Optimal monetary policy in a regime-switching economy

Fabrizio Zampolli[*]

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

Preliminary and incomplete. COMMENTS WELCOME

First draft: 1 September 2003; This draft: 31 January 2004. Some graphs require colour printing.

This paper represents the views and analysis of the author and should not be thought to represent those of the Bank of England or Monetary Policy Committee members.

This paper is work in progress and should not be cited or quoted from without the permission of the authors.

Contents

Abstract

This paper is of interest for two reasons. First, it provides a simple algorithm for solving an optimal control problem in which the law of motion of the economy is a Markov regime-switching vector autoregression. This allows, among other things, to model asymmetric multiplicative uncertainty. Second, it applies this algorithm to study optimal monetary policy in a stylised small open economy model modified to incorporate a boom-bust cycle in the exchange rate. The economy is assumed to alternate randomly between two states: a 'no-bubble' regime, in which the exchange rate fluctuates, in a stationary way, around its long-run equilibrium; and a 'bubble' regime, in which the exchange rate (absent any offsetting impact of policy or exogenous shocks) increasingly deviates from it. We compute the optimal policy rule for this economy, as opposed to an optimised reaction function. This rule is regime-contingent in that policy response varies according to whether the economy is experiencing a bubble or not.

The main results are as follows. First, while the optimal weights on output and inflation do not vary much between regimes, the optimal reaction to the asset price is highly dependent on the regime as well as the stochastic properties of the bubble. Second, uncertainty about the regime makes policy more cautious. Third, a policymaker uncertain about the true stochastic properties of the asset price tends to obtain a 'robust' performance (i.e. minmax outcome) by responding little to the asset price. Finally, over-estimating the probability of an incipient bubble is generally more costly than under-estimating it.

Key words: monetary policy, asset prices, asymmetric risk, regime switching, policy rules

JEL classification: C6, E5

Summary

Your summary here, please.

1 Introduction

Sudden changes in economic behaviour are one of the key uncertainties facing the policymaker. Even if the policymaker were to know the correct economic model today, she would face the risk that tomorrow some of the behavioural equations as well as the properties (mean, persistence, etc.) of the exogenous shocks could dramatically change. Indeed, in contrast to standard models of monetary policy, in which uncertainty is typically modelled as the dispersion of a symmetric distribution, actual policymaking often involves the discussion of risks that are skewed or one-sided.[1] Some of these risks involve large or extreme events, such as the collapse of an asset price, an oil price hike, and abrupt changes in some key aspects of their econometric models (e.g. the degree of exchange rate pass-through). For example, one of the major risks that has worried some MPC members in recent years has been the possibility that sterling could suddenly fall by a material amount. See e.g. the minutes of 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*" (p.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 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.*" (p. 5). Indeed, a common concern amongst 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.[2]

How to respond to potential large asymmetric shocks or abrupt shifts in economic behaviour is very much an open question. Should the policymaker try to pre-empt the shock or adopt a "wait and see" approach (i.e. respond to the shock only if and when it materialises)? In this regard,

---

(1) See e.g. the discussion of skews and asset prices in Goodhart (2001), pp. 178-180.
(2) A proper account of large risks of this type 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 (2002), estimated Taylor rules for a closed economy like the US leaves approximately one third of the variance unexplained.

there are at least two difficulties that the policymaker faces in reaching a decision. The first one is to assess the relative lags with which a given shock and policy affects the economy. If policy operates with longer lags, there is clearly a trade-off between "not taking insurance" against the shock and having the inappropriate policy in place in case the shock does not occur. On the contrary, if a shock affects demand or inflation with the same delay as policy, then preemption may not be necessary, provided the shock is observable when it materialises. The second difficulty concerns the possible endogeneity to policy of the risk that the policymaker might want to pre-empt. This is particularly relevant when the risks concern the behaviour of asset prices. In this case, even the sign of a pre-emptive action is not clear-cut. Assuming that an asset price bubble is fuelled by cuts in the interest rate, two responses are conceivable: the interest rate can be kept lower than otherwise to fend off the deflationary effects of the bubble bursting tomorrow; or, if there is no certainty that the bubble bursts tomorrow, an increase in interest rate may be deemed necessary to prevent the bubble from growing further and, therefore, causing greater damage further down the road.

The literature on monetary policy has produced a number of papers that deals with the issue of whether simple rules should include asset prices or asset price misalignments, but has unfortunately been relatively silent about the more general question of how policy should optimally react to asymmetric or one-sided risk. Al Nowaihi and Stracca (2003) provide a general discussion of optimal monetary policy in response to exogenous asymmetric shocks, when the policymaker's behaviour is affected by the types of cognitive biases commonly considered in the behavioural finance literature. More recently, Svensson (2003) investigates, in a simple model, how low-probability extreme events should be responded to under various types of loss functions, which have been proposed by various authors as a realistic description of central bankers' preferences.[3] In the present paper we also tackle the issue, but from a different angle. We maintain the assumption of a quadratic loss and focus, instead, on incorporating the possibility of regime shifts in the policymaker's model. Allowing for regime shifts is a general and convenient way to model the risks of a sudden change in economic behaviour, whether this involves changes in the model parameters, the model structure or the properties of the exogenous shocks. More precisely, we propose an optimal control problem in which the constraint is a Markov regime-switching vector autoregression (with any finite number of regimes), rather than a

---

(3) Under a quadratic loss the weight given to the extreme event is, *ceteris paribus*, equal to its probability. Under some simplifying assumptions, other loss functions - combined quadratic/absolute deviations, absolute deviation, combined quadratic/constant, and perfectionist - imply, in decreasing order, a lower weight on the extreme event.

stationary vector autoregression. Providing a simple algorithm that solves this problem, which extends the standard linear regulator problem, is the first contribution of this paper. To the best of the author's knowledge, this algorithm and its application (to be discussed below) are new in the economic literature. One of its advantages is its relative simplicity: the problem solved by the policymaker is effectively non-linear/quadratic but maintains all the advantages of the linear quadratic framework.

The second contribution of this paper is to analyse optimal monetary policy in the context of a stylised small open economy model, in which the real exchange rate is assumed to alternate between enduring booms (or bubbles) and sudden corrections (busts) towards equilibrium. Optimal policy is 'asymmetric' or non linear in the sense that it generally responds to the state of the economy (e.g. inflation, output, etc.) in a different way depending on whether the economy is experiencing a bubble or not. Optimal policy is also affected by the uncertainty about future regime shifts, that is whether there will be a bubble or not. The recent literature on monetary policy and asset prices - e.g. Bernanke and Gertler (2000,2001), Cecchetti et al. (2000,2003), Batini and Nelson (2000), Filardo (2001) - look 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 reaction functions (optimised or not) and ranking them according to the computed losses. By contrast, the current paper computes a regime-dependent (and hence time-variant) policy rule as the solution of an optimal control problem.

Before moving to the description of the application, it is worth noticing that the approach presented in this paper can be regarded as one of at least two ways of dealing with *model uncertainty*. One option is to think that there exist a number of possible candidate models that can be the true model at any given time, for the empirical evidence is typically consistent with more than one theory or model specification. A *bayesian* policymaker assigns a probability to each model and then works out the optimal policy by averaging across different models in every period. A more sophisticated version of this approach is for the policymaker to learn from any recent observations about the likelihood of each model and to adjust their probabilities accordingly.[4] The approach adopted in this paper is related but different. The policymaker believes that her model is a good approximation of available data today, but that it might not necessarily continue to

---

(4) For recent applications of Bayesian model averaging see Brock et al. (2003), Cogley and Sargent (2003), and Milani (2003).

be a valid approximation tomorrow. Hence, the policymaker takes into account the possibility that tomorrow the economy can be governed by different candidate models. For example, MPC members could be worried that the behaviour of the exchange rate (e.g. mean, persistence, sensitivity to interest rate, etc.) as well as its relationship with other endogenous variables in the model (e.g. the degree of pass-through) could change radically. [5] Furthermore, in each period the policymaker is assumed - unlike in the model averaging approach - to observe the current regime, and therefore her uncertainty is always about how the model will evolve in the subsequent period. [6]

The application is based on the small open economy model of Ball (1999). Ball's framework is modified to incorporate a boom-bust cycle in the exchange rate by assuming that the economy alternates randomly between two regimes: the 'bubble' regime, in which the exchange rate diverges increasingly from its long-run equilibrium (if not held in check by policy); and the 'no bubble' or 'normal' regime, in which the exchange rate fluctuates, in a stationary way, around its long-run equilibrium. The evolution over time of these two regimes is described by a Markov chain so that the times at which the boom begins and abruptly ends are stochastic. [7] 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 (and shocks).

A number of results emerge from the analysis. The first main result is that, while optimal policy in principle responds to all state variables, including the asset price, the optimal reaction to the exchange rate is not robust to different assumptions about its own behaviour, and may also significantly differ between the 'bubble' and the 'no-bubble' regimes. By contrast, the optimal weights that the policymaker should put on output and inflation do not, according to our numerical analysis, differ markedly between the two regimes (although they do in principle).

Examining this result in more detail reveals that the optimal response to the exchange rate is

---

(5) One of the reasons that could cause parameters to drift is a change in the way private agents' expectations are formed, especially as a result of an unusual changes in the economic environment. Allowing for regime shifts in the structural equations describing private sector's behaviour can be seen as one way of mitigating the Lucas critique.
(6) This is a plausible assumption in some circumstances (e.g. an asset price collapse can be observed), but not in all (e.g. permanent improvement in productivity is not visible unless with several years delay). Extending the current solution algorithm to incorporate the possibility of a hidden state should be feasible but not straightforward. It would involve solving a non linear filtering problem together with the maximisation problem. This is left to future research.
(7) Note that the bubble can restart after bursting.

highly dependent on the probability of the bubble bursting (indicated with p) and the probability of the bubble emerging (indicated with q). Indeed, the optimal response often switches from positive to negative and vice versa. In the 'bubble' regime, it is optimal to pre-empt a future crash if p exceeds a cut-off value (p^*) (i.e. if it is relatively likely that the bubble will collapse tomorrow). [8] By contrast if $p < p^*$ (i.e. if the bubble is relatively unlikely to burst tomorrow), it is optimal to lean against the bubble by responding negatively (positively) to an appreciation (depreciation). If $p = p^*$, then optimal policy involves no response to the asset price. Likewise, in the 'no-bubble' regime, if q exceeds a cut-off value (q^*) (i.e. if it is relatively likely that a bubble is about to develop), then optimal policy counters any shock that could develop into a prolonged future bubble by responding negatively (positively) to an appreciation (depreciation). But if $q < q^*$, then optimal policy responds positively (negatively) when the currency appreciates (depreciates) to offset the impact on inflation of shocks that are relatively likely to be of transitory nature. Again, optimal policy does not respond to the asset price when $q = q^*$. The cut-off values of the probabilities, p^* and q^*, are found in the numerical analysis to depend on the relative weight on output in the policymaker's loss function. In particular, p^* tends to be 'low' (i.e. less than 0.5) under strict inflation targeting and to rise with the relative loss weight. [9]

A further interesting result is the fact that policy becomes more cautious when there is uncertainty about whether or not there is a bubble, with cautiousness increasing with the 'strength' of the bubble (i.e. the rate at which past asset prices translate into current asset prices). The degree of cautiousness is usually strongest for intermediate values of the probabilities. This reflects a form of Brainard (1967) uncertainty.

Finally, a key result of the paper concerns the selection of the transition probabilities (p and q) which, as we have seen, critically affect the optimal response to the asset price. Given the difficulty involved in identifying bubbles and the fact that previous episodes may contain little information about current and future episodes, the policymaker may be highly uncertain about these probabilities. For this reason, we also compute the relative costs of policy mistakes, namely the cost involved in selecting the wrong probabilities, under different scenarios. We normally find that the robust (minmax) value of the probability is close to a half. This means that if a

---

(8) This involves keeping interest rates higher (lower), ceteris paribus, to offset the future inflationary (deflationary) impact of a currency depreciation (appreciation).

(9) p^* and q^* are approximately 1/2 when the policymaker's loss function attaches equal weight to output and inflation.

policymaker has a strong preference for an early resolution of uncertainty, she would attach approximately equal probability to being in a 'bubble' or being in a 'normal' regime. For bubbles of longer expected duration (i.e. five years), the minmax value tends to become smaller (around 25%). More interestingly, we find that overestimating the probability of an incipient bubble can be significantly more costly than underestimating it under most scenarios.

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 policy making. For instance, while, according to some commentators, the UK 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.

The paper is organised as follows. Section 2 describes the optimal control problem with structural or regime shifts, and Section 3 illustrates how to compute its solution. Section 4 describes the model used in the application. Section 5 analyses the sensitivity of the optimal interest rate to the model's state variables, namely output, inflation and lagged exchange rate. Section 6 performs a robustness analysis on the choice of the probabilities. Section 7 concludes.

2 The optimal control problem with regime shifts

The policymaker's problem is to choose a decision rule for the control u_t to minimise the intertermporal loss function: [10]

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

subject to x_0, s_0 given, and the following model of the economy

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

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. 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, which is observable

---

(10) An introduction to dynamic programming and the optimal linear regulator problem can be found in Ljungqvist and Sargent (2000), Ch. 2-4.

at t,[11] is assumed to be a Markov chain with probability transition matrix[12]

$$P = \left[p_{ij} \right]_{i,j=1,\dots,N} \tag{2}$$

in which $p_{ij} = prob\{s_{t+1} = j | s_t = i\}$ is the probability of moving from state i to state j at $t+1$; and $\sum_{j=1}^{N} p_{ij} = 1, i = 1, \dots, N$. These probabilities are assumed to be time-invariant and exogenous. The formulation **(1)** is general enough to capture different types of jumps or extreme changes in the economic system.

Note that 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 ($i = 1, \dots, N$) are identical, regardless of P.

3   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 current regime s_t.[13] 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:

$$v(x_t, i) = \max_{u_t} \left\{ r(x_t, u_t) + \beta \sum_{j=1}^{N} p_{ij} E_t^\varepsilon \left[v(x_{t+1}, j) \right] \right\} \qquad i = 1, \dots, N \tag{3}$$

where $v(x_t, i)$ 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, $s_t = i$; 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 payoff $r(\cdot)$ as well as the discounted sum of all future payoffs. 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 $t+1$ weighted by the transition probabilities **(2)**. Given the

---

(11) This means that the uncertainty faced by the policymaker is about where the system will be at t+1, t+2, and so forth.

(12) For an introduction to Markov chain and regime switching vector autoregressive models see e.g. Hamilton (1994).

(13) Henceforth, regime or state of the world will be used interchangeably.

infinite horizon of the problem, the continuation values (conditioned on a particular regime) have the same functional forms.

In principle, one could solve the above problem for the case in which $r\,(\cdot)$ is not a quadratic form. Techniques for discretising the state space or approximating the Bellman equations could be used for this purpose (see e.g. Judd, 1998). However, these techniques can easily run into the 'curse of dimensionality', thereby limiting the size of the state and control spaces to an unsatisfactory number of variables. We therefore prefer to proceed under the assumption that r is a quadratic form:

$$r\,(x_t, u_t) = x_t' R x_t + u_t' Q u_t \tag{4}$$

where R is a $n \times n$ positive definite matrix and Q is $m \times m$ positive semi-definite matrix. Given the linear-quadratic nature of the problem, let us assume that

$$\upsilon\,(x_t, i) = x_t' V_i x_t + d_i \qquad\qquad i = 1, .., N \tag{5}$$

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 **(5)** into the Bellman equations **(3)** (after using **(4)** and **(5)**) and compute the first-order conditions, which give the following set of decision rules:

$$u\,(x_t, i) = -F_i \cdot x_t \qquad\qquad i = 1, .., N \tag{6}$$

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 **(3)**, 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

$$[V_1 ... V_N] = T\,([V_1 ... V_N]) \tag{7}$$

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:

$$V_i = R + \beta G\left[A'VA|_{s=i}\right] - \beta^2 G\left[A'VB|_{s=i}\right]\left(Q + \beta G\left[B'VB|_{s=i}\right]\right)^{-1} G\left[B'VA|_{s=i}\right] \tag{8}$$

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

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

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 **(8)**, the matrices F_i in the closed-loop decision rules **(6)** are given by:

$$F_i = \left(Q + \beta G\left[B'VB|_{s=i}\right]\right)^{-1} \left(\beta G\left[B'VA|_{s=i}\right]\right) \qquad i = 1, .., N \qquad \textbf{(9)}$$

Solving for the constant terms in the Bellman equations **(3)** - after substitution of **(6)** - 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 **(5)** is given by

$$d = (I_N - \beta P)^{-1}\, \beta P\Gamma \qquad \textbf{(10)}$$

where $\Gamma = [tr\,(V_i \Sigma_\varepsilon)]_{i=1,...,N}$ ($N \times 1$ vector).[14] [15]

The decision rules **(6)** depend on the uncertainty about which state of the world will prevail in the future, as reflected in the transition probabilities **(2)**. Yet, the response coefficients (i.e. the entries in F_i) do not depend on the variance-covariance matrix Σ_ε of the zero-mean shock ε in **(1)**. Thus, with respect to ε *certainty equivalence* holds in that the policy rules **(6)** are identical to the ones obtained by assuming that *within* each regime the system behaves in a completely deterministic fashion. The noise statistics, as is clear from **(10)**, affect the objective function.

It is interesting to note that the above solutions incorporate the standard linear regulator solutions as two special cases. First, by setting the transition matrix $P = I_N$ (i.e. N-dimensional identity matrix), 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 (i.e. $A_i = A$, $B_i = B$), the solution obtained is trivially that of a standard linear regulator problem with a time-invariant law of transition.[16]

---

(14) 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.

(15) The law of transition **(1)** can be generalised to make the variance of the noise statistics vary across states of the world: i.e.

$$x_{t+1} = A\,(s_{t+1})\,x_t + B\,(s_{t+1})\,u_t + C(s_{t+1})\varepsilon_{t+1}$$

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 **(10)**: $\Gamma = [tr\,(V_i \Sigma_i)]_{i=1,...,N}$.

(16) In this case **(9)** boils down to

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

13

4 Exchange rate booms and busts 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:

$$y_{t+1} = \alpha y_t - \beta (i_t - \pi_t) - \chi a_t + \eta_{t+1} \tag{11}$$

$$\pi_{t+1} = \delta \pi_t + \gamma y_t - f (a_t - a_{t-1}) + \varepsilon_{t+1} \tag{12}$$

$$a_t = \theta (i_t - \pi_t) + v_t \tag{13}$$

(11) 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.[17] 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. **(12)** 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. **(13)** is a *reduced-form* equation that relates the real exchange rate to the current level of the real interest rate and a transitory shock v. The positive sign of θ 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 v, with variance σ_v^2, reflects all other factors (future expectations, foreign interest rate, etc.) that can influence the exchange rate.

The use of the reduced-form **(13)** in place of the uncovered interest parity (UIP) condition can be realistically justified with the empirical failure of the latter. As stressed e.g. 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

where V is the solution to the single Riccati equation

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

and **(10)** boils down to the constant

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

See e.g. Ljungqvist and Sargent (2002), pp. 56-58.

[17] All variables are in logs except i and π, and are expressed as deviations from its long-run equilibria.

a good empirical approximation of the actual exchange rate behaviour.[18] 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.

In our application we replace **(13)** with:

$$a_t = \rho\,(s_{t+1})\,a_{t-1} + \theta\,(s_{t+1})\,(i_t - \pi_t) + v_t \tag{14}$$

where the term ρa is an *ad-hoc*, but not unrealistic way, of capturing a bubble-like behaviour, that is the fact that the exchange rate can grow out of line of its long-run equilibrium, whatever the underlying factors. More precisely, **(14)** is modelled as a regime-switching autoregression, in which the autoregressive coefficient $\rho\,(s_{t+1})$ takes on different values depending on the regime in which the system is:

$$\begin{aligned}\rho\,(s_{t+1}) \quad &> 1 \quad &&\text{if} \quad s_{t+1} = 1 \\ &= 0 \quad &&\text{if} \quad s_{t+1} = 2\end{aligned}$$

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 long-run equilibrium 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 behaviour abruptly collapses towards its long-run equilibrium (again assuming that the real interest rate is also at its long-run equilibrium and that there is no shock). Likewise, **(14)** also allows for the possibility that policy can have a different contemporaneous impact on the exchange rate depending on the regime in which the system is (i.e. θ is contingent on s_{t+1}).[19] 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 $p = \Pr\{s_{t+1} = 2 | s_t = 1\}$ and $q = \Pr\{s_{t+1} = 1 | s_t = 2\}$, $(t = 0, 1, 2, ...)$. 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 the bubble and of the policy decision. At the time policy is chosen the policymaker knows s_t, but not s_{t+1}. Therefore, we assume that the state of the asset price during period t is revealed only towards the end of it, after policy has been decided. This assumption is meant to capture, realistically, the uncertainty about the bubble faced by the policymaker.

---

(18) Wadhwani (1999) discusses an empirical model of the exchange rate for the UK. Beechey et al. (2000) provides a model of Australia. Both models are similar to the reduced form adopted in Ball (1999).

(19) More generally, other structural parameters could change across regimes. For example, following a large depreciation of the exchange rate, the Phillips curve could become steeper. The algorithm described in Section 2 allow to compute optimal policies for these scenarios.

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.[20] Indeed, there is no consensus in the literature on the definition of a bubble and its underlying causes.[21] Bubbles may be rational and unrelated to fundamentals (e.g. 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 compensate 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 under-valued. However, bubbles do not have to rely only on self-fulfilling expectations: they may arise from manias or irrational exuberance (Kindleberg, 1978; Minsky, 1982; 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.[22] More recently, De Grauwe and Grimaldi (2003a, 2003b) 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 past values, a feature confirmed by the empirical evidence cited by the authors.[23]

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

$$\sum_{t=0}^{\infty} \phi^t \left(\pi_t^2 + \lambda y_t^2 \right) \tag{15}$$

where $\phi \in (0, 1)$ is the discount factor, and λ is the penalty on output stabilisation. It is important

---

(20) The 'bubble interpretation' is not the only possible one. An alternative could be to assume $0 < \rho\,(1) < 1$ and $\rho\,(2) = 0$. In this case the shock is either a stationary but persistent one, reflecting the influences of short- or medium-term factors not explicitly captured by the stylised model, or a purely transitory one. 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.

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

(22) As suggested e.g. 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 Knigthian uncertainty rather than 'well-behaved' probability distributions.

(23) The use of simple trading rules represents the agents' rational response to their own limited ability to process information. Because agents are rational, the relative use of these rules amongst agents varies over time reflecting how well they do in terms of risk-adjusted profit.

to note that the policymaker only targets output and inflation. As stressed by e.g. Bean (2003), there is no need to specify a target for asset prices when the targeting rule is a statement about the loss function. [24] However, the instrument rule - as will be shown below - is in principle a function of all state variables, including the asset price, for the policymaker uses all available information to forecast inflation and output. One could argue that optimal rules are difficult to implement in practise, because of the technical and political difficulties involved in forecasting for *all* future periods, and difficult to communicate if the model is more complex than the one analysed here. Therefore, simple rules could provide useful guidance. The crucial question then is not whether in principle one ought or not to include asset prices as an argument but which weight the policymaker ought to put on them. Below, we will see that if the decision rule arising from the control problem is interpreted as a rule of thumb, the choice of a 'robust' coefficient on the exchange rate could be problematic.

Given the simple dynamic structure of the model, a period is normally interpreted as a year. The chosen parameterisation is summarised in Table A. The parameter α, β and γ are calibrated using the estimated values found in Bean (1998). The equations are estimated individually using GLS over the sample 1950-1996 (δ is kept fixed in the estimation). The estimates are similar to those obtained with annual data for the Euro Area and the US in similar papers (e.g. Orphanides and Wieland, 1999). The standard deviations of the demand and inflation shocks, σ_η and σ_ε are also taken from Bean (1998). These are not the actual standard errors of the estimation but those of the Treasury forecasting record during the period 1985-92. The parameters describing the existence of a foreign exchange rate channel, χ, f and θ, are taken from Ball (1999). We have chosen a discount factor $\phi = .96$ in accordance with annual data. The bottom of the table reports the parameters chosen for the equation governing the exchange rate. In addition to θ, whose value is the same as in Ball (1999), we calibrate the existence of a bubble by assuming that the exchange grows by a rate comprised between 10 and 30 percent a year (in the absence of any shock and any offsetting policy action). The choice of σ_v is somewhat arbitrary but probably not implausible.

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

---

(24) In the terminology of Svensson (e.g. 2002) the loss **(15)** is a general targeting rule. For a definition of instrument rule, see *ibidem*.

Bean (1998)		loss	
α	.72	ϕ	.96
β	.47	λ	[0, 3]
γ	.49		
δ	1		
σ_η	.0167		
σ_ε	.0108		

Ball (1999)		exchange rate	
χ	.2	$\theta_{(1)} = \theta_{(2)}$	2
f	.2	$\rho_{(1)}$	[1.1,1.3]
θ	2	$\rho_{(2)}$	0
		σ_v	.05

Table A: Parameters for the small open economy

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 interest rate when the lagged exchange rate has appreciated. A simple Taylor rule, not augmented with the exchange rate, is generally suboptimal when the economy is described by **(11)-(13)**. The intuition of this result can be explained as follows. For simplicity, assume that the policymaker is targeting inflation only (i.e. $\lambda=0$ in **(15)**). In Ball's model, shocks to the real exchange rate are transitory. An appreciation at time t causes inflation and output to fall at t+1, the time at which the appreciation is reversed. If the policymaker behaves as predicted by a simple Taylor rule, thus ignoring the exchange rate, she would cut the interest rate. As a result, policy would have a positive impact on inflation in t+2 through the import price channel, and a further positive impact in t+3 through the output channel. By ignoring the transitory nature of the shocks to the exchange rate, policy would then cause excessive and suboptimal fluctuations in inflation. By contrast, a positive response to a lagged appreciation would tend to reduce the response of policy to the observed current change in inflation. Equivalently, optimal policy in his model involves targeting a measure of long-run inflation, which filters out any transitory variations due to the exchange rate.

In Ball (1999) optimal policy is computed and analysed with reference to a time-invariant model in which there is no parameters uncertainty. A general 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 regarded as a good approximation, it may turn out to be inadequate in some aspects tomorrow (e.g. 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. But the behaviour of the exchange rate is certainly one of the most uncertain structural feature of **(11)-(13)**.

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? And, if there is a bubble, how should the policymaker deal with it? Leaning against the bubble is one possibility, which involves keeping the interest rate lower than warranted by the current levels of output and inflation, in order to limit the growth of the bubble and, consequently, the detrimental effects of its possible crash going forward. On the other hand, the prospect of an immediate crash calls for a rise in the interest rate today in order to offset the inflationary consequences of a large fall tomorrow. Clearly, the optimal policy must depend on the probability q that a bubble is emerging as well as its average duration, as reflected in the probability p of the bubble collapsing. Unfortunately, these probabilities can be difficult to assess. The policymaker can end up making serious mistakes whatever her decision. If she leans against the bubble and the bubble bursts, policy would be too loose to counter its inflationary effects. By contrast, if the policymaker is bracing for a crash and this does not occur, her policy would be too tight and could further fuel the bubble, with even more unpleasant consequences further ahead. Below we set out to investigate how policy should respond in the presence of a bubble like the one in **(14)**. In Section 6 we examine how the wrong choice of the probabilities affects the policymaker's loss.

5   Analysis of optimal policy responses

Optimal policy is computed using the algorithm described in Section 2 after casting **(11)**, **(12)** and **(14)** in matrix form. This requires setting $\beta \equiv \phi$ (discount factor in **(15)**), $x = [y, \pi, a_{-1}]'$, $u = [i]$, and $\varepsilon = [\eta, \varepsilon, \upsilon]'$ in **(1)** and **(4)**, and constructing the matrices A, B, Q, and R accordingly. The solution gives the following optimal decision rule:

$$i_t = f_y(s_t) y_t + f_\pi(s_t) \pi_t + f_a(s_t) a_{t-1} \tag{16}$$

19

There are two notable aspects. First, the optimal rule depends on all state variables. Therefore, a simple Taylor rule, in which the interest rate responds only to output and inflation, is generally suboptimal. In practice, whether the asset price is given a large or small weight depends on the particular model at hand. We will return on this below. Second, the optimal response coefficients generally take different values in different states of the world. This means that the optimal response is not invariant but must adapt to the changing circumstances. Clearly, a state-contingent rule may no longer be simple and many of the advantages (mainly transparency) of a simple rule could then be lost. This can help explain why, in the face of the considerable degree of uncertainty involved in assessing the existence and impact of asset price bubbles - besides of course other reasons - a policymaker might not want to commit to, or be seen as following the prescriptions of a simple decision rule. Below we set out to investigate the sensitivity of the optimal interest rate to the properties of **(16)** under different scenarios about the characteristics of the bubble, namely its size and its transition probabilities.

Table B reports the optimal response coefficients in **(16)**. The left-upper part shows the optimal response coefficients in the absence of the bubble, that is when **(14)** holds, for different values of the weight on output stabilisation in **(15)**. In the case of strict inflation targeting, $\lambda = 0$, policy responds more aggressively to all variables. As the loss places more weight on output stabilisation, the responses decrease. In particular, the response to inflation as well as the positive response to a lagged exchange rate appreciation is much less strong . These responses are compared with those implied by the explicit consideration of a bubble or non-linear behaviour in the exchange rate. The table reports only the coefficients computed for $\lambda = 1$. On the left hand-side of the table are the response coefficients for a bubble that grows ten percent per period, other things equal; on the right hand-side are the response coefficients for a bubble that grows thirty percent per period. For each pair of probabilities (q, p), we report the responses in both states: the 'bubble' state (s=1) and the 'normal' state (s=2). For example, (.25,.5) 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 $1/p$).

There are a number of qualitative features that emerge from the inspection of Table B. We begin by examining the responses to output and inflation separately from those to the real exchange rate. First, the responses to output and inflation are generally smaller relative to the case in which there is no bubble uncertainty. Second, policy attenuation is generally stronger, other things equal, the

greater the growth rate of the bubble (as can be seen by comparing the left-hand side to the right-hand side of the table). Both findings have the interesting implication that policy should respond less not only when there is a switch between regimes but also when the economy is hit by ordinary shocks to demand and inflation. Third, the responses to output and inflation are very similar in the 'bubble state' and the 'normal state' in most cases. This means that most of the time when the economy moves from one state to the other, a failure by the policymaker to adjust policy to the new circumstances, bar the treatment of the asset price, should cause relatively minor losses.

By contrast, the size and the sign of the responses to the exchange rate vary greatly with the transition probabilities as well as the state of the world. In particular, the coefficients generally decrease with the probability of a bubble emerging q and with its average duration $(1/p)$. Let's look first at the effects of p on a_{-1}. Normally, the higher p (the lower $1/p$), the greater the positive response to an appreciation when the system is in a 'bubble state'. For $p=.75$, the responses are the closest to the ones that would prevail in a stationary, no bubble world, as the policymaker is anticipating the inflationary (deflationary) effects of a fall (rise) in the exchange rate. However, the response in the 'bubble state' becomes almost negligible when $p=.5$ and slightly negative when $p=.33$, for $\rho_{(1)} = 1.1$; a bubble that grows faster tends to make these figures slightly more negative (as can be seen by looking at the left-hand side of the table). Note that, the responses in the bubble state are remarkably similar across different values of q. In other words, they appear to depend only on the expected duration of the bubble. The same values becomes smaller and more negative when the bubble grow faster ($\rho_{(1)}=1.3$), but again the values are remarkably similar for different values of q. Overall, for bubble of expected duration of two or more periods the coefficients appear to be quite small and biased towards a negative response, which corresponds to a policy of (gently) leaning against the bubble.

As to the effects of q on a_{-1}, the response to the exchange rate is positive for low values of the probability, and becomes negative as the probability rises. Unlike in the bubble state, the response to the exchange rate in the normal state is positive for low probability q, close to zero for probability of around 50%, and negative for high probability q. The responses are remarkably similar across bubbles of different expected duration. In other words, they do not seem to depend on p. When a bubble is not likely, the responses are closer to the ones that would prevail in a no bubble world, though they are attenuated. When a bubble is likely, the responses are negative, indicating that the adequate response to an exchange rate shock is opposite to the one required by

a transitory shock. In this case the policymaker wants to limit the bubble from its inception. However, for intermediate probabilities, which reflects the highest degree of uncertainty about the nature of an exchange rate disturbance, the policymaker does not attach any weight to the exchange rate. A fast growing bubble ($\rho_{(2)}$=1.3) changes the picture slightly in that the responses become smaller and more negative.

To sum up, policy is more cautious in responding to output and inflation when the economy alternates between normal periods and period of inflating bubbles, compared to an economy in which bubbles never emerge. This also means that policy responds more cautiously to ordinary shocks to demand and inflation, not only to shocks to the exchange rate. The response to the exchange rate, instead, is highly dependent, perhaps unsurprisingly, on the stochastic properties of the bubble.

There is not much more that can be said about the optimal response coefficients by looking only at Table B. We resort to a graphical illustration to see if more can be learnt. Figures 1-3 show the contour lines for the optimal response coefficients in the decision rule **(16)**. The vertical axes indicate the probability p of a bubble bursting assuming one exists; the vertical axes indicate the probability q of moving into a bubble. The columns correspond to different coefficients (in order f_y, f_π, f_a); the rows correspond to different states of the world (the bubble state is in the upper part; the normal state is in the lower part). Each figure refers to a different weight λ on output stabilisation in the loss **(15)**, and they all refer to a bubble growing 20% a year (ceteris paribus).

We initially focus on the sensitivity of policy to output and inflation. Chart 1 corresponds to strict inflation targeting $\lambda = 0$. The first thing that one notices is that as both the probabilities q and p rise, there is an initial attenuation of responses to output and, more strongly, to inflation, in both states of the world, which tends to disappear as these probabilities both continue to rise. Indeed, the locus where policy is the most cautious is where the probabilities of moving into one regime or remaining in the other are more or less equally balanced. In the normal state, there is a steep decrease of the inflation coefficient as q rises. For bubble expected to last for about two periods (p=.5) the decline in the coefficient is much steeper as q rises towards .5 but then increases more gently after that. When the system is in a bubble state, the coefficient on inflation declines steeply for low values of the probability p and appears to flatten as the probability further rises.

	q	p		y	π	a_{-1}		y	π	a_{-1}
$\lambda=0$	0	0	s=2	1.23	3.50	.50				
1			s=2	1.16	1.88	.18				
3			s=2	1.06	1.55	.11				
$\rho(1)=1.1$							$\rho(1)=1.3$			
$\lambda=1$	.25	.75	s=1	1.08	1.83	.09	s=1	1.05	1.80	.07
			s=2	1.08	1.83	.09	s=2	1.05	1.80	.07
	.50	.75	s=1	1.06	1.82	.09	s=1	1.01	1.78	.06
			s=2	1.07	1.82	.01	s=2	1.02	1.78	-.03
	.75	.75	s=1	1.09	1.85	.09	s=1	1.05	1.82	.07
			s=2	1.08	1.82	-.07	s=2	1.04	1.79	-.13
	1	.75	s=1	1.16	1.92	.11	s=1	1.15	1.92	.09
			s=2	1.06	1.79	-.16	s=2	1.03	1.75	-.23
	.25	.50	s=1	1.07	1.82	0	s=1	1.03	1.79	-.04
			s=2	1.08	1.82	.08	s=2	1.03	1.79	.06
	.50	.50	s=1	1.05	1.81	0	s=1	1.00	1.77	-.04
			s=2	1.05	1.81	0	s=2	1.00	1.77	-.04
	.75	.50	s=1	1.07	1.83	.01	s=1	1.03	1.80	-.04
			s=2	1.06	1.82	-.08	s=2	1.02	1.78	-.14
	1	.50	s=1	1.12	1.88	.02	s=1	1.10	1.87	-.02
			s=2	1.05	1.80	-.16	s=2	1.01	1.77	-.23
	.25	.33	s=1	1.07	1.83	-.05	s=1	1.03	1.80	-.10
			s=2	1.07	1.82	.08	s=2	1.03	1.79	.06
	.50	.33	s=1	1.05	1.82	-.05	s=1	1.00	1.78	-.10
			s=2	1.05	1.82	0	s=2	1.00	1.78	-.04
	.75	.33	s=1	1.07	1.83	-.05	s=1	1.02	1.80	-.10
			s=2	1.06	1.83	-.07	s=2	1.02	1.80	-.13
	1	.33	s=1	1.10	1.87	-.04	s=1	1.08	1.85	-.09
			s=2	1.07	1.83	-.15	s=2	1.03	1.80	-.23

Table B: Optimal response coefficients

Under flexible targeting (Figure 2), and in particular when equal weights are placed on output and inflation stabilisation (i.e. $\lambda = 1$), the attenuation result described above is much more milder for inflation but more noticeable for output. Both the coefficients on output and inflation are quite sensitive to initial rises in the probability q in the normal state, whereas they seem to be less sensitive to initial rises of both p and q in the bubble state. As output versus inflation stabilisation becomes more important (see e.g. Figure 3), the responses to inflation and output become very insensitive to the transition probabilities in both states of the world.

We now turn our attention to the optimal response of policy to the real exchange rate. There are a number of features that emerge from the joint examination of all the charts. First, the response largely depends on: the probability p of the bubble collapsing if in the bubble state (top of the figures); on the probability q of a new bubble forming if in the normal state (bottom). The contour lines, however, are not completely flat, meaning that in some cases both transition probabilities matter, especially under strict inflation targeting (see Fig. 1).

Second, there is a cut-off value for the probabilities, at which the weight put on the exchange rate is zero. As this value is crossed, the policy response switches sign. For example, consider the bubble state in Chart 1 (the plot in the top-right corner). The cut-off value is around 30%. For low probabilities of the bubble crashing (meaning a longer expected duration), the optimal response consists in keeping the interest rate lower, ceteris paribus, in order to limit the bubble growth. As the probability p crosses the cut-off point, the response becomes positive: policy turns pre-emptive, that is the interest rate is kept higher (lower), ceteris paribus, in anticipation of the inflationary (or deflationary) impact of a downward (upward) jump in the real exchange rate. Now, imagine that the bubble has burst. Depending on p and q the weight on the exchange rate normally jump to a different values and can even switch sign. For example, for low p and low q, the weight will move from a negative to a positive one. The cut-off point for q in the normal state is approximately between 70% and 80% (depending on p). Hence, unless the probabilities of a bubble emerging next period are high, optimal policy will put a positive weight on the real exchange rate, though less than in the case where the shocks are purely transitory.

An interesting aspect of this cut-off point, which emerge on examination of all figures, is that it varies with the weight λ placed on output stabilisation: the cut-off point for p in the bubble state tends to rise with λ, whereas the cut-off point for q in the normal state tends to fall with λ. The

cut-off point for p raises from around 30% when $\lambda=0$ to just over 50% and to around 75% in the bubble state when $\lambda=1$ and $\lambda=5$, respectively. The cut off for q falls from around 70-80% under inflation targeting to about 50% and 25% when $\lambda=1$ and $\lambda=5$, respectively. Therefore, in the bubble state an increasing concern for output stabilisation widens the range of probabilities for which it is optimal to lean against the bubble, at the expense of the pre-emption motive. Likewise, in the normal state an increasing concern for output stabilisation widens the range of probabilities for which it is optimal to 'take insurance' against the inception of a bubble rather than offsetting a transitory shock.

To sum up, whether a policymaker decides to lean against the bubble or not depends not only on the probabilities of moving from one state to the other, but also on how flexible is her inflation targeting. A greater degree of flexibility tends to make a negative response to an appreciation of the real exchange rate more 'attractive' in both states of the world, in the sense that this would be the optimal thing to do for a wider range of probabilities.

Given the crucial importance of the transition probabilities in determining the optimal policy, the next section sets out to investigate more closely how choosing the 'wrong' probabilities affect the policymaker's loss.

6   Robust selection of transition probabilities

The stochastic properties of the bubble are undisputedly the most uncertain ingredient in the policy decision process, which of course requires a lot of judgement by the policymaker. Because of the difficulties involved, the policymaker might be reluctant or, most likely, unable to specify a single probability value. Instead, she would probably form an opinion about a probability interval, whose length would be a reflection of her degree of uncertainty. The policymaker is effectively in a world of Knightian uncertainty. How can she deal with it? And how can we reconcile this uncertainty with an optimal control exercise which requires the assignment of unique values to those probabilities? A heuristic way to deal with it, which is the one adopted here, is to analyse how different decisions affect the policymaker's loss under different scenarios.[25] In this regard, the natural questions one can ask are: is there a decision which delivers relatively better outcomes than others in most or all scenarios? Or, is there a decision which delivers a particularly bad

---

[25] For a formal approach to robustness in macroeconomics see Hansen and Sargent (2003).

outcome in some cases, and which should perhaps be avoided (assuming the policymaker is not a risk lover)? How asymmetric are the losses for a particular choice? To answer these questions, we will carry out the following exercise. We assume that the policymaker, who does not know q, chooses a probability $\hat{q}$. For a given expected duration of the bubble, we then compute the losses associated with all the pairs $(\hat{q}, q)$. We want to see how a given choice $\hat{q}$, and hence the optimal policy based on it, performs across different realisations of the true value of the probability q.

It is worth noticing that exercises similar to the one carried out here, we believe, may help to inform the judgement of the policymaker in choosing the optimal policy in the face of uncertainty. There is no need to base this informal robustness or sensitivity exercises on specific decision criterion, like for example the *maxmin*. Indeed, there is no need to adopt any criteria at all for decision under uncertainty, nor to know the policymaker's preferences towards risk. Instead of incorporating risk preferences in the formal optimisation problem and work out one optimal policy, the approach followed here - in a sense - works in the reverse. It presents the policymaker with a menu of options from which she could choose, thereby letting her reveal her priors about the likelihood of specific events or her taste for risk, elements which are in any case very difficult to elicit and formalise.

Chart 4-6 report the results of the experiment described above for a bubble assumed to grow 10% a year (keeping other factors constant). Each chart corresponds to a different weight λ on output stabilisation, and each plot in a given chart refers to a bubble of different expected duration ($=1/p$). Each line in a plot corresponds to a true value of the probability q. Clearly, this has a minimum for $q = \hat{q}$. All losses are computed using an algorithm similar to the one presented in Section 2.[26] All losses are normalised so that the loss corresponding to $\hat{q} = q = 0$ in each chart equals one.

Consider first Chart 4, in which the policymaker has an equal concern for output and inflation stabilisation (i.e. λ=1), which we will take as the benchmark. Three features are immediately apparent from this chart. First, there is a value of $\hat{q}$ which guarantees a minimum loss under all realisations of q. This is the minmax value of $\hat{q}$, which we indicate with $\hat{q}^{MM}$. For bubbles expected to last between 1 and 3 periods, this value is comprised between 40 and 50%, which

---

(26) More precisely, the decision rules $u_t = -F(s_t)x_t$ are taken as given, instead of being derived from the first order conditions of the optimisation problem. These are then replaced into the Bellman equations. Going through the same steps as described in Section 2, one obtains a system of interrelated Lyapunov equations, which can be solved by iteration in the same manner as the Riccati equations therein.

roughly corresponds to attach no (or very small) weight to the exchange rate in the decision rule. For a bubble expected to last 5 periods, the robust value of q falls to around 25%. Second, the losses are asymmetric: they tend to be higher when q rises above $\hat{q}^{MM}$ than vice versa. To put it differently, overestimating the probability of a bubble can lead to worse outcomes than underestimating it. Note that a missing value on the chart indicates an infinite loss.[27] Hence, assuming that there is a bubble with a high probability $\hat{q}$ when in fact there is none (i.e. $q = 0$) can lead to a catastrophic outcome. Third, as the expected duration of the bubble increases there appear to be ever smaller variation in the losses for different realisations of q, but it remains true that overestimating the probability of moving into a bubble is considerably more costly than underestimating it. The losses tend to flatten for all values of $\hat{q}$ but especially for the initial ones.

Chart 5 refers to the case in which the concern for inflation stabilisation dominates (i.e. $\lambda=.1$).[28] Again, the 'robust' value of $\hat{q}$ is just below 50% for bubbles of expected duration up to three periods, decreasing for bubble of more extended length. And, losses tend to flatten as the expected length of the bubble rises. The most striking feature is how immediately steep the loss becomes as $\hat{q}$ increases above $\hat{q}^{MM}$ when there is in fact no bubble, $q = 0$. Therefore, in the case of quasi-strict inflation targeting it is never 'safe' to assume $q > .5$. Chart 6 show the same plots for the case in which the policymaker has a greater concern about output stabilisation than inflation stabilisation (i.e. $\lambda=3$). Again, putting a large weight on the existence of a bubble can lead to a very bad outcome when a bubble in fact does not arise. Another thing to notice is that as output stabilisation becomes more important, the losses get smaller compared to the two previous cases. Chart 7 refers to the case in which $\lambda=1$ and the bubble grows by 30% a year, instead of 10% unlike the previous charts. Even for faster growing bubbles, the general picture emerging from the other charts remains unaltered.

To recap, it appears that the a mistake in the choice of the probability has worse consequences when the policymaker is targeting inflation strictly. Indeed, the dispersion of relative losses is generally smaller as more weight is put on output stabilisation. Unless the bubble is expected to last for long, the safest strategy for a policymaker who has a strong preference for the early resolution of uncertainty is to select a probability q of between 40% and 50%. It is better to underestimate than overestimate the probability of moving into a bubble. And finally, the size of

---

[27] In this case the algorithm that solves for the value function for a given policy produces an infinite norm.
[28] This is almost strict inflation targeting. We do not report the case of strict inflation targeting because it produces too many extreme outcomes.

the bubble does not appear to alter these conclusions.

The above results about the minmax probability can be compared to those normally found in the robust control literature, in which uncertainty concerns the set of feasible parameters or models (see e.g. Svensson (2000) and the literature cited therein). As stressed by Svensson (2000), a preference for robustness often leads to choose a parameter which is on the boundary of the feasible set of models. By contrast, a desire for robustness in the present context leads, interestingly, to pick an intermediate value of the probability.

7 Conclusion

The current paper is meant to be an initial step towards understanding how policy should be set in economic environment characterised by important non linearities, such as when asset prices undergo periods of sustained and prolonged booms followed by sudden and large adjustment towards their long-run equilibrium. To the author's knowledge the proposed optimal control problem presented in the paper is new in the economic literature. One of its advantages is its relative simplicity: the problem solved by the policymaker is effectively non-linear/quadratic but maintains all the advantages of the linear quadratic framework. A general implications for monetary policy is that, at least in principle, the decision rules are state-contingent and therefore must adapt to changing economic circumstances. We apply this tool to study how monetary policy should be set in a stylised small open economy model in which the real exchange rate alternates between periods of booms and busts, namely periods in which it tends to move away from its long-run equilibrium (bubbles) and periods in which abruptly returns to it (crashes).

One limitation of the current tool is that it applies to backward looking models. Extending it to models that allow for forward-looking agents is clearly desirable as such models are the natural setting for modelling asset prices. A further extension is to make the state of the world a latent variable, which therefore need to be estimated through some filtering problem. These are the next steps in our research agenda.

References

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

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

Bank of England, 2002, *Minutes of February 2002 MPC meeting*.

Batini, N. and E. Nelson, 2000, "When the bubble bursts: monetary policy rules and foreign exchange market behaviour", mimeo, Bank of England.

Bean, C., 1998, "The new UK monetary arrangements: a view from the literature", *Economic Journal* 108 (November): 1795-1809.

Bean, C., 2003, "Asset prices, financial imbalances and monetary policy: are inflation targets enough?", BIS Working Paper 140.

Beechey, M., N. Bharucha, A. Cagliarini, D. Gruen, and C. Thompson, 2000, "A small model of the Australian macroeconomy", Reserve Bank of Australia Research Discussion Paper 2000-05.

Bernanke, B. and M. Gertler, 2000, "Monetary policy and asset price volatility", NBER Working Paper 7559.

Bernanke, B. and M. Gertler, 2001, "Should central banks respond to movements in asset prices?", *American Economic Review* 91(2), pp. 253-257.

Blanchard, O. and M. Watson, 1982, "Bubbles, rational expectations and financial markets", in P. Wachtel (ed), *Crises in the Economic and Financial Structure*, Lexington Books.

Brainard, W., 1967, "Uncertainty and the effectiveness of policy", *American Economic Review* (May) 57: 411-425.

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

Cogley, T., and T.J. Sargent, 2003, "The conquest of US inflation: learning and robustness to model uncertainty", mimeo, UC Davis and New York University.

Cecchetti, S.G., H. Genberg, J. Lipsky, and S. Wadhwani, 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., H. Genberg, and S. Wadhwani, 2003, "Asset prices in a flexible inflation targeting framework", in W.C. Hunter, G.G. Kaufman, and M. Pomerleano (eds), *Asset Price Bubbles: the Implications for Monetary, Regulatory, and International Policies*, MIT Press, pp. 427-444.

De Grauwe, P. and M. Grimaldi, 2003a, "Heterogeneity of agents, transaction costs and exchange rates", forthcoming in *Journal of Economic Dynamics and Control*.

De Grauwe, P. and M. Grimaldi, 2003b, "A theory of bubbles and crashes", mimeo, Riksbank.

Filardo, A.J., 2001, "Should monetary policy respond to asset price bubbles? Some experimental results", in G. Kaufman (ed), *Asset Price Bubbles: Implications for Monetary and Regulatory Policies*, Elsevier Science, Amsterdam, Vol. 13, pp. 99-123.

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

Froot, K., and M. Obstfeld, 1991, "Intrinsic bubbles: the case of stock prices", *American Economic Review* 81(5), pp. 1189-1214.

Goodhart, C., 2001, "Monetary transmission lags and the formulation of the policy decision on interest rates", *Federal Reserve Bank of St. Louis Review* 83, no.4 (July/August), pp. 165-182.

Hamilton, 1994, *Time Series Analysis*, Princeton University Press.

Hansen, L.P. and T.J. Sargent, 2003, *Robust Control and Model Uncertainty in Macroeconomics*, unpublished monograph (October 2003).

Judd, K.L., 1998, *Numerical Methods in Economics*, MIT Press.

Kindleberg, C., 1978, *Manias, Panics, and Crashes. A History of Financial Crises.* John Wiley and Son, New York.

Ljungqvist, L. and T.J. Sargent, 2000, *Recursive Macroeconomic Theory*, MIT Press.

Meltzer, A.H., 2003, "Rational and irrational bubbles", in W. Hunter, G. Kaufman, and M. Pomerleano (eds), *Asset Price Bubbles: the Implications for Monetary, Regulatory, and International Policies*, MIT Press, Cambridge.

Milani, F., "Monetary policy with a wider information set: a Bayesian model averaging approach", mimeo, Princeton University.

Minsky, H., 1982, "Can 'it' happen again?", *Essays on Instability and Finance.*

Orphanides, A. and V. Wieland, 1999, "Inflation zone targeting", *European Central Bank* Working Paper 8.

Sargent, T., 1999, "Policy rules for open economies: Comment", in J. Taylor (ed.), *Monetary Policy Rules*, The University of Chicago Press.

Shiller, R., 2000, *Irrational Exuberance,* Princeton University Press.

Svensson, L.E.O., "Robust control made simple", unpublished paper, October 2000.

Svensson, L.E.O., 2002, "What is wrong with Taylor rules? Using judgement in monetary policy through targeting rules?", forthcoming in *Journal of Economic Literature*.

Svensson, L.E.O., 2003, "Monetary policy and low-probability extreme events", *Festschrift for Charles Freedman*, Bank of Canada, 19-20 June 2003.

Taylor, J.B., 1993, "Discretion versus policy rules in practice", *Carnagie-Rochester Conferences Series on Public Policy* 39: 195-214.

Wadhwani, S.B., 1999, "Currency puzzles", lecture delivered at LSE, 16 September 1999.

Chart 1: Contour lines for optimal response coefficients. Strict inflation targeting $\lambda=0$**; bubble growing 20% a year. 'Bubble state' on top; 'normal state' below.**

Chart 2: Contour lines for optimal response coefficients: λ=1; bubble growing 20% a year. 'Bubble state' on top; 'normal state' below.

Chart 3: Contour lines for optimal response coefficients: λ=5; bubble growing 20% a year. 'Bubble state' on top; 'normal state' below.

4a:

4b:

4c:

4d:

Relative losses as a function of the probability chosen by the policymaker, $\hat{q}$, for different values of the true probability q. Note: λ=1. Bubble grows 10% a year. Missing values indicate ∞. Losses conditioned on 'no bubble' regime.

5a:
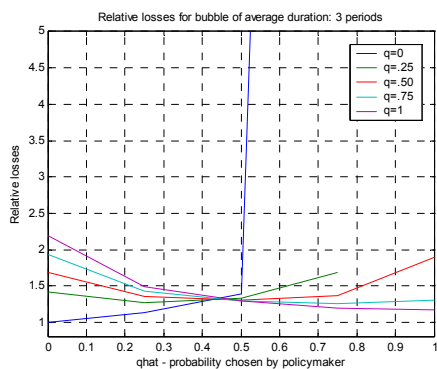
Relative losses for bubble of average duration: 1.25 periods

5b:

Relative losses for bubble of average duration: 2 periods

5c:

Relative losses for bubble of average duration: 3 periods

5d:

Relative losses for bubble of average duration: 5 periods


Relative losses as a function of the probability chosen by the policymaker, $\hat{q}$, for different values of the true probability q. Note: λ=0.1. Bubble grows 10% a year. Missing values indicate ∞. Losses conditioned on 'no bubble' regime.

6a:

6b:

6c:
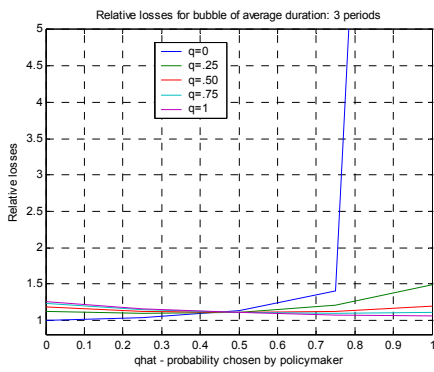
6d:

Relative losses as a function of the probability chosen by the policymaker, $\hat{q}$, for different values of the true probability q. Note: λ=3. Bubble grows 10% a year. Missing values indicate ∞. Losses conditioned on 'no bubble' regime.

7a:

Relative losses for bubble of average duration: 1.25 periods

7b:

Relative losses for bubble of average duration: 2 periods

7c:

Relative losses for bubble of average duration: 3 periods

7d:

Relative losses for bubble of average duration: 5 periods

Relative losses as a function of the probability chosen by the policymaker, $\hat{q}$, for different values of the true probability q. Note: λ=1. Bubble grows 30% a year. Missing values indicate ∞. Losses conditioned on 'no bubble' regime.