Working Paper no. 297

Optimal monetary policy in a regime-switching economy: the response to abrupt shifts in exchange rate dynamics

Fabrizio Zampolli

June 2006

Bank of England

Optimal monetary policy in a regime-switching economy: the response to abrupt shifts in exchange rate dynamics

Fabrizio Zampolli[*]

Working Paper no. 297

[*] Monetary Strategy and Assessment Division, Bank of England.
Email: fabrizio.zampolli@bankofengland.co.uk

The views expressed in this paper are those of the authors, and not necessarily those of the Bank of England. I am grateful to Andrew Blake for several conversations during the preparation of this paper, and to John Frain for acting as a discussant of an earlier version presented at the ECB workshop on Asset Prices and Monetary Policy, 11-12 December 2003, with the title 'Optimal policy and abrupt shifts in economic behaviour'. I also thank Peter Andrews, Mark Astley, Tobias Broer, Tim Cogley, Andrew Filardo, Richard Harrison, Marianna Grimaldi, Lavan Mahadeva, Joe Pearlman, Arnulfo Rodriguez, Mark Salmon, Peter Sinclair, Lars Svensson, Tim Taylor, Tony Yates, participants at the ECB workshop, 2004 CEF conference in Amsterdam, and seminar participants at the Bank of England, and Catholic University of Milan for helpful comments on previous drafts. All mistakes, omissions, and obscurities are my own. This paper was finalised on 21 December 2005.

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

Information on the Bank's working paper series can be found at www.bankofengland.co.uk/publications/workingpapers/index.htm.

Publications Group, Bank of England, Threadneedle Street, London, EC2R 8AH; telephone +44 (0)20 7601 4030, fax +44 (0)20 7601 3298, email mapublications@bankofengland.co.uk.

ISSN 1749-9135 (on-line)

Contents

Abstract

This paper examines the trade-offs that a central bank faces when the exchange rate can experience sustained deviations from fundamentals and occasionally collapse. The economy is modelled as switching randomly between different regimes according to time-invariant transition probabilities. We compute both the optimal regime-switching control rule for this economy and optimised linear Taylor rules, in the two cases where the transition probabilities are known with certainty and where they are uncertain. The simple algorithms used in the computation are also of independent interest as tools for the study of monetary policy under general forms of (asymmetric) additive and multiplicative uncertainty. An interesting finding is that policies based on robust (minmax) values of the transition probabilities are usually more conservative.

Key words: Monetary policy; regime-switching; uncertain probabilities

JEL classification: C6, E5

Summary

A common concern among central bankers is that the true or perceived existence of financial imbalances or asset price misalignments could at some point in time lead to sudden and large adjustments in asset prices, with potentially adverse consequences for inflation and output. For instance, one of the major risks that has worried some members of the Bank of England's Monetary Policy Committee (MPC) in the past has been the possibility that sterling could suddenly fall by a material amount. Other risks routinely debated by actual policymakers, including oil price hikes or abrupt changes in key econometric relationships, may also be asymmetric – that is, a given change may be more likely to occur in one direction than in the opposite. Nevertheless, modelling of asymmetric risks is not very common in the monetary policy literature, possibly because of the lack of readily-applicable technical tools.

In this paper we examine the trade-offs that the policymaker faces when the exchange rate can experience sustained deviations from its fundamental value (ie the value implied by interest rates absent any economic shock) and occasionally collapse. To do so we use a simple method which has rarely been applied in the economics literature. The method allows us to solve for the optimal monetary policy in an economy subject to regime shifts, while retaining the flexibility and simplicity of more commonly applied methods. The method could be applied in other ways that are not considered in this paper and can be considered as a general tool for studying uncertainty in monetary policy. In particular, it provides an example of how policymakers can incorporate judgemental information about a potential misalignment (and the uncertainties associated with it) into their macroeconomic model, and work out the best policy response based on that judgement.

Our analysis is based on a small open economy model, comprising a demand equation, a Phillips curve which determines prices, and an equation linking the real exchange rate to the domestic real interest rate. We modify this model to incorporate regime switching in the exchange rate. In one regime, which we call the bubble regime, any shock can lead the exchange rate to increasingly deviate from its fundamental value. Depending on the sign of the shock, the exchange rate can continue to rise above its fundamental value, or it can continue to fall below it. In the other regime, which we call the no-bubble regime, the exchange rate displays transitory fluctuations around its fundamental value. The times at which the bubble begins and ends are uncertain to the policymaker. Moreover, the size of the correction in the exchange rate, which occurs when the

economy switches from the bubble to the no-bubble regime, will vary over time as it depends on the past behaviour of the exchange rate as well as the interest rate.

Analysis of the optimal regime-switching policy rule shows the existence of an intuitive link in the bubble regime between the optimal response of the interest rate to the exchange rate and the expected duration of a bubble. When the bubble is expected to last for at least two years, the optimal interest rate is negatively related to movements in the real exchange rate and becomes more responsive as the expected duration of the bubble lengthens (an increase in the exchange rate being an appreciation). Similarly, in the no-bubble regime there is an intuitive link between the response to the exchange rate and the probability of the bubble emerging: for lower probabilities of bubbles the interest rate is positively correlated with exchange rate fluctuations (reflecting the likely transitory nature of exchange rate movements) but becomes less responsive as the probability of a bubble increases. For high probabilities of the bubble the interest rate responds negatively and becomes more reactive to exchange rate fluctuations as the probability rises further (reflecting the likely onset of a bubble). Another characteristic of the optimal regime-switching interest rate rule is that in both regimes the interest rate is for the most part less responsive to inflation and output fluctuations than in the absence of regime uncertainty, with the degree of caution increasing as both transition probabilities approach a half.

A key result of the paper concerns the assumptions that the policymaker makes about the (unknown) probabilities of moving between bubble and no-bubble regimes. These probabilities could be highly uncertain since historical experience might provide little or no help in quantifying them. We find that there are ‘robust’ values of the probabilities corresponding to more muted policy responses, where by ‘robust’ we mean values of the probabilities which can be assumed by the policymaker without fear of causing unnecessary volatility in output and inflation, were they to prove wrong in hindsight. This result is interesting as in the robust control literature uncertainty is often found to lead to more reactive policy responses than in the absence of uncertainty.

1 Introduction

A common concern among central bankers is that the true or perceived existence of financial imbalances or asset price misalignments could at some point in time lead to sudden and large adjustments in asset prices, with potentially adverse consequences for inflation and output stability. For instance, one of the major risks that has worried some members of the Bank of England's Monetary Policy Committee (MPC) in the past has been the possibility that sterling could suddenly fall by a material amount.(1) Other risks routinely debated by actual policymakers, including oil price hikes or abrupt changes in key econometric relationships, may also be asymmetric.(2) Nevertheless, modelling of asymmetric risks is not very common in the monetary policy literature, possibly because of the lack of readily-applicable technical tools.(3)

In this paper we examine the trade-offs that the policymaker faces when the exchange rate can experience sustained deviations from fundamentals and occasionally collapse. To do so we use a simple algorithm which has rarely been applied in the economics literature. Our analysis is based on the small open economy model of Ball (1999), which comprises a demand equation, a Phillips curve and an equation linking the real exchange to the real interest rate. We modify this model to incorporate regime switching in the exchange rate. In one regime, which we call the bubble regime, any shock can lead the exchange rate to deviate increasingly from fundamentals. In the other regime, which we call the no-bubble regime, the exchange rate displays transitory fluctuations around its fundamental value. The evolution over time of these two regimes is described by a Markov chain so that the times at which the boom begins and ends are stochastic.

(1) See eg the *Minutes* of the February 2002 MPC meeting: "*(...) Some members placed weight on upside risks to the inflation outlook. Two main risks to inflation were emphasised: from the possibility of a depreciation of sterling's exchange rate and from the possibility that consumption would not slow as much as projected*" (page 10). At the same meeting some members were also worried about potential financial imbalances: "*Persistently rising debt levels potentially increased the probability that any adjustment to household balance sheets would be abrupt rather than smooth, with an attendant risk of a fall in asset prices and, thus, in the value of collateral. (...) In the view of some members, therefore, rising debt levels risked increasing the volatility of output and so of inflation in the medium term, potentially making future inflation outturns more uncertain. Other members placed little or no weight on this.*" (page 5). On the risks posed by financial imbalances see also Borio and Lowe (2002).
(2) See eg the discussion of skews and asset prices in Goodhart (2001), pages 178-80. A proper account of asymmetric risks could also help explain part of the (large) deviations often observed between the actual policy rate and that implied by various versions of estimated simple rules. As pointed out by Svensson (2003a), estimated Taylor rules for a closed economy like the US leave approximately one third of the variance unexplained.
(3) The literature on monetary policy has produced a number of papers on whether simple rules should include asset prices or asset price misalignments but has unfortunately been relatively silent on the more general question of how policy should optimally react to asymmetric or one-sided risk. Skewed risks and policymakers' cognitive biases are discussed by Al Nowaihi and Stracca (2003). Svensson (2003b) investigates, in a simple model, the optimal response to low-probability extreme events under various types of loss functions.

Moreover, the size of the correction in the exchange rate, which occurs when the economy switches from the bubble to the no-bubble regime, is endogenous for it depends on the lagged exchange rate as well as the policy instrument (plus any additive shocks). We compute the optimal control rule for such a regime-switching economy, which is itself regime switching, and compare it to the responses implied by an optimised linear Taylor rule. Optimal monetary policy in any given regime, whether implemented through the optimal regime-switching rule or through a (suboptimal) optimised regime-invariant rule, crucially depends on the probabilities of moving from one regime to the other.[(4)]

The algorithm for computing the optimal regime-switching policy is a modification of the standard linear quadratic regulator problem, in which the constraint is given by a Markov regime-switching vector autoregression (with any finite number of observable regimes) rather than a stationary vector autoregression. The regime-switching model belongs to a class of models that have been studied in the engineering literature at least since Aoki's contribution (1967).[(5)] This formulation is sufficiently general to allow the modelling of a large range of different asymmetric (and symmetric) risks, either concerning changes in the economy's dynamics or additive disturbances, and to accommodate different models. The solution algorithm is also sufficiently simple to be amenable to further interesting developments. For these reasons the algorithm's applicability extends beyond the particular application considered in this paper and can be considered as a general tool for the study of multiplicative and additive uncertainty in monetary policy.[(6)]

The small open economy model used in the analysis is meant to capture the main effects of monetary policy in the simplest possible way. In particular, it does not incorporate rational or other forms of forward-looking expectations. Nevertheless, this simple model is close in spirit to

(4) There are some interesting papers which introduce regime switching in monetary policy. Andolfatto and Gomme (2003) assume that the growth rate of money switches periodically between high and low growth regimes. Schorfheide (2005) assumes that monetary policy follows an interest rate rule that is subject to regime switches in the target inflation rate. Both papers study the macroeconomic implications of these assumptions in environments where private agents cannot observe policy but have to learn about it. Our work differs from these studies because the regime-switching monetary policy rule is not assumed but derived optimally from a model in which regime switches characterise the behaviour of the exchange rate. In other words, here we focus on the decision problem faced by the policymaker, whereas the cited papers deal with the decision problems faced by private agents.

(5) In the engineering control literature models with regime shifts are currently referred to as Markov Jump Linear Systems (MJLS). Several results on the control of MJLS can be found in Mariton (1990) and, more recently, in Costa, Fragoso and Marques (2005).

(6) To the best of the author's knowledge, the only application to macroeconomics of MJLS is do Val and Başar (1999), who extend the macroeconometric model of Pindyck (1973) to allow time-variant parameters and apply receding horizon control problems to the resulting MJLS. Within the engineering literature, see also Blair and Sworder (1975) for an application to an econometric model in continous time.

the larger macroeconometric models that recently were or are still in use at several central banks. It is not unusual that in versions of these models the uncovered parity condition is, due to its empirical failure, replaced by reduced forms which are good empirical approximations of actual exchange rate behaviour.(7) The introduction of regime switching in the exchange rate equation in our model is meant to capture the complex behaviour of financial market agents, which arguably it would not be possible to characterise explicitly in current state-of-the-art fully microfounded general equilibrium models.(8) Thus, the application in this paper can also be thought of as an example of how a policymaker can incorporate judgemental information about a potential misalignment and the uncertainties associated with it into her macro model and work out the best policy response based on that judgement. Most importantly, there are clear advantages in terms of intuition and transparency of working with a backward-looking model. The trade-offs which a policymaker faces can often be more clearly analysed without the additional layer of complexity represented by any expectation-formation mechanism.(9) Besides, we believe that the main insights and conclusions in this paper are likely to carry over to a forward-looking model.(10)

The recent literature on monetary policy and asset prices – eg Bernanke and Gertler (2000, 2001), Cecchetti *et al* (2000, 2003), Batini and Nelson (2000), Filardo (2001), Tetlow (2003) – looks at whether simple rules should give weight to asset prices, usually over and above their predictive power for inflation and output. These papers derive their conclusions mainly from simulating some model under different time-invariant linear reaction functions (optimised or not) and ranking them according to the computed losses. However, the presence of a non-linearity (eg a bubble or misalignment) makes an otherwise linear model non-linear and calls in principle for a non-linear reaction function. In this paper we compute a regime-dependent (and hence time-variant) policy rule as the solution of an optimal control problem. Bordo and Jeanne (2002) have also pointed out

(7) See also the discussion in Section 3 below.

(8) Popular linear rational expectation models used for monetary policy analysis are usually derived from general equilibrium models based on representative infinitely-lived agents (see eg Woodford (2003)). We believe that these models are hardly the appropriate vehicles for studying the role of non-fundamental deviations of asset prices. General equilibrium models featuring overlapping generations and aggregate uncertainty would be more appropriate but are currently difficult to solve. Hence, *ad hoc* modelling of asset price misalignments or bubbles in linear macro models seems difficult to avoid.

(9) Extending the solution algorithm to solve for rational expectations models of private sector's behaviour involves non-trivial technical issues, which are the object of current research.

(10) An essential ingredient of the model is the endogeneity of the exchange rate misalignment to the interest rate which, among other things, creates a trade-off between the need to offset the current and expected future consequences of a booming exchange rate on one side, and the need to offset the consequences of a potential sizable correction on the other side. The inclusion of rational expectations or other forms of forward-looking behaviour on the part of private agents would not eliminate such a trade-off. In a forward-looking model the exchange rate risk premium featuring in the uncovered interest parity equation could, for instance, be modified to include regime switching and a dependency on the interest rate (as well as other state variables).

that in reality the optimal monetary rule is unlikely to take the form of a linear time-invariant rule, even if augmented by a linear term in asset prices.[(11)] To prove their point, they consider a stylised New Keynesian model of the economy in which monetary policy can affect the probability of a credit crunch. This model is assumed to have only three periods and is solved by backward induction. The optimal interest rate is shown to be a function of the probability that private agents attach to being in a 'new economy'. In our application we compute optimal policy for an infinite horizon and show how the optimal rule's feedback coefficients depend on the regime-switching transition probabilities. Unlike Bordo and Jeanne, however, while policy affects the size of the misalignment, it does not affect the probabilities. Assuming that the probabilities are exogenous is not an unreasonable assumption if one considers the high degree of uncertainty concerning both the knowledge of the stochastic properties of an asset price and their relationship with monetary policy.[(12)] Given such uncertainty, we also complement the computation of optimal policy with the welfare analysis of the incorrect assumptions about the transition probabilities.[(13)]

It is important to note that the optimal regime-switching rule is computed in this paper under the assumption that the regime (bubble or no-bubble) is observable by the policymaker, albeit with a delay. *Ex-post* identification of a regime is not implausible in several situations but more generally it might be regarded as a limitation. If the regime is not identifiable, even with a delay, then one obvious alternative for the policymaker is to adopt a regime-invariant policy rule optimised to take into account the regime-switching nature of the economy. For this reason, we also include an analysis of optimised regime-invariant rules and see how they compare with the optimal regime-switching rule. In particular, we consider a Taylor rule which includes a response to the

(11) For a critique of the simple rule approach also see Bean (2003).

(12) Assuming the wrong relationship between the probabilities of a regime shift and monetary policy is not necessarily better than assuming the probabilities to be exogenous. Clearly, endogenising the probabilities is in principle possible, but even for relatively small models the solution algorithms may in practice encounter the curse of dimensionality. Therefore, an additional advantage of the proposed method is that it can be applied to models with a relatively large number of state variables such as those in use at several central banks.

(13) Another closely related paper is Gruen *et al* (2003). Gruen *et al* considers the optimal response to a bubble in a two-equation closed economy model. Their bubble affects aggregate demand only, unlike in our paper in which the misalignment in the exchange rate directly affects both output and inflation. Most importantly, their method of computation differs from ours. In their paper optimal policy is computed by backward induction, assuming that the bubble must burst within 14 periods and will never form again. In our paper optimal policy is computed for an infinite horizon and it is therefore the solution to a fixed point problem. Our algorithm, therefore, enables us to analyse the optimal feedback coefficients across different scenarios. Another difference is that we also provide an analysis of the losses which the policymaker would incur when basing her policy on the incorrect probabilities governing the evolution of the bubble. A judgement over such probabilities is clearly a key, albeit highly uncertain, input in her policy decision.

exchange rate as well as a simple Taylor rule which does not.[14]

We briefly anticipate the main results which emerge from our application. The analysis of the optimal regime-switching control rule shows an intuitive link in the bubble regime between the response coefficient associated with the exchange rate and the (unconditional) expected duration of a bubble: when the bubble is expected to last for at least two years, the optimal interest rate is negatively correlated with real exchange rate fluctuations and becomes more responsive as the expected duration of the bubble lengthens (an increase in the exchange rate being an appreciation). In the no-bubble regime there is an intuitive link between the response to the exchange rate and the probability of the bubble emerging: for lower probabilities the interest rate is positively correlated with exchange rate fluctuations (reflecting the likely transitory nature of exchange rate movements) and becomes less responsive as the probability of a bubble increases; for higher probabilities the interest rate responds negatively and becomes more reactive to exchange rate fluctuations as the probability rises further (reflecting the likely onset of a bubble). Another characteristic of the optimal regime-switching interest rate rule is that in both regimes the interest rate is for the most part less responsive to inflation and output fluctuations than in the absence of regime uncertainty, with the degree of caution increasing as both transition probabilities approach their intermediate values of a half.

As stressed above, the difficulty or impossibility in identifying the regime might in practice lead the policymaker to follow a regime-invariant rule. A natural question then is how costly the latter rule would be relative to the optimal one or, equivalently, how costly it would be not to be able to observe the regime. In this regard, we find that the losses from adopting the latter rule are small relative to the optimal regime-switching rule except when the bubble has both a low probability and a long expected duration. Indeed, as both transition probabilities approach 1/2 the optimal regime-switching rule becomes ever more similar and eventually converge to the optimised

(14) Bullard and Schaling (2001) and Filardo (2001) are previous studies that have analysed optimal linear Taylor rules in a regime-switching economy. Bullard and Schaling (2001) consider a closed economy where there are unobserved regime shifts in productivity. They provide an analytical derivation of the optimal linear policy rule and find that the optimal Taylor rule includes additional terms which are leading indicators of future supply shocks. Apart from these additional terms the rule is unchanged: the feedback coefficients on inflation and current output deviations are not affected by regime uncertainty. In our paper regime uncertainty affects one of the model's parameters instead of an exogenous driving process such as productivity. As a result, all the feedback coefficients in the optimised linear rules are affected by regime uncertainty. Filardo (2001) uses a Markov process to model an asset price bubble but in a closed-economy model. His paper differs from the current one for two reasons. First, the non-fundamental component of the asset price does not depend on the interest rate or other shocks, whereas in the present paper the interest rate can affect the deviation from fundamentals. Second, the evaluation of policy rules is obtained through a Monte Carlo simulation, whereas we evaluate the interest rate rules by solving a system of Lyapunov equations.

regime-invariant rule. The differences between the two rules tends to be large when both transition probabilities are low.

The optimised regime-invariant rule (which includes the exchange rate) also shows intuitive links between the response to the exchange rate and the transition probabilities, with the obvious difference that in this case the optimal responses are heavily affected by both probabilities. In addition, policy is also found to be for the most part less responsive to output and inflation fluctuations. By contrast, an optimised simple Taylor rule (without the exchange rate) entails stronger responses to inflation and output fluctuations than in the absence of regime uncertainty. For bubbles of high expected duration the responses increase with the probability of the bubble arising. These stronger responses possibly compensate for the lack of a negatively correlated response to the exchange rate which is found to be appropriate in the unconstrained regime-invariant rule.

Finally, a key result of the paper concerns the assumptions that the policymaker makes about the (unknown) transition probabilities. These probabilities could be highly uncertain since historical experience might provide little or no help in quantifying them. We find that there are robust (minmax) values of the probabilities not falling on the boundaries of the feasible set of values. These robust values generally correspond to more muted policy responses. This result is interesting as in the robust control literature uncertainty is often found to lead to more reactive policy responses than under certainty equivalence.(15)

The paper is organised as follows. Section 2 describes the quadratic optimal control problem with regime shifts and its solution. It also explains how to evaluate regime-invariant and simple rules in a regime-switching economy. Section 3 describes the model used in the application. Section 4 analyses the responses of monetary policy to movements in output, inflation, and the lagged exchange rate. Section 5 examines the choice of the optimal assumptions about the uncertain transition probabilities. Section 6 concludes indicating possible future avenues for research.

---

(15) It is important to notice that the application provides some useful insights, but abstracts from the possible existence of more than one asset price misalignment, which in practice considerably complicates actual policymaking. For instance, while, according to some commentators, the United Kingdom apparently had an overvalued exchange rate in the late 1990s, suggesting lower interest rates to lean against it, it has also had rapidly rising house prices, which might have suggested the opposite response.

2 The quadratic optimal control problem with regime shifts

2.1 Formulation

The policymaker's problem is to choose a decision rule for the control u_t to minimise the expected value of the intertemporal loss function:[16]

$$\sum_{t=0}^{\infty} \beta^t r\left(x_t, u_t\right) \qquad \beta \in (0, 1) \tag{1}$$

subject to x_0, s_{-1} given, and the following model of the economy

$$x_{t+1} = A\left(s_t\right) x_t + B\left(s_t\right) u_t + \varepsilon_{t+1} \qquad t \geq 0 \tag{2}$$

where $r\left(x, u\right)$ is the period loss, β is the discount factor, x is the $n \times 1$ vector of state variables, u is the $m \times 1$ vector of control variables and ε is the $n \times 1$ vector of mean-zero shocks with variance-covariance matrix Σ_ε. We assume that the period loss is a quadratic form:

$$r\left(x_t, u_t\right) = x_t' R x_t + u_t' Q u_t \tag{3}$$

where R is a $n \times n$ positive definite matrix and Q is $m \times m$ positive semi-definite matrix. The matrices A and B are stochastic and take on different values depending on the regime or state of the world $s_t \in \{1, ..., N\}$. The regime s_t is assumed to be a Markov chain with probability transition matrix

$$P = \left[p_{ij}\right]_{i,j=1,..,N} \tag{4}$$

in which $p_{ij} = prob\left\{s_t = j | s_{t-1} = i\right\}$ is the probability of moving from state i to state j at t; and $\sum_{j=1}^{N} p_{ij} = 1$, $i = 1, ..., N$.[17] These probabilities are assumed to be time-invariant and exogenous. The policymaker's information set is assumed to be

$$I_t = \left\{x^t, \varepsilon^t, s^{t-1}, P, A, B, \Sigma_\varepsilon\right\}$$

where $t \geq 1$; $A = \left(A_1, ..., A_N\right)$, $B = \left(B_1, ..., B_N\right)$ are set of values that the stochastic matrices A and B can take; $x^t = \{x_r\}_{k=0}^{t}$, $\varepsilon^t = \{\varepsilon_k\}_{k=0}^{t}$ and $s^{t-1} = \{s_k\}_{k=0}^{t-1}$ are the 'histories' of the variables. Therefore, at t the policymaker knows the realisation of s_{t-1}, but not the realisation of the current regime s_t. This means that the uncertainty faced by the policymaker is about where the system is at time t and where it will be at $t+1$, $t+2$, and so forth. Note the difference with the stochastic specification in Hamilton (1989). In the latter the regime is a latent variable.

(16) An introduction to dynamic programming and the optimal linear regulator problem can be found in Ljungqvist and Sargent (2000), Chapters 2-4. Kendrick (2002) provides a more comprehensive and advanced treatment.
(17) For an introduction to Markov chain and regime-switching vector autoregressive models see eg Hamilton (1994).

The above control problem boils down to N *separate* optimal control problems (with generic period loss $r(\cdot)$) when $P = I_N$ (N-dimensional identity matrix), each corresponding to a different regime. Obviously, the above problem also reduces to a unique one in the case in which the matrices A_i and B_i ($i = 1, ..., N$) are identical, regardless of P.

2.2 Solution

Solving the problem means finding a state-contingent decision rule, that is a rule which tells how to set the control u_t as a function of the current vector of state variables, x_t, and the time-$(t-1)$ regime, s_{t-1} (henceforth, regime or state of the world will be used interchangeably). Associated with each state of the world is a Bellman equation. Therefore, solving the model requires jointly solving the following set of *intertwined* N Bellman equations:

$$\upsilon(x_t, i) = \max_{u_t} \left\{ r(x_t, u_t) + \beta \sum_{j=1}^{N} p_{ij} E_t^{\varepsilon} \left[\upsilon(x_{t+1}, j) \right] \right\} \qquad i = 1, ..., N \tag{5}$$

where $\upsilon(x_t, i)$ is the continuation value of the optimal dynamic programming problem at t written as a function of the state variables x_t as well as the state of the world at $t-1$, $s_{t-1} = i$; E_t^{ε} is the expectation operator with respect to the martingale ε, conditioned on information available at t, such that $E_t^{\varepsilon}\left[\varepsilon_{t+1}\right] = 0$. The policymaker has to find a sequence $\{u_t\}_{t=0}^{\infty}$ which maximises her current pay-off $r(\cdot)$ as well as the discounted sum of all future pay-offs. The latter is the expected continuation value of the dynamic programming problem and is obtained as the average of all possible continuation values at time t weighted by the transition probabilities **(4)**. Given the infinite horizon of the problem, the continuation values (conditioned on a particular regime) have the same functional forms.

Given the linear-quadratic nature of the problem, let us assume that

$$\upsilon(x_t, i) = x_t' V_i x_t + d_i \tag{6}$$

($i = 1, .., N$), where V_i is a $n \times n$ symmetric positive-semidefinite matrix, and d_i is a scalar. Both are undetermined. To find them, we substitute **(6)** into the Bellman equations **(5)** (after using **(3)** and **(6)**) and compute the first-order conditions, which give the following set of decision rules:

$$u(x_t, i) = -F_i \cdot x_t \tag{7}$$

($i = 1, .., N$), where the set of F_i depend on the unknown matrices V_i, $i = 1, .., N$. By substituting these decision rules back into the Bellman equations **(5)**, and equating the terms in the quadratic forms, we find a set of *interrelated* Riccati equations, which can be solved for V_i

$(i = 1, .., N)$ by iterating jointly on them, that is

$$[V_1 ... V_N] = T\left([V_1 ... V_N]\right) \quad \textbf{(8)}$$

The set of interrelated Riccati equations defines a contraction T over $V_1, ..., V_N$, the fixed point of which, $T(\cdot)$, is the solution being sought. After lengthy matrix algebra, the resulting system of Riccati equations can be written in compact form as:

$$V_i = R + \beta G\left[A'VA|_{s=i}\right] - \beta^2 G\left[A'VB|_{s=i}\right]\left(Q + \beta G\left[B'VB|_{s=i}\right]\right)^{-1} G\left[B'VA|_{s=i}\right] \quad \textbf{(9)}$$

where $i = 1, .., N$, and $G(\cdot)$ is a conditional operator defined as follows:

$$G\left[X'VY|_{s=i}\right] = \sum_{j=1}^{N} X'_j \left(p_{ij} V_j\right) Y_j$$

where $X \equiv A, B$; $Y \equiv A, B$. Written in this form the Riccati equations contain 'averages' of different 'matrix composites' conditional on a given state i.

Having found the set of V_i which solves **(9)**, the matrices F_i in the closed-loop decision rules **(7)** are given by:

$$F_i = \left(Q + \beta G\left[B'VB|_{s=i}\right]\right)^{-1}\left(\beta G\left[B'VA|_{s=i}\right]\right) \quad \textbf{(10)}$$

$(i = 1, .., N)$. Solving for the constant terms in the Bellman equations **(5)** – after substitution of **(7)** – gives $(I_N - \beta P)\, d = \beta P \Gamma$. The vector of scalars $d = [d]_{i=1,...,N}$ ($N \times 1$ vector) in the value functions **(6)** is given by[(18)]

$$d = (I_N - \beta P)^{-1} \beta P \Gamma \quad \textbf{(11)}$$

where $\Gamma = [tr\,(V_i \Sigma_\varepsilon)]_{i=1,...,N}$ ($N \times 1$ vector).[(19)]

The decision rules **(7)** depend on the uncertainty about which state of the world will prevail in the future, as reflected in the transition probabilities **(4)**. Yet, the response coefficients do not depend on the variance-covariance matrix Σ_ε of the zero-mean shock ε in **(2)** or the initial conditions (x_0, s_0). Thus, certainty equivalence holds with respect to the additive shock ε. The noise statistics, as is clear from **(11)**, affect the objective function.[(20)]

(18) We assume that $(I_N - \beta P)$ is invertible. Given that P is a stochastic matrix, a necessary condition for its invertibility is that $\beta < 1$. If $(I_N - \beta P)$ is not invertible, other methods can be used to find the solution d.
(19) The law of transition **(2)** can be generalised to make the variance of the noise statistics vary across states of the world, ie $x_{t+1} = A(s_t)\, x_t + B(s_t)\, u_t + C(s_t)\varepsilon_{t+1}$. Assuming $E^\varepsilon\left(\varepsilon_t \varepsilon'_t\right) = I$, then the covariance matrix of the white-noise additive shocks would be $\Sigma(s_t) = C(s_t)\, C(s_t)'$ or, to simplify notation, $\Sigma_i = C_i C'_i$ $(i = 1, .., N)$. The introduction of a state-contingent variance for the noise process does not affect the decision rules $u_t = -F_i x_t$ but affects the value functions through Γ in **(11)**: $\Gamma = [tr\,(V_i \Sigma_i)]_{i=1,...,N}$.
(20) The solution method described above also applies, *mutatis mutandis*, to one problem in which the *timing of uncertainty* is slightly different than assumed in Section 2. This problem differs from the original one in that the

The above solutions incorporate the standard linear regulator solutions as two special cases. First, by setting the transition matrix $P = I_N$, one obtains the solution of N separate linear regulator problems, each corresponding to a different regime on the assumption that each regime will last forever (and no switching to other regimes occurs). This case could be useful as a benchmark to see how the uncertainty about moving from one regime to another impacts on the state-contingent rule. In other words, by setting $P = I_N$, we are computing a set of rules which will differ from ones computed with $P \neq I_N$, in that the latter will be affected by the chance of switching to another regime. Second, by choosing identical matrices A_i and B_i, the solution obtained is trivially that of a standard linear regulator problem with a time-invariant law of transition. In this case **(10)** boils down to

$$F = \left(Q + \beta B'VB\right)^{-1} \beta B'VA$$

where V is the solution to the single Riccati equation

$$V = R + \beta A'VA - \beta^2 A'VB \left(Q + \beta B'VB\right)^{-1} B'VA$$

and **(11)** boils down to the constant

$$d = (1-\beta)^{-1} \beta \cdot tr\left(V\Sigma_\varepsilon\right)$$

The above algorithm can be derived having in mind the dynamic programming solution to a simple precautionary savings model (see eg Chapter 3 in Ljungqvist and Sargent (2000)), which includes both continuous and discrete value state variables, and imposing the assumption of a quadratic loss function common in the optimal monetary policy literature. Similar solutions exist in the engineering literature. The algorithm in fact can be regarded as a variation of those introduced by Aoki (1967). His general approach can be adapted to include the use of a Markov chain over an infinite horizon. More recent work by Costa *et al* (2005) has provided several results on the control of models such as **(2)**, which in the engineering literature have been given the name of Markov Jump Linear Systems (MJLS). Despite its simplicity and potential usefulness in the analysis of symmetric and asymmetric uncertainty, the above algorithm or similar ones have apparently not been applied to optimal monetary policy problems, with the exception of do Val and Başar (1999).[21] Beside the application considered in this paper, there is a couple of other

stochastic matrices A and B in **(2)** are assumed to be contingent on s_{t+1} instead of s_t. Accordingly, the information set is assumed to include s_t (whereas in the original formulation the information set includes s_{t-1} but not s_t). The solution procedure goes through the same steps as described above yielding solutions of the same form. The only difference is that the value functions and the closed-loop decision rules **(7)** are now contingent on the latest realisation of the regime variable: that is $\upsilon(x_t, s_t)$ replaces $\upsilon\left(x_t, s_{t-1}\right)$ in **(5)** and **(6)**; $F(s_t)$ replaces $F\left(s_{t-1}\right)$ in **(7)**; and the initial conditions (x_0, s_0) replace $\left(x_0, s_{-1}\right)$.

(21) Another related paper is Shupp (1972). It applies general dynamic programming techniques to solve a macro

relevant applications that come to mind. One is to use the above algorithm in conjunction with quadrature methods. The latter would allow to discretise various types of continuous distributions imposed over the model's parameters to reflect eg the judgemental information possessed by the policymaker. Another application involves using mixture models to approximate non-normal distributions (eg t-distribution), which could be accommodated within the current framework with few modifications.(22)

2.3 Evaluation and optimisation of regime-invariant rules and simple rules

There may be situations in which implementing the optimal regime-switching (ORSC) control rule is not appropriate or feasible. Therefore we consider the case in which the policymaker wants to restrict the policy rule to respond to a subset of the state variables. In what follows such a suboptimal (or simple) rule can be written as:

$$u_t = -F_i x_t \qquad \textbf{(12)}$$

where i indicates the regime at $t-1$ (ie $s_{t-1} = i$) and where F may or may not have zero entries. One important special case is a rule whereby the control variable responds to all variables in the state vector x but not to the regime s. In this case, $F_i = F \quad (i = 1, ..., N)$ in **(12)**. Such regime-invariant rule can be further simplified by imposing zero restrictions on F.

The first step in finding the matrix of coefficients F which minimise the expected value of the loss function **(1)** is to evaluate the performance of a generic policy rule. To this purpose, we replace **(12)** into the set of Bellman equations **(5)** and equate the coefficients of the resulting quadratic forms to obtain a system of interrelated Lyapunov equations:

$$V_i = R + F_i' Q F_i + \beta H \left[D'VD|_{s=i} \right] \qquad \textbf{(13)}$$

with $i = 1, ..., N$, and $H[\cdot]$ is a conditional operator defined as follows:

$$H \left[D'VD|_{s=i} \right] \equiv \sum_{j=1}^{N} D_{ij}' \left(p_{ij} V_j \right) D_{ij}$$

$$D_{ij} \equiv A_j - B_j F_i$$

with $i, j = 1, ..., N$. Note that i indicates the state of the world at $t-1$ (ie $s_{t-1} = i$) and j the state of the world at t (ie $s_t = j$). Equating the constant terms yields the same expression for

model in which the autonomous components of aggregate consumption and investment are modelled as Markov processes and where the welfare function is assumed to be non-quadratic. This paper, however, does not consider multiplicative uncertainty.

(22) Good introductions to numerical quadrature methods can be found in Judd (1998) and Miranda and Fackler (2002). On mixture models see eg Kim *et al* (1998).

$d = (d_1, .., d_N)'$ as in **(11)**.[23] The system of Lyapunov equations can be solved by value iteration, in the same fashion as the system of Riccati equations. When all matrices are identical or $N = 1$, the system boils down to a single equation in V:

$$V = R + F'QF + \beta D'VD$$

Having found the set of matrices $(V_1, ..., V_N)$ that solve **(13)**, one can compute the vector of scalars d and the conditional losses associated with **(12)**:

$$\upsilon(x_0, i) = x_0' V_i x_0 + d_i \qquad \textbf{(14)}$$

$(i = 1, ..., N)$. These losses are conditioned on an initial condition for the continuous-value state variables x_0 and a particular regime i.[24] Note that in general, these losses depend not only on the matrix of transition probabilities P and the initial conditions, but also on the variance-covariance matrix Σ_ε. Hence, certainty equivalence with respect to the additive shocks ε is a special case which holds for the optimal unrestricted control rule.[25] Once a routine for computing **(14)** has been created an appropriate numerical optimiser can be employed to find F.

3 Exchange rate regime switching in a stylised small open economy model

In this section we consider Ball's model (1999) of a small open economy. The model is meant to capture the main effects of monetary policy in the simplest possible way. It consists of three equations:

$$y_{t+1} = \alpha y_t - \beta(i_t - \pi_t) - \chi a_t + \eta_{t+1} \qquad \textbf{(15)}$$

$$\pi_{t+1} = \delta\pi_t + \gamma y_t - f(a_t - a_{t-1}) + \varepsilon_{t+1} \qquad \textbf{(16)}$$

$$a_t = \theta(i_t - \pi_t) + \upsilon_t \qquad \textbf{(17)}$$

---

(23) The vector of scalars d is algebraically identical to **(11)**. Numerically, however, its value will generally be different, for Γ depends on the set $(V_1, .., V_N)$ computed from the system of Lyapunov equations rather than from the set of Riccati equations. Of course, if the set of feedback matrices $(F_1, ..., F_N)$ characterises the optimal (unconstrained) solution, then iterating on **(13)** gives the same solution as the system of Riccati equations.

(24) To obtain the average loss not conditioned on the initial regime s_0, we need the unconditional probabilities π_i, which can be computed as the eigenvector of P corresponding to the unit eigenvalue. The average loss is then $E(\upsilon(x_0)) = \sum_{i=0}^{N} \pi_i \cdot \upsilon_i(x_0)$. In some exercises it is appropriate to judge different policies not on the basis of how well they perform across regimes, or on average, but how well the policy performs given the initial regime the economy is in and an initial condition for x, x_0. For example, if the economy is in a world without a bubble but the policymaker believes one may well develop in the future, then policies are probably better assessed conditional on the current regime, which is the one in which a policy action needs to be taken.

(25) For a proof in the context of a linear time-invariant model see also Currie and Levine (1987).

(15) is an open-economy IS equation in which the real interest rate $i - \pi$ and the real exchange rate a affects the output gap y with one period delay.[26] An increase in a is an appreciation of the domestic currency and, therefore, tends to depress spending on domestic goods. η is a white-noise shock with variance σ_η^2. **(16)** is an open-economy Phillips curve, in which the output gap as well as the change in the real exchange rate affects inflation with one period lag. With $\delta = 1$, the Phillips curve is an accelerationist one, the assumption we make throughout. ε is a white-noise shock with variance σ_ε^2. **(17)** is a *reduced-form* equation that relates the real exchange rate to the current level of the real interest rate and a transitory shock υ. The positive sign of θ ($\theta > 0$) captures the idea that a rise in the real interest rate makes domestic assets more attractive relative to foreign ones, leading, other things equal, to an appreciation. The white-noise shock υ, with variance σ_υ^2, reflects all other factors (future expectations, foreign interest rate, etc) that can influence the exchange rate.

The use of the reduced-form **(17)** in place of the uncovered interest parity (UIP) condition can be realistically justified with the empirical failure of the latter. As stressed in eg Batini and Nelson (2000) (see also references cited therein), there are two responses to this failure. One is to assume that deviations from UIP takes the form of a structural shock, which can be thought of as a time-varying risk premium. The other is to replace the UIP with a reduced-form equation which is a good empirical approximation of the actual exchange rate behaviour.[27] In this sense, assuming a reduced form linking the exchange rate to interest rates is not unrealistic or unreasonable, especially if one makes the assumption – as we will do below – that this relationship is not stable but can change over time or across regimes.

As illustrated by Ball (1999), a key feature of the above model is the fact that the exchange rate affects inflation directly through the costs of imported inputs, and indirectly through the effect on demand. The import price channel operates with one period lag, while the output channel takes two periods: one period to affect output, and an additional period for output to affect inflation through the Phillips curve. Ball shows how the existence of an exchange rate channel requires policy to respond not only to output and inflation, as prescribed by the simple version of the Taylor rule (Taylor (1993)) but also to the lagged level of the real exchange rate. More precisely, the optimal response is positive, requiring a rise in the real interest rate when the lagged real

(26) All variables are in logs except i and π, and are expressed as deviations from its long-run equilibria.
(27) Wadhwani (1999) discusses an empirical model of the exchange rate for the UK. Beechey *et al* (2000) provides a model for Australia. Both models are similar to the reduced form adopted in Ball (1999).

exchange rate has appreciated. A simple Taylor rule, not augmented with the exchange rate, is generally suboptimal.

In our application we replace **(17)** with:

$$a_t = \rho(s_t)\, a_{t-1} + \theta(i_t - \pi_t) + \upsilon_t \qquad \textbf{(18)}$$

where the term ρa is an *ad hoc*, but not unrealistic way, of capturing *bubble-like behaviour*, that is the fact that the exchange rate can grow out of line with its long-run equilibrium, whatever the underlying factors. More precisely, **(18)** is modelled as a regime-switching autoregression, in which the autoregressive coefficient $\rho(s_t)$ takes on different values depending on the regime in which the system is:

$$\begin{aligned} \rho(s_t) \quad &> 1 \quad \text{if} \quad s_t = 1 \\ &= 0 \quad \text{if} \quad s_t = 2 \end{aligned}$$

We assume there are two regimes: in one regime (indicated with $s = 1$) $\rho(1) > 1$, that is the exchange rate tends to grow away from its fundamental value (assuming an unchanged real interest rate and no shock); in the other regime (indicated with $s = 2$), $\rho(2) = 0$, that is the exchange rate abruptly collapses towards its fundamental value (again assuming that the real interest rate is also at its neutral level and that there is no shock). The variable s is assumed to evolve as a Markov chain with the following transition probability matrix:

$$P = \begin{bmatrix} 1-p & p \\ q & 1-q \end{bmatrix}$$

where $p = \Pr\{s_t = 2 | s_{t-1} = 1\}$ and $q = \Pr\{s_t = 1 | s_{t-1} = 2\}$, $(t = 0, 1, 2, ...)$. Hence, p is the probability that the bubble crashes when one exists and q is the probability that a new bubble starts.

Note the timing of uncertainty and of the policy decision. At the time policy is chosen the policymaker knows s_{t-1}, but not s_t. That is, the policymaker only knows the past regime but not the current one. Consistently with this assumption, the shock υ_t is also assumed to be observed with a period delay.[28] Thus, the policymaker cannot tell whether a given exchange rate movement is due to a temporary shock or is instead the start of a boom. These assumptions are meant to capture, realistically, the uncertainty that the policymaker faces about the nature of any movements in the exchange rate.

(28) This represents an additional departure from the model of Ball (1999) in that Ball assumes that the shock is observable at the time policy is chosen.

The above specification captures the existence of a bubble or misalignment and should be thought of as reflecting the existence of various phenomena that can cause the exchange rate to grow out of line from its long-run equilibrium. In other words, it does not reflect any particular definition of bubble or theory, but the bubble-like behaviour common to several theories or models.[(29)] Indeed, there is no consensus in the literature on the definition of a bubble and its underlying causes.[(30)] Bubbles may be rational and unrelated to fundamentals (eg Blanchard and Watson (1982)). A reason for holding an asset even if the price is above that suggested by fundamentals is that there is a chance that it will continue to rise, generating an expected capital gain that compensates the asset holder for the risk of a price collapse. Froot and Obstfeld (1991) argue that rational bubbles may be intrinsic, that is reflecting the excessive reaction of market participants to fundamentals. Thus, persistent changes in fundamentals could lead an asset price to be persistently over or undervalued. However, bubbles do not have to rely only on self-fulfilling expectations: they may arise from manias or irrational exuberance (Kindleberger (1978), Minsky (1982) and Shiller (2000)). They may reflect expectations of future higher productivity or earnings, often associated with important technological innovations; when these expectations are subsequently disappointed, asset prices collapse.[(31)] For example, models of learning with endogenous information, in which public information varies with the level of economic activity, may explain the alternation of slow booms and sudden crashes exhibited by several asset markets (Veldkamp (2004)). Bubbles may also emerge and persist because of a co-ordination failure among rational arbitrageurs. Recent work by Abreu and Brunnermeier (2003) show that a synchronisation problem, together with the individual incentive to time the market, results in persistence of a bubble over a sustained period of time. In their set-up news events can have a disproportionate impact relative to their intrinsic information content by enabling synchronisation.[(32)] More recently, de Grauwe and Grimaldi (2003a, b) have shown that, when market participants use different trading rules and transaction costs are important, exchange rates can alternate between periods in which they tightly follow fundamentals and periods in which they appear to be 'disconnected' and more related to their own

---

(29) The 'bubble interpretation' is not the only possible one. An alternative could be to assume $0 < \rho(1) < 1$ and $\rho(2) = 0$. Accordingly, the uncertainty faced by the policymaker would be about whether any observed shock to the exchange rate is going to persist for some time or to vanish after just one period.

(30) See Filardo (2003) for an interesting discussion of macroeconomic bubbles and their implications for monetary policy.

(31) As suggested eg by Meltzer (2003), alternative explanations to rational and irrational bubbles can be formulated by assuming that agents have only imperfect knowledge of the underlying fundamentals and face Knightian uncertainty rather than 'well-behaved' probability distributions.

(32) For an excellent survey of models of asset price behaviour under asymmetric information see Brunnermeier (2001).

past values, a feature confirmed by the empirical evidence cited by the authors.[33]

The rest of the model is standard. The policymaker is assumed to minimise the following standard quadratic loss function:

$$\sum_{t=0}^{\infty} \phi^t \left(\pi_t^2 + \lambda y_t^2 + \varsigma i_t^2\right) \quad \textbf{(19)}$$

where $\phi \in (0, 1)$ is the discount factor, and λ and ς are the penalties on output and the interest rate stabilisation (relative to inflation stabilisation), respectively. In the benchmark case we will assume that $\varsigma = 0$, but we will consider $\varsigma > 0$ in our sensitivity exercises.[34] It is important to notice that the policymaker does not target the exchange rate. As stressed by eg Bean (2003), there is no need to specify a target for asset prices when the targeting rule is a statement about the loss function.[35] However, the instrument rule – as will be shown below – is a function of all state variables, including the exchange rate, for the policymaker uses all available information to forecast inflation and output. From the point of view of optimal control theory, the crucial question then is not whether or not an instrument rule should include an asset price but what weight the policymaker should put on it. We will see that such weight could be not sufficiently 'robust'.

Given the simple dynamic structure of the model, a period is interpreted as a year. The chosen parameterisation is summarised in Table A. The parameters α, β and γ are calibrated using the values estimated for the UK economy in Bean (1998). The equations are estimated individually using GLS over the sample 1950-96 (δ is kept fixed in the estimation). The estimates are similar to those obtained with annual data for the euro area and the United States in similar papers (eg Orphanides and Wieland (1999)). The standard deviations of the demand and inflation shocks, σ_η and σ_ε are taken from the February *Inflation Report* of the Bank of England (2005). These are the standard deviations used in the construction of the fan charts and constitute the subjective assessment of the Monetary Policy Committee. They are, however, based on forecast errors of the past ten years and hence reflect more closely the uncertainty faced by the policymaker in more

(33) de Grauwe and Grimaldi's model is based on high-frequency behaviour.
(34) Several arguments have been put forward in the literature to rationalise a concern for interest rate stabilisation: eg reducing the risks to the financial sector posed by maturity mismatch (Cukierman (1996)); exerting greater influence on long-term bond rates and thus over future economic activity and inflation (Goodfriend (1991)); reducing transaction frictions and the frequency with which a linear policy rule would violate the zero bound on the nominal interest rate (Woodford (2003)). The reason we add a concern for interest rate stabilisation is purely for comparison purposes with the existing literature.
(35) In the terminology of Svensson (2003a) the loss function **(19)** is a general targeting rule. For a definition of instrument rule, see *ibidem*.

Table A: Parameters for the small open economy

Economy		Loss	
α	.72	ϕ	.96
β	.47	λ	1
γ	.49	ς	0
δ	1		
χ	.2		
f	.2		
θ	2		
$\rho_{(1)}$	1.3		
$\rho_{(2)}$	0		
σ_{ε}	.0076		
σ_{η}	.0040		
σ_{υ}	.0250		

recent times. The parameters describing the existence of a foreign exchange rate channel, χ, f and θ, are taken from Ball (1999). We calibrate the existence of a bubble by assuming that the exchange grows by 30 % a year (in the absence of any shock and any offsetting policy action). Together with ρ (1), the standard deviation of the real exchange rate σ_{υ} is chosen so as to provide responses of plausible magnitude. We have chosen a discount factor $\phi = .96$ in accordance with annual data, and we assume that the policymaker cares equally about inflation and output stabilisation. Even though some of the calibrated values are taken from a model of the UK economy, it is not intended that the line of work set out in this paper applies particularly to the United Kingdom: the analysis is meant to be of general relevance.

In Ball (1999) optimal policy is analysed with reference to a time-invariant model in which there is no parameter uncertainty. A criticism levelled against Ball's analysis is that the policymaker fails to take into account (at least) some of the uncertainties that real-world policymakers face, and even though the model can be currently regarded by the policymaker as a good approximation, it may turn out to be inadequate in some aspects in the future (eg Sargent (1999)). In particular, in Ball's model the shock to the real exchange rate is known by the policymaker to be transitory. This is probably critical to the finding that a modified Taylor rule should incorporate a role for the real exchange rate and that the sign of the response should be positive. By contrast, in the model considered here, there is uncertainty about the nature of a given change in the real exchange rate: this can be either a transitory shock or the inception of a bubble. How does this uncertainty alter the above conclusions about the optimal response to the exchange rate? Below we set out to investigate how policy should respond in the presence of a bubble like the one in **(18)**.

4 Analysis of optimal policy responses

Optimal policy is computed using the algorithm described in Section 2.[36] The solution gives the following optimal decision rule:

$$i_t = f_y(s_{t-1})\, y_t + f_\pi(s_{t-1})\, \pi_t + f_a(s_{t-1})\, a_{t-1} \quad \textbf{(20)}$$

There are two notable aspects. First, the optimal control rule depends on all state variables. Second, the optimal response coefficients generally take different values in different regimes. Therefore, a regime-invariant rule, or a rule in which the interest rate responds only to a subset of the state variables x, is generally suboptimal.[37] Below we will examine the properties of the optimal regime-switching control rule (ORSC) and of a (regime-invariant) Taylor rule with and without the exchange rate. We want to understand, in particular, how optimal policy depends on the probabilities governing the dynamics of the real exchange rate.

4.1 Optimal regime-switching policy

Table B reports the optimal response coefficients in **(20)** for the benchmark parameterisation (Table A). The left-upper part shows the optimal response coefficients when the probability of moving in the bubble regime is zero. This corresponds to the standard optimal control rule computed in the model without regime switching or regime uncertainty. These responses can be compared with those implied by the existence of regime uncertainty. We report the coefficients for different values of the transition probabilities p and q and different values of the autoregressive coefficient $\rho(2)$. The reader is reminded that p is the probability of moving from the bubble regime into the no-bubble regime, while q is the probability of moving into the bubble regime from the no-bubble regime. The parameter $\rho(2)$ characterises the growth rate of the real exchange rate in the absence of any disturbances or changes in policy and hence indicates the 'strength' of the bubble. For each pair of probabilities (q, p), we report the responses in both regimes: the 'bubble' regime (s=1) and the 'no-bubble' regime (s=2). For example, (.25,.5)

(36) To use the algorithm, the model equations **(15)**, **(16)** and **(18)** are cast in matrix form. This requires setting $\beta \equiv \phi$ (discount factor in **(19)**), $x = [y, \pi, a_{-1}]'$, $u = [i]$, and $\varepsilon = [\eta, \varepsilon, \upsilon_{-1}]'$ in **(2)** and **(3)**, and constructing the matrices A, B, Q and R accordingly.

(37) To compute the optimised coefficients of the (regime-invariant) Taylor rule we use the algorithm described in Section 2.3.

indicates that there is a 25% probability that a bubble with an expected duration of two periods can emerge (the expected duration of a bubble being $1/p$). Given the great number of coefficients in the table, graphical illustrations are probably the most effective tools for analysing the relationship between transition probabilities and feedback coefficients. Chart 1 plots the contour lines for the optimal response coefficients against the transition probabilities. The columns refer to different feedback coefficients (in order f_y, f_π, f_a). The upper panel corresponds to the bubble regime while the lower panel corresponds to the no-bubble regime.

There are a number of interesting features that emerge from the inspection of Table B and Chart 1. We begin by examining the response coefficients associated with the lagged real exchange rate separately from those to output and inflation. Its sign and magnitude varies with the transition probabilities and with the observed regime. To understand how the response coefficient varies according to the transition probabilities, it is helpful first to consider what would be the dilemmas faced by the policymaker if she did not know such probabilities. If the economy happened to be in the bubble regime in the previous period, the dilemma for the policymaker today would be between offsetting the effects of the bubble bursting and leaning against the bubble. For example, if the bubble bursts (causing the exchange rate to depreciate), the optimal choice is to raise interest rates (other things equal) to counteract the effects on output and inflation of the depreciation. On the contrary, if the (appreciating) bubble persists, the optimal choice for the policymaker is to lower interest rates (other things equal) to limit the appreciation. This action has two effects: it offsets the negative effects that the (appreciating) exchange rate currently has (other things equal) on output and inflation, and reduces the detrimental impact that a larger bubble might cause when bursting in some future period. If the economy happened to be in the no-bubble regime in the previous period, the dilemma today would be between restraining an incipient bubble and a less active policy. For example, if a bubble arises, the optimal choice for the policymaker is to restrain the movement in the lagged real exchange rate (other things equal) by moving interest rates in the opposite direction. This action, loosely speaking, prevents the bubble from gaining momentum (absent any additive shock) and limits its effects on inflation and output. On the contrary, if the economy remains in the no-bubble regime, the optimal choice for the policymaker is to move interest rates positively with any lagged exchange rate fluctuation. For example, an appreciation of the lagged exchange rate causes a fall in inflation and output prompting the policymaker to lower interest rates. However, in the no-bubble regime such an appreciation is temporary, being reversed at the time policy is chosen. A positive response to the lagged exchange rate, therefore,

Table B: Optimal response coefficients

q	p		y	π	a_{-1}		y	π	a_{-1}		y	π	a_{-1}
0	0	s=2	1.16	1.88	.18								
		ρ (1)=1.1				ρ (1)=1.3				ρ (1)=1.5			
.1	.50	s=1	1.10	1.85	.01	s=1	1.06	1.82	-.04	s=1	1.02	1.79	-.09
		s=2	1.12	1.85	.14	s=2	1.10	1.83	.13	s=2	1.07	1.81	.11
.25	.50	s=1	1.07	1.82	0	s=1	1.03	1.79	-.04	s=1	.97	1.74	-.09
		s=2	1.08	1.82	.08	s=2	1.03	1.79	.06	s=2	.98	1.74	.03
.50	.50	s=1	1.05	1.81	0	s=1	1.00	1.77	-.04	s=1	.93	1.71	-.09
		s=2	1.05	1.81	0	s=2	1.00	1.77	-.04	s=2	.93	1.71	-.09
.75	.50	s=1	1.07	1.83	.01	s=1	1.03	1.80	-.04	s=1	.97	1.75	-.09
		s=2	1.06	1.82	-.08	s=2	1.02	1.78	-.14	s=2	.96	1.73	-.21
1	.50	s=1	1.12	1.88	.02	s=1	1.10	1.87	-.02	s=1	1.08	1.85	-.07
		s=2	1.05	1.80	-.16	s=2	1.01	1.77	-.23	s=2	.97	1.72	-.31
.1	.20	s=1	1.09	1.86	-.08	s=1	1.06	1.85	-.14	s=1	1.02	1.82	-.21
		s=2	1.12	1.85	.14	s=2	1.10	1.84	.13	s=2	1.07	1.82	.12
.25	.20	s=1	1.08	1.85	-.08	s=1	1.04	1.83	-.14	s=1	.99	1.79	-.21
		s=2	1.07	1.83	.09	s=2	1.04	1.80	.07	s=2	.99	1.76	.04
.50	.20	s=1	1.07	1.84	-.09	s=1	1.02	1.81	-.15	s=1	.96	1.76	-.21
		s=2	1.05	1.82	.01	s=2	1.01	1.79	-.03	s=2	.94	1.74	-.07
.75	.20	s=1	1.07	1.85	-.09	s=1	1.03	1.82	-.15	s=1	.98	1.78	-.21
		s=2	1.07	1.85	-.07	s=2	1.03	1.82	-.13	s=2	.98	1.78	-.19
1	.20	s=1	1.09	1.87	-.08	s=1	1.07	1.85	-.14	s=1	1.03	1.83	-.21
		s=2	1.08	1.85	-.15	s=2	1.05	1.83	-.22	s=2	1.01	1.80	-.30
.1	.10	s=1	1.10	1.88	-.11	s=1	1.07	1.87	-.17	s=1	1.03	1.85	-.24
		s=2	1.12	1.85	.14	s=2	1.10	1.84	.13	s=2	1.07	1.82	.12
.25	.10	s=1	1.09	1.87	-.11	s=1	1.06	1.85	-.17	s=1	1.01	1.83	-.24
		s=2	1.07	1.83	.09	s=2	1.04	1.80	.07	s=2	.99	1.77	.05
.50	.10	s=1	1.08	1.86	-.11	s=1	1.04	1.84	-.18	s=1	1.00	1.81	-.25
		s=2	1.05	1.83	.01	s=2	1.01	1.80	-.02	s=2	.95	1.76	-.06
.75	.10	s=1	1.08	1.87	-.11	s=1	1.05	1.85	-.18	s=1	1.00	1.82	-.25
		s=2	1.08	1.86	-.06	s=2	1.04	1.84	-.12	s=2	.99	1.81	-.17
1	.10	s=1	1.09	1.88	-.11	s=1	1.07	1.86	-.17	s=1	1.03	1.85	-.24
		s=2	1.09	1.87	-.14	s=2	1.06	1.86	-.21	s=2	1.03	1.84	-.29

ensures that the overall response of policy is small.

Chart 1 provides a quantitative resolution of the above policy dilemmas by showing how the transition probabilities map into the optimal policy response to the lagged exchange rate. Consider first the response in the bubble regime (top-right corner of the chart). The first thing to notice is that there is a threshold value of the probability p of the bubble bursting which implies no response to the lagged exchange rate. The policymaker chooses not to respond to the lagged exchange rate for a value of the probability of roughly 60%. As the transition probability of a bubble bursting deviates from roughly 60%, the optimal monetary policy becomes more responsive to exchange rate fluctuations. For higher probabilities, the response coefficient is positively correlated with the exchange rate fluctuations, and for lower probabilities it is negatively correlated. Loosely speaking, for higher probabilities the motive for offsetting the effects of the bubble bursting is stronger than the motive for leaning against it and offsetting its effects. The probability q of moving into the bubble regime has little or no impact on the responses. This is due to the fact that the optimal policy is anticipated to adjust to the no-bubble regime after a switch occurs.

Consider now the response coefficients associated with the exchange rate in the no-bubble regime. Again, there is a threshold value for the transition probability q of a bubble emerging such that the policymaker does not respond to the exchange rate fluctuations. As the transition probability of a bubble emerging deviates from roughly 50%, the optimal monetary policy becomes more responsive to the exchange rate fluctuations. For higher probabilities, the response is negatively correlated with the exchange rate fluctuations, and for lower probabilities it is positively correlated. Thus, for lower probabilities, the possibility of any fluctuation being transitory has more weight than the possibility of a bubble emerging. For higher probabilities, the opposite is true.

We now examine the responses to output and inflation. First, the responses to output and inflation do not differ much across regimes. Second, the responses to output and inflation are *for the most part* smaller relative to the case in which there is no bubble uncertainty. In Chart 1 the exception is given by some small regions corresponding to high values of either p and q in which policy is more active. High values of p and q mean that the regime is almost certain to switch. So it is not surprising that more aggressive responses are optimal.[(38)] The differences in the coefficients are however rather small.

Regime switching in a model's parameters – as in the model analysed here – is a particular form of parameter uncertainty. As shown by Brainard (1967), parameter uncertainty causes policy to be more cautious (under certain conditions).[(39)] In other terms, attempting to hit the target immediately is generally not optimal if there is uncertainty about the policy multiplier, because with a quadratic loss function a policymaker cares about not only the deviation of a variable from its target level but also about its volatility. The existence of a trade-off between closing the gap and the volatility induced by policy makes it optimal to do less than under certainty equivalence. Unlike in Brainard's static model, however, the policy multiplier in the present model does not change when a regime change occurs. But the basic insights of Brainard's analysis about the trade-off between closing the gap and volatility carries over to the current model. The point is that here a policy which is appropriate in one regime might not be appropriate in the other regime. Hence fluctuations in inflation and output might be exacerbated by a more active policy if switching between regimes occurs relatively frequently, as is the case when p and q approach their intermediate values.

The probability intervals $0 \leq p \leq 1/2$ and $0 \leq q \leq 1/2$ correspond to a world of persistent regimes, in which bubbles are expected to arise with an interval of at least two years and once started they are expected to last a minimum of two or more years. It is interesting to notice that within these intervals optimal monetary policy has two relevant characteristics. First, the optimal policy in the bubble regime requires responding negatively to fluctuations in the lagged real exchange rate, whereas it requires responding positively, albeit less than in a no-uncertainty world, to the lagged exchange rate in the no-bubble regime. Second, looking at the entire set of feedback coefficients, optimal monetary policy is generally less responsive in both regimes as the transition probabilities p and q move away from zero towards intermediate values (ie $p = q = 1/2$).[(40)]

(38) In the limit case in which $p = q = 1$ the economy will undergo entirely predictable fluctuations in inflation and output. In this case it is obvious that a policy reacting strongly will reduce the volatility of both variables.
(39) In Brainard's analysis a sufficient condition is that the policy instrument is uncorrelated with the additive disturbance. Craine (1979) and Söderström (2002) have shown that in a dynamic context the standard result of Brainard uncertainty can be overturned. In the context of minmax analysis Giannoni (2001, 2002) and Tetlow and von zur Muehlen (2001), among others, have also found that parameter uncertainty may lead to more aggressive policy.
(40) Note, however, that while the intervals $0 \leq p \leq 1/2$ and $0 \leq q \leq 1/2$ are presumably the most relevant for policymaking, one cannot rule out the possibility that $q > 1/2$. This is an environment in which the real exchange rate undergoes sustained deviations from fundamentals most of the time and in which once a crash occurs a new boom might begin soon after. Because of the uncertainty about these probabilities we have decided not to confine our

An important feature of the optimal regime-switching control rule is that the responses in the two regimes become more and more similar as both p and q approach 1/2. In the limit case in which $p = q = 1/2$ the responses are identical across regimes. In other words, the optimal policy converges to a regime-invariant rule,(41) as can be clearly seen from Table B. This suggests that the farther away the probabilities are from their intermediate values, the greater are the losses from adopting a (suboptimal) regime-invariant rule relative to the optimal regime-switching control rule (ORSC). In the special case in which $p = q = 1/2$ the losses are indeed identical as the optimal regime-invariant rule is identical to the ORSC rule.(42) It follows that the regime-switching nature of the policy is crucial when one or both regimes are persistent. We will return to this below.

Does the degree of flexibility in inflation targeting matter? Charts 2 and 3 show the optimal monetary policy reaction for $\lambda = 0$ and $\lambda = 5$, respectively. It can be seen that a stronger preference for output stabilisation (relative to inflation stabilisation) raises the threshold value of the probability p that determines whether, in the bubble regime, the optimal policy response is positively or negatively correlated with a given exchange rate fluctuation. A stronger preference for output stabilisation also lowers the threshold value of the probability q that determines whether, in the no-bubble regime, the optimal policy response is positively or negatively correlated with a given exchange rate fluctuation.(43) In a small open economy, the quickest channel from policy to inflation is the exchange rate. Stabilising inflation requires larger fluctuations in interest rates and the exchange rate, and ultimately in output compared with a closed economy (in a closed economy a one-off change in output is sufficient). Note that for this reason all response coefficients become smaller as λ increases in the absence of the bubble regime (ie $q = 0$); the same holds true for any given positive value of q. Similarly, in the bubble regime for a given probability p of the bubble bursting the response coefficients become smaller.

Before moving to the analysis of the Taylor rules in the next section, we examine the sensitivity of

---

analysis above to subintervals of the probabilities.

(41) As noted above, this is also the point at which optimal policy is generally less active.

(42) This result resembles that obtained by Bullard and Schaling (2001) in a closed economy subject to random productivity shifts. In their model the optimal linear rule is a Taylor rule augmented with terms whose role is to offset the potential impact of future supply shocks. When the transition probabilities governing the productivity shifts are both equal to 1/2, the optimal rule collapses to the one which is optimal in the absence of productivity shifts. By contrast, in our model, when $p = q = 1/2$ uncertainty continues to matter, so that the rule is different from the one that is optimal when eg $q = 0$.

(43) Equivalently, for lower probabilities, a stronger preference for output stabilisation makes policy more responsive for any given p in the bubble regime. Likewise, for lower probabilities, a stronger preference for output stabilisation makes policy less responsive for any given q in the no-bubble regime.

the above results to changes in the benchmark parameter. Table B and Chart 1 report the optimal response coefficients for the benchmark economy in which we assume that the autoregressive coefficient of the bubble $\rho(2) = 1.3$. This means that any initial fluctuation in the real exchange rate will grow on average at 30% a year absent any change in the real interest rate. What happens if this growth rate is different? Table B shows that as $\rho(2)$ increases the responses to output and inflation, for the most part, become smaller relative to an economy in which bubble uncertainty is absent (this can be seen for any given pair (p, q) by moving along the rows from left to right). Furthermore, as $\rho(2)$ increases the response to lagged exchange rate gets smaller, and if negative, policy becomes more reactive to the lagged exchange rate. Intuitively, a stronger bubble growth makes the motive for leaning against the bubble even stronger for any given pair of the transition probabilities.[(44)]

One important caveat of the above analysis is the assumption that the regime is observable by the policymaker, albeit with a delay of one year. *Ex-post* identification of a regime is not implausible in several situations, but it can arguably be considered a limitation more generally. If we think that the policymaker cannot observe the regime, even with a delay, the analysis of regime-invariant Taylor rules in the next section acquires greater importance. The optimal regime-switching control rule remains helpful in two ways. First, it helps develop intuition about the appropriate policy response and offers a useful benchmark for assessing other types of rules. Second, it is a starting point for building more sophisticated models of the policymaker's choice under regime uncertainty.

4.2 Optimised Taylor rules

Here we analyse optimal monetary policy under the assumption that policy does not change as the economy switches between regimes. The optimal feedback coefficients are selected by minimising the loss function conditional on the economy being in the no-bubble regime.

(44) We also checked the sensitivity of the results to changes in other relevant parameters. To economise on space we briefly summarise the findings (details are available from the author on request). A reduction in χ (the elasticity of demand with respect to the real exchange rate) lowers the threshold value for p and raises the threshold value for q. Policy is also generally more responsive to output and inflation fluctuations. A reduction in the degree f of exchange rate pass-through raises the threshold for p and lowers the threshold for q, with similar effects to an increase in the preference for output stabilisation λ. It also makes policy more responsive to output and inflation. A fall in θ (the sensitivity of the real exchange rate to the real interest rate) leaves the threshold values for p and q roughly unchanged but makes policy more responsive as the probability p and q deviate from their respective thresholds. The responsiveness of policy to output and inflation increases. Finally, including the nominal interest rate in the policymaker's loss function (eg $\varsigma = .5$) causes policy to be less responsive to output and inflation as one would expect, but does not significantly alter its responsiveness to the exchange rate.

Optimised response coefficients are plotted in Chart 4. On the upper panel is a Taylor rule which features the lagged real exchange rate,[45] while on the lower panel is the more familiar Taylor rule in which no response to the lagged exchange rate is allowed.

We first comment on the Taylor rule with the exchange rate. The response to the lagged exchange rate displays an intuitive link with the probability of a bubble's onset and with its expected duration (given by $1/p$). Unlike the optimal regime-switching control rule, both transition probabilities have now a significant impact on the optimal response. This means, for instance, that the threshold value that 'divides' positive from negative responses to the lagged exchange rate cuts the probability space obliquely. However, the behaviour of optimal policy in response to the lagged exchange rate has the same intuitive explanations discussed in the previous section.[46] There a couple of features which are worth noticing. First, the response coefficients are quite sensitive to small changes in q close to the origin. This is interesting because it means that even a small probability of a costly event (ie the high-duration bubble) can have significant implications for policy.[47] Second, the responses to output and inflation are generally more muted than in the absence of regime uncertainty. In particular, policy is less responsive to output and especially to inflation in the bottom-left corner, which corresponds roughly to low-probability high-duration bubbles.[48] The pronounced weakening of the optimal responses is presumably related to the fact that a regime-invariant policy is much more inefficient for low-probability high-duration bubbles than for bubbles characterised by higher values of q and p. Low-probability high-duration bubbles requires the optimal unconstrained policy to respond very differently to the lagged exchange rate in the bubble and no-bubble regimes (recall the discussion in the previous section).[49] For other values of the transition probabilities, this problem does not arise in that the regime-invariant rule is much more similar to the ORSC rule. As noticed in the previous section,

(45) Note that in the absence of regime switching the optimal interest rate rule takes the form of a time-invariant Taylor rule augmented with the exchange rate.

(46) For a given probability p, policy becomes less responsive to positive exchange rate fluctuations as q rises. It then becomes more responsive to negative fluctuations as q rises above the threshold, reflecting the greater importance the policymaker gives to either leaning against a bubble (bubble regime) or to offset any movements that can cause a bubble to gain momentum (no-bubble regime). The time-invariant response must be, so to speak, a 'good' compromise between what would have been the appropriate responses in each regime, were policy allowed to switch. Similar considerations hold for the effects of the probability p.

(47) Note, however, that there is also a tiny region of the probability space corresponding to low values of p and q where the feedback coefficient associated with inflation is a bit larger than the one prevailing in the absence of regime uncertainy.

(48) The response coefficient associated with inflation displays a trough also in the upper-right corner, in correspondence of high values of both transition probabilities.

(49) Because of the low probability of a bubble's onset, the optimal response should be positive in the no-bubble regime; but due to the high expected duration, it should turn negative in the bubble regime.

the latter becomes ever closer to the optimal regime-invariant rule as the probabilities approach $1/2$. When the probabilities are close to $1/2$, there is greater expected volatility caused by the more frequent switches of regimes. In this case policy becomes less responsive to all variables than in the absence of regime uncertainty.

How does a simple Taylor rule without the exchange rate compare? The contrast is quite striking. Monetary policy's responsiveness increases with the expected duration of the bubble as well as with the probability of it arising. Presumably, policy needs to be more active for a given p as q rises in order to compensate for the lack of a negatively correlated response to exchange rate fluctuations.[50]

How inefficient is it to adopt a Taylor rule in our regime-switching small open economy? Charts 5a-5b show the percentage difference between the losses implied by the Taylor rules with and without the exchange rate and the losses implied by the ORSC rule. The common feature displayed by the two rules is that the relative losses are mainly concentrated where the transition probabilities are both low or high. As stressed above, this is where a regime-switching rule is most useful as the optimal policy responses in the two regimes differ the most. By contrast, intermediate values of the transition probabilities are where the losses are the smallest. The Taylor rule with the exchange rate is identical to the ORSC rule when $p = q = 1/2$. This is obviously not true in general for the Taylor rule without the exchange rate, although the difference in the losses are tiny for intermediate values of the probabilities. Another common feature, perhaps surprising, is that in the corners the simple Taylor rule entails loss differences which are only slightly larger than the regime-invariant rule. For a large duration bubble, however, the Taylor rule without exchange rate fares less well as the probability of such a bubble arising increases.[51]

---

(50) Note that the coefficients of the simple Taylor rule in the absence of regime uncertainty (ie $q = 0$) are smaller. They are clearly of similar magnitude in correspondence of the loci of p and q where the response to exchange rate is zero.

(51) We have also computed the optimised coefficients of the Taylor rules for $\lambda = 0$ and $\lambda = 5$ (these are not reported here but are available on request). Under strict inflation targeting monetary policy is much more responsive to inflation fluctuations, as one would expect, and to a minor extent, is also more responsive to output fluctuations relative to the benchmark case. The qualitative pattern is, however, unchanged for the most part. The losses implied by the Taylor rules relative to the ORSC rule are now much larger than in the benchmark case, reflecting the fact that the responses to the exchange rate are much bigger. For $\lambda = 5$, policy with both Taylor rules is generally less responsive to output and inflation fluctuations relative to the benchmark case, though showing a similar pattern with respect to changes in the transition probabilities. Moreover, in the Taylor rule that includes the exchange rate the area of the probabilities over which it is optimal to respond negatively to exchange rate movements is now much larger, which is unsurprising given the findings about the ORSC rule. A simple Taylor rule which does not include the exchange rate does much worse than the unrestricted regime-invariant rule for most values of the probabilities, reflecting the crucial importance that responding to exchange rate fluctuations have in stabilising output. The relative

Given the crucial importance of the transition probabilities, the analysis in the next section takes into full account the consequences of assuming the incorrect probabilities and the beliefs the policymakers might hold on them.

5 Optimal assumptions about uncertain transition probabilities

Historical experience might provide little or no guidance as to how to calibrate the transition probabilities p and q, thus effectively putting the policymaker in a world of Knightian uncertainty. The policymaker is more likely to specify intervals for p and q, in which case the tools outlined in the previous sections cannot directly be applied to achieve a policy decision. In this section we will examine the losses incurred by the policymaker in making the incorrect assumptions about the transition probabilities as well as the selection of the best assumptions using the most common criteria for choice under uncertainty: the minmax and Bayesian averaging.

To start with, it is useful to define the notation that we will use. Let us indicate the policymaker's assumptions about the true probabilities (p, q) with $(\hat{p}, \hat{q})$. Note that making an assumption about (p, q) is equivalent to selecting the policy that minimises the loss function **(19)** under that assumption.[52] The performance of this policy, or equivalently the validity of the assumption, can be assessed across different realisations of the true but unknown probabilities. Each pair (p, q) indeed corresponds to a different model. So, in what follows it is easier to think of a pair $(\hat{p}, \hat{q})$ as a 'policy' and of a pair (p, q) as a 'model'. To evaluate how a given policy $(\hat{p}, \hat{q})$ fares in a given model (p, q) we compute the loss using the algorithms outlined in Section 2.3. Let us indicate with $L\left(p, q, \hat{p}, \hat{q}\right)$ the value of the loss function **(19)** under policy $(\hat{p}, \hat{q})$ and model (p, q).[53] Using this function we can select the optimal policy (or optimal probability assumptions) under uncertainty using two different criteria. The first is the *minmax*, which amounts to choosing the policy (or assumption) which delivers the best performance among all possible worst-case scenarios. In other words, for each policy there is a worst-case scenario represented by a pair (p, q) which maximises the loss. Policy is then selected by solving the following problem:

$$\min_{(\hat{p},\hat{q})} \max_{(p,q)} L\left(p, q, \hat{p}, \hat{q}\right) \quad \textbf{(21)}$$

loss differences are, however, of similar size compared with those in the benchmark case.

(52) We will consider optimal unconstrained policies or constrained policies such as a time-invariant Taylor rule. By optimal policy here we mean the optimal policy based on the probability assumption $(\hat{p}, \hat{q})$. It is not in general the optimal policy under uncertain probabilities.

(53) In the computations this is a four-dimensional matrix.

This criterion for choice under uncertainty may be appropriate when the policymaker has no priors on p and q other than an interval.(54) The second criterion which we consider is the minimisation of the *Bayesian aggregate loss*, which takes into account the policymaker's priors over p and q.(55) Formally, the policymaker's problem is:

$$\min_{(\hat{p},\hat{q})} L_B(\hat{p},\hat{q}) \tag{22}$$

where

$$L_B(\hat{p},\hat{q}) = \int_0^1 \int_0^1 L(p,q,\hat{p},\hat{q})\, f(q)\, g(p)\, dp dq$$

$f(\cdot)$ and $g(\cdot)$ are distinct density functions which represent the policymaker's beliefs over q and p respectively (assumed to be independent from each other). Without loss of generality we will characterise such beliefs with two Beta distributions, ie $p \sim B(a_p, b_p)$, $q \sim B(a_q, b_q)$. The Beta distribution is continuous over the interval [0, 1] and encompasses a very ample set of beliefs (as will be seen below). To compute the Bayesian loss we employ Gaussian quadrature methods.(56)

Before embarking on the analysis of the optimal policy under uncertainty, a useful exercise consists in assessing the asymmetry of the losses implied by assuming the incorrect probabilities. This will help the interpretation of the results of the optimisation exercises spelled out above. We do this in the simplest possible way: we keep one of the probabilities fixed (which we also assume to be true), and look at how 'bad' policy does when the value assumed for the other probability deviates from the true one.(57) This exercise is similar in spirit to the analysis of models' 'fault tolerance' in Levin and Williams (2003) and Jääskelä (2005). Charts 6-11 report the outcome of this exercise for different types of policy rules and different assumptions about which probability is kept fixed. All losses are computed using the benchmark values in Table A and are conditioned on the economy being in the no-bubble regime.

Let us start with examining Chart 6, which shows the losses associated with the ORSC rule. Each plot in this chart refers to a bubble of different expected duration $(1/\bar{p})$. All losses are normalised

(54) Gilboa and Schmeidler (1989) provide the axiomatic foundation for the minmax as a criterion for choice under uncertainty (dating back at least to the statistician Wald (1945)). Applications of the minmax criterion in economics include: von zur Muehlen (1982), Rustem (1994), Hansen and Sargent (2003), Giannoni (2001, 2002), Onatski and Stock (2002), Tetlow and von zur Muehlen (2001, 2004).

(55) Recent examples of Bayesian aggregation of the losses implied by competing reference models are Levin and Williams (2003) and Levin, Wieland and Williams (2003).

(56) See eg Judd (1998) and Miranda and Fackler (2002).

(57) More precisely, when we fix $p = \hat{p} = \bar{p}$, L collapses to a two-dimensional matrix $L_{\bar{p}}(q,\hat{q}) \equiv L(\bar{p},q,\bar{p},\hat{q})$. Likewise, when we fix $q = \hat{q} = \bar{q}$, L collapses to $L_{\bar{q}}(p,\hat{p}) \equiv L(p,\bar{q},\hat{p},\bar{q})$.

so that the loss corresponding to $\hat{q} = q = 0$ in each plot equals one. Each line in a plot corresponds to a true value of the probability q, plotted against $\hat{q}$, the value assumed by the policymaker. Clearly, each line has a minimum at $\hat{q} = q$. Three features are immediately apparent from this chart.[58] First, there are values of $\hat{q}$ such that the losses are insensitive or robust to realisations of the true value of q.[59] These are also the minmax values of q (given $\bar{p}$), as it is easy to see. It is interesting to note that the 'robust' values of $\hat{q}$ correspond to less activist policies in the no-bubble regime. Indeed, Chart 1 shows that the optimal responses to the real exchange rate are close to zero. This is not so in the bubble regime. Optimal policy corresponding to the selected $\hat{q}$ displays different degrees of responsiveness to the exchange rate. However, this is to expect as we are fixing and assuming to know the true values of p with certainty, which is the mostly relevant transition probability in the bubble regime. Besides, policy is less responsive to output and inflation fluctuations in both regimes. A second feature of Chart 6 is that the losses are asymmetric: they tend to get larger as $\hat{q}$ rises above the cross point. To put it differently, overestimating the probability that the economy moves into a bubble can lead to worse outcomes than underestimating such probability. Note that a missing value on the chart indicates an infinite loss.[60] Hence, assuming that the economy moves into a bubble regime with a high probability $\hat{q}$ when in fact there is no such regime (ie $q = 0$) can lead to a catastrophic outcome. Finally, as the expected duration of the bubble increases the losses flatten (especially to the left-hand side) although they remain asymmetric.

Chart 7 shows the outcome of a similar experiment as in Chart 6 for the ORSC rule. The difference now is that the probability being fixed and assumed to be known is q, rather than p. Hence, each plot in this chart refers to a bubble which can arise with different probabilities q (or, equivalently, at intervals of expected length 2, 4 or 10 periods, respectively). Chart 7 displays the same qualitative features, *mutatis mutandis*, of Chart 6. There is a robust value of the assumed probability $\hat{p}$ which makes the loss insensitive to the realisations of the true values of p.[61] As shown in Chart 1, for robust values of p, policy is not very responsive in both regimes, unless q is known to be low: eg for $q = .10$ policy is more responsive, although it is negatively correlated in

(58) Note that for readability the charts are produced using a coarse grid of probability values and that the graphical tools automatically interpolate the values inbetween. Using a finer grid would produce similar graphs. In particular, the cross points obtained using a coarse grid are roughly identical to the ones obtained using a finer grid.
(59) For bubble regimes expected to last 2 and 3 periods, this value is 50%; for bubble regimes expected to last 5 periods it is slightly less than 50%; and for bubbles expected to last for 10 periods it is roughly 37.5%.
(60) In this case the algorithm that solves for the value function for a given policy produces an infinite norm.
(61) This value is 50% for bubbles arising with a 50% and 25% probability, and falls just below 40% for bubbles of 10% probability.

the bubble regime and positively correlated in the no-bubble regime. Furthermore, the losses show similar asymmetry and flattening as in Chart 6.

Charts 8 and 9 report the results of similar experiments as reported in Charts 6 and 7 but for a Taylor rule which includes the exchange rate. Again, the losses are asymmetric with very large values generally located to the right of the loci where the lines cross. In Chart 8 the robust value of $\hat{q}$ becomes smaller as the expected duration of the bubble regimes increases. Likewise, in Chart 9 the robust value of $\hat{p}$ tends to become smaller as the expected interval at which bubbles arise lengthens. The interesting thing to notice here is that the robust values of $\hat{q}$ (together with the fixed values of p) and $\hat{p}$ (together with the fixed values q) correspond to policies which are very similar to each other. These policies involve a very small response to the real exchange rate (see upper panel of Chart 4).[62] They are also among the least responsive to output and inflation fluctuations.

Charts 10 and 11 finally report the results of the same type of experiments as above for a Taylor rule which does not include a response to the exchange rate. The robust values of $\hat{q}$ (given p) and $\hat{p}$ (given q) decrease as q and p, respectively, fall. The optimal response coefficients corresponding to such robust values of $\hat{q}$ and $\hat{p}$ are, however, very similar to each other. In this case policy is more responsive in general than in the case in which there is no regime uncertainty (ie $q = 0$), but less responsive than under higher values of q or lower values of p.

So far we have examined the consequences of making incorrect assumptions about the transition probabilities by looking at one probability at a time. To complete the analysis we need to look jointly at both probabilities, assuming that we know none of them. The top of Table C reports the minmax values of both probabilities (the solution to problem **(21)** above) for different policy rules using the benchmark parameterisation of Table A.[63] For the ORSC rule it is optimal to assume that p and q are 50%. It is easy to see from Chart 1 that these probabilities correspond to a policy which is the least responsive to output and inflation fluctuations and respond very little to the exchange rate in both the bubble and no-bubble regimes. Considering the findings in Chart 6 and 7 this result is not surprising. For the Taylor rule which includes the exchange rate (Taylor 1 in Table C), it is optimal to assume that p and q are both 70%. This result might seem a bit odd at a first glance. However, a careful examination of the worst-case losses associated with all possible

(62) These policies lie very close to the oblique contour line corresponding to -0.04 in Chart 4.

(63) We compute the minmax values $(\hat{p}, \hat{q})$ by performing a search over a grid of step 0.1. A finer grid would give more accurate results at the cost of much larger computational time but would not give any new insights.

policies (ie the matrix of losses $L_{wc}\left(\hat{p},\hat{q}\right) \equiv \max_{(p,q)} L\left(p,q,\hat{p},\hat{q}\right)$, not reported here) reveals that for policies along the main diagonal $\hat{p}=\hat{q}$ the differences are small,[(64)] and indeed the feedback coefficients of the corresponding optimal policies are very similar. This can be seen clearly in Chart 4. In other words, policies based on the assumption that $p=q$ are generally quite robust for all values bar those on or close to the boundaries. Moreover, such policies are generally more conservative than for other values of the probabilities and relative to certainty equivalence. Finally, for the Taylor rule which does not include the exchange rate (Taylor rule 2 in Table C) it is optimal to assume that p is 20% and q is 30%. Again, when the assumed values of p and q are similar (as when moving obliquely in the matrix of losses L_{wc}) policies are also very similar as can be seen in the lower panel of Chart 4.

The above results about the minmax choice of probability assumptions can be compared to those normally found in the robust control literature. As stressed by Svensson (2000), a preference for robustness often leads to the choice of a parameter which is on the boundary of the feasible set of models. By contrast, a desire for robustness in the present context, interestingly, leads the policymaker to pick an intermediate value of the probability. Furthermore, in the literature a robust policy is normally found to be more activist. By contrast, here both the regime-switching and the regime-invariant policies based on their respective minmax probability assumptions are found to be less responsive than in the case of no uncertainty. The exception is given by the simple Taylor rule which does not include a response to the exchange rate. Minmax assumptions about the probabilities correspond to policies which are in general more activist than in the absence of uncertainty.[(65)]

What are the optimal assumptions about the transition probabilities if the policymaker does not rely on the minmax criterion? Table C shows the solution to problem **(22)** under different sets of relevant priors. We consider first the case in which the policymaker has uniform priors over both p

(64) Excluding the extreme cases $\hat{p}=\hat{q}=0$ and $\hat{p}=\hat{q}=1$, the differences in the losses along the main diagonal of L_{wc} range from 0 to 5.5% relative to $\hat{p}=\hat{q}=.70$. The difference of $\hat{p}=\hat{q}=.50$ is only 1%. The matrix L_{wc} is not reported here for reasons of space but is available from the author.
(65) The preference for output stabilisation (versus inflation stabilisation) does not change the results for minmax $\hat{p}$ and $\hat{q}$ in the case of the ORSC rule (for $\lambda=0$ and $\lambda=5$ they are both $\hat{p}=\hat{q}=.5$). Note for the case of strict inflation targeting (not reported here) there is a greater number of policies $\left(\hat{p},\hat{q}\right)$ associated with an infinite loss, normally concentrated near the boundary values of the probabilities. The weight on output stabilisation somewhat affects the results for the optimised linear rules. For the unrestricted regime-invariant rule, we obtain $\hat{p}=\hat{q}=.4$ for $\lambda=0$, and $\hat{p}=\hat{q}=1$ for $\lambda=5$. Again, it is important to notice that the feedback coefficients of different policy rules for which $p=q$ are extremely similar to each other, and therefore entail tiny loss differences. For the simple Taylor rule the results are $\hat{p}=\hat{q}=.1$ for $\lambda=0$, and $\hat{p}=.4$, $\hat{q}=.7$. The general conclusions about the degree of responsiveness remain valid.

Table C: Optimal assumptions about transition probabilities

	ORSC		Taylor 1		Taylor 2	
	$\hat{p}$	$\hat{q}$	$\hat{p}$	$\hat{q}$	$\hat{p}$	$\hat{q}$
Minmax	0.5	0.5	0.7	0.7	0.2	0.3
Priors:						
$p : a = b = 1$ $q : a = b = 1$	0.43	0.45	0.42	0.42	0.44	0.44
$p : a = 2, b = 8$ $q : a = 2, b = 8$	0.19	0.19	0.17	0.16	0.18	0.17
$p : a = b = 1$ $q : a = 2, b = 8$	0.42	0.21	0.34	0.17	0.36	0.18

Note: ORSC is the optimal regime-switching control rule.
Taylor 1 is a Taylor rule which includes the exchange rate.
Taylor 2 is a Taylor rule which does not. Minmax probabilities were obtained performing a grid search over [0 : .1 : 1]. Bayesian probabilities were obtained by combining a grid search and optimisation by the Nelder-Mead algorithm. Policymaker's beliefs over p and q: $p \sim Beta\,(a, b), q \sim Beta\,(a, b)$.
All parameters set to their benchmark values unless stated otherwise.

and q. For $a = b = 1$ the density function of the Beta distribution corresponds to that of the uniform distribution, thus putting the same weight on all values between 0 and 1 (see Chart 12). Both $\hat{p}$ and $\hat{q}$ turn out to be approximately equal to 40% across all types of policy rules. This result is perhaps not surprising given the asymmetry of the losses shown in Charts 6-11. Next we consider the case in which the policymaker has a 'strong' belief that bubbles arise with a small probability q but have high expected duration (low p). In particular, we set $a = 2$ and $b = 8$ which implies an average value of p and q of 20% and a mode of around 13% (see Chart 12). The optimal assumptions $\hat{p}$ and $\hat{q}$ turn out to be quite close to the the priors' mean values.(66) Finally, we consider the case in which the policymaker has a strong belief that the bubble regime emerges with a low probability but no strong prior about its expected duration. The optimal assumptions about p and q turn out to be very close to the values computed under previously considered priors. The only notable difference is perhaps the lower values of $\hat{p}$ in the Taylor rules. This aspect is presumably related to the fact – which we have previously noticed examining Charts 8-11 – that the variability of the losses is generally very small when p and q are not too distant from each other.

6 Conclusion

The current paper has discussed a method for analysing how policy should respond when the exchange rate experiences sustained booms followed by occasional corrections towards fundamentals. It would be interesting to extend the analysis to models with richer dynamics than the one used in this paper. Furthermore, the algorithm used to compute the optimal regime-switching control rule has some limitations. One limitation is that it applies to backward-looking models. Extending it to models that allow for forward-looking agents involves non-trivial technical issues and is clearly desirable.(67) Another limitation is the assumption that the policymaker is able – unlike in the Bayesian model averaging approach(68) – to identify the regime, albeit with a delay. So the policymaker's uncertainty is always about how the model will evolve in subsequent periods. This is a plausible assumption in some circumstances (eg an asset price collapse can be observed) but not in all (eg permanent improvement in productivity is not

(66) Varying the preference for output stabilisation (versus inflation stabilisation) does not alter the conclusions of the analysis. The differences in the results are tiny. Details are available from the author.
(67) This is the object of current research by the author.
(68) For recent applications of Bayesian model averaging see Brock et al. (2003), Cogley and Sargent (2003), and Milani (2003). Cogley and Sargent (2003) considers a policymaker that is uncertain about a number of models equally valid from an empirical point of view. Such a policymaker continuously updates the probabilities of these models in a Bayesian fashion, but such uncertainty does not enter her optimisation problem.

visible except with several years' delay). For this reason we also looked at the performance of regime-invariant rules whose implementation does not require the policymaker to classify the regime. A further possible extension is to make the state of the world a latent variable as well as to require the policymaker to learn about the transition probabilities. A policymaker who does not know the regime could learn about it and estimate the transition probability matrix of the Markov chain by application of the Bayes rule. Optimal policy would then be computed taking into account the uncertainty about the current regime as well as the uncertainty of a future switch to a different regime. The type of learning we envisage is passive because in each period the policymaker would assume that the transition probabilities would not change in all subsequent periods, while in fact they will be changing.

Charts

Chart 1: Optimal regime-switching control rule: benchmark parameterisation

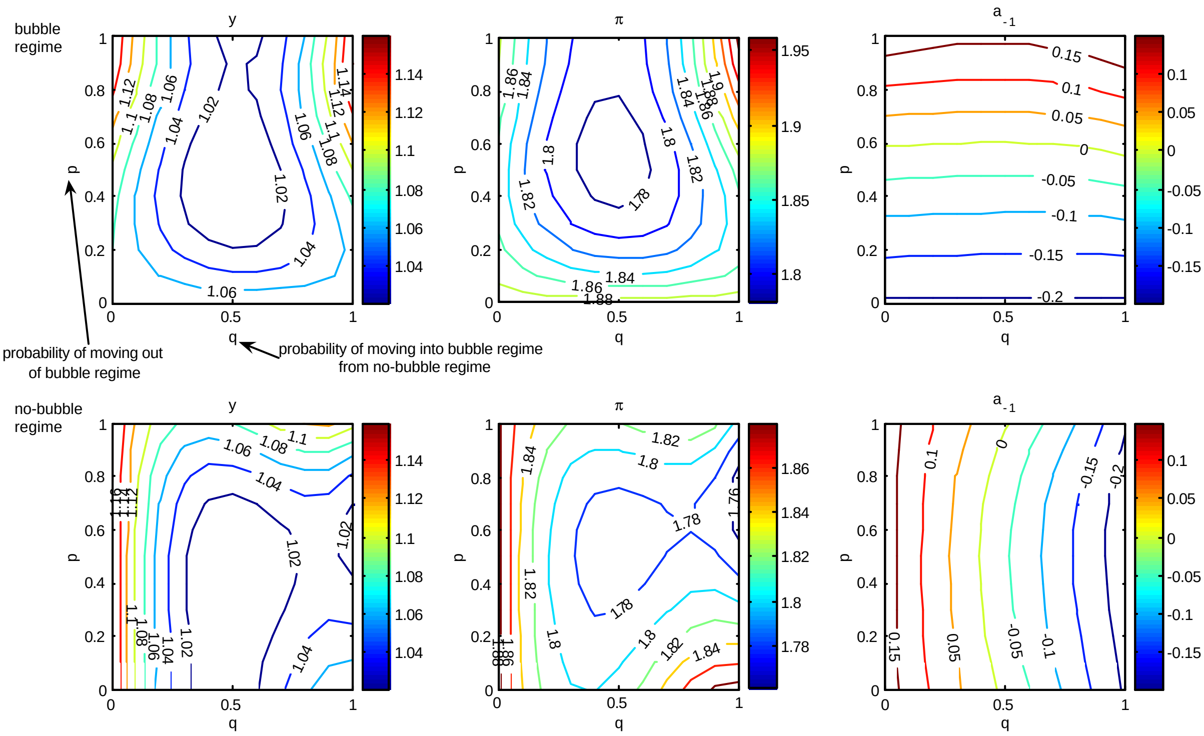

Chart 2: Optimal regime-switching control rule: strict inflation targeting $\lambda = 0$
(other parameters set to their benchmark values)

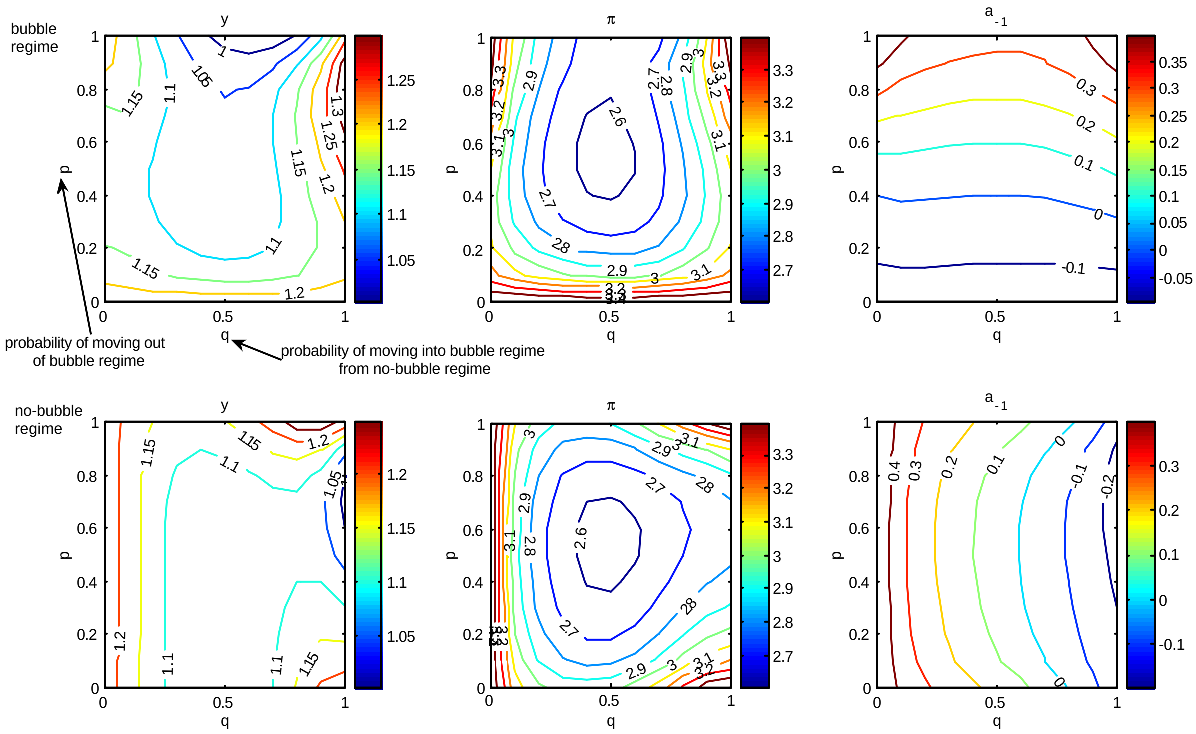

Chart 3: Optimal regime-switching control rule: $\lambda = 5$
(other parameters set to their benchmark values)

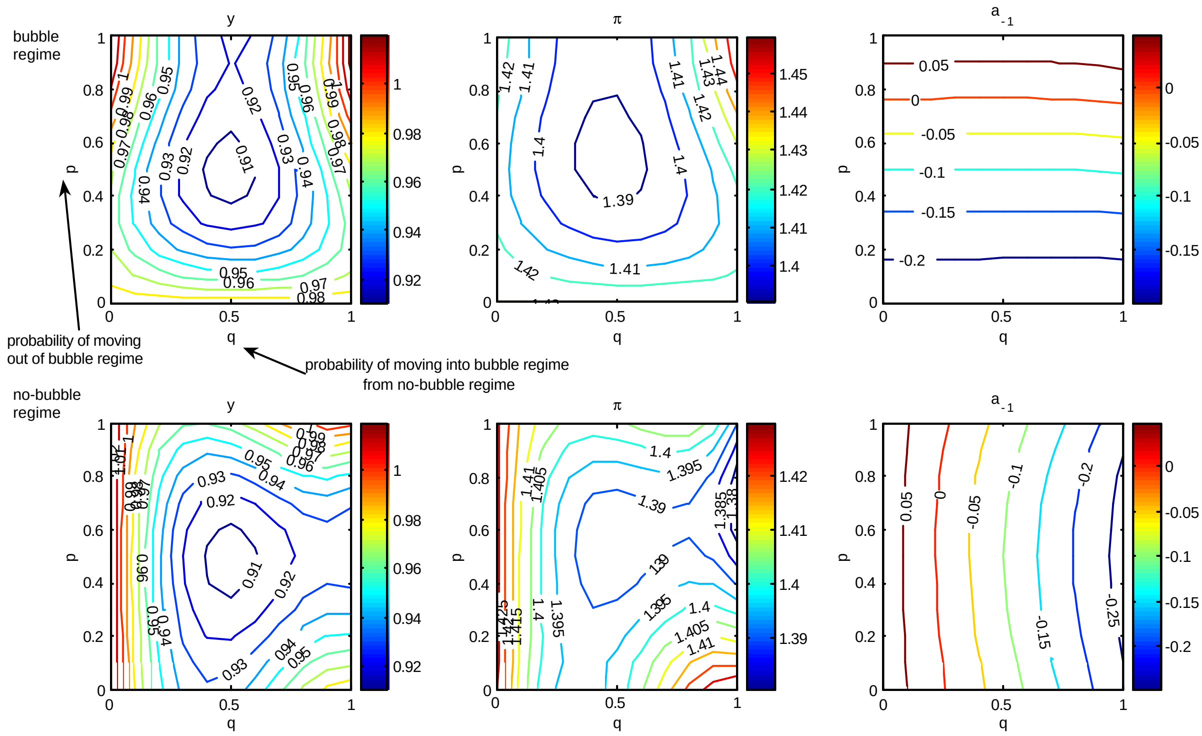

Chart 4: Optimised Taylor rules: benchmark parameterisation

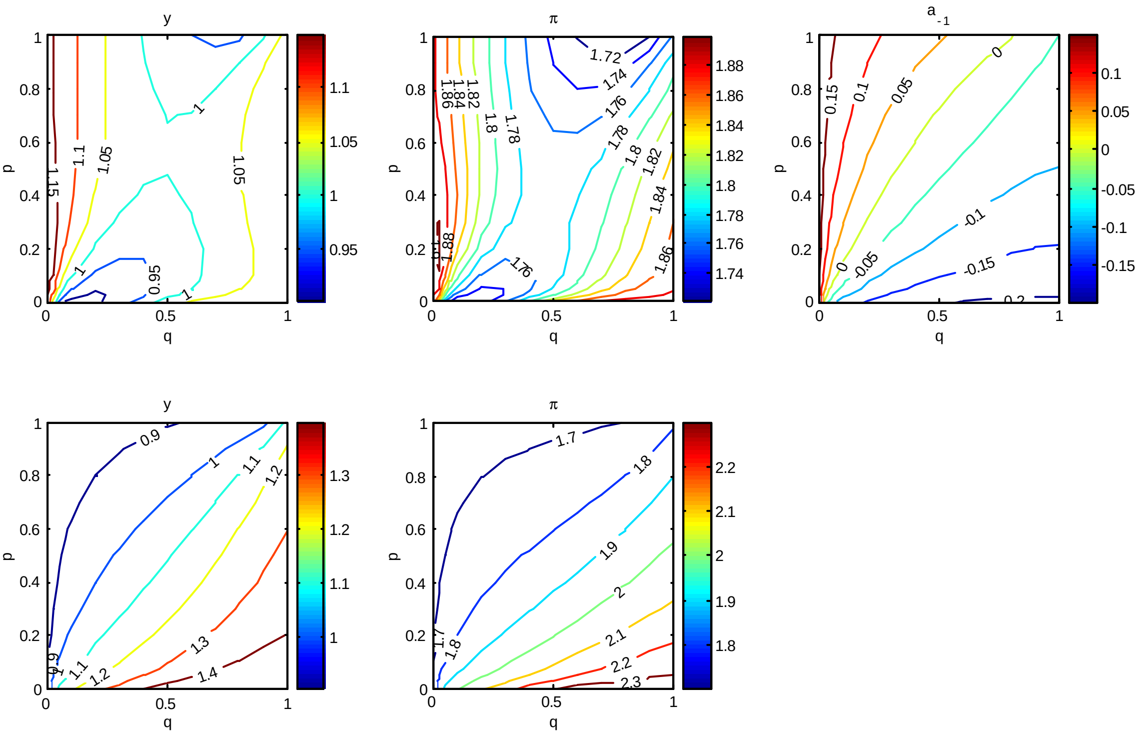

Chart 5: Taylor rules vs optimal regime-switching control rule (ORSC)

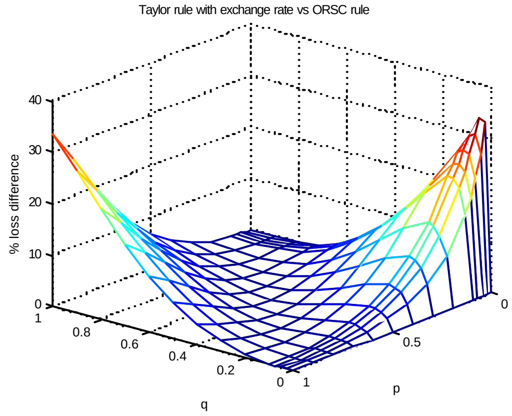

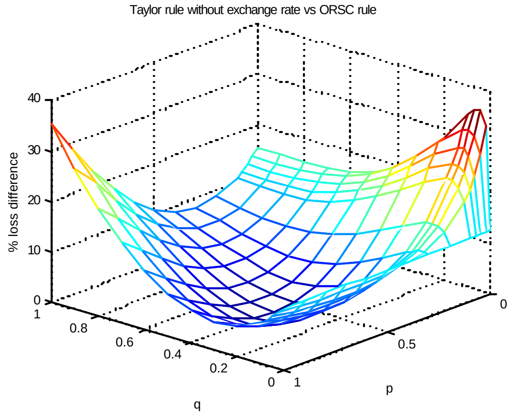

Chart 6

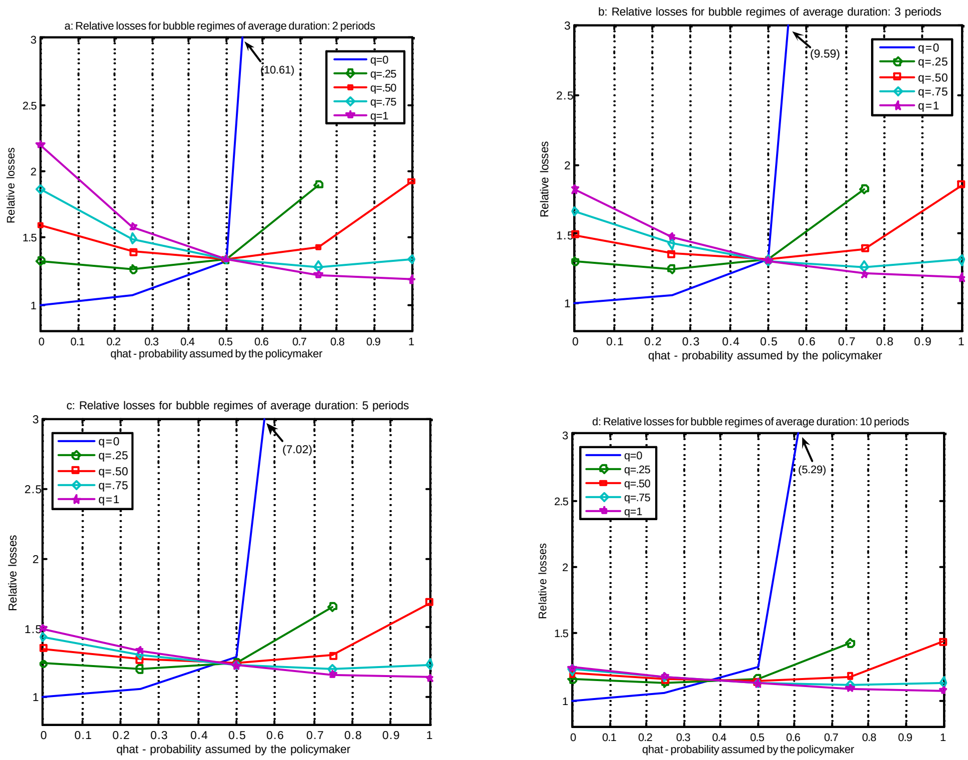

Relative losses of ORSC rule as a function of $\hat{q}$ and q, given p. Benchmark parameterisation.

Note: Missing values indicate 1. Losses conditioned on no-bubble regime.

Chart 7

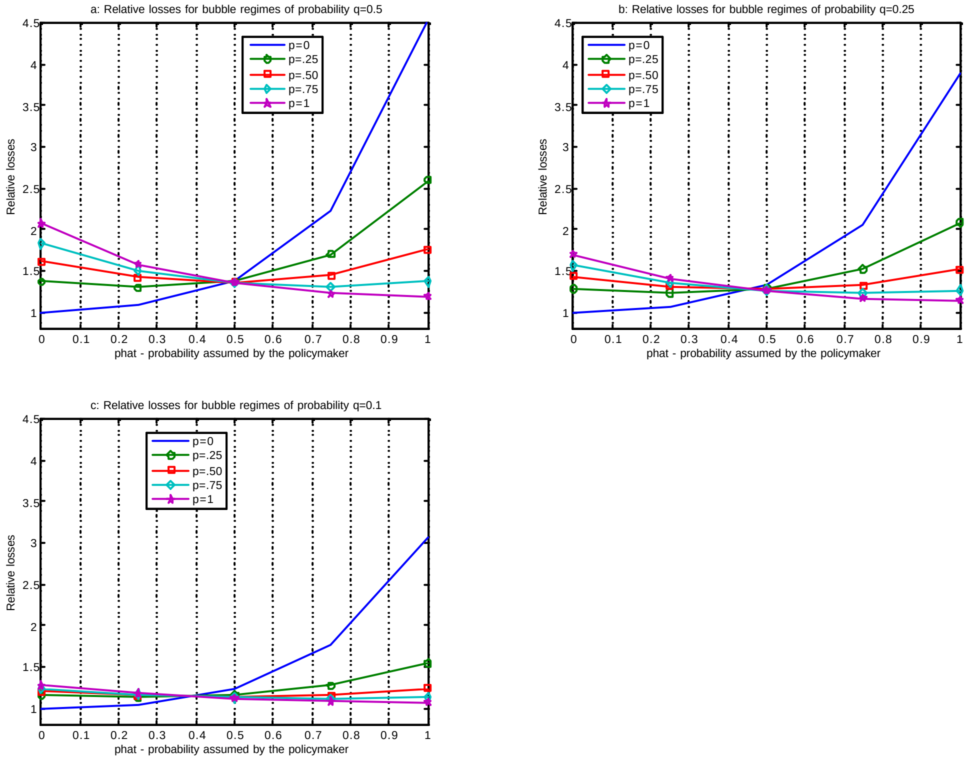

Relative losses of ORSC rule as a function of $\hat{p}$ and p, given q. Benchmark parameterisation.

Note: Losses conditioned on no-bubble regime.

Chart 8

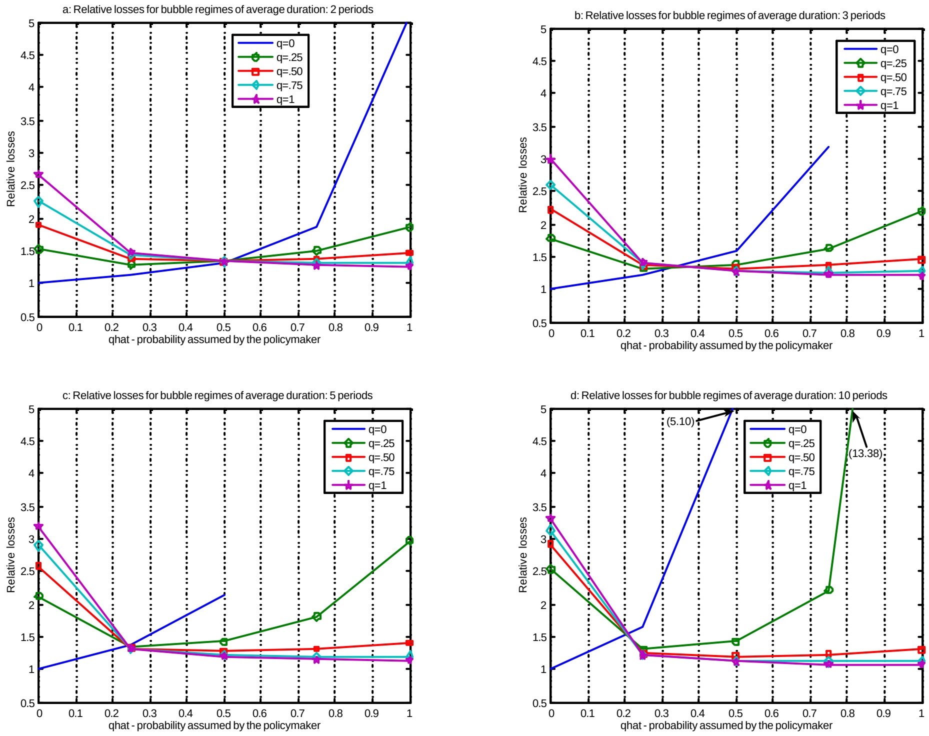

Relative losses of Taylor rules with exchange rate as a function of $\hat{q}$ and q, given p. Benchmark parameterisation. Note: Missing values indicate 1. Losses conditioned on no-bubble regime.

Chart 9

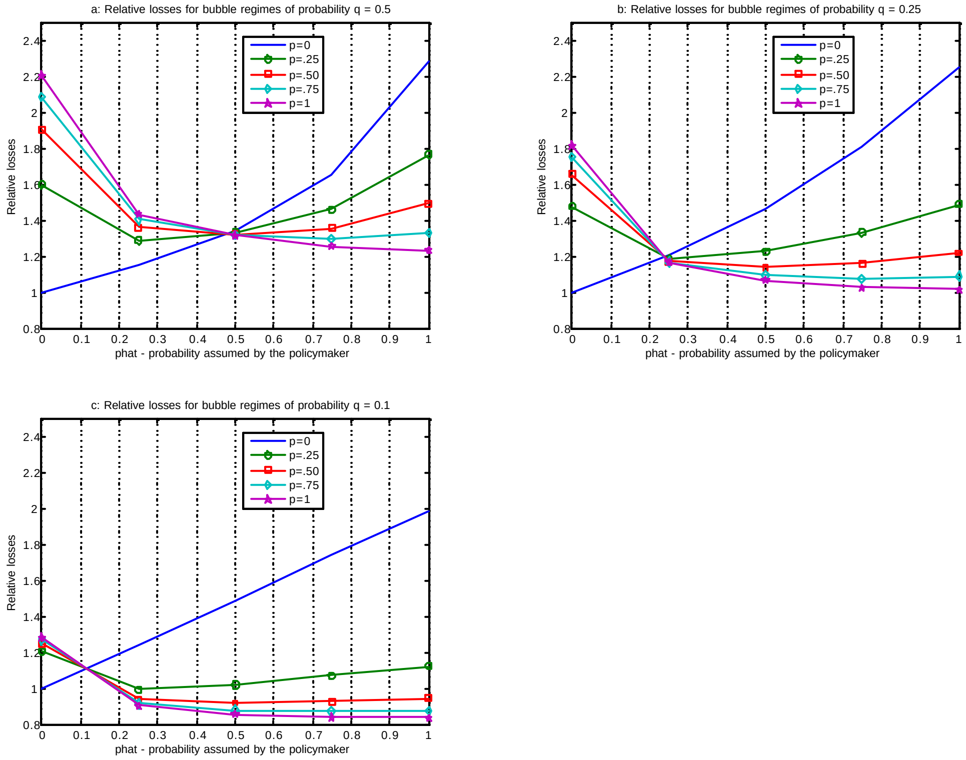

Relative losses of Taylor rules with exchange rate as a function of $\hat{p}$ and p, given q.

Benchmark parameterisation. Note: Losses conditioned on no-bubble regime.

Chart 10

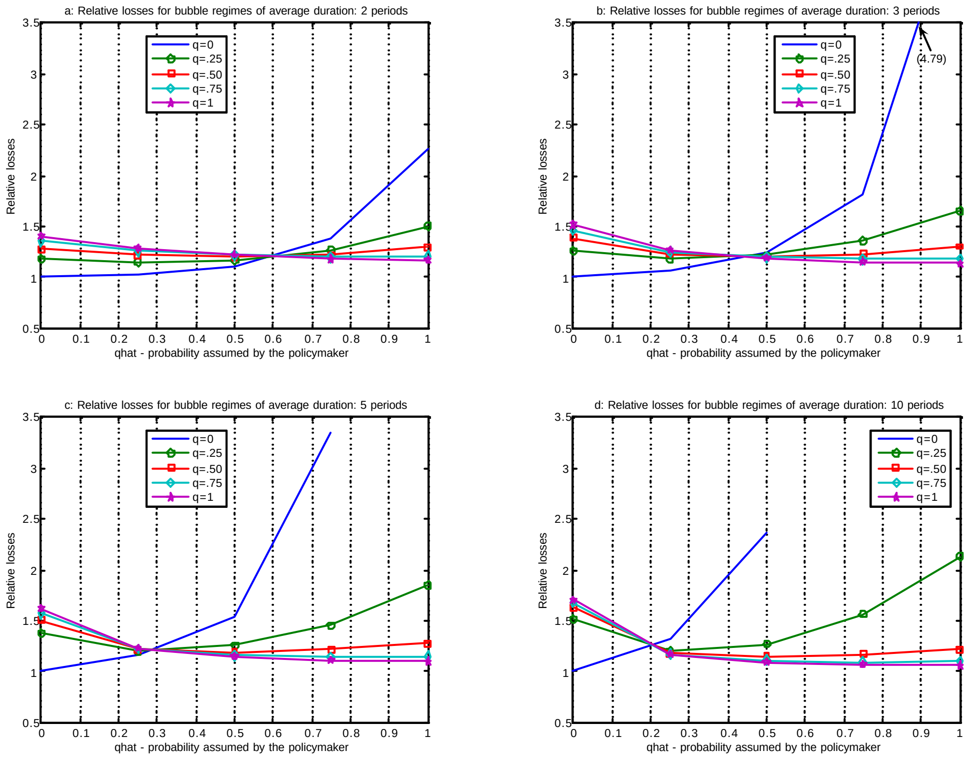

Relative losses of Taylor rules without exchange rate as a function of $\hat{q}$ and q, given p.

Benchmark parameterisation. Note: Missing values indicate 1. Losses conditioned on no-bubble regime

Chart 11

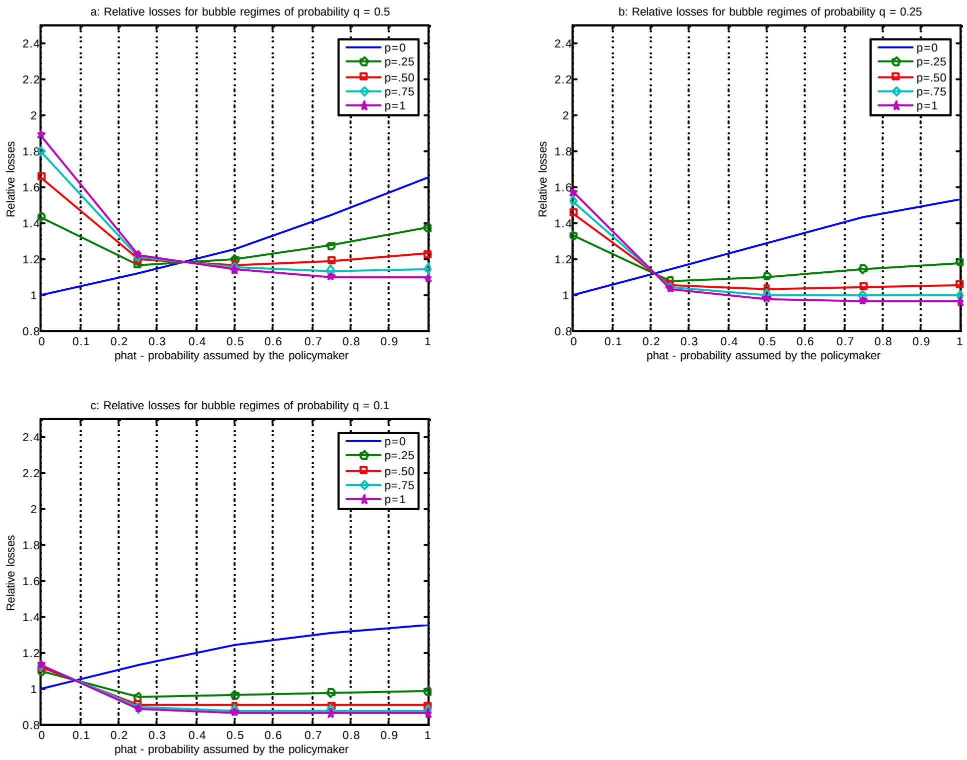

Relative losses of Taylor rules without exchange rate as a function of $\hat{p}$ and p, given q.

Benchmark parameterisation. Note: Losses conditioned on no-bubble regime.

Chart 12: Beliefs of the policymaker over transition probabilities

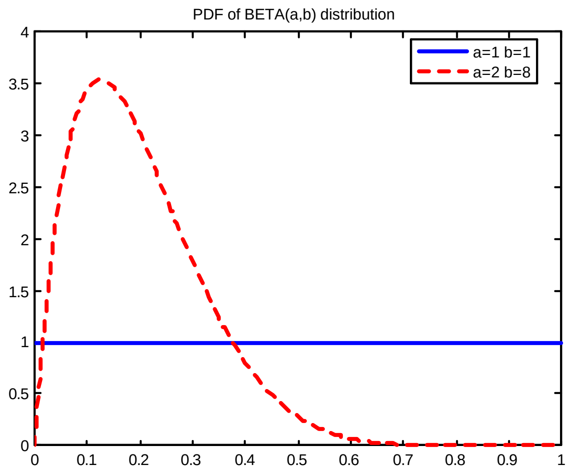

References

Abreu, D and Brunnermeier, M K (2003), 'Bubbles and crashes', *Econometrica*, Vol. 71, No. 1, pages 173-204.

Al Nowaihi, A and Stracca, L (2003), 'Behavioural central bank loss functions, skewed risks and certainty equivalence', *Manchester School*, Vol. 71 (Supplement), pages 21-38.

Andolfatto, D and Gomme, P (2003), 'Monetary policy regimes and beliefs', *International Economic Review*, Vol. 44, No. 1, pages 1-30.

Aoki, M (1967), *Optimization of stochastic systems*, Academic Press.

Ball, L (1999), 'Policy rules for open economies', in Taylor, J (ed), *Monetary policy rules*, The University of Chicago Press.

Bank of England (2002), *Minutes of the MPC*, February meeting.

Bank of England (2005), *Inflation Report*, February.

Batini, N and Nelson, E (2000), 'When the bubble bursts: monetary policy rules and foreign exchange market behaviour', unpublished, available at http://research.stlouisfed.org/econ/nelson/bubble.pdf.

Bean, C (1998), 'The new UK monetary arrangements: a view from the literature', *Economic Journal*, Vol. 108, pages 1,795-809.

Bean, C (2003), 'Asset prices, financial imbalances and monetary policy: are inflation targets enough?', *BIS Working Paper no. 140*, Bank for International Settlements.

Beechey, M, Bharucha, N, Cagliarini, A, Gruen, D and Thompson, C (2000), 'A small model of the Australian macroeconomy', *Research Discussion Paper no. 2000-05*, Reserve Bank of Australia.

Bernanke, B and Gertler, M (2000), 'Monetary policy and asset price volatility', *NBER Working Paper no. 7559*, National Bureau of Economic Research.

Bernanke, B and Gertler, M (2001), 'Should central banks respond to movements in asset prices?', *American Economic Review*, Vol. 91, No. 2, pages 253-57.

Blair, W P and Sworder, D D (1975), 'Continuous-time regulation of a class of econometric models', *IEEE Transactions on Systems Man and Cybernetics*, Vol. 5, pages 341-46.

Blanchard, O and Watson, M (1982), 'Bubbles, rational expectations and financial markets', in Wachtel, P (ed), *Crises in the economic and financial structure*, Lexington Books.

Bordo, M D and Jeanne, O (2002), 'Boom-busts in asset prices, economic instability, and monetary policy', *NBER Working Paper no. 8966*, National Bureau of Economic Research.

Borio, C and Lowe, P (2002), 'Asset prices, financial and monetary stability: exploring the nexus', *BIS Working Paper no. 114*, Bank for International Settlements.

Brainard, W (1967), 'Uncertainty and the effectiveness of policy', *American Economic Review*, Vol. 57 (May), pages 411-25.

Brock, W A, Durlauf, S N and West, K D (2003), 'Policy evaluation in uncertain economic environments', *Brooking Papers on Economic Activity*, pages 235-302.

Brunnermeier, M K (2001), *Asset pricing under asymmetric information*, Oxford University Press.

Bullard, J and Schaling, E (2001), 'New economy - new policy rules?', *Federal Reserve Bank of St. Louis Review*, Vol. 83, Issue 5, pages 57-66.

Cecchetti, S G, Genberg, H, Lipsky, J and Wadhwani, S (2000), 'Asset prices and central bank policy', *Geneva Reports on the World Economy no. 2*, International Centre for Monetary and Banking Studies and Centre for Economic Policy Research, Geneva.

Cecchetti, S G, Genberg, H and Wadhwani, S (2003), 'Asset prices in a flexible inflation targeting framework', in Hunter, W C, Kaufman, G G and Pomerleano, M (eds), *Asset price bubbles: the implications for monetary, regulatory, and international policies*, MIT Press, pages 427-44.

Cogley, T and Sargent, T J (2003), 'The conquest of US inflation: learning and robustness to model uncertainty', unpublished, University of California at Davis and New York University.

Costa, O L V, Fragoso, M D and Marques, R P (2005), *Discrete-time markov jump linear systems*, Springer, London.

Craine, R (1979), 'Optimal monetary policy with uncertainty', *Journal of Economic Dynamics and Control*, Vol. 1, pages 59-83.

Cukierman, A (1996), 'Why does the Fed smooth interest rates', in Belongia, M (ed), *Monetary policy on the 75th anniversary of the Federal Reserve System*, Kluwer Academic Publishers, Norewell, Massachusetts, pages 111–47.

Currie, D and Levine, P (1987), 'The design of feedback rules in linear stochastic rational expectations models', *Journal of Economic Dynamics and Control*, Vol. 11, pages 1-28.

de Grauwe, P and Grimaldi, M (2003a), 'Heterogeneity of agents, transaction costs and exchange rates', *Journal of Economic Dynamics and Control*, Vol. 29, Issue 4, pages 691-719.

de Grauwe, P and Grimaldi, M (2003b), 'A theory of bubbles and crashes', unpublished, Sveriges Riksbank.

do Val, J B R and Başar, T (1999), 'Receding horizon control of jump linear systems and a macroeconomic policy problem', *Journal of Economic Dynamics and Control*, Vol. 23, pages 1,099-131.

Filardo, A J (2001), 'Should monetary policy respond to asset price bubbles? Some experimental results', in Kaufman, G (ed), *Asset price bubbles: implications for monetary and regulatory policies*, Elsevier Science, Amsterdam, Vol. 13, pages 99-123.

Filardo, A J (2003), 'Monetary policy and asset price bubbles: calibrating the monetary policy tradeoffs', unpublished, Bank for International Settlements.

Froot, K and Obstfeld, M (1991), 'Intrinsic bubbles: the case of stock prices', *American Economic Review*, Vol. 81, Issue 5, pages 1,189-214.

Giannoni, M (2001), 'Robust optimal monetary policy in a forward-looking model with parameter and shock uncertainty', unpublished manuscript, Federal Reserve Bank of New York.

Giannoni, M (2002), 'Does model uncertainty justify caution? Robust optimal monetary policy in a forward-looking model', *Macroeconomics Dynamics*, Vol. 6, Issue 1, pages 111-44.

Gilboa, I and Schmeidler, D (1989), 'Maxmin expected utility with a nonunique prior', *Journal of Mathematical Economics*, Vol. 18, pages 141-53.

Goodfriend, M (1991), 'Interest rates and the conduct of monetary policy', *Carnegie-Rochester Conference Series on Public Policy*, Vol. 34, pages 7-30.

Goodhart, C (2001), 'Monetary transmission lags and the formulation of the policy decision on interest rates', *Federal Reserve Bank of St. Louis Review*, Vol. 83, Issue 4, pages 165-82.

Gruen, D, Plumb, M and Stone, A (2003), 'How should monetary policy respond to asset-price bubbles?', *Research Discussion Paper no. 2003-11*, Reserve Bank of Australia.

Hamilton, J D (1989), 'A new approach to the economic analysis of nonstationary time series and the business cycle', *Econometrica*, Vol. 57, Issue 2, pages 357-84.

Hamilton, J D (1994), *Time series analysis*, Princeton University Press.

Hansen, L P and Sargent, T J (2003), *Robust control and model uncertainty in macroeconomics*, unpublished monograph.

Jääskelä, J P (2005), 'Inflation, price level and hybrid rules under inflation uncertainty', *Scandinavian Journal of Economics*, Vol. 107, pages 141-56.

Judd, K L (1998), *Numerical methods in economics*, MIT Press.

Kendrick, D A (2002), *Stochastic control for economic models*, II edition.

Kim, S, Shephard, N and Chib, S (1998), 'Stochastic volatility: likelihood inference and comparison with ARCH models', *Review of Economic Studies*, Vol. 65, pages 361-93.

Kindleberger, C (1978), *Manias, panics, and crashes. A history of financial crises*, John Wiley and Son, New York.

Levin, A, Wieland, V and Williams, J C (2003), 'The performance of forecast-based monetary policy rules under model uncertainty', *American Economic Review*, Vol. 93, pages 622-45.

Levin, A and Williams, J C (2003), 'Robust monetary policy with competing reference models', *Journal of Monetary Economics*, Vol. 50, Issue 5, pages 945-75.

Ljungqvist, L and Sargent, T J (2000), *Recursive macroeconomic theory*, MIT Press.

Mariton, M (1990), *Jump linear systems in automatic control*, Marcel Dekker, New York.

Meltzer, A H (2003), 'Rational and irrational bubbles', in Hunter, W, Kaufman, G and Pomerleano, M (eds), *Asset price bubbles: the implications for monetary, regulatory, and international policies*, MIT Press.

Milani, F (2003), 'Monetary policy with a wider information set: a Bayesian model averaging approach', unpublished, Princeton University.

Minsky, H (1982), 'Can 'it' happen again?', *Essays on instability and finance.*

Miranda, M J and Fackler, P L (2002), *Applied computational economics and finance*, The MIT Press.

Onatski, A and Stock, J H (2002), 'Robust monetary policy under model uncertainty in a small model of the U.S. economy', *Macroeconomic Dynamics*, Vol. 6, No. 1, pages 85-110.

Orphanides, A and Wieland, V (1999), 'Inflation zone targeting', *ECB Working Paper no. 8*, European Central Bank.

Pindyck, R S (1973), *Optimal planning for stabilization*, North-Holland, Amsterdam.

Rustem, B (1994), 'Robust min-max decisions with rival models', in Belsley, D (ed), *Computational techniques for econometrics and economic analysis*, Kluwer Academic Publishers.

Sargent, T (1999), 'Policy rules for open economies: comment', in Taylor, J (ed), *Monetary policy rules*, University of Chicago Press.

Schorfheide, F (2005), 'Learning and monetary policy shifts', *Review of Economic Dynamics*, Vol. 8, pages 392-419.

Shiller, R (2000), *Irrational exuberance*, Princeton University Press.

Shupp, F R (1972), 'Uncertainty and stabilization policies for a nonlinear macroeconomic model', *Quarterly Journal of Economics*, Vol. 86, No. 1, pages 94-110.

Söderström, U (2002), 'Monetary policy with uncertain parameters', *Scandinavian Journal of Economics*, Vol. 104, Issue 1, pages 125-45.

Svensson, L E O (2000), 'Robust control made simple', unpublished, Princeton University.

Svensson, L E O (2003a), 'What is wrong with Taylor rules? Using judgement in monetary policy through targeting rules?', *Journal of Economic Literature*, Vol. 41, Issue 2, pages 426-77.

Svensson, L E O (2003b), 'Monetary policy and low-probability extreme events', *Festschrift for Charles Freedman*, Bank of Canada.

Taylor, J B (1993), 'Discretion versus policy rules in practice', *Carnagie-Rochester Conferences Series on Public Policy*, Vol. 39, pages 195-214.

Tetlow, R J (2003), 'Monetary policy, asset prices, and misspecification: the robust approach to bubbles with model uncertainty', unpublished, ECB workshop on 'Monetary Policy and Asset Prices', December.

Tetlow, R J and von zur Muehlen, P (2001), 'Robust monetary policy with misspecified models: does model uncertainty always call for attenuated policy?', *Journal of Economic Dynamics and Control*, Vol. 25, Issue 6/7, pages 911-49.

Tetlow, R J and von zur Muehlen, P (2004), 'Avoiding Nash inflation: Bayesian and robust responses to model uncertainty', *Review of Economic Dynamics*, Vol. 7, Issue 4, pages 869-99.

Veldkamp, L L (2004), 'Slow boom, sudden crash', *Journal of Economic Theory*, forthcoming.

von zur Muehlen, P (1982), 'Activist versus nonactivist monetary policy: optimal rules under extreme uncertainty', manuscript reprinted as *Finance and Economics Discussion Series paper no. 2001-02*, Federal Reserve Board.

Wadhwani, S B (1999), 'Currency puzzles', Monetary Policy Committee, Bank of England.

Wald, A (1945), 'Statistical decision functions which minimize maximum risk', *Annals of Mathematics*, Vol. 46, pages 265-80.

Woodford, M (2003), *Interest and prices: Foundations of a theory of monetary policy*, Princeton University Press.