21.0%
$\sum_{s=0}^{500} \beta^s \pi_s^2$	0.0034	0.0025	0.0025		
$\sum_{s=0}^{500} \beta^s \lambda_y (y_s - y_s^n)^2$	0.0001	0.0004	0.0004		

^(a)ONP: Optimal non-inertial plan; OHDP: Optimal history-dependent plan; PLT: Price level targeting.

display the expected welfare when the central bank and private agents have to use recursive least squares in order to learn about the parameters of the model, while the last three rows of the table (under the Full Information heading) contain numbers derived under the assumption that the central bank and the private agents have perfect knowledge.

The results in the lower part of the table are now fairly standard in the literature (Clarida et al. (1999), Woodford (1999b)). The optimal history-dependent rule delivers a lower expected value of the loss function than the optimal non-inertial rule. This is because the history-dependent rule optimally internalises the effects of predictable policy on private agent expectations. This affects inflation expectations and limits the impact effect of cost-push shocks on inflation, thereby improving the output/inflation trade-off faced by the monetary authority. The central bank controls private sector expectations by promises of future action, which are embodied in the lagged output gap term in the optimal history-dependent rule (18). In contrast, the optimal non-inertial rule is restricted to respond only to current variables. It, therefore, cannot optimally internalise its effect on private sector expectations. Consequently, this worsens the inflation/output trade-off caused by cost-push shocks, reducing expected welfare.

The results in the last three rows of Table A confirm another well-known result in the literature.²⁸ The optimal history-dependent rule and the price level targeting rules deliver identical welfare losses under full

²⁸ See Clarida et al. (1999), Giannoni (2000) and Vestin (1999).

information, because both of them are timelessly optimal in our set-up. Because correcting deviations of the price level from a deterministic path involves a strong commitment against inflation, the price level targeting rule is a good policy for a central bank that wants to improve its inflation/output trade-off by stabilising inflation expectations.

However, the first three rows of Table A shows that the advantages of the history-dependent rule over the optimal non-inertial rule are reversed under learning. This is mainly due to a substantial deterioration of welfare under the optimal history-dependent rule.²⁹ By contrast, under learning, the price level targeting rule performs much better than the optimal history-dependent rule. It seems, therefore, that imperfect knowledge of the model parameters breaks down the link between the price level targeting rule and the history-dependent rule as the performance of the latter deteriorates markedly. This is a feature of our results we will explore further in the next subsection.

4.1.3 Impulse responses

The impulse responses of our model to different shocks offers another way of evaluating and explaining the performance of our three candidate rules. This exercise can show exactly which shocks cause the marked deterioration of welfare under the OHDP rule when the central bank is learning about parameter values. Below we show the model's impulse responses to a cost-push shock as well as a demand (IS curve) shock.³⁰

Cost-push shock Figure 1 plots the dynamic response of our model economy to a cost-push shock under the three rules we consider.³¹ It shows that, even under learning, both the price level targeting rule and the optimal history-dependent rule stabilise inflation better than the optimal non-inertial rule under learning. So the standard result from the monetary policy literature holds in our set-up too (Clarida et al. (1999),

²⁹ Because the benefit of the optimal history-dependent rule derives from its ability to improve the output/inflation trade-off, the volatility of cost-push shocks will matter greatly for welfare comparisons. Therefore, even under learning, sufficiently volatile cost-push shocks will mean that OHDP dominates non-inertial policy.

³⁰ Those are impulse responses to a 10 basis point innovation to cost-push shock and demand shock.

³¹ In Figures 1-2, the economy starts with initial values of $\kappa(t)$ and $\sigma(t)$ used in the stochastic simulation. Those charts also show the 'policy mistakes' that are discussed in Section 4.2.

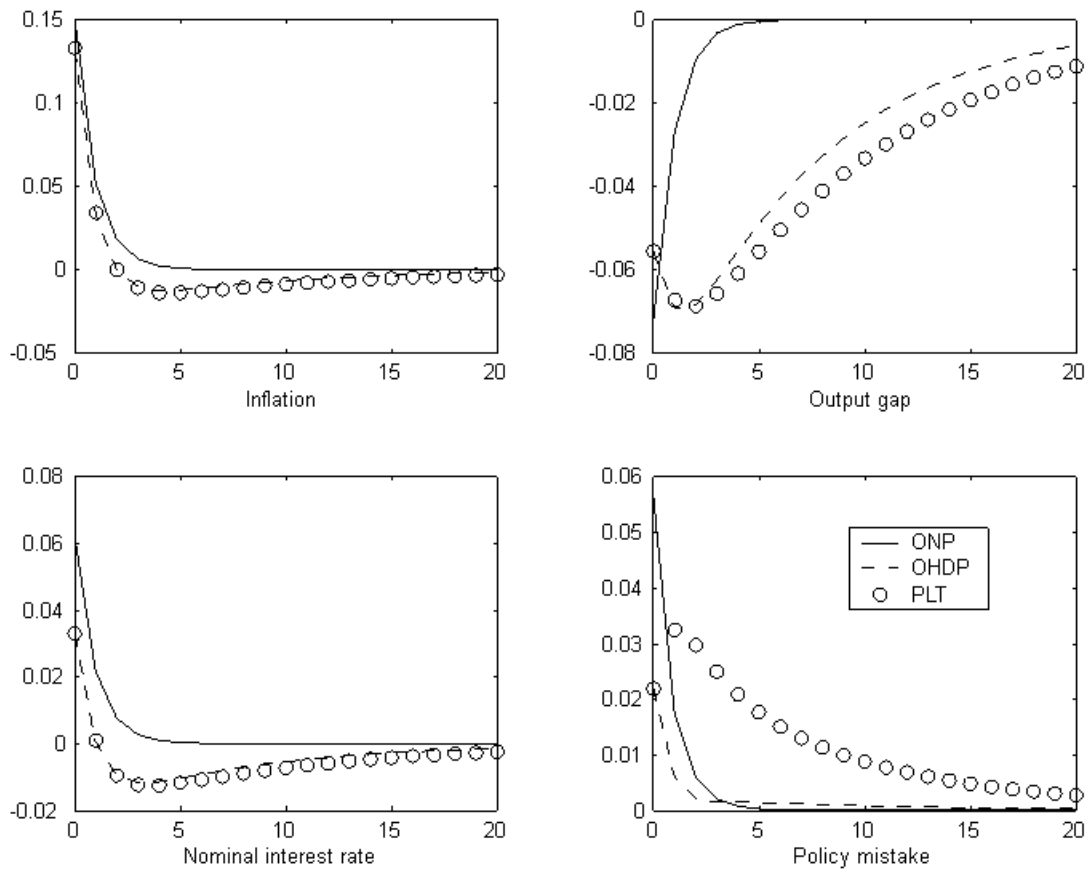


Figure 1: Responses to cost-push shock (benchmark case)

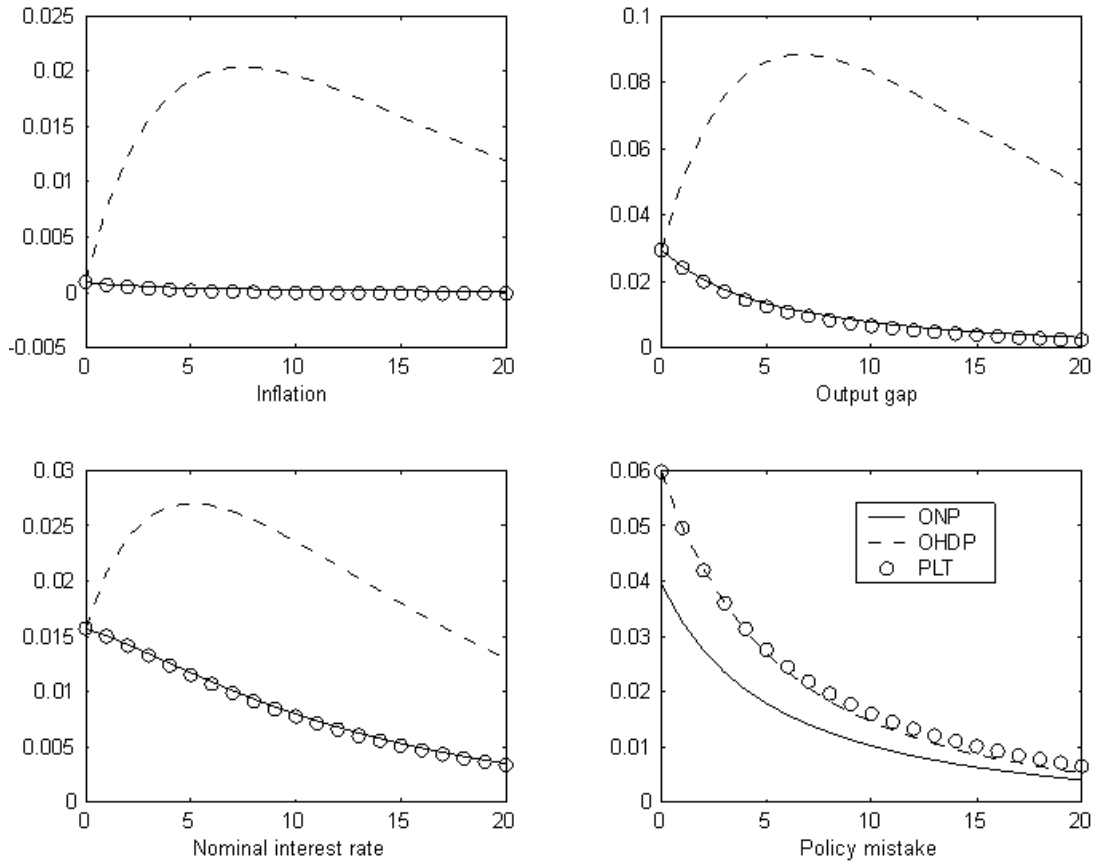


Figure 2: Responses to demand shock (benchmark case)

Giannoni (2000), Vestin (1999)).

Demand shock Figure 2 plots the dynamic response of our model economy to a demand shock under the three rules we consider.³² It is a well-known result in the literature (for example Clarida et al. (1999)) that under full information and rational expectations, the central bank would stabilise demand and supply shocks perfectly. Figure 2 shows that this is not the case under learning. Because of wrong parameter estimates, demand and supply shocks affect output and inflation in the resulting equilibrium. And interestingly, the optimal non-inertial rule and the price level targeting rule dominate the optimal history-dependent rule. Indeed, the performance of the optimal history-dependent rule is particularly poor

³² Properties of the impulse responses to supply shock are similar to those of demand shock. Here we focus our discussion on demand shock.

- demand shocks lead to larger and more persistent fluctuations in inflation and output. It is this poor performance in the face of demand (and supply) shocks that is at the heart of the superior performance of the price level targeting rule.

4.2 History dependence and policy mistakes

In the previous section we argued that, under perfect knowledge, the price level targeting rule and the optimal history-dependent rules deliver a better performance than the non-inertial rule because they affect inflation expectations in a desirable way. In our basic model, those two rules deliver identical equilibria, even though they embody different mechanisms of affecting private expectations. But our simulation results showed that certain ways of implementing history dependence (such as the optimal history-dependent rule) may have undesirable ‘side effects’ under imperfect knowledge, which can be avoided by adopting a different way of implementing history dependence (such as the price level targeting rule). In this section we show how imperfect knowledge of parameters under learning can lead to imprecise estimates of the natural real rate of interest. This results in policy mistakes for all the three rules we consider. However, we will show that where the three rules crucially differ is in how they propagate the effect of the policy mistake through time.

4.2.1 The source of policy mistakes: wrong estimates of the natural real rate of interest

The natural or Wicksellian interest rate is a key concept in New Keynesian models such as the one we study.³³ It is the equilibrium real interest rate, which would prevail under fully flexible prices. In our model, the natural interest rate, r_t^n , is given by

$$r_t^n = \sigma^{-1} [g_t + (\delta_{y_n} - 1) y_t^n]$$

The IS equation (1) can be rewritten in terms of the output gap as

$$y_t - y_t^n = E_t(y_{t+1} - y_{t+1}^n) - \sigma [i_t - E_t\pi_{t+1} - r_t^n]$$

³³ For the recent discussion of this concept, see Blinder (1998) and Woodford (1999b).

This shows that demand and supply shocks only affect the determination of inflation and the output gap through the ‘interest rate gaps’ in our model. In our setting, the central bank could perfectly insulate the effect of demand and supply shocks under all the three targeting rules if it were able to estimate the natural rate accurately.

But in our set-up, the central bank can only observe the primitive demand and supply shocks; it must calculate the natural real rate of interest using its estimate of the slope of the IS curve. Consequently, wrong estimates of σ lead to biased estimate of r^n . This, in turn, means that the central bank cannot ensure that its targeting rule always holds. In fact, as long as the monetary authority is still updating its parameter estimates, these optimality conditions will fail to hold *ex post*. This will act very much like a ‘monetary policy mistake’. Wrong parameter estimates will imply the wrong estimates of the natural real interest rate, leading to biased projections of inflation and the output gap, and, consequently, to the wrong level of policy rates. So, under learning, all shocks will have an effect on economic activity and inflation. Below we make some simplifying assumptions, which allow us to derive analytical expressions for such ‘monetary policy mistakes’. We then explain the intuition of how the different rules propagate these mistakes through time.

The optimality conditions (12), (11), and (13) imply that implementing the optimal plan requires that (15) holds for $s \geq t$. In order to implement (15), the OHDP rule commits to satisfy target (18), while the price level targeting rule commits to satisfy target (19). In both cases, the bank bases its decision on its projections using the estimated model (3) and (4).

Since the estimated model is parameterised imprecisely, the target criteria do not hold exactly. In order to focus on *ex-post* policy mistakes, let us abstract from learning, so that $\sigma(t)$ and $\kappa(t)$ are constant over time and denoted respectively by $\hat{\sigma}$ and $\hat{\kappa}$.³⁴ ³⁵

Notice that the policy rate calculated using the estimated model satisfies

$$\dot{i}_t = \hat{\sigma}^{-1} (y_{t+1|t} - y_{t|t} + g_t) + \pi_{t+1|t} \quad (23)$$

³⁴ We find that the speed of learning about parameters is rather slow, although the estimated parameters converge to their true values. Therefore, the assumption of constant $\sigma(t)$ and $\kappa(t)$ is not a bad approximation when we are thinking about the implications of quarterly monetary policy mistakes for economic fluctuations. In order to save space, we do not report the speed of learning, but it is available upon request.

³⁵ Also, we continue to assume $\varepsilon_{y,t} = \varepsilon_{\pi,t} = 0$.

Substituting this into the true IS (1) and noticing that expectations are symmetric under our assumption of symmetric information, we can calculate the deviation of actual output from the bank's projection as

$$y_t - y_{t|t} = \left(\frac{\sigma}{\hat{\sigma}} - 1\right) [y_{t|t} - y_{t+1|t} - g_t] \quad (24)$$

Similarly, by noticing that

$$\pi_{t|t} = \hat{\kappa} (y_{t|t} - y_t^n) + \beta E_t \pi_{t+1|t} + u_t$$

and using (24), we have

$$\pi_t - \pi_{t|t} = \left(\kappa \frac{\sigma}{\hat{\sigma}} - \hat{\kappa}\right) y_{t|t} + \kappa \left(1 - \frac{\sigma}{\hat{\sigma}}\right) (y_{t+1|t} + g_t) - (\kappa - \hat{\kappa}) y_t^n \quad (25)$$

Substituting (24) and (25) into the OHDP rule (18), we obtain³⁶

$$\pi_t + \frac{\hat{\lambda}_y}{\hat{\kappa}} [(y_t - y_t^n) - (y_{t-1} - y_{t-1}^n)] = \nu_t \quad (26)$$

where

$$\begin{aligned} \nu_t \equiv & \left[\left(\kappa \frac{\sigma}{\hat{\sigma}} - \hat{\kappa}\right) + \frac{\hat{\lambda}_y}{\kappa} \left(\frac{\sigma}{\hat{\sigma}} - 1\right) \right] y_{t|t} + \left(1 - \frac{\sigma}{\hat{\sigma}}\right) \left(\kappa + \frac{\hat{\lambda}_y}{\kappa}\right) y_{t+1|t} \\ & + \left(1 - \frac{\sigma}{\hat{\sigma}}\right) \left(\kappa + \frac{\hat{\lambda}_y}{\kappa}\right) g_t - (\kappa - \hat{\kappa}) y_t^n \end{aligned} \quad (27)$$

The right-hand side of (26) represents the policy mistake when the bank follows the OHDP rule (18). Since the OHDP rule directly corresponds to the optimality condition (15), ν_t also represents the degree to which the optimality conditions fail to hold. Since $y_{t|t}$ and $y_{t+1|t}$ are model-consistent expectations based on the estimated model, those can be expressed in terms of the structural shocks. So the policy mistake (27) is a function of the structural shocks, the parameter estimates, and the true parameter values. It is clear that ν_t converges to zero as $\hat{\sigma} \rightarrow \sigma$ and $\hat{\kappa} \rightarrow \kappa$.

Similarly, we can calculate the policy mistake when the bank commits to targeting rule (19). Using the IS and Phillips curves, we can write:

$$p_t - p_{t|t} = \left(\kappa \frac{\sigma}{\hat{\sigma}} - \hat{\kappa}\right) y_{t|t} + \kappa \left(1 - \frac{\sigma}{\hat{\sigma}}\right) (y_{t+1|t} + g_t) - (\kappa - \hat{\kappa}) y_t^n \quad (28)$$

³⁶ Here $\hat{\lambda}_y \equiv 16\hat{\kappa}/\theta$.

And using (27) we show that the policy mistake is identical to that under the OHDP rule when the central bank follows the price level targeting rule:

$$p_t + \frac{\hat{\lambda}_y}{\hat{\kappa}} [(y_t - y_t^n)] = \nu_t$$

However, even though the central bank is committing itself to a different rule, its preferences and its optimality conditions remain the same. Therefore, although the implementation error, ν_t is the same across the price level targeting and the OHDP rules, the degree to which the optimality condition (15) fails to hold is different and is given by

$$\pi_s = -\frac{\hat{\lambda}_y}{\hat{\kappa}} [(y_s - y_s^n) - (y_{s-1} - y_{s-1}^n)] + \nu_t - \nu_{t-1} \quad (29)$$

In the next section we will see that this is crucial for the relative performance of the two rules under learning.

4.2.2 Intuition: the two benefits of integral control in forward-looking models

In the previous subsection we showed that when the central bank is learning about the parameter values of its model, it cannot implement the optimal equilibrium exactly, leading to higher welfare losses and a breakdown in the equivalence between different methods of implementing history-dependent monetary policy. In this section we offer a ‘classical control’ theory explanation for why this is the case. We first recall a well-known result in the engineering literature that steady-state errors can be corrected by feeding back from the integral of past target misses.³⁷ We then argue that the inertial nature of the optimal equilibrium in forward-looking models brings additional benefits of following integral control policies compared to the backward-looking dynamic models normally studied in the classical control theory literature.

The traditional ‘engineering’ argument for integral control

In general the implementation error, ν_t will be a persistent process. As (27) shows, this persistence comes from two sources - the persistence of the structural shocks, g_t and y_t^n , but also the persistence of y_t ,

³⁷ See, for example, Franklin et al. (2002).

which arises due to the history-dependent nature of the optimal policy. Consequently, the first difference of the monetary policy mistake ($\nu_t - \nu_{t-1}$) is likely to be substantially smaller than its level, ν_t . As equations (26) and (29) show this certainly seems to be part of the reason why the price level targeting rule performs better than the OHDP rule under learning. Under the OHDP rule the mistake enters in levels whereas under price level targeting rule, the policy mistake enters in first differences - it is 'undone' one period later.

This feature of the price level targeting rule is an implication of targeting the 'integral' of deviations of target variables from their target values, and as Franklin et al. (2002) show, it can substantially improve the performance of feedback rules in mechanical systems. In particular, when the rule is potentially subject to errors in estimating the steady state of the system, an integral control term can help stabilise the system at the target value, while a proportional control rule may lead to convergence to the wrong steady-state value for the target variables. These arguments are certainly not new and Phillips (1954) offers an early economic application using a simple dynamic model.

The price level targeting rule has elements of proportional and integral control because it targets the price level (the integral of past inflation deviations from target) and the level of the output gap. The OHDP rule, on the other hand, combines elements of proportional control (the level of inflation) and derivative control (the change in the output gap). This makes the OHDP rule vulnerable to persistent control errors as Figure 2 showed in the previous section. This vulnerability has important effects on welfare in our model because the demand shocks are highly persistent (their autocorrelation coefficient is equal to 0.92) and this causes persistent monetary policy mistakes.

The benefits of integral control under forward-looking behaviour

We showed above that integral control carries substantial benefits, which arise out of persistent (or permanent) unobservable shocks. In our case, persistent demand shocks (although observable) led to persistent errors in estimating the natural rate of interest, which led to difficulties in implementing the optimal equilibrium. We then found that the integral control terms in the price level targeting rule helped to reduce the impact of the errors on the first-order conditions of the central bank's maximisation problem and this improved welfare.

However, our model is forward-looking and this brings additional benefits of implementing history dependence by means of integral control methods. This additional benefit arises out of the persistence of output and inflation, which is induced by history-dependent monetary policy in the optimal equilibrium. This allows the central bank to manage private sector expectations and improves the output/inflation trade-off. However, we will show below that this policy-induced persistence, which is unambiguously good for welfare under cost-push shocks, can have undesirable effects when implemented by proportional-derivative control (the OHDP rule) by a monetary policy authority that is unable to implement the first-order conditions of its maximisation problem exactly. These undesirable effects can be reduced when history dependence is implemented through proportional-integral control terms, which is what the price level targeting rule does.

To demonstrate analytically the benefits of implementing history dependence through integral control methods, here we consider a simple model that abstracts from learning and instead assumes that the central bank as well as the private sector have full information about the model parameters.³⁸ However, we will assume that the central bank can only achieve its target up to a white-noise control error. Although the assumptions are somewhat different from those in the benchmark model used in Section 4.1, this model helps to obtain clearer intuitions behind our simulation results.³⁹ So, consider the following simplified model which consists of the Phillips curve and the optimality condition that now includes the white-noise disturbance term. In the case of the OHDP rule, the corresponding system is:

$$\pi_t = \kappa (y_t - y_t^n) + \beta E_t \pi_{t+1} + u_t \quad (30)$$

$$\pi_t + \frac{\lambda_y}{\kappa} [(y_t - y_t^n) - (y_{t-1} - y_{t-1}^n)] = \varepsilon_t \quad (31)$$

Equation (31) corresponds to (26). In the case of the price level targeting rule, equation (31) is replaced by

$$\pi_t + \frac{\lambda_y}{\kappa} [(y_t - y_t^n) - (y_{t-1} - y_{t-1}^n)] = \varepsilon_t - \varepsilon_{t-1} \quad (32)$$

³⁸ Also, we keep assuming that $\varepsilon_{y,t} = \varepsilon_{\pi,t} = 0$.

³⁹ As $\kappa(t)$ and $\sigma(t)$ become close to their true values, the behaviour of the benchmark economy become close to this simplified model.

which corresponds to (29). If we solve the system (30) and (31), we have

$$y_t - y_t^n = \lambda_1 (y_{t-1} - y_{t-1}^n) + \lambda_1 \frac{\kappa}{\lambda_y} \left(\varepsilon_t - \frac{u_t}{1 - \beta\lambda_1\delta_u} \right) \quad (33)$$

where λ_1 is the stable root of the characteristic equation

$$\lambda^2 - (\beta^{-1} + 1 + \beta^{-1}\kappa^2/\lambda_y) \lambda + \beta^{-1} = 0$$

For the system (30) and (32), the solution is

$$y_t - y_t^n = \lambda_1 (y_{t-1} - y_{t-1}^n) + \lambda_1 \frac{\kappa}{\lambda_y} \left(-\varepsilon_{t-1} + [1 + \beta(1 - \lambda_1)] \varepsilon_t - \frac{u_t}{1 - \beta\lambda_1\delta_u} \right) \quad (34)$$

Note that the responses of the output gap to cost-push shock are identical between (33) and (34). Under the OHDP rule, impulse responses of the output gap to disturbance ε_t is⁴⁰

$$y_{t+i} - y_{t+i}^n = \lambda_1^i \frac{\kappa}{\lambda_y} \varepsilon_t, \quad i \geq 0 \quad (35)$$

while under the price level targeting rule, it is given by

$$\begin{aligned} y_{t+i} - y_{t+i}^n &= \lambda_1^i \frac{\kappa}{\lambda_y} [1 + \beta(1 - \lambda_1)] \varepsilon_t, \quad i = 0 \\ &= \lambda_1^{i-1} \frac{\kappa}{\lambda_y} (1 - \lambda_1)(\beta\lambda_1 - 1) \varepsilon_t, \quad i \geq 1 \end{aligned} \quad (36)$$

Looking at our solutions for the path of the output gap under our two rules we can see that the path of the output gap exhibits persistence under both rules even when shocks are white noise. This persistence arises due to the history dependence of the optimal policy, and the rate of decay of output gap fluctuations is governed by λ_1 . And as we can see from (35) when this degree of optimal policy-induced persistence is large, monetary policy mistakes will be propagated through time by the OHDP rule, reducing welfare.

However closer examination of (36) reveals that matters are more complicated under the price level targeting rule. Because of the integral control properties of the rule, as λ_1 gets large,⁴¹ the equations actually imply that the impact of ε_t on the output gap (and therefore on inflation) becomes smaller, not

⁴⁰ Here we assume that the economy is in the steady state at time $t - 1$.

⁴¹ In our calibration, λ_1 is given by 0.87, implying substantial policy-induced persistence. Persistence can be even higher under higher output-gap weights in the welfare function. For example λ_1 tends to unity when λ_y tends to infinity.

larger. In the limiting case in which $\lambda_1 \rightarrow 1$, the effect of policy mistakes under the price level targeting rule vanishes completely after one period. This feature of the price level targeting rule again arises out of the integral control terms in the rule. This means that the rule will offset most of the endogenous propagation of monetary policy mistakes, while still preserving the benefits of history dependence for the stabilisation of cost-push shocks. Note that this benefit of the price level targeting rule arises entirely out of the inertial nature of the optimal equilibrium in forward-looking models. It is entirely independent of the time series properties of the control error, which is the benefit of integral control methods usually emphasised by the engineering literature.⁴²

This result has one important implication. It can easily be shown that λ_1 is increasing in λ_y , the weight on output gap stabilisation in the central bank's loss function. So a central bank with a strong preference for output gap stabilisation would be more vulnerable to policy errors under the OHDP rule. In contrast, we can see from (36) that the effects of policy mistakes under price-level targeting vanish as $\lambda_y \rightarrow \infty$ and $\lambda_1 \rightarrow 1$. This means that a central bank that is concerned with output gap stabilisation can still get very close to the rational expectations benchmark as long as it implements the optimal policy through integral rather than derivative control methods. This is in contrast with the findings of Orphanides and Williams (2004) who show that, under learning, a central bank with a strong preference for output gap stabilisation will induce near unit root behaviour for output and inflation.⁴³ Of course, their results are based on a model, in which agents form expectations through least squares learning, whereas in our framework, agents form model-consistent expectations conditional upon their latest parameter estimates. One extension, which we leave for future research, could be to check whether integral control methods can prove robust to the effects of adaptive learning in a model such as the one used by Orphanides and Williams (2004).

⁴² In fact, our numerical results suggest that this channel is more important than the traditional channel which relies on exploiting the time-series properties of the monetary policy mistakes. Our simulated policy mistake series had relatively similar variances across the three policy rules. From this we concluded that in our framework, it is the forward-looking nature of the model that is driving the benefits of implementing the optimal policy by integral control methods.

⁴³ Orphanides and Williams (2004) show that econometric learning makes a strong response to inflation deviations much more important than under rational expectations. The paper shows that a weak response to inflation when agents are learning can lead to near unit root behaviour in inflation and the output gap in a small macroeconomic model.

5 Extensions

The previous section showed that integral control representations of timelessly optimal policies were more robust to small model misspecifications than derivative control representations. In this section we depart from the stylised model of Section 3 and consider two empirically motivated extensions in order to demonstrate the generality of our results. Following Woodford (2003), we introduce an interest rate variability term in the central bank's objective function and add inflation indexation into the Phillips curve.

These modifications lead to a framework in which stabilising the price level is no longer timelessly optimal.⁴⁴ However, we still show that integral representations of the optimal rule dominate derivative ones when the central bank is learning about the model's parameter values.

5.1 Interest rate variability objective

Introducing an interest rate variability objective leads to a period loss function as follows.

$$L_t = \frac{1}{2} \left[\pi_t^2 + \lambda_y (y_t - y_t^n)^2 + \lambda_i (i_t - \bar{i})^2 \right]$$

Woodford (2003) shows that this can be justified as a welfare-based loss function when monetary transactions frictions are not negligible. In this case, the first-order conditions for the central bank's maximisation problem are as follows (assuming for simplicity that $\bar{i} = 0$):

$$\pi_s - \beta^{-1} \sigma \phi_{1s-1} + \phi_{2,s} - \phi_{2s-1} = 0 \quad (37)$$

$$\lambda_y (y_s - y_s^n) + \phi_{1s} - \beta^{-1} \phi_{1s-1} - \kappa \phi_{2s} = 0 \quad (38)$$

$$\lambda_i i_t + \sigma \phi_{1s} = 0 \quad (39)$$

The main difference between (39) above and (13) lies in the fact that now the presence of an interest rate variability term constrains the ability of the central bank to stabilise all demand and supply shocks. Because varying the nominal interest rate is costly in terms of welfare, the central bank will optimally trade off some inflation and output variability against the costs of interest rate variability. What this implies is

⁴⁴ See Giannoni (2000) for this point.

that the benefit of commitment is even greater than under the welfare function we considered in Section 3. This is because affecting private sector expectations now allows the central bank to stabilise demand and supply shocks at a lower cost in terms of interest rate variability - an added benefit of commitment relative to the case of no interest rate variability objective.

We again focus on the timelessly optimal policy. From (37), (38), and (39), the optimal history-dependent rule in this case is given by:

$$i_t = \rho_1 i_{t-1} + \rho_2 (i_{t-1} - i_{t-2}) + \rho_\pi \pi_t + \rho_y [(y_t - y_t^n) - (y_{t-1} - y_{t-1}^n)] \quad (40)$$

where $\rho_1 = 1 + \kappa/\beta\sigma$, $\rho_2 = 1/\beta$, $\rho_\pi = \kappa/\lambda_i\sigma$ and $\rho_y = \lambda_y/\lambda_i\sigma$.

A natural extension of our analysis could be a rule of the form:

$$i = \theta_\pi p_t + \theta_y (y_t - y_t^n)$$

But as shown by Giannoni (2000), when there is an interest rate variability objective, price level targeting no longer implements the optimal equilibrium. However, following our results in Section 3 we want to explore price level targeting methods of implementing the optimal degree of history dependence. In other words, we want to find an integral control formulation for our optimal rule.

We can rewrite our optimal history-dependent rule in a way that resembles a price level targeting rule.

Using lag operators we can write

$$i_t = \rho_1 i_{t-1} + \rho_2 (i_{t-1} - i_{t-2}) + \rho_\pi (1 - L) p_t + \rho_y (1 - L) (y_t - y_t^n)$$

This, in turn, implies that

$$R_t = \rho_1 R_{t-1} + \rho_2 (R_{t-1} - R_{t-2}) + \rho_\pi p_t + \rho_y (y_t - y_t^n) \quad (41)$$

where

$$R_t = \frac{i_t}{1 - L} = \sum_{s=0}^{\infty} i_{t-s}$$

is a sum of all past nominal interest rates. We can now rewrite (41) in a more intuitive interest rate rule form:

$$i_t = \rho_2 i_{t-1} + (\rho_1 - 1) R_{t-1} + \rho_\pi p_t + \rho_y (y_t - y_t^n) \quad (42)$$

Table B: Welfare loss (The results with interest rate variability)^(a)

Learning (sample average of 1,000 simulations)			
	OHDP	PLT	PLT/OHDP
$W = \sum_{s=0}^{500} \beta^s L_s$	0.0112	0.0035	-68.8%
$\sum_{s=0}^{500} \beta^s \pi_s^2$	0.0056	0.0026	
$\sum_{s=0}^{500} \beta^s \lambda_y (y_s - y_s^n)^2$	0.0048	0.0006	
$\sum_{s=0}^{500} \beta^s \lambda_i i_s^2$	0.0008	0.0003	
Full information (sample average of 1,000 simulations)			
	OHDP	PLT	PLT/OHDP
$W = \sum_{s=0}^{500} \beta^s L_s$	0.0033	0.0033	0%
$\sum_{s=0}^{500} \beta^s \pi_s^2$	0.0026	0.0026	
$\sum_{s=0}^{500} \beta^s \lambda_y (y_s - y_s^n)^2$	0.0004	0.0004	
$\sum_{s=0}^{500} \beta^s \lambda_i i_s^2$	0.0004	0.0004	

^(a)OHDP: Optimal history-dependent plan; PLT: Price level targeting.

Both (41) and (42) are optimal from a timeless perspective under rational expectations.

Whereas (40) only has elements of proportional control (the terms in inflation and the interest rate) and derivative control (the terms in the change in interest rates and the change in the output gap), rule (42) also has elements of integral control. These elements are introduced by the terms in R_{t-1} and also the price level targeting term. Note that under full information this ‘quasi-price-level targeting rule’ is just a transform of (40) and, therefore, it implements the optimal equilibrium in exactly the same way as (40). Without control errors there is no benefit of an integral control formulation of our rule. However as we know from Section 3, imperfect knowledge of parameter values introduces policy mistakes into the central bank’s decision-making. So again we expect to find that (42) delivers a superior equilibrium to (40) under learning. Below we again use stochastic simulations to test our hypothesis in the case when the central bank has an interest rate variability objective.

Table B presents the value of the welfare loss under the derivative control (40) and the integral control representations of the timelessly optimal rule (42). We calibrate the weight on the interest rate term of the central bank’s loss function by appealing to the transactions frictions discussed in Woodford (2003)⁴⁵.

⁴⁵ Woodford (2003) argues that the threat of hitting the zero lower bound may be an alternative motivation for an interest rate variability objective for the central bank.

We assume that utility is separable between money balances and consumption,⁴⁶ so as Woodford (2003) shows, the optimal interest rate weight is given by:

$$\lambda_i = \frac{\eta_i}{v\epsilon_{mc}} \frac{16\kappa}{\theta} \quad (43)$$

where η_i is the elasticity of money demand with respect to the nominal interest rate, ϵ_{mc} is the elasticity of real marginal cost with respect to the output gap, and v is the steady state velocity of circulation of high powered money. We assume that the only parameter in (43), which is unknown to the central bank is the slope of the Phillips curve κ . Consequently, the monetary authority will be changing the interest rate weight in its loss function only to the extent that it learns about κ .

We calibrate the parameters in (43) in order to match the initial interest rate weight of the central bank to Woodford (2003)'s value for λ_i of 0.077. We set η_i equal to 28, ϵ_{mc} equal to 0.65 and θ equal to 8. The initial belief about κ is 0.0238. To match Woodford's interest rate weight calibration we set v equal to 26.63, implying a ratio of high powered money to nominal GDP of 3.75%. So our central bank will begin with an interest rate weight of 0.077, but eventually it will learn the true value of κ - 0.0308 and increase its λ_i to the true interest rate weight of 0.10.

We get very similar results to the ones we obtained in Section 3. Under the true parameter values, the two rules deliver identical performances. However, the table shows that the performance of the proportional-derivative control rule deteriorates substantially relative to the quasi-price level targeting rule when we allow for imperfect knowledge of the model's parameter values.⁴⁷ So, our simulation results confirm the intuition from Section 3. When the central bank's policy rule is potentially subject to persistent policy mistakes, the presence of integral control terms helps to offset the effects of these errors, improving welfare. This is the reason why our quasi price level targeting rule (41) (which has elements of integral control) delivers higher welfare than the optimal history-dependent rule (42) (which does not).

⁴⁶ This value is based on a model of non-separable utility between money balances and consumption. Such an assumption will complicate the model by introducing an interest rate term in the marginal cost function without changing the qualitative nature of our results. So, for the sake of simplicity, we decided to maintain our assumption of separable utility.

⁴⁷ Indeed, our simulation results show that imperfect knowledge of the model parameter affects the price level targeting rule only marginally. The welfare loss becomes marginally larger.

5.2 Inflation inertia

We now consider another empirically motivated extension, which adds inflation inertia into the simple Calvo (1983) Phillips curve we considered in Section 3. Since there is lively debate on forward-lookingness of inflation dynamics, this exercise would be of particular interest in practice.⁴⁸ To preserve the clarity and simplicity of our analysis we now set the interest rate variability term equal to zero in order to focus on the effects of inflation inertia.

Following Giannoni and Woodford (2002b) we allow firms that cannot re-optimize to index their price to an average of past inflation and steady-state inflation. Let γ denote the weight on lagged inflation in the firm's indexation rule. Giannoni and Woodford (2002b) show that under such an indexation rule, the firm's log-linearised optimal pricing condition is given by:

$$\pi_t - \gamma\pi_{t-1} = \kappa(y_t - y_t^n) + \beta(E_t\pi_{t+1} - \gamma\pi_t) + u_t \quad (44)$$

In this case, Giannoni and Woodford (2002b) show that the utility-based loss function of the central bank is given by:

$$L_t = \frac{1}{2} \left[(\pi_t - \gamma\pi_{t-1})^2 + \lambda_y (y_t - y_t^n)^2 \right]$$

and the corresponding timeless optimal rule is:

$$\pi_t - \gamma\pi_{t-1} = -\frac{\lambda_y}{\kappa} [(y_t - y_t^n) - (y_{t-1} - y_{t-1}^n)] \quad (45)$$

The integral representation of this rule is then given by

$$(1 - \gamma)p_t - \gamma\pi_t = -\frac{\lambda_y}{\kappa} (y_t - y_t^n) \quad (46)$$

So the timeless optimal policy with inflation inertia is to target a weighted average of the price level and the rate of inflation. This hybrid rule converges to the price level targeting rule when price-setters do not index their prices ($\gamma = 0$), while it only aims to stabilise the inflation rate when indexation is full ($\gamma = 1$).

Table C displays the simulation results in the model with inflation inertia. We calibrate the starting value of γ at 0.75 while we set the true value of γ at 0.5. The results confirm our earlier intuition. Integral

⁴⁸ See, for example, Fuhrer (1997) and Galí and Gertler (1999).

Table C: Welfare Loss (The results with inflation inertia)^(a)

Learning (sample average of 1,000 simulations)			
	OHDP	PLT	PLT/OHDP
$W = \sum_{s=0}^{500} \beta^s L_s$	0.0109	0.0031	-71.6%
$\sum_{s=0}^{500} \beta^s (\pi_s - \gamma\pi_{s-1})^2$	0.0062	0.0026	
$\sum_{s=0}^{500} \beta^s \lambda_y (y_s - y_s^n)^2$	0.0048	0.0005	
Full information (sample average of 1,000 simulations)			
	OHDP	PLT	PLT/OHDP
$W = \sum_{s=0}^{500} \beta^s L_s$	0.0029	0.0029	0%
$\sum_{s=0}^{500} \beta^s (\pi_s - \gamma\pi_{s-1})^2$	0.0025	0.0025	
$\sum_{s=0}^{500} \beta^s \lambda_y (y_s - y_s^n)^2$	0.0004	0.0004	

^(a)OHDP: Optimal history-dependent plan; PLT: Price level targeting.

and derivative formulations of the optimal policy perform equally well under rational expectations but the integral control formulation is more robust to model misspecification.

6 Conclusion

In this paper, we combine insights from three kinds of literature: optimal monetary policy rules; learning; and feedback control. When agents are forward-looking, the optimal policy involves history dependence, as advocated by Woodford (1999b). History dependence can internalise the effects of predictable policy on private agent expectations. In general, there is no unique way of implementing optimal policy. When the central bank has perfect knowledge, many different policy rules can implement the optimal equilibrium. However we find that this equivalence breaks down when there is imperfect knowledge about the structure of the economy. In particular, the way a policy rule incorporates history dependence has important implications for stabilisation policy when the central bank's knowledge of the economy is not perfect. We use the performance of policy rules under learning as a selection criterion for narrowing down optimal rules. We find that rules that involve integral terms are more robust to policy mistakes caused by imprecise parameter estimates. For example, this class of policy includes rules that involve price level, which is the integral of inflation. Those rules automatically 'undo' past policy mistakes, while keeping the advantage of history-dependent monetary policy.

It is worth comparing this robustness to another kind of robustness proposed by the literature on robust monetary policy, such as Hansen and Sargent (2003) and Giannoni (2001). Their approach examines the choices facing a policymaker who is concerned that his model may be misspecified but cannot put a probability distribution over the possible alternative models. But crucially, the policy problem examined in Hansen-Sargent is one in which the policymaker can implement his policy rule precisely because he does not have to base his policy on model-based forecasts. In contrast, the policymaker in our paper can only achieve the optimal policy up to an expectational error, which is persistent due to the non-rational nature of expectations in our framework. Such errors are absent from the world of Hansen and Sargent (2003) and Giannoni (2001) and, therefore, their policymakers do not need to be concerned with designing policy in ways that are robust against them. In practice, of course, most central banks use forecasts from some sort of wrong model to inform policy decisions and we believe that the kind of robustness we discuss in our paper is also very important.

It would also be interesting to compare our results with a Bayesian optimal policy under parameter uncertainty for forward-looking models. The analysis of optimal policy rules under Bayesian model averaging is left for future research. However, based on the results shown in Aoki (2005), we expect that a rule that involves integral terms could be a good approximation to the optimal Bayesian policy in forward-looking models. Using a simple New Keynesian model, Aoki (2005) considers an optimal commitment policy when the bank's information about the state of the economy is imperfect, while the agents have perfect information. It is shown that an optimal plan makes the current policy slightly expansionary if it turns out that the policy in the previous period was too contractionary, and *vice versa*. In other words, the optimal policy has the overshooting property, whereby the monetary authority commits to undo past policy mistakes caused by information problems. When private agents understand this overshooting, they adjust their expectations and react less to monetary policy mistakes. In our setting, the rules that involve integral terms have this overshooting property.

There are a number of ways to extend our analysis. First, we assumed that all of the variables are observable except $\varepsilon_{y,t}$ and $\varepsilon_{\pi,t}$. It would be more realistic to assume more serious uncertainty about

some of the variables, such as potential output. Second, we assumed that the central bank's model is correctly specified, and the central bank learns the structural parameters by recursive least squares. It would be interesting to change those assumptions. We could consider a case in which the bank's model is misspecified. Investigation of robustness of policy rules under model misspecification would be an interesting direction to future research. For example, we could consider a case in which the bank's model includes a backward-looking Phillips curve, while the true inflation dynamics involves forward-looking behaviour, or *vice versa*. Since there is lively debate on forward-lookingness of inflation dynamics, this exercise would be of particular interest in practice.⁴⁹ Third, it would be interesting to consider a case in which the bank estimates the parameters by constant-gain learning, as in Sargent (1999) and Cho et al. (2001). Compared with least squares learning, constant-gain learning puts more weight on most recent data, and is suitable for a central bank that suspects there is some kind of structural change in the economy. In this case, when the bank's economic model is misspecified, there would be a potential for escape dynamics, which may behave differently under different policy rules. One could also study how different rules alter the frequency and character of the escape dynamics.

⁴⁹ See, for example, Fuhrer (1997) and Galí and Gertler (1999).

References

- Aoki, K., 2005. Optimal commitment plan under noisy information. *Journal of Economic Dynamics and Control* forthcoming.
- Bernanke, B., Woodford, M., 1997. Inflation forecasts and monetary policy. *Journal of Money, Credit, and Banking* 29 (4), 653–684.
- Blinder, A., 1998. *Central Banking in Theory and Practice*. MIT Press, Cambridge.
- Bullard, J., Mitra, K., 2002. Learning about monetary policy rules. *Journal of Monetary Economics* 49, 1105–1129.
- Calvo, G., 1983. Staggered prices in a utility-maximizing framework. *Journal of Monetary Economics* 12 (3), 383–398.
- Cho, I.-K., Williams, N., Sargent, T., 2001. Escaping nash inflation. *Review of Economic Studies* 69 (1), 1–40.
- Clarida, R., Gali, J., Gertler, M., 1999. The science of monetary policy: A new keynesian perspective. *Journal of Economic Literature* 37 (4), 1661–1707.
- Currie, D., Levine, P., 1987. The design of feedback rules in linear stochastic rational expectations models. *Journal of Economics Dynamics and Control* 11, 1–28.
- Evans, G., Honkapohja, S., 2001. *Learning and Expectations in Macroeconomics*. Princeton University Press, New Jersey.
- Evans, G., Honkapohja, S., 2003. Expectations and the stability problem of optimal policies. *Review of Economic Studies* 70, 807–824.
- Evans, G., Honkapohja, S., Mitra, K., 2003. Notes on agents' behavioral rules under adaptive learning and recent studies of monetary policy. Working Paper, University of Oregon.
- Franklin, G., Powell, D., Emami-Naeini, A., 2002. *Feedback Control of Dynamic Systems*, 4th Edition. Prentice Hall, New Jersey.
- Fuhrer, J., 1997. The (un)importance of forward-looking behavior in price specifications. *Journal of Money, Credit, and Banking* 29 (3), 338–350.
- Gali, J., Gertler, M., 1999. Inflation dynamics: A structural econometric analysis. *Journal of Monetary Economics* 44 (2), 195–222.
- Giannoni, M., 2000. Optimal interest-rate rules in a forward-looking model, and inflation stabilization versus price-level stabilization. Working Paper, Columbia University.
- Giannoni, M., 2001. Does model uncertainty justify caution? robust optimal monetary policy in a forward-looking model. *Macroeconomic Dynamics* 6 (1), 111–144.
- Giannoni, M., Woodford, M., 2002a. Optimal interest rate Rules I: General theory. NBER Working Paper Series No. 9419.
- Giannoni, M., Woodford, M., 2002b. Optimal interest rate rules: II. applications. NBER Working Paper Series No. 9420.
- Hansen, L. P., Sargent, T., 2003. Robust control and filtering of forward-looking models. *Journal of Monetary Economics* 50 (3), 581–604.

- Kerr, W., King, R., 1996. Limits on interest rate rules in the IS model. *Economic Quarterly*, Federal Reserve Bank of Richmond, 47–56.
- Kreps, D., 1998. Anticipated utility and dynamic choice. In: *Frontiers of Research in Economic Theory: The Nancy L. Schwartz Memorial Lectures*. Cambridge University Press.
- Marcet, A., Sargent, T., 1989. Convergence of least squares learning mechanisms in self referential linear stochastic models. *Journal of Economic Theory* 48, 337–368.
- McCallum, B., Nelson, E., 1999. Performance of operational policy rules in an estimated semi classical structural mode. In: Taylor, J. (Ed.), *Monetary Policy Rules*. University of Chicago Press.
- Orphanides, A., Williams, J., 2004. Imperfect knowledge, inflation expectations, and monetary policy. In: Bernanke, B., Woodford, M. (Eds.), *Inflation Targeting Debate*. University of Chicago Press, Chicago.
- Phillips, A., 1954. Stabilisation policiess in a closed economy. *Economic Journal* 64, 290–323.
- Preston, B., 2005. Learning about monetary policy rules when long-horizon expectations matter. *International Journal of Central Banking*, forthcoming.
- Rotemberg, J., Woodford, M., 1997. An optimization-based econometric framework for the evaluation of monetary policy. In: Bernanke, B., Rotemberg, J. (Eds.), *NBER Macroeconomics Annual*. MIT Press, Cambridge.
- Sargent, T., 1999. *The Conquest of American Inflation*. Princeton University Press, New Jersey.
- Svensson, L., 1997. Inflation forecast targeting: Implementing and monitoring inflation targets. *European Economic Review* 41, 1111–1146.
- Svensson, L., 2002. Inflation targeting: Should it be modeled as an instrument rule or a targeting rule? *European Economic Review* 46, 771–780.
- Svensson, L., 2003. What is wrong with taylor rules? using judgment in monetary policy through targeting rules. *Journal of Economic Literature* 41, 426–477.
- Svensson, L., Woodford, M., 2004. Implementing optimal policy through inflation-forecast targeting. In: Bernanke, B., Woodford, M. (Eds.), *Inflation Targeting Debate*. University of Chicago Press, Chicago.
- Taylor, J. (Ed.), 1999. *Monetary Policy Rules*. Chicago University Press, Chicago.
- Vestin, D., 1999. Price-level targeting versus inflation targeting in a forward- looking model. Working Paper, Stockholm University.
- Wieland, V., 2000. Monetary policy, parameter uncertainty, and optimal learning. *Journal of Monetary Economics* 46, 199–228.
- Woodford, M., 1996. Control of the public dept: A requirement for price stability? NBER Working Paper Series 5684.
- Woodford, M., 1999a. Commentary: How should monetary policy be conducted in an era of price stability? In: *New Challenges for Monetary Policy*. Federal Reserve Bank of Kansas City.
- Woodford, M., 1999b. Optimal monetary policy inertia. NBER Working Paper Series 7261.
- Woodford, M., 2003. *Interest and Prices*. Princeton University Press, Princeton.