First draft: November 2002
This draft: May 2003

Rule-Based Monetary Policy under Central Bank Learning *

Abstract

We evaluate the performance of three kinds of rule-based monetary policy under central bank learning about the parameter values of a simple New Keynesian model. The central bank and the private sector learn the slopes of the IS and the Phillips curves by recursive least squares estimation, and form expectations and set policy based on their estimated model. The three policies we evaluate are: (1) the optimal non-inertial rule; (2) the optimal history-dependent rule; (3) the Wicksellian rule. It is shown that the optimal Wicksellian rule delivers the highest welfare, by improving the inflation-output gap variability trade-off, without introducing undesirable feedback from past policy mistakes that are caused by imprecise parameter estimates.

Kosuke Aoki
CREI, Universitat Pompeu Fabra and CEPR
kosuke.aoki@crei.upf.es

and

Kalin Nikolov
Bank of England
kalin.nikolov@bankofengland.co.uk

---

*We would like to thank Tim Cogley, George Evans, Jordi Galí, Albert Marcet, Kaushik Mitra, Katharine Neiss, Thomas Sargent, and Michael Woodford for helpful suggestions. We would also like to thank Peter Andrews, Luca Benati, Jan Groen, Jan Vlieghe, David Vines, and Peter Westaway for helpful comments and discussions. Aoki thanks to the MCYT and FEDER (project No. SEC2002-03816) for financial support. The views expressed here are those of the authors and do not reflect those of the Bank of England.

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 not accurate but improves through learning. Much of the recent literature on monetary policy has emphasised the role and importance of the systematic component of monetary policy. For example, Woodford (1999a) strongly argues the case for rule-based monetary policy, and the papers in Taylor (1999) extensively examine the performance of simple policy rules in small macroeconomic models. It has been 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.

However, the usefulness of rule-based policies is sometimes criticised from a practical point of view, because their performance depends 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 addition, there may be structural change in the economy that affects the underlying model parameters. In this case, it may not even be appropriate for a central bank to commit to following a fixed policy rule since the optimal rule itself may be changing over time. In light of these issues, the question of how to think about policy rules under central bank learning is a key area of research.

The fact that the central bank is continually learning does not necessarily imply that there are no gains from conducting monetary policy systematically. The advantage of rule-based policy stems from the fact that the central bank can internalise its effects of predictable policy on private sector expectations. Rather than completely abandoning rule-based policy, it makes more sense for the monetary authority to commit to a rule and follow that rule until its 'best guess' of the structure of the economy is revised.

In order to analyse monetary policy rules under central bank learning, we consider a central bank which learns structural parameters of a simple New Keynesian model. We assume that both the central bank and private agents do not know the true slopes of the IS and Phillips curve when they form expectations and set policy.[1] The central bank and the private agents are assumed to update their parameter estimates by recursive least squares learning. We assume that they adjust expectations and policy in light of the most recent parameter estimates. Thus, following Sargent (1999) we assume 'anticipated utility' behaviour of the central bank.[2] Under these settings, we evaluate the performance of several policy rules.

The three policy rules we evaluate are (1) optimal non-inertial rule, (2) optimal history dependent rule, and (3) optimal Wicksellian rule. In our basic model, the latter two rules are optimal from a 'timeless' perspective, in the sense defined by Woodford (1999b) and Giannoni and Woodford (2002a). Rules (2) and (3) involve history-dependence, as a result of optimally internalising the effects of the predictable part of policy on private expectations formation.[3] Under perfect knowledge about model parameters, these two rules result in identical (and optimal) equilibria, although they introduce history-dependence in different ways. As Clarida, Galí, and Gertler (1999) have shown, history dependence delivers higher

[1] Thus we assume symmetric information throughout the paper. See section 3 for more discussion.

[2] See Kreps (1998)

[3] This feature of optimal monetary policy is emphasised in Woodford (1999b).

welfare because it improves the output-gap inflation variability trade-off.

However, imperfect knowledge about the structure of the economy may adversely affect the performance of certain policy rules that deliver good economic outcomes in a world of perfect knowledge. This paper shows that imprecise parameter estimate leads 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 Wicksellian rule maintains the benefit of commitment without generating undesirable feedback from past policy mistakes that worsens the performance of the optimal history dependent rule. It is argued that the optimal Wicksellian rule 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 miss-measurement 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 adaptive learning 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 to include an interest rate variability term in the central bank's loss function. Section 6 concludes.

2 Related literature

There is large literature on monetary policy that shows the desirability of commitment over discretion in small macroeconomic models that emphasise forward looking behaviour of agents, followed by Barro and Gordon (1983). More recently, Woodford (1999b) has shown that optimal monetary policy should involve a degree of history dependence, which, for example, implies the dependence on lagged interest rates in the case of optimal interest rate rules. This strand of literature argues the case for rule-based monetary policy (Woodford (1999a)).

Svensson (1999), Woodford (1999a), and Giannoni and Woodford (2002a) have argued that the once-and-for-all commitment implicit in the optimal equilibrium is not a practically useful concept, because the optimal plan derived in this way in general depends on the initial state of the economy, and the implied policy rule is not time-invariant. Therefore, Woodford (1999a) proposed the concept of 'timeless optimality' which involves a time-invariant policy that asymptotically implements the optimal equilibrium. 'Timeless optimality' is an important concept in our paper too, because it provides us with a convenient and practically realistic way of characterising optimal policy with uncertain parameter values under central bank learning.

Our paper is also related to a growing number of studies that have examined the performance of rules when either private agents or the central bank are learning. The literature on adaptive learning has grown rapidly in the 1990s, following the seminal paper by Marcet and Sargent (1989). The literature, which is thoroughly covered by Evans and Honkapohja (2001b), explores the consequences of expectations formation by econometric learning rather than model-consistent expectations. The main question for this literature is whether (and under what circumstances) models with boundedly rational agents produce asymptotically identical behaviour to rational expectations models. In other words, the literature investigates whether agents that use simple econometric models to form expectations eventually 'learn' to make the 'right' forecasts. Rational expectations equilibria, which are asymptotically reached by boundedly rational agents, are known to be 'stable under learning' or 'E-stable'. Bullard and Mitra (2002) and Evans and Honkapohja (2001a) consider adaptive learning about rational expectations equilibrium by private agents in New Keynesian models. They show that a central bank that directly responds to private expectations can achieve a determinate and E-stable equilibrium. In contrast, an instrument rule defined in terms of the underlying economic shocks cannot achieve a learnable equilibrium. Orphanides and Williams (2002) 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.

Cho, Williams, and Sargent (2001) and Wieland (2000) are two examples of papers that examine the consequences of central bank learning. Cho, Williams, and Sargent (2001) studies the behaviour of inflation when the central bank is estimating a misspecified Phillips curve and finds that the economy is characterised by recurrent inflationary episodes followed by rapid 'disinflations'. These are driven by the central bank 'learning' and then 'forgetting' that there is no long-term output-inflation trade-off. Wieland (2000) is an example of a paper that studies optimal Bayesian learning in a simple dynamic model. He finds that the optimal policy involves experimentation whereby the monetary authority deliberately creates economic instability in order to learn about the parameter values.

This paper assumes adaptive learning about structural parameters by both the central bank and private agents. A difference from the previous literature is that we focus on the welfare evaluation of alternative policy rules under learning, rather than long-run learnability of rational expectations equilibria. Unlike Evans and Honkapohja (2001a) and Bullard and Mitra (2002), agents form model-consistent expectations given most recent parameter estimate. Similar to Cho, Williams, and Sargent (2001), we are primarily interested in implications of the central bank learning about model parameters. But we focus on the welfare implications of the way in which policy rules introduce history dependence when a central bank does not know the true model parameter values but estimate them in real time.

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, Galí, and Gertler (1999), McCallum and Nelson (1999)).

The expectational IS equation is given by

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

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 intertemporal elasticity of substitution of expenditure, and the exogenous disturbance g_t represents a demand shock, which may arise from autonomous variations in spending not motivated by intertemporal substitution in response to the real interest rate.

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

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

where y_t^n is an exogenous supply shock at time t. It represents the potential 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 or an inefficient supply shock that represents time variation in the markup on the goods market or time variation in distortionary tax rates.[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 sticky price models, such as 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. Equations (1) and (2), together with a monetary policy rule, determine the equilibrium path for inflation, output, and the nominal interest rate.

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

$$y_t^n = \delta_y y_{t-1}^n + e_{yt}, \quad 0 < \delta_y < 1, \tag{3}$$

$$g_t = \delta_g g_{t-1} + e_{gt}, \quad 0 < \delta_g < 1, \tag{4}$$

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

[5] See Clarida, Galí, and Gertler (1999) and Giannoni (2000).

$$u_t = \delta_u u_{t-1} + e_{ut}, \quad 0 < \delta_u < 1, \tag{5}$$

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

3.2 Adaptive Learning about Structural Parameters

Next we discuss uncertainty. We assume that the central bank knows the structure of the economy but does not know precisely some of the structural parameters. Our modelling strategy is to assume adaptive learning and anticipated utility behaviour in the sense of Kreps (1998) and Sargent (1999). Following the literature on adaptive learning, we assume that the central bank updates those parameters by recursive least squares.

Adaptive learning, which is extensively studied in Sargent (1999) and Evans and Honkapohja (2001b), enables us to introduce central bank learning in the design of optimal policy in a very simple way. Under adaptive learning, the central bank recursively updates the parameters of the model 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. Thus we deviate slightly 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. Our modelling strategy has some advantages over other alternatives such as setting up an optimal policy problem under parameter uncertainty and Bayesian learning. This approach assumes full rationality of economic agents to model optimal policy under parameter uncertainty and learning. Compared to Bayesian learning, our adaptive approach seems more appropriate since our focus is on the evaluation of policy rules under learning rather than optimal learning per se.[6]

Yet another approach may be to derive robust monetary policy under learning. Robust monetary policy is a policy that minimises the worst possible outcome under model uncertainty (Hansen and Sargent (2000), Giannoni (2001)). The robustness approach is conceptually different from adaptive learning, and is suitable for analysing situations in which it is not possible to distinguish or indeed learn about alternative models. In contrast, the adaptive approach assumes that there is one known model but the central bank needs to improve its estimates of the model parameters.

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

$$y_t = E_t y_{t+1} - \sigma(t)\left[i_t - E_t \pi_{t+1}\right] + g_t + \varepsilon_{y,t}, \tag{6}$$

$$\pi_t = \kappa(t)\left(y_t - y_t^n\right) + \beta E_t \pi_{t+1} + u_t + \varepsilon_{\pi,t}. \tag{7}$$

Here κ(t) and $\sigma(t)$ are respectively the bank's most recent estimates of κ and σ. For simplicity, we focus on implications of uncertainty about κ and σ by assuming that β and the stochastic process (3), (4), (5) are known. The bank's model also includes unobservable white noise disturbance terms $\varepsilon_{y,t}$, $\varepsilon_{\pi,t}$. This assumption prevents the bank from being able to identify the true values of κ and σ after observing data for one period.[7]

We assume that the bank estimates κ(t) and $\sigma(t)$ recursively. Following the literature on adaptive learning, we assume that, at time t, the central bank updates its parameter

[6] Also, the existing literature on optimal Bayesian learning has shown that it is not possible in general to derive an analytical solution, although numerical approximations can be obtained for smaller models. See, for example, Wieland (2000). Compared to this approach, adaptive learning is much simpler to solve.

[7] The bank believes that those shocks are hitting the economy, although there are no such shocks in (1) and (2). In this sense, the bank's models is slightly misspecified. As the bank eventually learns the true values of κ and σ, it also learns that the variance of $\varepsilon_{y,t}$, $\varepsilon_{\pi,t}$ are equal to zero. Therefore, this small misspecification vanishes as the bank learns over time.

estimates $(\sigma(t), \kappa(t))$ based on time $t-1$ data.[8] The estimation equation is therefore given by

$$y_{t-1} - E_{t-1}y_t - g_{t-1} = -\sigma\left[R_{t-1} - E_{t-1}\pi_t\right] + \varepsilon_{y,t-1}, \tag{8}$$

$$\pi_{t-1} - \beta E_{t-1}\pi_t - u_{t-1} = \kappa\left(y_{t-1} - y^n_{t-1}\right) + \varepsilon_{\pi,t-1}. \tag{9}$$

In order to focus our analysis on the implications of learning about structural parameters, we assume for simplicity that expectations and the structural shocks are observable variables. In each period, the updated estimates $(\sigma(t), \kappa(t))$ are given by the usual recursive least squares estimators.

We also assume symmetric information between the private agents and the central bank. This implies that the private agents use (6) and (7) to make projections about future inflation and output, and the policy rate. This in turn 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 discussed in 3.3, 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, although the policy will eventually be revised based on future improvements in the knowledge of the economy. 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, and simplifies our analysis

[8] This assumption is commonly used in the literature on adaptive learning to avoid circularity between learning and equilibrium. See, for example, Evans and Honkapohja (2001a).

significantly.[9]

3.3 Modelling central bank and policy rules

We turn to the analysis of monetary policy. The welfare loss of the central bank at time t is the expected discounted sum of period loss functions:

$$W_t \equiv E_t \left[\sum_{s=t}^{\infty} \beta^s L_s \right], \tag{10}$$

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

$$L_s = \frac{1}{2} \left[\pi_s^2 + \lambda_y \left(y_s - y_s^n \right)^2 \right]. \tag{11}$$

Rotemberg and Woodford (1997) show that the loss measure (11) is a quadratic approximation to the expected utility of the representative household in the Calvo model, when $\lambda_y > 0$ is appropriately chosen.[10] However, our analysis would be valid for arbitrary values of λ_y.

We focus on commitment to policy rules that are optimal from a 'timeless perspective' defined in Giannoni and Woodford (2002a). In our setting, a timelessly optimal policy rule is a time-invariant policy rule, conditional on the bank's current knowledge of the economy, that implements the optimal state-contingent plan, which minimises expected welfare losses, 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.

---

[9] Hansen and Sargent (2000) also assume that private agents and policy makers 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, 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 to be robust under this alternative scenario.

[10] See, also, Woodford (1999b). Our loss measure corresponds to his when the central bank does not face the zero nominal interest rate bound. In Section 5 we extend the analysis to the case in which the central bank's period loss function involves an interest rate variability term.

As a result, the central bank commits to follow a certain fixed rule indefinitely, given its best knowledge of the economy, and the bank is expected to follow this rule unless its best knowledge of the economy improves in subsequent periods. Combined with our assumption of anticipated utility behaviour, the policy rule is expected to be time-invariant in each period, even though it would be eventually adjusted through learning.

Now consider a central bank in period t whose most recent estimates of the two structural parameters are $\kappa(t)$ and $\sigma(t)$. As we discussed in Section 3.2, we assume that both private agents and the central bank regard those parameter values as constant over time. Based on those parameters, the central bank at time t seeks to minimise (10) subject to the constraints

$$y_s = E_s y_{s+1} - \sigma(t)\left[R_s - E_s \pi_{s+1}\right] + g_s, \tag{12}$$

$$\pi_s = \kappa(t)\left(y_s - y_s^n\right) + \beta E_s \pi_{s+1} + u_s. \tag{13}$$

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

$$\pi_s - \beta^{-1}\sigma(t)\phi_{1s-1} + \phi_{2,s} - \phi_{2s-1} = 0, \tag{14}$$

$$\lambda(y_s - y_s^n) + \phi_{1s} - \beta^{-1}\phi_{1s-1} - \kappa(t)\phi_{2s} = 0, \tag{15}$$

$$\sigma(t)\phi_{1s} = 0, \tag{16}$$

for each $s \geq t$. Here ϕ_{1s} and ϕ_{2s} respectively represents the Lagrange multiplier associated with (12) and (13). We also need to impose the initial conditions ϕ_{1t-1} and ϕ_{2t-1}, which will be discussed below.

It follows from the first order conditions that, in the case of optimal commitment plan that has been in implemented since $s-1$, we can infer ϕ_{1s-1} and ϕ_{2s-1} from the past observable variables π_{s-1} and $y_{s-1} - y^n_{s-1}$. Those are given by $\phi_{1s-1} = 0$ and $\phi_{2s-1} = \lambda_y/\kappa(t)(y_{s-1} - y^n_{s-1})$. In this case, the first order conditions imply

$$\pi_s = -\frac{\lambda_y}{\kappa(t)} \left[(y_s - y^n_s) - \left(y_{s-1} - y^n_{s-1} \right) \right] \tag{17}$$

for $s \geq t + 1$.

Next we turn to the initial conditions on ϕ_{1t-1} and ϕ_{2t-1}. A once-and-for-all commitment from period t on typically stipulates the initial conditions $\phi_{1t-1} = 0$ and $\phi_{2t-1} = 0$, implying that the bank exploits the existing expectations at the time the policy is chosen. However, this is not desirable 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, this could easily result in discretionary policy, leading to a suboptimal equilibrium as in the analysis of Barro and Gordon (1983). We instead restrict our policy regime on policies that are optimal from a timeless perspective, following Giannoni and Woodford (2002a) and Woodford (2002). Optimality from a timeless perspective imposes a restriction on the initial evolution of the endogenous variables in such a way that the bank does not exploit the existing expectations at the time the policy is chosen. 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 $\phi_{2t-1} = \lambda/\kappa(t)(y_{t-1} - y^n_{t-1})$. Then we obtain a time-invariant targeting criteria

$$\pi_s = -\frac{\lambda_y}{\kappa(t)} \left[(y_s - y^n_s) - \left(y_{s-1} - y^n_{s-1} \right) \right], \quad s \geq t. \tag{18}$$

As is shown in Woodford (2002), commitment to the rule (18) from period t is optimal from a timeless perspective in the sense defined in Giannoni and Woodford (2002a). In other words, the equilibrium associated with commitment to the policy of the form (18) involves the same responses to exogenous shocks in all periods $s \geq t$ as the optimal plan, given that the model of the economy is described by (6) and (7). Notice that the rule involves the lagged output gap $\left(y_{t-1} - y_{t-1}^n\right)$. 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 (18) does not involve the nominal interest rate, it is a targeting rule in the sense of Svensson (2001) and Giannoni and Woodford (2002a). Targeting rules are specified in terms of a criterion that certain target variables should satisfy. The central bank commits itself to choose the interest rate so as to satisfy this criterion, and it must use its estimated model of the economy in setting the interest rate.[11] When implementing (18), the bank needs to make its decisions on the projections of the current endogenous variables, such as $\pi_{s|t}$ and $y_{s|t}$ using its model. Here a variable $x_{s|t}$ denotes the projection of x_s based on the Bank's model at time t. Therefore, in these policy regimes, the bank commits itself to choose the nominal interest rate in order to satisfy

$$\pi_{s|t} = -\frac{\lambda_y}{\kappa(t)} \left[\left(y_{s|t} - y_s^n\right) - \left(y_{s-1} - y_{s-1}^n\right) \right], \quad s \geq t. \tag{19}$$

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

$$p_s = -\frac{\lambda_y}{\kappa(t)} \left(y_s - y_s^n\right), \quad s \geq t. \tag{20}$$

[11] When the target criterion involves the nominal interest rate, the target can be interpreted as an interest rate rule. So simple interest rate rules, such as the Taylor rule, can also be interpreted as a special class of targeting rules. But a target criterion can define a policy rule even if it does not involve the interest rate. As Svensson (2001) argues, it is not always possible to express 'optimal rules' in terms of instrument rules.

Notice that, since $\pi_s \equiv (1-L)p_s$, 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 an 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).[12] 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.[13] Compared with (18), (20) introduces history-dependence without explicitly committing to respond to the lagged output gap. Following Woodford (1998), we call such a rule a Wicksellian rule.

When implementing (20), the bank also bases its decisions on the projections of the current endogenous variables, such as $p_{s|t}$ and $y_{s|t}$ using its model. The bank commits itself to choose the nominal interest rate in order to satisfy

$$p_{s|t} = -\frac{\lambda_y}{\kappa(t)} \left(y_{s|t} - y_s^n \right), \quad s \geq t. \tag{21}$$

Policy rules (21) and (19) are closely related. In the special case in which the true structural parameters values are known and the economy starts from its steady state, commitment to rule (19) and commitment to rule (21) result in the same rational expectations equilibrium, as is shown in Clarida, Galí, and Gertler (1999) and Vestin (1999).[14]

[12] Currie and Levine (1987) show that an optimal commitment policy can be expressed in terms of integral of past predetermined variables. Here we interpret (20) as a targeting rule which involves an integral term of one of the goal variables, namely, π.

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

[14] See Woodford (2002) for further discussion.

As a comparison, we also consider a rule that does not involve history dependence. This is called the optimal non-inertial rule:

$$\pi_{t|t} = -\frac{\lambda_y}{\kappa(t)} \left(1 - \beta\delta_u\right) \left(y_{t|t} - y_t^n\right). \qquad (22)$$

Conceptually, this is similar to a simple Taylor rule that involves only contemporaneous feedback from inflation and the output gap. The coefficient $\frac{\lambda_y}{\kappa(t)} \left(1 - \beta\delta_u\right)$ is chosen to implement the optimal non-inertial plan defined in Woodford (2002).[15] Since our structural equations are purely forward-looking and the rule does not involve lagged endogenous variables, 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 $\kappa(t)$.[16] After they update their best estimates of the parameters, they observe time-t shocks. Based on their updated model, they form expectations and set policy. Expectations are model-consistent given the most recent parameter estimates. The policy rate is chosen on the basis of the central bank's projection of the current and future endogenous variables, given its updated model. After expectations and policy rates are set, those are put into the true model and inflation and the output gap are realised.[17] In the

[15] More specifically, (22) implements the optimal non-inertial plan discussed in Section 3.1 of Woodford (2002). 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.

[16] As stated above, we do not assume that time t endogenous variables are used for time t estimation, as this would complicate the learning procedure without adding anything to the analysis.

[17] Recently, Preston (2003) argues that, when agents are forward looking and therefore long term expectations matter, it may not be appropriate to use as a representation of aggregate dynamics the log-linearised IS

next section, we provide some numerical examples to evaluate the performance of the three rules under learning.

4 Numerical examples

4.1 Simulation exercises

In this section we conduct stochastic simulations using our model and 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 1000 simulated paths.

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). 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 per cent. The autocorrelation coefficients of the shocks are set as follows: δ_g is 0.92, $\delta_{y_n}$is 0.15, δ_u is 0.35.[18] The 'true' value of σ and κ in the model

and Phillips curve equations with expectations terms replaced by subjective expectations given by reduced-form expectation equations implied by rational expectations equilibrium. He argues that this 'Euler equation approach' is consistent with optimising behaviour of agents only when expectations are rational. He found some examples in which his approach and the 'Euler equations approach' give different E-stability conditions. In our model, expectations are rational and therefore optimal forecast given the estimated model. So we can safely interpret (1) and (2) as determining aggregate dynamics after expectations are determined by the estimated model. See Evans, Honkapohja, and Mitra (2003) for further discussion.
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. We confirmed that the main results given in the subsequent sections remain the same. This implies that our results mainly stem from central bank's imperfect knowledge and are robust to changes in assumptions on private sector information.

[18] 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 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. The standard deviations of the innovations in the demand and supply disturbances in the Rotemberg-Woodford model are calculated as 1.022 percent and 1.906 percent, respectively.[19] The standard deviation of the innovation in the cost push shock is set 0.37 per cent, cited from Giannoni (2000).

4.1.2 Welfare results

Table 1 presents the expected values of the welfare function under the optimal non-inertial rule (ONP) and the optimal history dependent rule (OHDP). The first row of the table displays 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 second row contains numbers derived under the assumption that the central bank and the private agents know the true parameter values of the model.

The results in the second row are now fairly standard in the literature (Clarida, Galí, and Gertler (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

with the same autoregressive root.

[19] I thank Marc Giannoni for providing these values. Specifically, these are the standard deviations of the innovations in the processes $E_{t-2}\left[G_{t+1} - G_t\right]$ 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.

(19). 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.

However, the first row of Table 1 shows that the advantages of the history dependent rule over the optimal non-inertial rule are reversed under adaptive learning. This is mainly due to a substantial deterioration of welfare under the optimal history dependent rule.[20] This is one feature of our results we will explore further in the next subsection.

Table 1 also compares expected welfare under the optimal history dependent rule and the optimal Wicksellian rule (WP). The second row of the table presents expected welfare when the bank and private agents know the true parameter values of the model. The two rules deliver identical welfare losses under full information, because both of them are timelessly optimal in our setup.[21] Because correcting deviations of the price level from a deterministic path involves a strong commitment against inflation, the Wicksellian rule is a good policy for a central bank that wants to improve its inflation-output trade-off by stabilising inflation expectations. Table 1 confirms this well-known result in the literature.[22] However, the first row of Table 1 reveals that, under learning, the Wicksellian rule performs much better than the optimal history dependent rule and this restores the advantage of history dependence.

[20] 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.

[21] See, also, Clarida, Galí, and Gertler (1999). In our model, the optimal commitment policy involves Wicksellian rule when there is no interest rate smoothing term in the central bank's loss function. Otherwise, the optimal commitment policy in general involves non-stationary path of the price level. See Giannoni (2000).

[22] See Giannoni (2000) and Vestin (1999) for more discussion of the links between price level targeting and inflation targeting under commitment.

It seems, therefore, that imperfect knowledge of the model parameters breaks down the link between the Wicksellian rule and the history dependent rule as the performance of the latter deteriorates markedly.

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.

Cost push shock Chart 1 below plots the dynamic response of our model economy to a cost push shock under the three rules we consider.

[Cost push shock chart here]

It shows that, even under learning, both the Wicksellian 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, Galí, and Gertler (1999), Giannoni (2000), Vestin (1999)).

Demand shock Chart 2 below plots the dynamic response of our model economy to a demand shock under the three rules we consider.[23]

[Demand shock chart here]

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

It is a well-known result in the literature (for example Clarida, Galí, and Gertler (1999)) that under full information and rational expectations, the central bank would stabilise demand and supply shocks perfectly. Chart 2 shows that this is not the case under learning. Because the wrong estimate of σ leads to biased estimates of the natural rate, demand and supply shocks affect output and inflation in the resulting equilibrium. And interestingly, the optimal non-inertial rule and the Wicksellian rule dominate the optimal history dependent rule. Indeed, the performance of the optimal history-dependent rule is particularly poor - demand shocks lead to very large and 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 Wicksellian rule.

4.2 History dependence and policy mistakes

In the previous section we argued that when there is no uncertainty about model parameters, the Wicksellian 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 imprecise parameter estimates, which can be avoided adopting a different way of implementing history dependence (such as the Wicksellian rule). In this section we show how imperfect knowledge of parameters during the learning phase 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.[24] It is the equilibrium real interest rate, which would prevail under fully flexible prices. In our model, the natural interest rate, r^n, is given by

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

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

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

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,

[24] For the recent discussion about this concept, see Blinder (1998)and Woodford (1999b)

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 (15), (14), and (16) imply that implementing the optimal plan requires that (18) holds for $s \geq t$. In order to implement (18), the OHDP rule commits to satisfy target (19), while the Wicksellian rule commits to satisfy target (21). In both cases, the bank bases its decision on its projections using the estimated model (6) and (7).

Since the estimated model is parameterised imprecisely, the target criteria do no 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 $\hat{\kappa}$. As we discuss in Section 4.3, the speed of learning about parameters is rather slow, so 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.

Notice that the policy rate calculated using the estimated model is given by

$$i_t = \hat{\sigma}^{-1} \left(y_{t+1|t} - y_{t|t} + g_t \right) + \pi_{t+1|t}. \tag{23}$$

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

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

Similarly, by noticing that

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

and using (24), we have

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

Substituting (24) and (25) into the OHDP rule (19), we obtain

$$\pi_t + \frac{\lambda_y}{\hat{\kappa}}\left[\left(y_t - y_t^n\right) - \left(y_{t-1} - y_{t-1}^n\right)\right] = \nu_t, \tag{26}$$

where

$$\begin{aligned} \nu_t \;\equiv\; & \left[\left(\kappa\frac{\sigma}{\hat{\sigma}} - \hat{\kappa}\right) + \frac{\lambda_y}{\kappa}\left(\frac{\sigma}{\hat{\sigma}} - 1\right)\right] y_{t|t} + \left(1 - \frac{\sigma}{\hat{\sigma}}\right)\left(\kappa + \frac{\lambda_y}{\kappa}\right) y_{t+1|t} \\ + \; & \left(1 - \frac{\sigma}{\hat{\sigma}}\right)\left(\kappa + \frac{\lambda_y}{\kappa}\right) g_t - (\kappa - \hat{\kappa})\, y_t^n. \end{aligned} \tag{27}$$

The right hand side of (26) represents the policy mistake when the bank follows the OHDP rule (19). Since the OHDP rule directly corresponds to the optimality condition (18), 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 $\hat{\sigma} \to \sigma$ and $\hat{\kappa} \to \kappa$.

Similarly, we can calculate the policy mistake when the bank commits to targeting rule (21). Using the IS and Phillips curves, we can write:

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

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

$$p_t + \frac{\lambda_y}{\hat{\kappa}} \left[(y_t - y_t^n) \right] = \nu_t,$$

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 Wicksellian and the OHDP rules, the degree to which the optimality condition (18) fails to hold is different and is given by

$$\pi_s = -\frac{\lambda_y}{\hat{\kappa}} \left[(y_s - y_s^n) - \left(y_{s-1} - y_{s-1}^n \right) \right] + \nu_t - \nu_{t-1}. \tag{29}$$

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.[25] We then argue that the inertial nature of the optimal equilibrium in forward-looking models brings additional benefits of following integral control policies

[25] See, for example, Franklin, Powell, and Emami-Naeini (2002)

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, 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, v_t. As equations (26) and (29) show this certainly seems to be part of the reason why the Wicksellian rule performs better than the OHDP rule under learning. Under the OHDP rule the mistake enters in levels whereas under Wicksellian rule, the policy mistake enters in first differences - it is 'undone' one period later.

This feature of the Wicksellian rule is an implication of targeting the 'integral' of deviations of target variables from their target values, and as Franklin, Powell, and Emami-Naeini (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 Wicksellian 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 Chart 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 Wicksellian 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 Wicksellian 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. 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.[26] 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:

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

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

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

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

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

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

where λ_1 is the stable root of the characteristic equation

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

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

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

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

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 is[27]

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

while under the Wicksellian rule, it is given by

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

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 under white noise shocks. This persistence arises due to history dependence, 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 Wicksellian rule. Because of the integral control properties of the rule, as λ_1 gets large,[28] the equations actually imply that the impact of ε_t on the output gap (and therefore on inflation) becomes smaller, not larger. In the limiting case in which $\lambda_1 \to 1$, the effect of policy

[27] Here we assume that the economy is in the steady state at time $t-1$

[28] In our calibration, λ_1 is given by 0.87, implying substantial policy-induced persistence. Persistence can be even higher under some parameter or welfare function values. For example λ_1 tends to unity when λ_y tends to infinity.

mistakes under the Wicksellian rule vanishes completely after one period. This feature of the Wicksellian 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 Wicksellian 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.

To give another example of the importance of our results, we change our calibration slightly in a way that increases substantially the degree of optimal policy-induced persistence. One interesting way of doing this is to increase the value of λ_y, the weight on the output gap in the welfare function, from our baseline value of 0.048 to 1. This implies that the monetary authority now cares equally about inflation and output gap deviations. Charts 3 and 4 compare the economy's dynamic responses to a white noise demand shock under the baseline calibration and under the alternative calibration with $\lambda_y = 1$. The increase in λ_y raises λ_1 from 0.87 to 0.97 - a considerable increase in policy-induced persistence. In addition, we now return to our set up in which the central bank and the private sector are learning. In those charts we also show policy mistake that is given by (27).

[Chart 3 here]

[Chart 4 here]

Firstly, Chart 3 clearly shows that the mechanism we have demonstrated using the simplified model is present in our benchmark model. The OHDP rule propagates monetary policy mistake while the Wicksellian rule does not.[29] Secondly, Chart 4 shows that the change in the social welfare function leads to a much more persistent effect of the monetary policy mistake on inflation and output when policy follows the OHDP rule. In contrast, there is no economically important difference between the behaviour of inflation and output when policy uses the Wicksellian rule. The reason for this is the integral control terms present in the Wicksellian rule. And although the increase in the output gap weight in the social welfare function raises the degree of policy-induced persistence increases under both rules, this only affects welfare under the OHDP rule. The integral control term in the Wicksellian rule 'offsets' most of the effect of the monetary policy mistake after one period and does not propagate its effect over time.

Interestingly, our results show that a central bank with a strong preference for 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 (2002) 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. 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 (2002).

[29] In charts 3 and 4, the responses under the Wicksellian rule from period 1 on are not exactly equal to zero, but very small numbers. For example, the nominal interest rate becomes slightly negative in period 1 and converges to zero.

4.3 The speed of learning

As we showed in the previous section, imprecise parameter estimate can lead to policy errors with substantial welfare consequences. This implies that monetary rules that increase the speed of learning are likely to deliver good welfare outcomes. Charts 5 and 6 below show the convergence of the central bank's estimates of σ and κ towards their true values. Both charts display the representative path of the estimates from 1000 simulations.

[Chart 5 here - sigma convergence]

[Chart 6 here - kappa convergence]

The parameter estimates always seem to converge to their true values. This implies that the central bank and the private agents will eventually learn the true model, and the equilibrium will converge to the rational expectations equilibrium with true parameter values. However, rigourous analysis of convergence, in line with Evans and Honkapohja (2001a) is left for future research.

The charts show that the speed of convergence of the estimated slope of the IS curve (σ) is relatively unaffected by the policy rule followed by the central bank. In contrast, the estimated slope of the Phillips curve (κ) converges more slowly under the optimal non-inertial rule than under the Wicksellian rule and the optimal history dependent rule. This is quite intuitive. Because the optimal non-inertial rule stabilises the output gap more than the other two rules, there is less data variation in the Phillips curve, and estimation is slower. Overall, the difference in the speed of convergence does not seem to be playing a major role in the difference of the performance of the three policy rules, compared with the effect on welfare of the precise way of implementing history dependence.

5 Consideration of interest rate variability

5.1 The Central Bank's policy rule

In the model used in the previous section, the optimal Wicksellian rule is timelessly optimal and delivers the same equilibrium as the optimal history dependent rule when the true structural parameter values are known. However, the Wicksellian rule is in general not optimal from a timeless perspective, because the optimal path of the price level under inflation targeting regime is in general not stationary.[30] In this section, therefore, we check our results for robustness to the optimality of the Wicksellian rule. We now assume that the loss function involves the interest rate variability as well as the output gap and inflation variability:

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

In this case, the first order conditions for the central bank's maximisation problem are as follows (assuming for simplicity that $\overline{R} = 0$):

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

$$\lambda(y_s - y_s^n) + \phi_{1s} - \beta^{-1}\phi_{1s-1} - \kappa(t)\phi_{2s} = 0, \tag{38}$$

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

The main difference between (39)above and (16) 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 with marginal social cost equal to $\sigma(t)\phi_{1s}$ against the costs of interest rate variability, the

[30] See Giannoni (2000) for this point.

marginal social cost of which is given by $\lambda_i i_t$. What this implies is 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.

Under once and for all optimal commitment, the solution of the central bank's maximisation problem is completed by setting the initial conditions for the shadow prices (ϕ_{1t-1} and ϕ_{2t-1}) associated with the Phillips curve and the IS curve equal to zero, which implies that the bank exploits the private sector's existing expectations prior to the commitment to the new policy. However, we again focus on the timelessly optimal policy as defined by Giannoni and Woodford (2002b). In this case the central bank will set the initial conditions for the two Lagrange multipliers equal to the values they would have had if a commitment to this rule were made in the past. The bank does not exploit the existing beliefs of the private sector. The values of the Lagrange multipliers are given by:

$$\phi_{1t-1} = \frac{\lambda_i i_t}{\sigma(t)} \tag{40}$$

$$\phi_{2t-1} = \frac{\lambda_y}{\kappa(t)} (y_{t-1} - y_{t-1}^n) \tag{41}$$

We can now use equations (40) and (41) to substitute the Lagrange multipliers out of the model. This allows us to solve for an explicit instrument rule in the same way as in (Giannoni and Woodford (2002b)). The optimal history dependent rule in this case is given by:

$$i_t = \rho_1 i_{t-1} + \rho_2 (i_{t-1} - i_{t-2}) + \rho_\pi \pi_t + \rho_y \left[(y_t - y_t^n) - (y_{t-1} - y_{t-1}^n) \right] \tag{42}$$

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

(2002b) also show, this rule is optimal from a timeless perspective.

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 re-write 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

$$R_t = \rho_1 R_{t-1} + \rho_2 (R_{t-1} - R_{t-2}) + \rho_\pi p_t + \rho_y (y_t - y_t^n) \tag{43}$$

where

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

is a sum of all past of nominal interest rates. We can now rewrite (43) in a more intuitive interest rate rule form:

$$i_t = \rho_2 i_{t-1} + (\rho_1 - 1) R_{t-1} + \rho_\pi p_t + \rho_y (y_t - y_t^n) \tag{44}$$

Both (43) and (44) are optimal from a timeless perspective.

Whereas (42) 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 (44) also has elements of integral control. These elements are introduced by the terms in R_{t-1} and also the price level targeting term (which is simply the sum of all past inflation deviations from target). Note that under full information this 'quasi-Wicksellian rule' is just a transform of (42) and, therefore, it implements the optimal equilibrium in exactly the same way as (42). 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 (44) delivers a superior equilibrium to (42). Below we again use stochastic simulations to test our hypothesis in the case when the central bank has an interest rate variability objective.

5.2 Numerical results

Table 2 below presents the value of the welfare loss under the optimal history dependent rule (42) and the Wicksellian rule (44). Following Rotemberg and Woodford (1997) we set λ_i, the weight on interest rate variability in the welfare function, equal to 0.236. All other parameter values are the same as in section 3.

[TABLE 2 HERE]

We get very similar results to the ones we obtained in section 3. Under the true parameter values, the optimal history dependent rule and the Wicksellian rule deliver identical performances. However, the table shows that the performance of the optimal history dependent rule deteriorates substantially relatively to the Wicksellian rule when we allow for imperfect knowledge of the model's parameter values.[31] 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-Wicksellian rule (43) (which has elements of integral control) delivers higher welfare than the optimal history dependent rule (44) (which does not).

6 Conclusion

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, we find that 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. A certain form of history dependence may propagate policy mistakes over time. History dependence in monetary policy may generate endogenous persistence in the responses of the economy to policy errors. We find that rules that involves integral terms are more robust when the bank's economic model is slightly misspecified. For example, this class of policy includes rules that involve price level, which is the integral of inflation.

[31] Indeed, our simulation results show that imperfect knowledge of the model parameter affects the Wicksellian rule only marginally. The welfare loss becomes marginally larger.

Those rules automatically 'undo' past policy mistakes, while keeping the advantage of history dependent monetary policy.

Throughout the paper, we assumed that information are symmetric between the central bank and the private agents. However, it would be of interest to extend our analysis to asymmetric information, particularly to the case in which the private sector has more accurate knowledge of the structure of the economy. We expect that our results would be strengthened further, by a similar argument in Aoki (2001). Aoki (2001) considers an optimal 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 tight, and vice versa. In other words, the optimal policy has overshooting property. Thus the optimal plan commits to undo the past policy mistakes caused by information problem. Since the private agents with perfect information anticipate this overshooting, the responses of economy to the current policy mistakes becomes smaller. In our setting, the rules that involve integral terms have this overshooting property.

References

Aoki, K. (2001): "Optimal Commitment Plan under Noisy Information," *CEPR Discussion Paper Series*, 3370.

Barro, R., and D. Gordon (1983): "A Positive Theory of Monetary Policy in a Natural Rate Model," *Journal of Policical Economy*, 91, 589–610.

Bernanke, B., and M. Woodford (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., and K. Mitra (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., N. Williams, and T. Sargent (2001): "Escaping Nash Inflation," *Review of Economic Studies*.

Clarida, R., J. Galí, and M. Gertler (1999): "The Science of Monetary Policy: A New Keynesian Perspective," *Journal of Economic Literature*, XXXVII, 1661–1707.

Currie, D., and P. Levine (1987): "The Design of Feedback Rules in Linear Stochastic Rational Expectations Models," *Journal of Economics Dynamics and Control*, 11, 1–28.

Evans, G., and S. Honkapohja (2001a): "Expectations and the Stability Problem of Optimal Policies," *Working Paper, University of Oregon*.

——— (2001b): *Learning and Expectations in Macroeconomics*. Princeton University Press, New Jersey.

Evans, G., S. Honkapohja, and K. Mitra (2003): "Notes on Agents' Behavioral Rules under Adaptive Learning and Recent Studies of Monetary Policy," *Working Paper, University of Oregon*.

Franklin, G., D. Powell, and A. Emami-Naeini (2002): *Feedback Control of Dynamic Systems*. Prentice Hall, New Jersey, 4th edn.

Giannoni, M. (2000): "Optimal Interest-Rate Rules in a Forward-Looking Model, and Inflation Stabilization versus Price-Level Stabilization," *Working Paper, Columbia University*.

——— (2001): "Does Model Uncertainty Justify Caution? Robust Optimal Monetary Policy in a Forward-Looking Model," *Macroeconomic Dynamics*.

Giannoni, M., and M. Woodford (2002a): "Optimal Interest Rate Rules I: General Theory," *Working Paper, Princeton University and Columbia University*.

——— (2002b): "Optimal Interest Rate Rules: II. Applications," *Working Paper, Princeton University and Columbia University*.

Hansen, L. P., and T. Sargent (2000): "Robust Control and Filtering of Forward-Looking Models," *Working Paper, New York Univerisy and University of Chicago*.

Kerr, W., and R. King (1996): "Limits on Interest Rate Rules in the IS Model," *Economic Quarterly, Federal Reserve Bank of Richmond*, pp. 47–56.

Kreps, D. (1998): "Anticipated Utility and Dynamic Choice," in *Froniters of Research in Economic Theory: The Nancy L. Schwartz Memorial Lectures*. Cambridge University Press.

Marcet, A., and T. Sargent (1989): "Convergence of Least Squares Learning Mechanisms in Self Referential Linear Stochastic Models," *Journal of Economic Theory*, 48, 337–368.

McCallum, B., and E. Nelson (1999): "Performance of Operational Policy Rules in an Estimated Semi Classical Structural Mode," in *Monetary Policy Rules*, ed. by J. Taylor. University of Chicago Press.

Orphanides, A., and J. Williams (2002): "Imperfect Knowledge, Inflation Expectations, and Monetary Policy," *Board of Governors of the Federal Reserve Board.*

Phillips, A. (1954): "Stabilisation Policiess in a Closed Economy," *Economic Journal*, 64, 290–323.

Preston, B. (2003): "Learning About Monetary Policy Rules When Long-Horizon Expectations Matter," *Working Paper, Princeton University.*

Rotemberg, J., and M. Woodford (1997): "An Optimization-Based Econometric Framework for the Evaluation of Monetary Policy," in *NBER Macroeconomics Annual*, ed. by B. Bernanke, and J. Rotemberg. MIT Press, Cambridge.

Sargent, T. (1999): *The Conquest of American Inflation.* Princeton University Press, New Jersey.

Svensson, L. (1999): "How Should Monetary Policy be Conducted in an Era of Price Stability?," in *New Challenges for Monetary Policy.* Federal Reserve Bank of Kansas City, Kansas City.

——— (2001): "Inflation Targeting: Should It Be Modeled as an Instrument Rule or a Targeting Rule?," *European Economic Review, forthcoming.*

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.*

——— (1998): "Doing Without Money: Controlling Inflation In A Post-Monetary World," *Review of Economic Dynamics*, 1, 173–219.

——— (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.

——— (1999b): "Optimal Monetary Policy Inertia," *NBER Working Paper Series 7261.*

——— (2002): "Gains From Commitment to a Policy Rule," *Working paper, Princeton University.*

Table 1: Welfare loss $\sum_{s=0}^{500} \beta^s L_s$

	ONP	OHDP	WP	ONP/OHDP	ONP/WP
Learning	0.0036	0.0090	0.0030	-60.0%	19.7%
Full Info	0.0034	0.0028	0.0028	23.7%	23.7%

The results under the benchmark calibration
Sample average of 1000 simulations

Table 2: Welfare loss $\sum_{s=0}^{500} \beta^s L_s$

	OHDP	WP	WP/OHDP
Learning	0.0085	0.0037	-57.2%
Full Info	0.0037	0.0037	0.0%

The results with interest rate variability
Sample average of 1000 simulations

Chart 1 RESPONSES TO COST-PUSH SHOCK

Benchmark Case

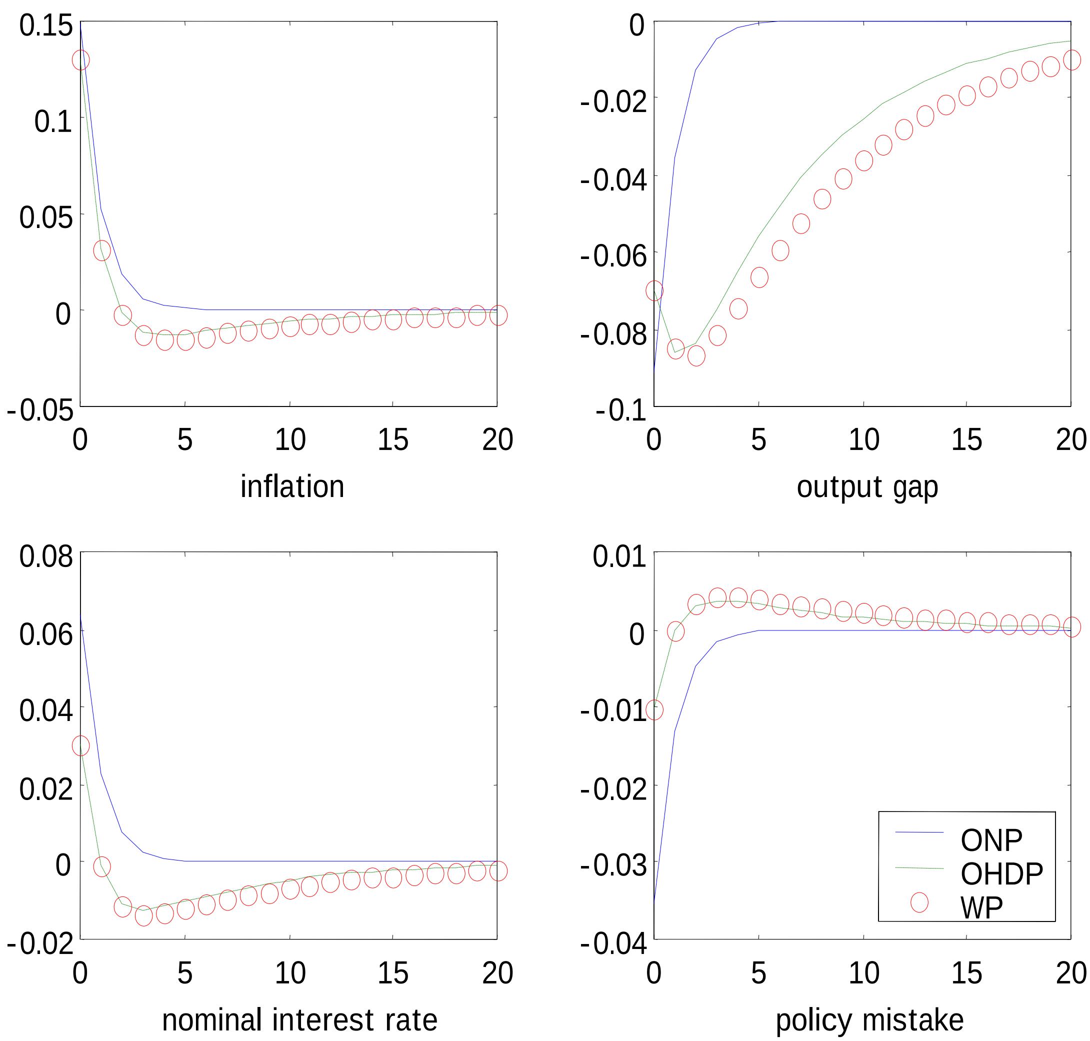

Chart 2 RESPONSES TO DEMAND SHOCK

Benchmark Case

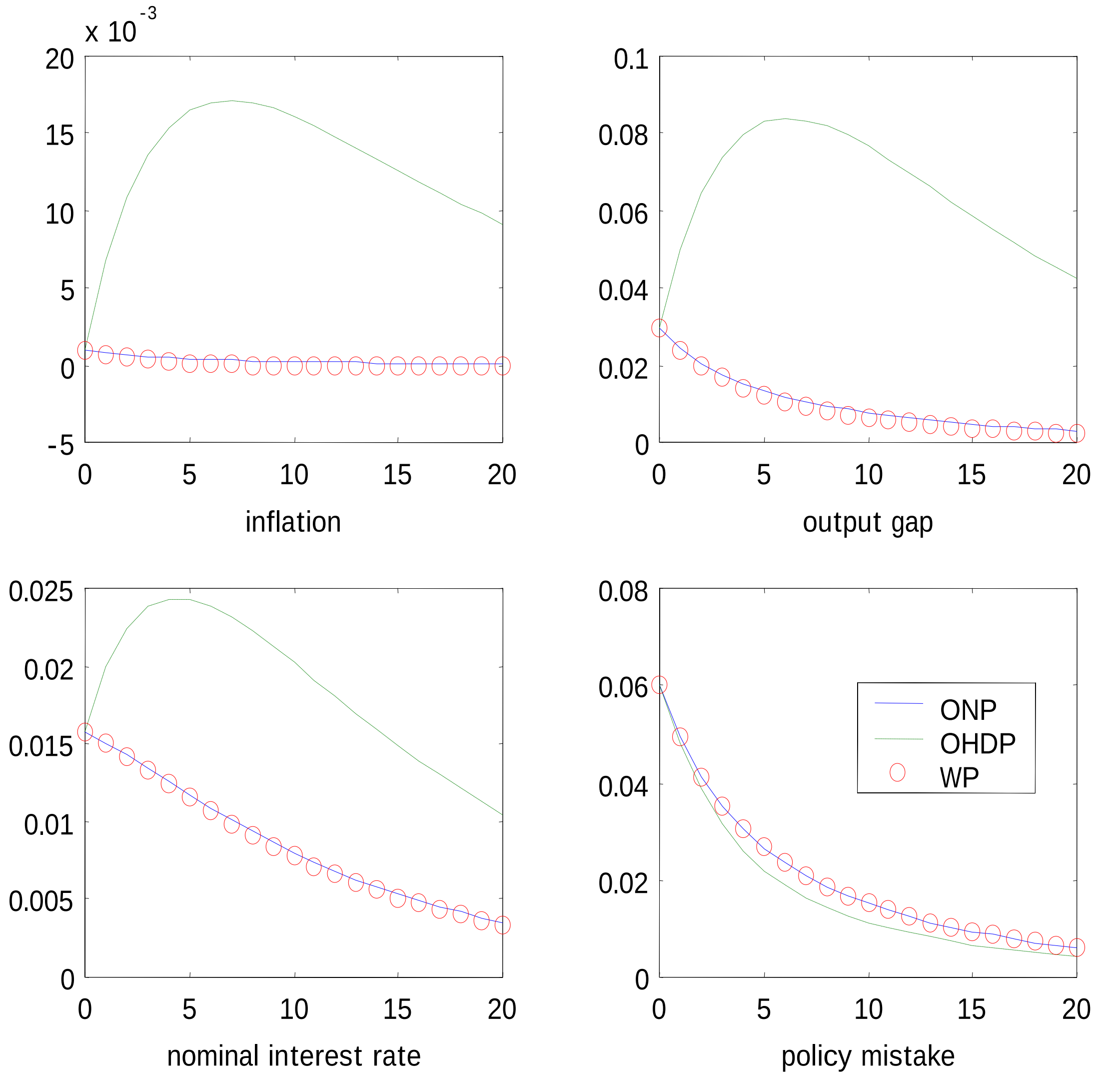

Chart 3 RESPONSES TO DEMAND SHOCK

White noise demand shock

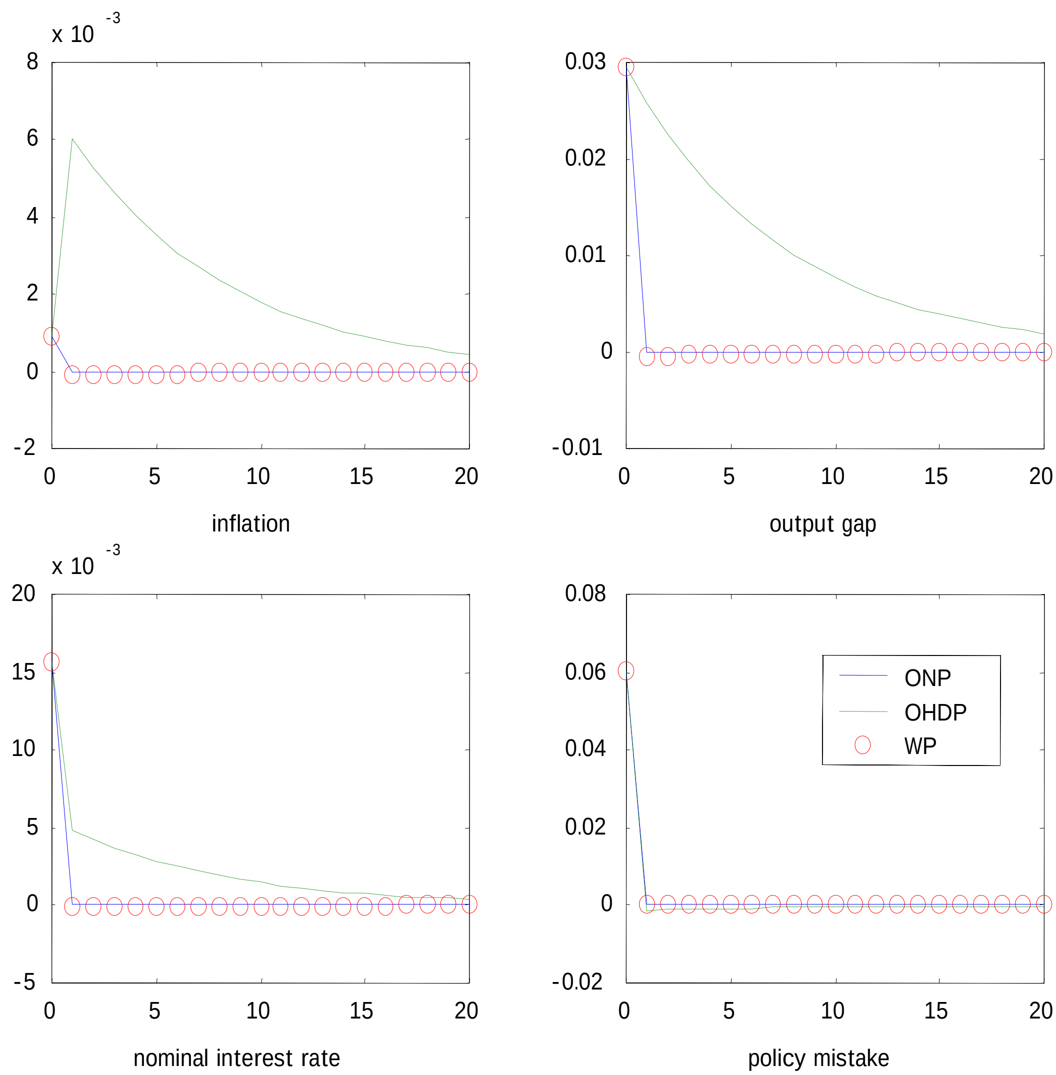

Chart 4 RESPONSES TO DEMAND SHOCK

White noise demand shock

Output weight in loss function = 1.0

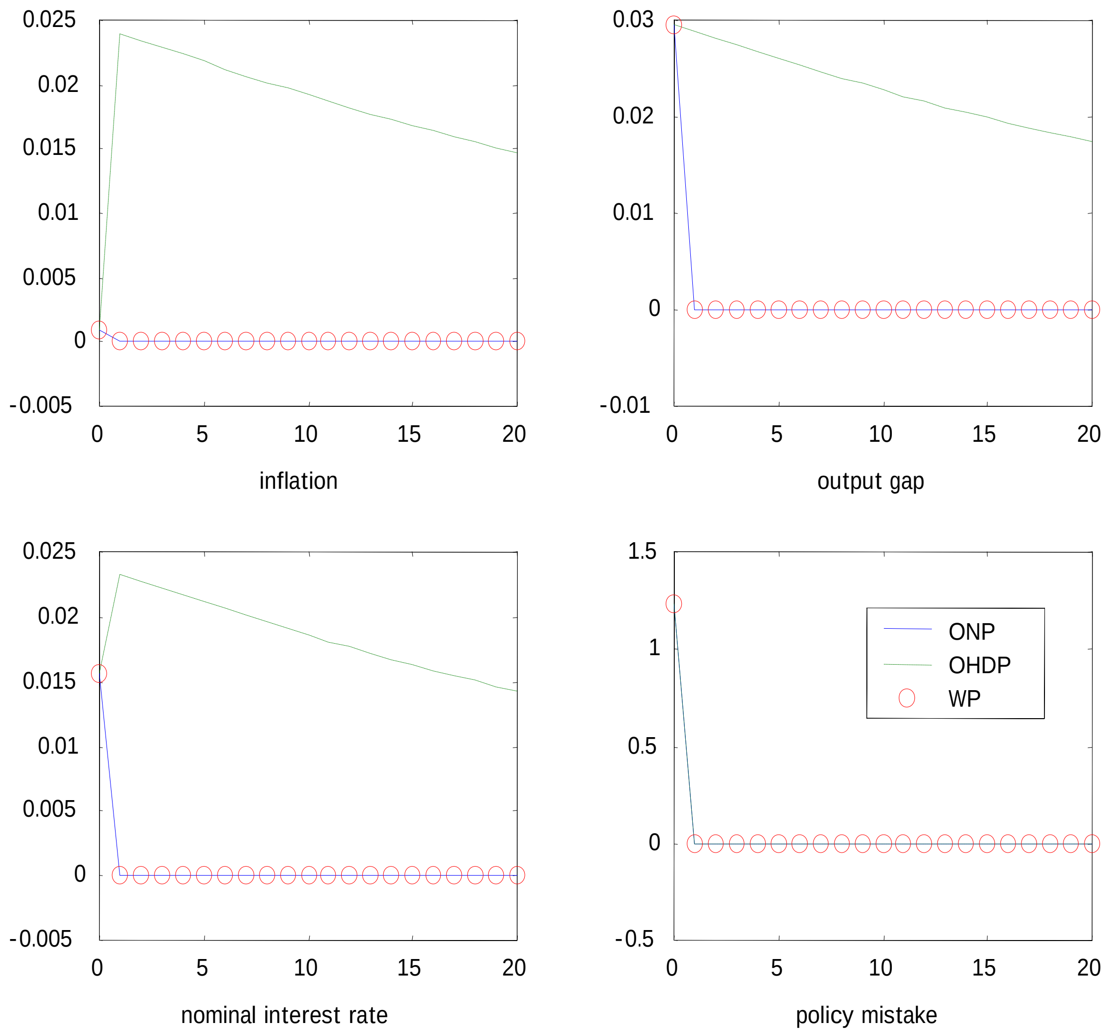

Chart 5: Slope of IS curve

Sample average of 1000 simulation

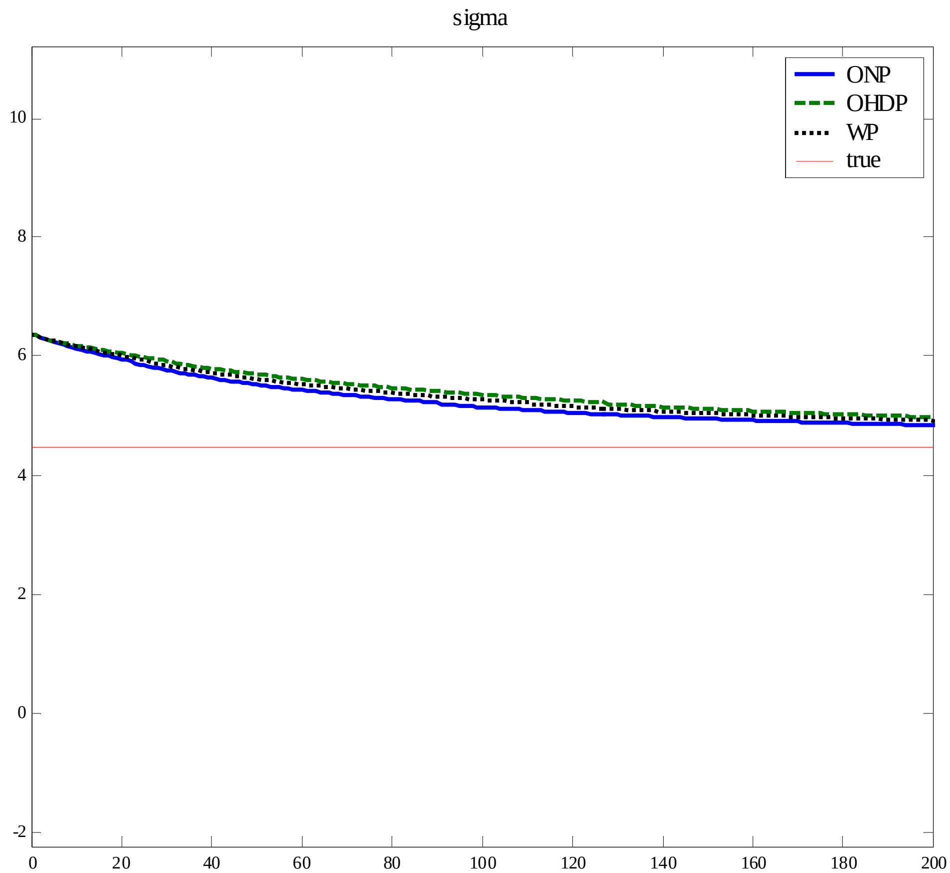

Chart 6: Slope of the Phillips Curve

Sample average of 1000 simulation