Elements of Robust Decision Theory Applied to U.K. Monetary Policy *

Juha Kilponen and Mark Salmon
Bank of Finland and Warwick Business School

August 18, 2005

Abstract

In this paper we provide an introductory discussion of several issues relating to robust policy design. We apply H_∞ methods to a standard empirical New Keynesian model of inflation and output gap and derive optimal LQG and H_∞ interest rate policy rules and compare them with the historical record in the UK over 1988 -2001. Both optimal rules are substantially more active than the historical policy record. We also investigate the importance of measurement errors on the output gap and inflation forecast. It is clear that implementing the *most* robust rule does not make economic sense but more robust rules than the LQG rule seem to coincide well with actually policy over the period when the MPC has been in place. However there is still a question as to why actual monetary policy has been less responsive than these optimal rules suggest and whether any preference for robustness is explicit within the MPC policy making process..

1 Introduction

The theory of macroeconomic policy has developed over many years from Klein and Tinbergen and through the Lucas Critique and the rational expectations revolution and yet one of the most fundamental issues facing policy makers, both then and now, that of how to cope with uncertainty, is still poorly understood theoretically. How exactly should the Monetary Policy Committee in the UK react to increased uncertainty regarding exchange rates, oil prices or measurement errors in the projected output gap or model uncertainty. Publishing fan charts showing the projected inflation uncertainty into the future does not

*We would like to thank Christos Roupakiotis for help while preparing this paper.We are also very grateful for comments received in seminar presentations at the Bank of England, The Evolving Macro Economy Conference at Warwick, Cambridge University and Cardiff Business School in particular. This research was generated under an ESRC Evolving Macro Economy grant L:138251018.

resolve how that uncertainty should be properly taken into account and reflected in the policy choice.

Economists have for many years adapted standard methods of control theory and these now enable us to numerically compute "optimal" (expectations consistent) policies in forward looking nonlinear models but this is with little or no formal recognition of the potential impact of misspecification in the econometric model and data uncertainty. The analysis of time consistency led to a strategic monetary policy games literature which largely exploited Linear Quadratic Gaussian methods and while this informed and indeed radically altered the policy debate in many ways the approach can formally only deal with additive uncertainty under a very specific class of errors following the Gaussian distribution. Given certainty equivalence, the policymaker should behave as if the *average* model were true and then "ignore" the additive uncertainty surrounding the model in the design of the optimal policy rule. Practical policy making is of course not so rigid or rule based as these theoretical constructions imply but models *are* used in the policy making process at least to generate forecasts and for scenario analysis and these models are necessarily only approximations. Policy makers are also clearly very aware that the current data they work with is likely to be inaccurate. Brainard (1967) using a Bayesian analysis, was one of the first to formally consider the effect of model uncertainty, in the particular form of parameter uncertainty, and showed how this should lead to a more cautious policy. However caution in the face of uncertainty is not always *a priori* intuitively sensible. Consider for instance the case of global warming; is it better to act cautiously and wait until the uncertainty is resolved or act more aggressively now given that if we do not there may simply be no time left, given the natural dynamics, to implement an appropriate policy?

The question of how to address uncertainty has been a major concern for policy makers, see for instance Charles Goodhart's Keynes Lecture at the British Academy (Goodhart(1999)) where he notes that "*One of the central problems is that uncertainty is far more complex, insidious and pervasive than represented by the additive error terms in standard models*". The Bank of England and the Federal Reserve have produced a number of papers (see, Batini et al (1999), Martin(1999) Martin and Salmon(1999) Sack(1998) and Wieland(1998) that have attempted to address these concerns. In this paper we apply the tools of Robust Decision Theory to reconsider these issues in the context of recent UK monetary policy. As Goodhart notes there has been a tendency for actual policy to smooth interest rate changes giving rise to the potential criticism that the MPC acts "too little too late". We are interested in seeing how this apparent preference for caution stands up against optimal policies designed with robustness in mind.

Several papers, notably Sargent (1999),Onatski and Stock (2002) and Gianonni (2002) have already explored the use of robust methods in the monetary policy context and shown how they can lead to more aggressive interest rate decisions than actually observed. This conforms with the insight of Hansen and Sargent (2000) that a preference for robustness is like a "discount factor", suggesting that the policy maker has an incentive to pre-empt the future consequences of current uncertainty by acting aggressively today. However, as we show below, this "aggressiveness" gets less pronounced when measurement uncertainty increases and is in fact problem specific as shown by Bernhard (2002). In other words there is no reason in theory why robust rules should always be more aggressive than non-robust rules.

In the next section we provide an introduction to the ideas and tools of robust decision making before turning to their application to UK Monetary Policy[1].

2 Robust Decision Theory

Major advances in Control Theory over the past ten years have led to the development of a range of new approaches to decision making under uncertainty which have typically been associated with what is known as $\mathcal{H}_\infty$ theory or more generally *Robust Decision Theory.* Critically $\mathcal{H}_\infty$ methods free us from the limitations of needing to assume additive Gaussian uncertainty and enable us to design policy rules that perform well under quite arbitrary misspecifications in the assumed nominal model and input data[2]. Essentially robust policy rules are designed to perform well across a set of alternative models *not just a single linear model* as in the classic LQG framework. The basic idea in this new approach is to search for a "safe" policy, in other words one which retains good performance under a presepcified set of potential misspecifications in the original decision problem. Zames (1981) recognized that the goal of guaranteeing some minimum performance in the presence of model uncertainty could be achieved by analysing traditional control problems in the $\mathcal{H}_\infty$ norm, rather than the standard linear quadratic norm and this observation induced the subsequent revolution in control theory.

We shall start by describing the familiar Linear Quadratic Gaussian (LQG) control problem and show how it can be reformulated, first into the minimisation of an H_2 norm problem and then be seen in a robust control setting. We then demonstrate how H_∞ control techniques can be used to solve for robust rules and compare the resulting robust rules with LQG rules.

The stability properties of different policy rules can be compared using the H_∞ norm which essentially tells us how close we are to instability for a given range of perturbations from a nominal model and then the so called small gain theorem states that the inverse of the maximum value of the H_∞ norm provides a conservative measure of (stability) robustness of the system under the policy rules. The original development of H_∞ theory was carried out in the frequency domain and this approach can still provide substantial intuition although the theory has now been fully developed in the time domain within the state space framework. We can for instance evaluate the performance of the different rules

[1] A considerably more detailed introduction to these methods can be found in the manuscript by Hansen and Sargent(2002).

[2] The relevant notion of robustness is simply that the policy should be designed so that the output is relatively insensitive to misspecification or disturbances while retaining reasonable performance and stability.

using a frequency decomposition of the expected losses. This frequency decomposition summarises how different attitudes towards model misspecification and measurement errors affect the average performance of the system. It also allows us to separate the contribution of the disturbance at each frequency to the total loss as well as assess the average expected loss. We shall use both time and frequency domain descriptions below.

2.1 Preliminaries[3]

Assume that the economy facing the policy maker evolves according to the law of motion given by

$$\dot{x}(t) = Ax(t) + B_1 v(t) + B_2 u(t) \tag{1}$$

where $x(t)$ is the vector of state variables, $v(t)$ a vector of disturbances (for the moment random), $u(t)$, the policy variables and $\dot{x}(t) = \frac{dx}{dt}$. The policymaker seeks to stabilize the economy using a state feedback rule, denoted by

$$u(t) = Fx(t) \tag{2}$$

where F contains the policy response parameters, yet to be determined[4]. Let $z_1(t)$ denote a vector of target variables and assume that there is a mapping from the state variables into observable target variables of the form,

$$z_1(t) = C_1 x(t) + D_{12} u(t) \tag{3}$$

C_1 and D_{12} are selection matrices which choose the appropriate vector of output variables from among the state and control variables. There is also a vector of measured output variables $z_2(t)$,

$$z_2(t) = C_2 x(t) + D_{21} v(t) \tag{4}$$

C_2 is a selection matrix and D_{21} is the impact matrix of measurement disturbances. Throughout we assume $B_1' D_{21} = 0$, that is, we assume that system disturbances and measurement disturbances are not correlated. We assume also that $D_{12}' C_1 = 0$. This implies that there are no cross-terms between state variables and control variables – but this can be relaxed. To make our assumptions clear the full state space system matrices are of the general form;

$$\begin{bmatrix} A & B \\ C & D \end{bmatrix} = \begin{bmatrix} A & B_1 & B_2 \\ C_1 & D_{11} & D_{12} \\ C_2 & D_{21} & D_{22} \end{bmatrix}$$

[3] In order to facilitate the further application of these robust methods we will follow the notation used in the MATLAB robust control toolbox and μ -synthesis and analysis toolbox.We discuss continuous time. $\mathcal{H}_\infty$ techniques for simplicity but discrete methods are almost entirely equivalent; a bilinear transformation can be used to transform discrete time systems to continuous time.

[4] The feedback rule could be driven by (robustly) estimated state variables when the full state vector needs to be reconstructed.

with the two inputs to the state evolution being the disturbances $\nu(t)$ and the control variables $u(t)$.

Under the assumption that $\{v(t)\}$ are *i.i.d.* disturbances, certainty equivalence holds in the standard LQG problem: in other words the optimal policy does not depend upon the matrix B_1 and hence the covariance matrix of the shocks. Moreover, the impact coefficients F in the optimal rule remain the same, whether D_{12} is singular or not due to separation principle. When D_{12} is different from zero, however, the state vector $x(t)$ is evaluated at an estimated state $\widehat{x}(t)$, where $\widehat{x}(t)$ can be calculated using Kalman filter, which forms recursive mean square error forecasts of the current state vector given current and past information. Adding uncorrelated measurement errors into the model has only a minor impact on the ordinary LQG problem.

On the contrary, in the robust control setting B_1 becomes the impact matrix for model *misspecifications* of an arbitrary form, contributing to the solution of the robust control problem. The disturbances $v(t)$ can be interpreted as *unstructured* model specification errors.

Certainty equivalence, in its standard sense, now breaks down as the robust control problem allows the disturbances $v(t)$ to feed back onto the state vector $x(t)$[5]. So in the robust control setting the nature of the measurement errors critically affects the design of the robust rule. A Robust Kalman filter can then used to recursively estimate the state vector.

Equations (1), (3) and (4) lead to an augmented state-space model $\mathcal{S}_a$ as follows

$$\mathcal{S}_a : \begin{cases} \dot{x}(t) = Ax(t) + B_1 v(t) + B_2 u(t) \\ z_1(t) = C_1 x(t) + D_{12} u(t) \\ z_2(t) = C_2 x(t) + D_{21} v(t) \end{cases} \tag{5}$$

When the control rule $u(t) = Fx(t)$ is embedded, the model can be written compactly in the reduced form as

$$\mathcal{S}_c : \begin{cases} \dot{x}(t) = A_c x(t) + B_c v(t) \\ z(t) = C_c x(t) \end{cases} \tag{6}$$

where A_c, B_c and C_c are reduced form counterparts of the A, B and C matrices. Under some specific, but not too restrictive circumstances, the matrix D does not in fact appear in the reduced form of the state space model. When measurement errors enter into the model, the states $x(t)$ are replaced with their estimated counterparts $\dot{\widehat{x}(t)}$.

The reduced form model can then be further represented by means of the closed loop transfer matrix $\mathcal{T}(s)$

$$\mathcal{T}(s) = C_c(sI - A_c)^{-1} B_c \tag{7}$$

[5] An extended notion of certainty equivalence can however be stated see Hansen and Sargent (2005) or the minimax certainty equivalence of Basar-Bernhard (1995).

where the transfer matrix $\mathcal{T}(s)$ compactly describes the mapping from the shocks $v(t)$ to the output variables $z(t)$ such that

$$z(t) = \mathcal{T}(s)v(t)$$

2.2 LQG and Robust Control

2.2.1 H2 -Norm

In this section we give a definition of the H_2 norm and the H_∞ norm. We discuss, how they can be computed both in the time domain and the frequency domain. We also discuss equivalence between the so called H_2 -problem and LQG problem.

Definition 1 *(H_2 system norm in the time domain) For a proper transfer matrix given in (7), the H_2 norm of the system is defined as*

$$\|\mathcal{T}\|_2 = \left(\frac{1}{2\pi} tr \int_{-\infty}^{\infty} g(t)g(t)'\right)^{\frac{1}{2}}$$

where tr is the trace operator and $g(t)$ is the system's impulse response defined as

$$g(t) = \begin{cases} C_c e^{A_c t} B_c, \ t \geq 0 \\ 0, \ t < 0 \end{cases}$$

Minimisation of the $\|\mathcal{T}\|_2$ norm, as presented above implies that the policymaker is only concerned with the magnitude of the impulse responses. In the robust control setting, as will been seen later, the policymaker is also concerned with the shape of the impulse response function. The H_2 norm can also be calculated directly from the state-space representation of the system (6). It can be shown that

$$\|\mathcal{S}\|_2^2 = tr(C_c' C_c L_c) \tag{8}$$

where $L_c = L_c' > 0$ is a unique solution to the Lyapunov equation

$$A_c L_c + L_c A' + B_c B_c' = 0 \tag{9}$$

Equation (8) can be best understood by noticing that L_c, which in the literature is often denoted as controllability grammian, is actually the correlation matrix of the state variables when the exogenous disturbances are assumed to be uncorrelated white noise errors. To see this, notice that controllability grammian is defined as

$$L_c = \int_0^{\infty} e^{A_c t} B_c B_c' e^{A_c t} dt \tag{10}$$

which resembles the correlation matrix of the stationary random vector process.

Given these, we are ready to develop the equivalence between the minimisation of the H_2 norm and more familiar LQG cost function, given as

$$J = E\int_0^\infty [z_1(t)'Qz_1(t)]\,dt = \int_0^\infty E\,[z_1(t)'Qz_1(t)]\,dt \tag{11}$$

The matrix Q in (11) is a positive definite weighting function, chosen to quantify the relative importance of the target variables. The cost can be calculated in steady state as

$$J = E\,(z_1(t)Qz_1(t)) = tr\{Q\Sigma_y\} \tag{12}$$

where Σ_y is a steady-state correlation matrix of the output. J can be further decomposed into

$$J = E[x(t)'\bar{Q}x(t) + u(t)'Ru(t)]$$

under the assumption that there are no cross-terms between the state variables and control variables. R is positive semi-definite weighting matrix, which attaches a penalty for the control variable u.

The steady-state correlation matrix of the output, Σ_y, can be calculated directly from the states of the reduced form system, using a diagonal selector matrix C_c

$$\Sigma_y = C_c\Sigma_x C_c' \tag{13}$$

where the state correlation matrix Σ_x is the unique solution to the following Lyapunov equation

$$A_c\Sigma_x + \Sigma_x A_c' + B_cB_c' = 0 \tag{14}$$

Substituting (13) into (12) then yields

$$J = tr\{QC_c\Sigma_x C_c'\} = tr\{QC_c'C_c\Sigma_x\} \tag{15}$$

due to the fact that the trace is invariant to cyclic permutations. Recall then that

$$\|\mathcal{S}\|_2^2 = tr(C_c'C_cL_c)$$

So the H_2 norm ($\|\mathcal{S}\|_2$) of the system $\mathcal{S}$ is a square root of J in equation (15) provided the weighting matrix $Q = I$ and $S_w = I$. In most cases, however, this is not the case. In order to state the LQ problem as a minimisation of an H_2 -norm problem, when in the original problem $Q \neq I$, $R \neq I$ and $S_w \neq I$, we simply need to use appropriate weighting matrices and partitioning.

The following H_2 problem statement is then equivalent to the minimisation of the linear quadratic loss function (11) subject to the law of motion of the economy (1).

Problem 2 ($\mathbf{H_2}$) *Find a stabilising feedback rule $u(t) = Fx(t)$ which minimizes the H_2 norm of the reduced form system*

$$\mathcal{S}_{H_2} : \begin{cases} \dot{x}(t) \;=\; Ax(t) + B_1\sqrt{S}v(t) + B_2u(t) \\ z_1(t) = \sqrt{Q}x(t) + \sqrt{R}u(t) \\ z_2(t) = C_2x(t) + \sqrt{S}v(t) \end{cases} \tag{16}$$

S refers to spectral density matrices of the errors. Throughout, we assume that errors have flat spectral densities, that they are uncorrelated and we set $S = I$.

Since squaring is monotonic, minimising the LQ cost function is equivalent to minimising the reduced form system H_2 norm. This equivalence is extremely useful as it implies that the H_2 norm can be used for performance analysis in the frequency domain as well[6].

Definition 3 *(H_2 system norm in the frequency-domain) For a stable transfer function given in (7),the H_2 norm of the system in the frequency domain is given by*

$$\|\mathcal{T}\|_2 = \left(\frac{1}{2\pi} tr \int_{-\pi}^{\pi} G(i\omega)G(-i\omega)' d\omega \right)^{\frac{1}{2}}$$

where ω is a point in the frequency range $(-\pi, \pi)$.

Plotting $tr\,(G(i\omega)G(-i\omega)')$ over the frequency range shows the contribution of the disturbance to the total loss at frequency ω. One fundamental difference between robust control and linear quadratic or H_2 control is that a robust decision maker does not assign equal weighting across frequencies, as would a conventional LQ policymaker. In the time domain this means that a robust policymaker is also concerned about shape of the impulse response function. In effect a robust policy maker assigns more weight to the "worst " frequency and it is this that makes H_∞ methods more robust but more delicate in design.

2.2.2 Robust Control

We are now in a position to set up the H_∞ norm as a cost function where the minimisation of the H_∞ norm can be seen as a problem of minimising the loss under some form of "worst-case" scenario. We first provide definitions of the H_∞ norm, both in the time domain and frequency domain assuming analyticity inisde the unit circle.

Definition 4 *(time domain) the H_∞ - norm can be defined as*

$$\|\mathcal{S}\|_\infty = \sup_{v \neq 0} \frac{\|g(t) \otimes v(t)\|_2}{\|v(t)\|_2}$$

where $g(t) \otimes v(t)$ is the convolution of the input with the impulse response matrix, yielding the time-domain system output.

Definition 5 *(frequency domain) The maximum gain over all frequencies is given by the H_∞ norm.*

$$\|\mathcal{T}\|_\infty = \sup_{\omega} \overline{\sigma} \left(\mathcal{T}\,(j\omega) \right)$$

where $\overline{\sigma}$ denotes maximum singular value over the frequency range ω.

[6] Sargent and Hansen (1999) (2005), Kasa (2000) and Tetlow and Muhlen (2000) discuss the frequency domain approach.

The H_∞ norm can then be calculated directly from the transfer function as follows

$$\sup_{\omega} \overline{\sigma} \left(\mathcal{T} \left(s \right) \right) = \sup_{\omega} \overline{\sigma} [C(sI - A]^{-1} B] \tag{17}$$

An alternative interpretation of the robust control problem is as a game between two participants; the policymaker who attempts to find a stabilising policy rule which minimises his loss while a malevolent fictitious opponent is choosing his inputs, the disturbances , so as to maximise the policymaker's loss.[7]

In order to state this min-max problem formally, let the economy evolve according to (1). The solution to the differential game consists of simultaneously finding an optimal trajectory for the policy variable, $\{u(t)\}_{t=0}^{\infty}$, and a trajectory for the worst-case disturbance input $\{v(t)_{worst}\}_{t=0}^{\infty}$. The solution to this problem can be found by searching for a saddle point of an objective function

$$J(u, \nu) \tag{18}$$

where u is the choice variable of the policymaker and ν is a choice variable of the malevolent opponent. If J is interpreted as a "loss function", we are interested in finding an optimal feedback rule u^* which minimises J and v^* which maximises J simultaneously. Choosing the loss function appropriately, it is possible to apply standard equilibrium concepts, notably a Nash equilibrium, to the solution of above game. A quadratic loss function will do, as long as we ensure that the loss function is convex in u and concave in v[8]. We next briefly develop this rule starting from the robust control problem itself.

In the robust control problem the policymaker seeks to minimise the supremum of the H_∞ norm with respect to the linear law of motion of the economy. This H_∞norm is defined as

$$\|\mathcal{S}\|_\infty = \sup_{\|v\|_2 \neq 0} \frac{\|z\|_2}{\|v\|_2} \tag{19}$$

In (19) $\|z\|_2$ and $\|v\|_2$ denote L_2 vector norms of real valued vectors z and v and sup denotes supremum. Vector z contains the linear (or non linear) combination of target variables that the policymaker seeks to stabilise based on his decision variables, u, and v contains the unobservable components of the disturbance.

The robust control, or $\mathcal{H}_\infty$ control problem is then to find a feedback policy rule, $u = u(z)$, which minimizes this H_∞ norm given some bound $\gamma > 0$,

$$\|\mathcal{S}\|_\infty = \sup_{\|v\|_2 \neq 0} \frac{\|z\|_2}{\|v\|_2} < \gamma \tag{20}$$

[7] See Basar and Bernhard (1995) or Green and Limebeer for books that develop the game theoretic interpretation to H_∞.

[8] Formally we require concavity in ν but because u is allowed to use state feedback the requirement in u is much less strict than convexity, See Basar and Bernhard (1995).

Essentially (20) describes the effect of the disturbance ν on the output z under the closed loop policy rule and by bounding the effect of ν on the output through γ we achieve a degree of robustness. This norm bounded uncertainty thus describes a set of alternative models (misspecifications) or disturbances of a completely unstructured form as long as they are contained within the γ bound. This unstructured description to the uncertainty facing the policy maker holds the real power of the H_∞ approach and notice that this is achieved in a deterministic formulation of the effect of the uncertainty on the output. Minimisation of the H_∞ norm implies that the policymaker minimises his loss $\|z\|_2$ recognising and explicitly taking into account the impact of the potentially destabilising (deterministic) inputs arising from misspecifications and disturbances to the system as captured in v. Remarkably, as demonstrated by Glover and Doyle (1988), it turns out that this optimal robust feedback policy rule based on the minimisation of the bounded H_∞ norm coincides with the optimal policy in a linear exponential risk-sensitive criterion LEQG (see Whittle(1981) and (2002)). If we square both sides of (20)

$$\|\mathcal{S}\|_\infty^2 = \sup_{\|v\|_2 \neq 0} \frac{\|z\|_2^2}{\|v\|_2^2} < \gamma^2$$

we can see that in order for a supremum to satisfy a strict inequality above, the term $\frac{\|z\|_2^2}{\|v\|_2^2}$ must be bounded away from γ^2 so for some $\varepsilon > 0$ we can write

$$\begin{aligned} \frac{\|z\|_2^2}{\|v\|_2^2} &\leq \gamma^2 - \varepsilon^2 \\ &\Leftrightarrow \\ \|z\|_2^2 - \gamma^2 \|v\|_2^2 &\leq -\varepsilon^2 \|v\|_2^2 \end{aligned} \tag{21}$$

When the inequality (21) holds, so does the strict equality (20) for all disturbances v and for some $\varepsilon^2 > 0$. Consequently, the left hand side of (21) can be used as an objective function in the dynamic game between the two players,

$$J_\gamma = \|z\|_2^2 - \gamma^2 \|v\|_2^2 \tag{22}$$

This objective function, in fact, provides the link to the stochastic risk-sensitive decision theory, developed and studied in Whittle (1981). The parameter γ^2 in (22) in effect describes the policymakers' attitude towards uncertainty or his desire for robustness- risk in the LEQG problem. It permits an interpretation of risk sensitivity parameter when the policy rule results from a saddle point solution to (22). That is, the optimal robust feedback policy rule (u^*) and the most de-stabilising deterministic input (v^*) are solutions to

$$\begin{aligned} u^* &\in \arg\min_u \{J_\gamma(u, \nu)\} \\ v^* &\in \arg\max_v \{J_\gamma(u, \nu)\} \end{aligned}$$

In particular, the solution pair (u^*, v^*) is a saddle point when it fulfils the following inequalities

$$J_\gamma(u^*, v) \leq J_\gamma(u^*, v^*) \leq J_\gamma(u, v^*),\ \forall\ v \in \mathcal{V},\ u \in \mathcal{U} \tag{23}$$

So $J_\gamma(u, v)$ has a minimum with respect to the policy variable u and a maximum with respect to v at the point (u^*, v^*).

All that has changed with respect to the original optimisation problem is that an unobserved disturbance input v is introduced into the objective function directly: v is simply another control variable which is penalised by an uncertainty preference factor $\gamma^2 > 0$. The policymaker now plays a mental game against the fictitious opponent: While the policymaker wants to minimise J_γ by choosing some u^*, the opponent wants J_γ to be maximised. The policy rule that results from this game is equivalent to the min-max policy rule. Standard methods of solving linear quadratic optimisation problems can now be applied; solving the optimisation problem using (22) as a loss function and a law of motion of the economy (1) as a constraint[9]. When the state space model is appropriately defined, as in (16), the resulting robust rules can be compared to outcomes of linear quadratic gaussian rules or equivalently the minimisation of the H_2 norm, as given above (2). MATLAB provides a variety of computational tools for solving for the H_∞ policy rule.

2.3 Stability and robustness

As discussed above, H_∞ techniques can be seen as providing an alternative to LQG methods or H_2, when policymakers want to design rules that perform well under unstructured uncertainty. A policy rule which *maximises robustness* can be found by searching for the infimum of γ such that the solution to the problem exists and the policy rule *stabilises* the reduced form system This is due to the *Small-Gain Theorem*, originally due to Sandberg and Zames (1981).

The Small-Gain Theorem states that the reduced form model, given for instance in (6) remains stable with respect to all possible unstructured perturbations Δ of size $\|\Delta\|_\infty < 1/\gamma$ as long as $\|\mathcal{T}\|_\infty < \gamma$. Therefore, the infimum of γ, or γ^*, effectively ensures that the policy rule is robust to the largest possible unstructured perturbations before destabilising the economy. If we interpret, γ as a preference for robustness we may decide that we are not really interested in the *maximum* degree of robustness that retains stability and in fact we may wish to move back from the optimal H_∞ rule provided by γ^* and thereby achieve less robustness. As we let γ increase to infinity we recover the LQG solution. By letting γ increase from γ^* we are reducing the set of alternative models/ misspecifications/disturbances the rule will be robust against. In this way we shape the degree of robustness desired and in fact the "worst case" misspecification may be very close to the nominal model. This enables the method to consider both behaviour and policy responses to minor deviations from say rational expectations equilibrium but it can also be used to construct policies

[9]See for instance Sargent (2000), Zhou et. al. (1996).

for dramatic uncertainties in which the relevant economic structure can lie a considerable distance from the nominal model. None of this of course obviates the need to construct models that are as accurate as possible since economic welfare will decline through the use of policy that is unecessarily robust.

2.4 Performance

Whilst stability criteria can be stated explicitly - conditional on the judgement of uncertainty of course - performance robustness must be evaluated with respect to some a priori performance criteria. If the H_∞ norm is used as a performance criterion, then it must be possible to express performance in terms of bounds on the target variables. This requires that the disturbance inputs and output variables in the state space model given in (16) are adequately normalized. If the target variable(s) fail the *a priori* criteria, then that policy rule does not perform robustly. However, such a policy rule can still be robustly stable. Many of the existing economic examples have focused on robust stability as a performance criterion (see Onatski and Stock (2002), Tetlow and von Muhlen (2000) etc). In economic models robust performance or performance in general is perhaps best analysed by using the original H_2 norm as an economic criterion function, rather than H_∞ norm. However assessment of relative performance is not straightforward, even when the H_2 norm is used. This is due to the fact that the reduced form state evolution equation in the robust control problem involves a state-feedback law for the "worst-case" disturbances. Further, when measurement errors are present, the standard separation principle breaks down and strictly speaking-does not allow us to separate the control and robust estimation problems.

To see this more formally and at the same time emphasise the differences between robust control and ordinary LQG problem, we can see [10] that the solution to the robust control problem yields the following laws of motion.

$$\widehat{\dot{x}}(t) = A\hat{x}_t + B_1\hat{v}_{worst}(t) + B_2u(t) + Z_\infty L_\infty(C_2\hat{x}(t) - z_2(t)) \quad (24)$$

$$u(t) = -B_2' X_\infty \hat{x}(t) \quad (25)$$

$$\hat{v}_{worst}(t) = \gamma^{-2} B_1' X_\infty \hat{x}(t) \quad (26)$$

where

$$L_\infty \; : \; = -Y_\infty C_2',$$
$$Z_\infty \; : \; = (I - \gamma^{-2} Y_\infty X_\infty')^{-1}$$

X_∞', and Y_∞ are the solutions to corresponding Ricatti equations[11]. Substituting the feedback rules for $u(t)$ and $v(t)$ into (24) yields.

$$\widehat{\dot{x}}(t) = \underbrace{(A - B_2B_2'X_\infty + \gamma^{-2}B_1B_1'X_\infty)}_{A_\gamma}\hat{x}(t) + Z_\infty L_\infty(C_2\hat{x}_t(t) - z_2(t)) \quad (27)$$

[10] See Zhou, Doyle and Glover (1996) for details.

[11] Solution involves following two Hamiltonian matrices

Equation (27) gives a law of motion for the economy under the "worst-case disturbances" and under the robust rule. Matrix A_γ therefore embeds both a feedback policy rule, as well as feedback rule for the worst-case disturbances.

The law of motion in the standard H_2 optimal control problem, in turn, can be written as

$$\begin{aligned}
\widehat{\dot{x}}(t) &= A\hat{x}(t) + B_2 u + B_1 v(t) + G_{LQ}(C_2\hat{x}_t(t) - z_2(t)) && (28)\\
u(t) &= F_{LQ}\hat{x}(t) && (29)\\
F_{LQ} &= -B_2' X_{LQ} && (30)\\
G_{LQ} &= Y_{LQ} C_2' && (31)
\end{aligned}$$

$\hat{x}(t)$ is the optimal estimate of x, obtained by using Kalman filter. X_{LQ} and Y_{LQ} are solutions to corresponding Riccati equations[12].

Substituting (29) into (28) delivers the reduced form state equation,

$$\widehat{\dot{x}}(t) = \underbrace{(A - B_2 B_2' F_{LQ})}_{A_{LQ}}\hat{x}(t) + B_1 v(t) + G_{LQ}(C_2\hat{x}_t(t) - z_2(t)) \qquad (32)$$

Given these, there are basically two different ways to assess a relative performance of LQG rule and robust rule[13]. The first possibility is to assess losses under an approximating model such that

$$\widehat{\dot{x}}(t) = (A - F_i)\hat{x}(t) + B_1 v(t) + G_i(C_2\hat{x}_t(t) - z_2(t)) \qquad (33)$$

where F_i corresponds either to the LQ rule (F_{LQ}) or the robust rule (F_γ). G_i corresponds either to the Kalman filter or the robust Kalman filter as given in (24) by the term $Z_\infty L_\infty$. At first sight this may seem an intuitive and natural way to proceed: we just set the worst case disturbances to zero and assume that the approximating model usually prevails. However, strictly speaking, in the robust control setting the stochastic separation principles does not allow us to separate control and estimation. This means that the robustified Kalman gain $Z_\infty L_\infty$ makes sense only under (27).

$$\begin{aligned}
H_\infty \quad &: \; = \begin{bmatrix} A & \gamma^{-2}B_1B_1' - B_2B_2' \\ -C_1'C_1 & -A' \end{bmatrix}\\
J_\infty \quad &: \; = \begin{bmatrix} A' & 0 \\ -B_1'B_1 & -A' \end{bmatrix} - \begin{bmatrix} C_2 \\ B_1'D_{21} \end{bmatrix} \begin{bmatrix} D_{21}'B_1' & C_2' \end{bmatrix}
\end{aligned}$$

such that $X_\infty = Ric(H_\infty)$ and $Y_\infty = Ric(J_\infty)$.

[12] Solution involves following two Hamiltonian matrices

$$\begin{aligned}
H_{LQ} \quad &: \; = \begin{bmatrix} A & -B_2B_2' \\ -C_1'C_1 & -A' \end{bmatrix}\\
J_{LQ} \quad &: \; = \begin{bmatrix} A' & 0 \\ -B_1'B_1 & -A' \end{bmatrix} - \begin{bmatrix} C_2 \\ B_1'D_{21} \end{bmatrix} \begin{bmatrix} D_{21}'B_1' & C_2' \end{bmatrix}
\end{aligned}$$

such that $X_{LQ} = Ric(H_{LQ})$ and $Y_{LQ} = Ric(J_{LQ})$.

[13] Criteria to assess performance is always quadratic loss or equivalently H_2 -norm.

A second possibility is to make a performance comparison directly between the approximating model under the LQG rule (33) and the worst case model under the robust rule as given in (27). We opt for the latter but note that we are actually comparing two "different models".

3 UK Monetary Policy

3.1 New Keynesian Model

We adopt a stripped down version of a New Keynesian form of policy model, extensively studied in the recent literature on monetary policy rules. The New Keynesian model consists of two aggregate relationships with direct foundations in a dynamic general equilibrium theory with temporary nominal price rigidities[14].

The first aggregate relationship links inflation to the real output gap as follows

$$\pi_t = \rho E_t \pi_{t+1} + (1-\rho)\pi_{t-1} + \lambda y_t + \epsilon_t \tag{34}$$

where π_t is the inflation rate, y_t is the output gap, $E_t \pi_{t+1}$ is expected inflation of period $t+1$ conditional on information available at time t and $0 \leq \rho \leq 1$. The $\{\epsilon_t\}$ are serially uncorrelated innovations.

The parameter ρ, describes the degree to which expectations are forward looking. Setting $\rho = 1$, implies that expectations are purely forward looking so that current inflation is determined with no reference to past inflation. Many authors assume that costs of adjustment and overlapping price and wage contracts generate some inertia in inflation, so that ρ will be less than one. At the other extreme, for instance Ball (1999) and Svensson (1997), it is assumed that expectations are purely backward looking. Empirically estimated values for ρ vary greatly. Rudebusch (1999) for instance suggests that based on the recent empirical studies with U.S. data a 90 % confidence interval for ρ would range from 0 to 0.6.

Although it would interesting to study the performance of robust and LQG rules under parametric uncertainty regarding ρ, we set $\rho = 0$ initially[15] Thus, our base model reduces to a backward looking model for inflation

$$\pi_t = \pi_{t-1} + \lambda y_{t-1} + \epsilon_t \tag{35}$$

The estimated counterpart for U.K. during the period of 1988:1-2000:2 delivers the following simple relationship

[14] See for instance Roberts (1995), Woodford (1996), Goodfriend and King (1997), Walsh (1998), Clarida, Gali and Gertler (1999).

[15] This is not completely implausible on empirical grounds. For instance Fair (1993) and Gruen, Pagan and Thompson (1999) obtain very small values for ρ.

$$\Delta\pi_t = \underset{(0.02)}{-0.02} + \underset{(0.013)}{0.05}\, y_{t-1} + \epsilon_t$$

$$R^2_{adj} = 0.50$$
$$BG = 5.59(0.04)$$
$$JB = 0.85(0.65)$$
$$S.E. = 0.20$$

$BG(4)$ and JB are the Breusch Godfrey and Jarque Bera statistics for heteroskedasticity and normality respectively. Standard errors are shown in parentheses. Potential output level is estimated using a Hodrick-Prescott filter with a smoothing parameter of 30,000 from monthly industrial production. The output gap was then calculated as a logarithmic difference between actual and its estimated potential level and is measured in percentages. Inflation is measured as annual percentage change in RPIX. Currently, the Bank of England uses RPIX as a preferred measure of underlying inflation where RPIX differs from RPI by excluding mortgage interest payments.

This relationship between inflation and the real output gap is very "approximate" and there is some evidence of serial correlation. However the explanatory power of the regression is reasonably high and the parameter of interest is determined accurately it would seem. This permits us to pin down a rough value for the λ parameter and allows us to keep the structure of the model as simple as possible in order to concentrate on the principal issue of robustness[16].

The second aggregate relationship in these New Keynesian models, links the real output gap to its expected value, nominal interest rates and an equilibrium real rate

$$y_t = E_t y_{t+1} - \beta(i_t - E_{t+1}\pi_t - \overline{r}_t) + \nu_t \tag{36}$$

where y_t is real output gap, $E_t y_{t+1}$ is expected real output gap at time $t+1$ conditional on information at time t, i_t is the short-term nominal interest rate and $\overline{r}$ is equilibrium or natural rate of interest. In the theoretical model of Woodford (1996), this natural rate of interest is interpreted as a steady-state value in the case of zero inflation and steady output growth. The ν_t are serially uncorrelated innovations. The empirical counterpart of (36) typically allows some form of costly adjustment or habit formation in order to generate inertia and lagged responses of output gap on the interest rate and inflation. Starting from the hypothesis that expectations are purely backward looking, and after experimenting with the U.K. data, we ended up into following specification

$$y_t = \delta_1 y_{t-1} + \delta_2 y_{t-2} - \beta(\overline{i}_{t-1} - \overline{\pi}_{t-1} - \overline{r}_{t-1}) + \nu_t \tag{37}$$

Real output gap depends upon its lagged values and the lagged real rate of interest less the equilibrium real rate of interest. The real rate of interest $\left(\overline{i} - \overline{\pi}\right)$ was calculated as a difference between a 15 month moving average of the Bank of England´s base rate and 15 month moving average of the periodic annualised

[16] Notice that to a degree we are not emphasising the need to develop a perfect econometric model since the robust methods used to construct policy below take the misspecification into account.

percentage change in RPIX. This suggests that monetary policy affects the real economy with a considerable lag. We experimented with different moving average lengths, but discovered that shorter moving averages would not deliver a significant value for β.

Estimating the equilibrium rate of interest $\overline{r}_t$ is somewhat more debatable. Some authors have argued that in the U.K, there is a strong evidence that the equilibrium rate of interest, that is, a rate of interest which keeps output growth constant and inflation at zero, has radically changed over the years. We have therefore "estimated" the equilibrium (real) rate of interest using a Hodrick-Prescott filter with a smoothing parameter (20,000). [17] The resulting estimated equation for the period of 1988:M1 to 2000:M2 yields the following relationship

$$y_t = 0.04 + \underset{(0.08)}{0.52}\, y_{t-1} + \underset{(0.08)}{0.28}\, y_{t-2} - \underset{(0.08)}{0.17}\, (\bar{\bar{i}}_{t-1} - \overline{\pi}_{t-1} - \overline{r}_{t-1}) + \epsilon_t$$

$$R^2_{adj} = 0.71$$
$$BG(4) = 1.75(0.14)$$
$$JB = 1.62(0.44)$$
$$S.E. = 0.73$$

The explanatory power of the regression is reasonably high and all the estimated parameters are relatively accurately determined. There is no evidence of serial correlation and the errors are normally distributed.

3.2 Measurement errors

As stressed at the outset monetary policymakers are not only concerned with model misspecification but with errors in the inflation forecast and estimates of the output gap. Apart from conceptual issues relating to the measurement of the output gap, setting monetary policy on the basis of the level of the estimated real output gap requires relying on a quantity that is difficult to measure accurately and the real-time measurement errors can be large, as noted by Rudebusch (1999) and many others[18]. As noted by Egginton et. al (2001), the size and number of revisions to UK macro data have also been large. A number of researchers have argued that the quality of UK macro statistics has deteriorated over time and showed a considerable downward bias in the initial measurement of key variables in the late 1980s. Policy-makers, at both HM Treasury and the Bank of England thought that this bias may have delayed the tightening of monetary policy in the late 1980s, and therefore, contributed to the inflationary Lawson boom.

In order to get some idea how large the real-time measurement errors may have been in the U.K during the estimation period, we have used a unique real-time macro data set from Eggington et. al. (2001). The data consists of a common set of macro economic variables with different vintages, which reflect

[17] Martin and Salmon (1999) detrended base rate directly.

[18] Orphanides (1999) found that real-time estimates of potential output severely overstated the economy's capacity relative to the revised estimates in the U.S. in 1970s.

the revisions and updates of the data over time. So we can use data that was actually available in 1988 rather than that which is presented today for 1988 in the revised national statistics.

Figure 1 shows the output gap, as calculated from the real-time data and from the revised data. Real-time output gap has been calculated from the first available estimates of the monthly industrial production and using the Hodrick-Prescot filter to estimate potential output. The revised output gap, in turn, has been calculated from the last revised figures of industrial production. Considerable differences can be seen in these output gap estimates throughout the whole sample period, particularly in $89-90, 91-92, 94-95, 97-98$, which emphasises the importance of this issue in practical policy making and th need for robustness.

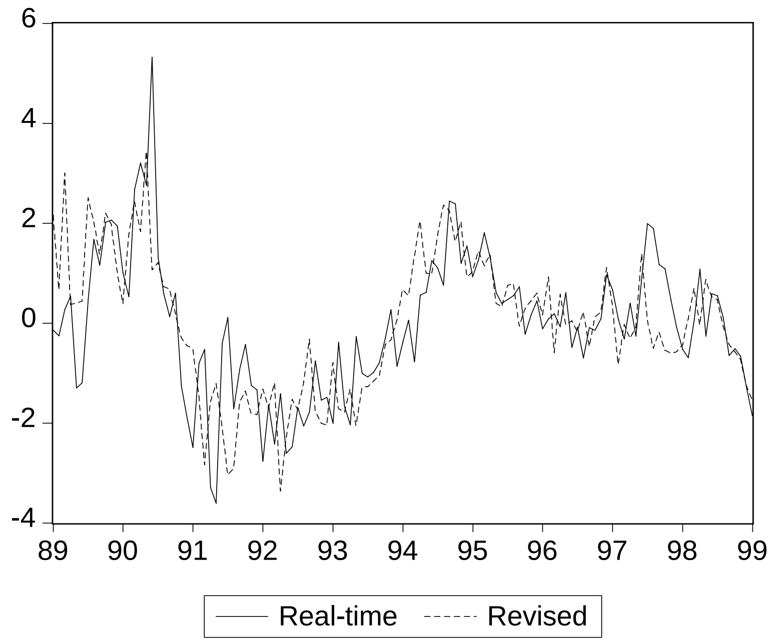

Figure 1: Output Gap in real time and after subsequent data revisions.

Figure (2) shows the distribution and summary statistics of the output gap measurement errors, calculated as a difference between real-time output gap and actual. The standard error of output-gap measurement errors during the whole sample period is 1.09.

There is also some evidence that output gap measurement errors persist over time, but we nevertheless make an assumption that the errors are purely transitory and can be modelled as white noise with variance $\sigma_{\eta t}$ and mean zero. Of course, at the time that the policymaker assesses the state of the

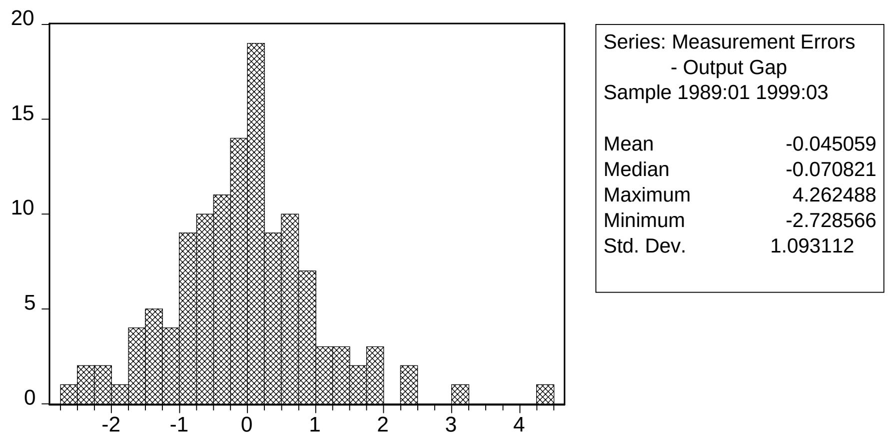

Figure 2: Distribution of Output Gap Measurement Errors and Summary Statistics
Source: Own calculations from Egginton et. al (2001).

economy, he has no indication of what the "true" measurement error actually is. However, over the course of time, we assume, that the policymaker has learned how large the measurement errors have been historically, and therefore uses that value as a proxy for "real-time-measurement error".Furthermore, we assume that measurement errors are uncorrelated with the innovations ϵ_t and ν_t; not an implausible assumption.

In practical implementation, diagonal elements of the impact matrix D_{21} contain the standard errors of the measurement errors and diagonal elements of the impact matrix B_1 contain the standard errors of innovations ϵ_t and ν_t. The assumption that $D_{21}B_1' = 0$ ensures that the innovations and measurement errors are uncorrelated.

3.3 Policy rules

We assume the monetary policymaker's task is to stabilise inflation *and* output and to achieve this goal with "minimum" effort. We also assume that the policymaker uses an interest rate feedback rule of a simple Taylor (1993) form, specifically

$$i_t = \theta_i i_{t-1} + \theta_\pi \pi_t + \theta_y y_t \tag{38}$$

where θ_i, θ_π and θ_y are the feedback coefficients, i_t is the nominal interest rate in annual terms, π_t is annual inflation rate and y_t denotes the output gap. Interest rate rules of this type (38) are perhaps only naive approximations to the actual policy making process, but as suggested initially by Taylor (1993), they do seem to describe actual outcomes relatively well. Several authors (for instance Clarida, Gali and Gertler (1999), Rudebusch (1998) and Nelson (2000) for the

U.K.) have found considerable inertia in interest rate rules and interpreted this as the policymaker's concern for interest rate-smoothing.[19]

In order to derive such a rule optimally in this setting, we assume that the policymaker's linear quadratic loss function can be expressed as

$$\begin{aligned} \|z_1(t)\|_2^2 &= \sum_{t=o}^{\infty} \beta^{t-1} L_t \\ L_t &= \lambda_\pi \left(\pi_t - \pi^*\right) + \lambda_y \left(y_t - y_n\right)^2 + \lambda_i (i_t - i_{t-1})^2 \end{aligned}$$

where β is the discount rate, L_t is the periodic loss function, λ_π, λ_y and λ_i are the preference weights associated with inflation, output-gap and the change in interest rate respectively. In the benchmark simulation, we set $\lambda_\pi = \lambda_y = \beta = 1$ and $\lambda_i = .25$. The policy maker's objective function is standard in that he is penalised by the deviation of output and inflation from their assumed target values and that the policymaker has an interest rate smoothing objective. When deriving the robust policy rule, the policymaker minimises the H_∞ norm, as defined in (19).

Several empirical studies have found that the optimal feedback coefficients from LQG optimisation are larger than those estimated empirically while others have shown that robust rules can be more aggressive than LQG rules. On the one hand, this challenges the original insight of Brainard (1967), who showed that parametric uncertainty should lead to less aggressive policy responses, when the policymakers are Bayesian. On the other, robust rules seem to be even further away from the realistic policy responses estimated from the observed data[20]. Whilst we have not estimated the policy response coefficients, we compare the performance of the different optimal rules with the actual policy visually below.

4 Results

4.1 LQ and Robust Interest Rate Rules

This section presents our initial results from applying both LQG and H_∞ to derive (optimal) policy rule for the model estimated in the previous section. When deriving the interest rate rule, we impose the condition that the rule will be a full state feedback rule in that it depends on all the state variables where our state variables are i_{t-1}, π_{t-1}, y_{t-1}, y_{t-2}.

Equations(39) and (39) provide a benchmark linear quadratic rule and the optimal robust rule ($\hat{\gamma}^* = 8.3552$).

$$i_t^{LQ} = .343 i_{t-1} + 1.28 \pi_{t-1} + .55 y_{t-1} + .17 y_{t-2} \tag{39}$$

$$i_t^{R} = .08 i_{t-1} + 7.80 \pi_{t-1} + 1.28 y_{t-1} + .34 y_{t-2} \tag{40}$$

[19] For the discussion of interest rate inertia in forward looking models, see for instance Woodford (1999)

[20] Tetlow and von zur Muehlen (2000) and Onatsi and Stock (2002) suggest, however, that rules that are robust to *structured* model uncertainty tend to be less aggressive than optimal LQG rules. They suggest also that the equilibrium responses of these rules are closer to empirically estimated rules.

where π_{t-1} is the15 month inflation moving average lagged one period, i_{t-1} is the one period lagged base rate and y_{t-1} and y_{t-2} are one and two period lagged output gaps respectively. These benchmark results are derived under the assumption that there are no measurement errors and that the policymaker attaches equal weight (unity) to inflation and the output gap as well as a small weight on interest rate smoothing (0.25). The optimal robust rule is clearly radically more aggressive and less autoregressive than linear quadratic rule.

This "optimal" robust rule would suggest an unrealistically strong inflation response, implying that the policymaker should respond to 0.1% percentage point increase in inflation rate with an approximately 80 basis points increase in the base rate, while 1 percentage point increase in output gap should trigger an approximately 150 basis point increase in the interest rate. In order to get some idea of how the linear quadratic interest rate rule performs with respect to the actual base rate we compare the two in Figure(3) and can see similar results to Sack(1998) in that the optimal LQG rule is much more active or variable than the actual policy. One potential explanation for apparent variation is that actual policy is discretised and rarely if at all moves in steps other than discrete multiples of 25 basis points. The optimal policy is continuous and the figure shows a discretised version in which the optimal rate is only moved to the nearest quarter percent.

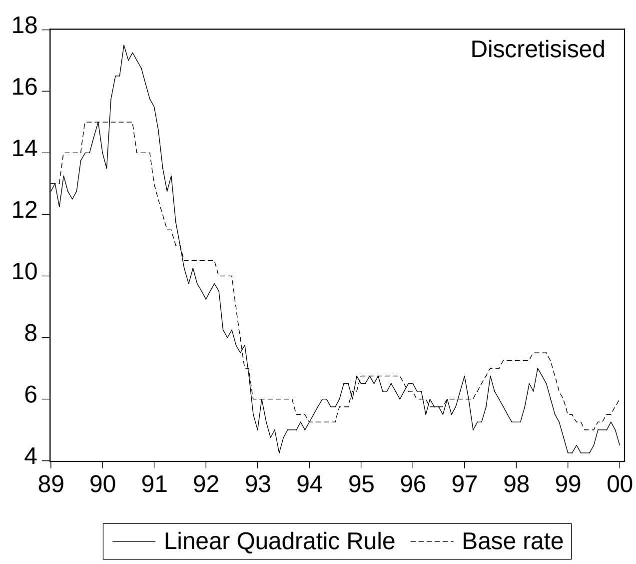

Figure 3: The discretised optimal LQG policy and the actual base.

So even the linear quadratic rule tends to be more aggressive than the actual policy (see Goodhart (1999) for more discussion as to why this may be the case). Contrasting these two policies with the optimal $\hat{\gamma}^*$ robust rule shown in (40) is meaningless since the implied responses are unrealistically high even with "realistic "measurement noises.

We next concentrate on modifying the optimal robust rule. We first experimented with changing the relative inflation preference weight, the inflation conservativeness parameter. Keeping everything else equal with respect to (40) but increasing the relative weight on inflation in the loss function to 2.25 delivers an even more aggressive inflation response. This, "inflation conservative robust rule" would suggest responding to a 0.1 percentage point increase in inflation by raising interest rates by approximately 160 basis points.

Less aggressive robust rules can be derived either by increasing a weight associated with interest rate smoothing or lowering the weight with respect to inflation. For instance, increasing the penalty on interest rate smoothing to1, and at the same time increasing inflation measurement errors would deliver an inflation response coefficient of around 1.50 (Taylor's nominal value). Whilst increasing inflation measurement errors seems reasonable, it would be hard to convince many of the need for such a high preference for interest rate smoothing.

We next study in some detail the effects of measurement errors. As argued above,the output gap is, in particular, difficult to estimate accurately in real-time. So we have assumed that the standard error of output gap measurement error is 1.09, as estimated from the real-time data above. We assume also a small standard error for inflation measurement errors (0.1). This produces the following optimal robust rule using the robust Kalman filter to estimate the output gap;

$$i_t^R = .32i_{t-1} + 7.48\pi_{t-1} + 0.5y_{t-1} + .5y_{t-2} \tag{41}$$

The introduction of the output gap measurement error quite radically reduces the output gap response coefficient and also brings down the inflation response coefficient slightly. Interestingly, the robust rule also shows more inertia as the measurement noise increases. The panels in Figure(4) display these effects in more detail and Figure (5) shows the effects of errors in the inflation forecast.

Figure(5) indicate the strong effect that recognising measurement errors in the inflation forecast has on the feedback coefficient on inflation in the robust rule. There appears to be very little effect on the degree of feasible attainable robustness as shown by the value of γ as the measurement error increases.

Nevertheless, the policy rule (41) is still unrealistically aggressive and we therefore turn to analyse rules where the preference parameter for robustness is allowed to depart from its lowest attainable value. In other words we look for less robust rules but those that still stabilise the economy.

4.2 Departing from Robustness

The optimal γ^* robust rules analysed in the previous section tend to yield unrealistically high inflation responses when compared with the actual behaviour of the base rate. We could of course try to search for a better fitting rule by optimising the inflation and output weights, but our motivation lies somewhere else

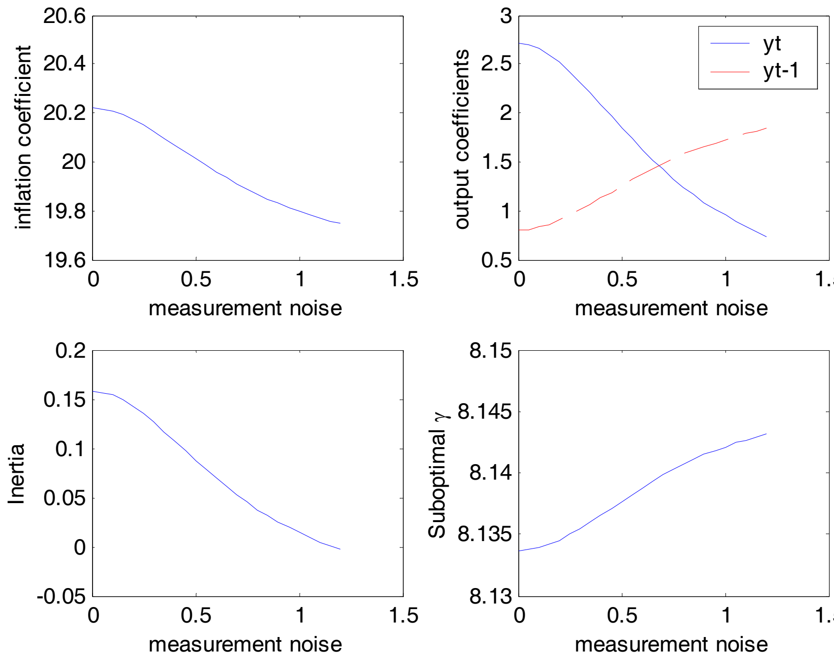

Figure 4: Effects of increasing measurement error on the output gap on the robust Taylor Rule.

(that approach is taken in Salmon and Weston (2001)). It might be unrealistic to assume that the policymaker would actually want to "maximise" robustness, in the sense of finding the smallest attainable γ, or γ^*, for which the stabilising solution exists. Essentially the difficulty we now face is how to specify a formal preference for robustness alongside the normal objectives of policy.

Without a utility maximising framework choosing a value for γ is rather arbitrary however it turns out that in practice γ values close to γ^*, even small departures can drastically change the optimal reaction coefficients in the implied robust policy rule. As we mentioned above at the another extreme, when γ approaches infinity, $\mathcal{H}_\infty$ delivers the LQG rule.

In order to demonstrate how the degree of robustness alters the behaviour of the policymaker we let the robustness parameter γ increase in small increments of .05, starting from γ^* and show how the policymaker's response coefficients change. Figure (6) shows how the inflation response coefficient and output gap response coefficient change as a function of γ. Preference parameters are as in the benchmark case, $\sigma_{\eta_\pi} = .1$ and $\sigma_{\eta_y} = 1.09$.

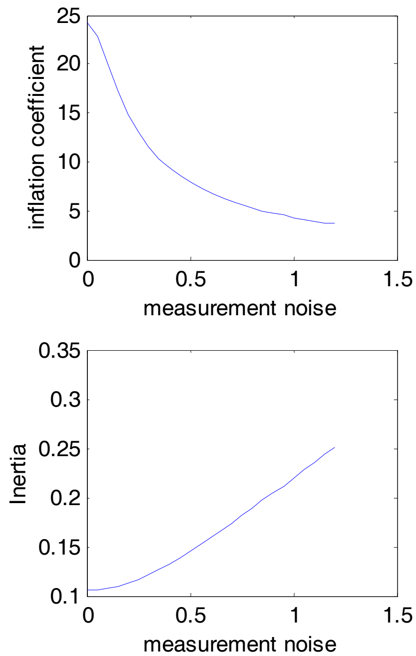

Figure 5: The effect of errors in the inflation forecast on the robust Taylor rule.

Small changes close to the neighbourhood of $\gamma^* = 8.43$, can make a substantial difference in the inflation response coefficient, whilst at higher values of γ inflation response coefficients settle down to values close to 1 corresponding to that in the LQG rule. Also, the output gap response coefficients decrease relatively fast close to the neighbourhood of γ^*.

Interestingly, Figure (7) also shows how inertia in the policy rule changes as the degree of robustness changes; the lagged interest rate coefficient becomes larger and hence the policy rule therefore more "backward looking " as the policymaker becomes less concerned with the robustness properties of his rule.

We can now explore the question raised initially about the degree of forward looking behaviour by the private sector in forming their expectations of inflation. As a simple calibration exercise we can increase the value of the coefficient, ρ, that determines the relative weight of forward expectations in equation (34). Up to this point we have set $\rho = 0$ so expectations have been purely backward looking and when $\rho = 1$ expectations are purely forward looking. Figure(8) below shows that the optimal value of γ decreases monotonically as ρ increases towards unity. Hence as there is more forward looking behaviour the policy maker is able to be more robust. This seems to suggest that the forward looking nature of expectations can absorb some of the concern the policy maker may have for robustness.

Finally, a moderately robust rule with γ approximately 30% higher than its lowest attainable value is calculated below. This implies γ value of 10.84 in the graphs above. .

$$i_t^R = .45i_{t-1} + 2.23\pi_{t-1} + .30y_{t-1} + .34y_{t-2} \tag{42}$$

This rule, after adjusting for mean, delivers decisions for the base rate, which by visual inspection are reasonably close to actual policy choices at least over the period 1993-2001. Figure(9) plots this moderately robust rule and base rate.

This rule would have suggested considerably more aggressive interest rate increases during the early period of the "Ken and Eddie" show. Then, from 1993 the rule seems to follow reasonably closely to actual base rate until about 1997. Towards the beginning of 2000 the robust rule would have again suggested somewhat more active interest rate movements. One interpretation of this observation that it appears impossible to capture both the preMPC and MPC periods with the same γ may be that the procedures adopted within the MPC for monetary policy making may be more robust than previous procedures. It is also noticable that even though we may be able explain part of the MPC's behaviour as a desire for robustness we still see more variability in this robust policy than the actual policy.

4.3 Average Performance

Finally, as discussed in section (2.4) we need somehow to compare the performance of these rules. We do this by plotting in Figure(10) the frequency decomposition of the expected losses under the robust rule, the moderately robust rule and the linear quadratic rule, corresponding to (40), (42) and (39). This corresponds to plotting the trace of a closed-loop transfer function $(tr(G(\theta)'G(\theta)'))$ over a frequency range θ.Each point on the plot shows the contribution of the shock process to the total loss at frequency θ. The area below the curve gives the average expected loss under the corresponding rule[21].Frequencies on the horizontal axis can be translated into a cyclical period of a shock in real time, as shown in the figure below. Consequently, the frequency decomposition of expected losses summarises how different attitudes towards model misspecifications and measurement errors affect the performance of the system.

One clear observation is that there is relatively little difference between the moderate robust rule and the Linear Quadratic rule, yet the figure shows that the linear quadratic rule is somewhat more vulnerable to the shocks that occur at normal "business cycle frequencies". The robust rule, instead delivers a flatter frequency response, suggesting that it insulates the economy across a wider range of disturbances at different frequencies. This is due to the fact that the

[21] See the discussion in the section(2.4).

robust policymaker assigns larger weights to frequencies to which the economy is most vulnerable as his rule is specifically designed to avoid the worst-case outcomes.

5 Conclusions

We have explored the use of robust control methods in the context of recent UK monetary policy in this paper and several insights have been gained. In the first place it is clear that the simple application of this approach in the case of unstructured uncertainty, where the set of misspecifications is determined simply by the choice of the γ parameter, does not immediately produce economically sensible solutions. This may imply that we should follow the alternative route of making specific assumptions as to the nature of the misspecifed model we wish our policy to be robust against (structured uncertainty) or alternatively that the minmax criterion is not always a suitable resolution to the problem of Knightian Uncertainty. Essentially each of the alternative misspecified models we could consider defines an alternative probability framework in which we could carry out our standard policy analysis and we are faced with resolving the multiple prior form of Knightian Uncertainty discussed by Gilboa and Schmedler(1989). The minmax criterion is one way of resolving the ambiguity in the decision problem and providing a unique solution but it is not the only route and perhaps more consideration should be given to considering alternatives.

However we found that relatively small adjustments to the size of the set of models we wish the policy to be robust against , determined by movements away from the optimal γ^*delivered sensible solutions. This emphasises that the concentration on robust *stability* within $\mathcal{H}_\infty$ methods serves only as a guide but it not necessarily the dominant concern. If policy makers are to make use of these powerful techniques we need to find formal justifications for the manner in which we move back from the most robust solution. Once we reduced the value of γ by some 30% we found economically sensible and interesting policy conclusions.

First, it is clear that even the non-robust LQG policy was more active or variable than the actual policy carried out over our sample period in common with all robust policies we computed. We need to understand better why actual monetary policy is relatively inactive when compared to these optimal robust and non-robust policies. Maybe the resolution here lies as suggested by Charles Goodhart (1999) in the need for the policy maker to retain credibility and rapid policy changes do not inspire much confidence. The framework we have developed here is not able to address this issue properly. We did however find that the policy maker's ability to construct more robust policy increased as the private sector became more forward looking when forming their inflation expectations which emphases the importance of addressing the strategic issues raised in the design of robust monetary policy.

We also found that the impact of errors in estimating the output gap and the inflation forecast have a strong impact on the robust policy design and are an important element to take into account. The robust Kalman filter approach we adopted can be easily used in practice.

From a political economy point of view we found that we were unable to explain the preMPC and MPC periods with the same γ value. In particular the MPC period seems to be explained by a considerable degree of robustness unlike the previous period in which the robust policy would have been much more aggressive than the monetary policy that was actually implemented from 1989-1992.

References

[1] Ball, L.(1999).'Policy Rules for Open Economies '.In J.B.Taylor (ed.), *Monetary Policy Rules* , University of Chicago Press,pp.127 –144.

[2] Basar, T and Bernhard, P. (1995).$\mathcal{H}_\infty$ *-Optimal Control and Related Minimax Design Problems, A Dynamic Game Approach*, Second edition, Birkhauser.

[3] Bernhard P.(2002), A Tutorial on Robust Control, *Macroeconomic Dynamics,* vol6,no.1.

[4] Brainard W.,(1967), Uncertainty and the effectiveness of policy, *American Economic Review*, Papers and Proceedings 57, 411-25.

[5] Clarida, R., Gali, J. and Gertler, M. (1999). The Science of Monetary Policy: A New Keynesian Perspective," *Journal of Economic Literature* 37, 1661-1707.

[6] Giannoni, M. (2002). Does Model Uncertainty Justify Caution? Robust Optimal Monetary Policy in a Forward-Looking Model. *Macroeconomic Dynamics*, vol6,no.1

[7] Gilboa, I. and Schmeidler, D. (1989). Max-min Expected Utility with Non-unique Prior. *Journal of Mathematical Economics*, 18, 141-153

[8] Glover, D. and Doyle, J. (1988). State-Space Formulae for All Stabilising Controllers that Satisfy an H_∞-Norm Bound and Relations to Risk Sensitivity, *System and Control Letters,* Vol 11(1), pp. 167-72.

[9] Goodhart C.,(1999), Central Bankers and Uncertainty, Keynes Lecture to the Royal Academy, reprinted i n the Bank of England Quarterly Bulletin, Feb , 102-114.

[10] Hansen, L. and Sargent, T. (2005). *Robustness . manuscript*, forthcoming Princeton University Press.

[11] Hansen, L. and Sargent, T. and Tallarini, T. (1999). Robust Permanent Income and Pricing, *Review of Economic Studies*, vol. 66, 873-907.

[12] Egginton, D., Pick, A., and Vahey, S. (2001), "Keep it Real!": A Real-time UK Macro Data Set, h*ttp://www.econ.cam.ac.uk/dae/keepitreal/*

[13] Jääskelä, J. and Kilponen, J. (2001). Risk, Uncertainty and Model Misspecification in a Monetary Policy Game, *manuscript.*

[14] Kasa, K. (2002). Model Uncertainty, Robust Policies, and the Value of Commitment. *Macroeconomic Dynamics,* vol6,no1.

[15] Knight, F. 1921. *Risk, Uncertainty and Profit.* Boston, New York, Houghton Mifflin Company.

[16] Martin B.,(1999), Caution and Gradualism in Monetary Policy under Uncertainty, *Bank of England Working Paper.No.105*

[17] Martin B. and C.Salmon,(1999), Should uncertain moneary policy makers do less?, *Bank of England Working Paper No.99*

[18] Onatski, A. (2000). Robust Monetary Policy Under Model Uncertainty: Incorporating Rational Expectations. *manuscript*, Harvard University.

[19] Onatski, A. (1999). Minimax Monetary policy: Comparison to Bayesian Approach, Worst Cases, and Exact Minimax Rules. *Robust Decision Theory in Economics and Finance*, Cambridge University Press, forthcoming.

[20] Onatski, A. and Stock, J. (2002) Robust Monetary Policy under Model Uncertainty in a Small Model of the U.S. Economy. *Macroeconomic Dynamics*, vol6,no.1

[21] Orphanides, A. (1998). Monetary Policy Rules Based on Real Time Data, *Federal Reserve Board Finance and Economics Discussion Series* 1998-3.

[22] Roberts, J. (1995). New Keynesian Economics and the Phillips Curve, *Journal of Money Credit and Banking*, 27 (4), pp. 975-84.

[23] Sack B.,(1998), Does the Fed act gradually?, A VAR analysis, *Board of Governors of the Federal Reserve System*, FED Working Paper No.17.

[24] Salmon M. and S.Weston (2001), Is the Monetary Policy Committee Robust?, Working Paper, City University Business theSchool.

[25] Sargent T.,(1999), Discussion of paper by L. Ball, Policy Rules for Open Economies, in J.Taylor, *Monetary Policy Rules*, NBER, Chicago University Press.

[26] Svensson, L. (1997). Inflation Forecast Targeting: Implementing and Monitoring Inflation Targets. *European Economic Review*, 41, 1111-1146.

[27] Svensson, L. (2001). What is wrong with Taylor Rules? Using Judgement in Monetary Policy through Targeting Rules", manuscript.

[28] Taylor, J. (1993)., Discretion Versus Policy Rules in Practice," *Carnegie-Rochester Conference Series on Public Policy*, 39, 1495-214.

[29] Tetlow R.J. and von zur Muehlen, P. (2000). Robust Monetary Policy with Misspecified Models: does model uncertainty always call for attenuated Policy. *Manuscript.* Federal Reserve Board.No.2000/28

[30] Whittle, P. (1981). Risk-Sensitive Linear Quadratic Gaussian Control, *Adv. Appl. Prob,* 13, 764-777.

[31] Whittle, P. (2002). Risk-Sensitivity - A Strangely Pervasive Concept, *Macroeconomic Dynamics,*vol6,no.1.

[32] Woodford, M. (1999). Optimal Monetary Policy Inertia. *NBER Working Paper* #7261.

[33] Zames, G. (1981). Feedback and Optimal Sensitivity: Model Reference Transformations, Multi

[34] Zhou, K. Doyle, J., Glover, K. (1996). *Robust and Optimal Control.* Prentice Hall, New Jersey.

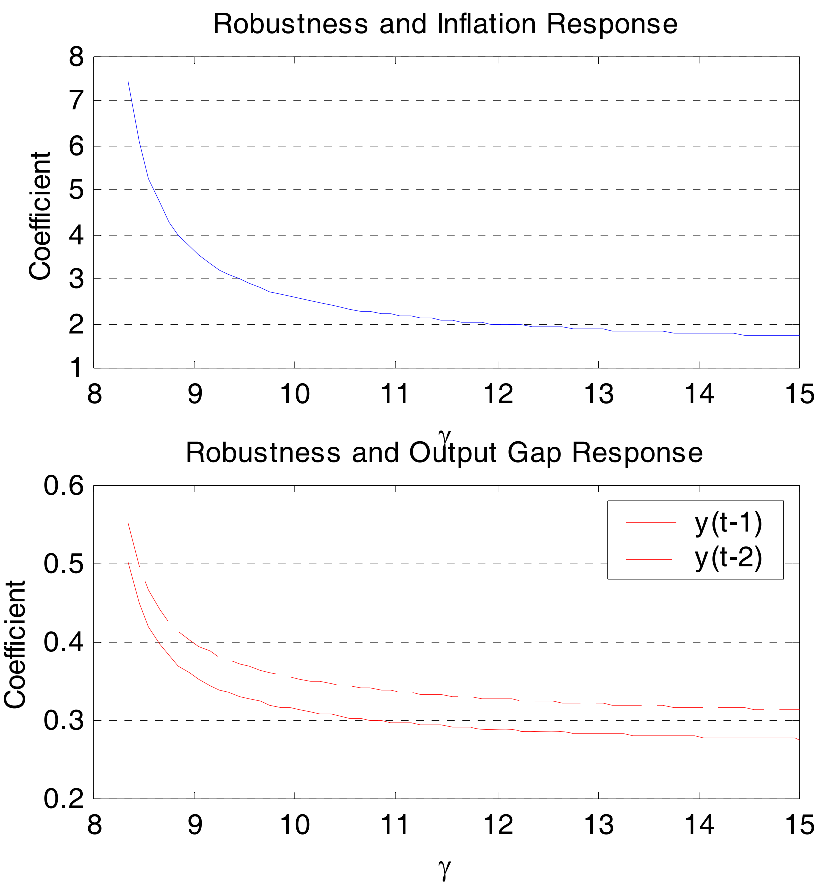

Figure 6: The effect on robust Taylor reaction coefficients as the degree of robustness is decreased

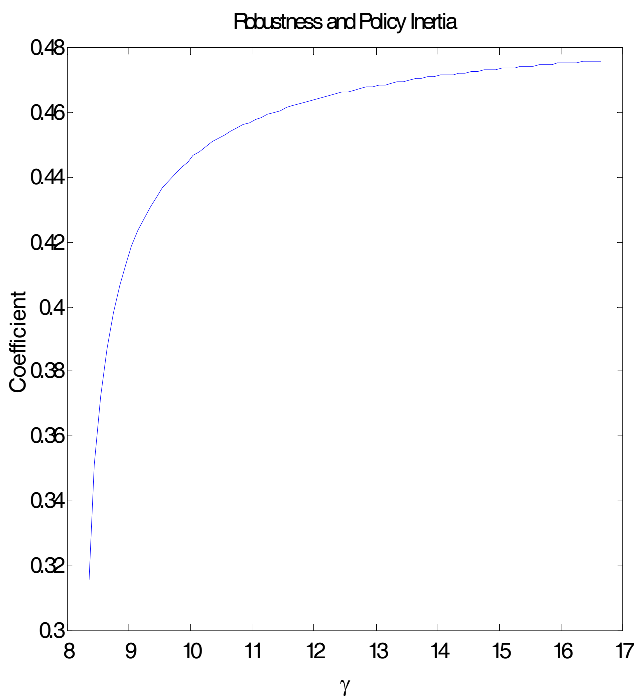

Figure 7: Robustness and Inertia in the Policy Rule

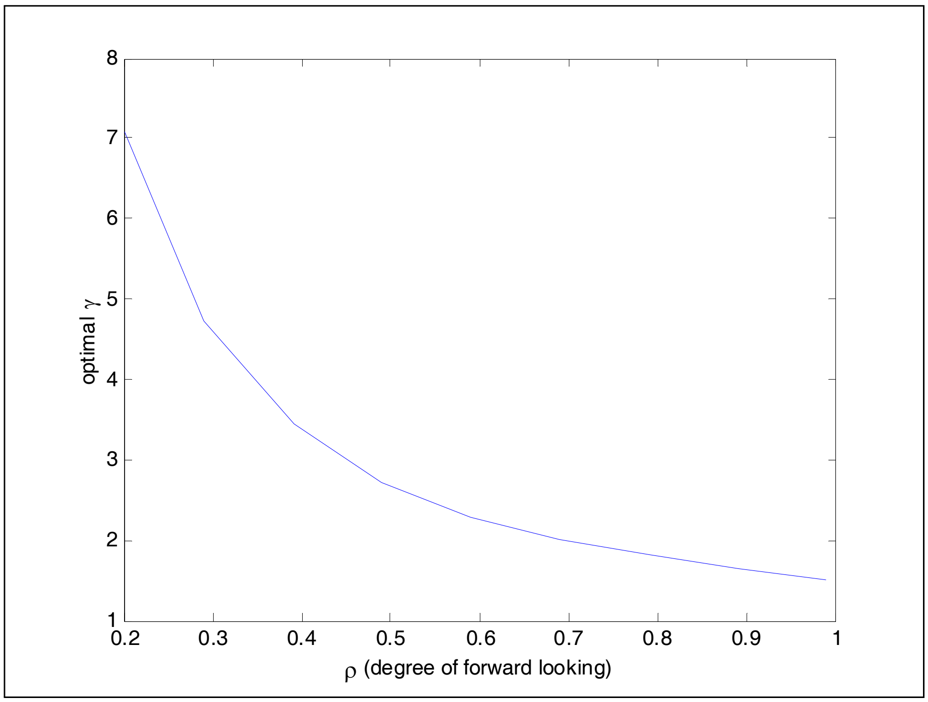

Figure 8: How a greater weight on forward looking expectations implies more robustness.

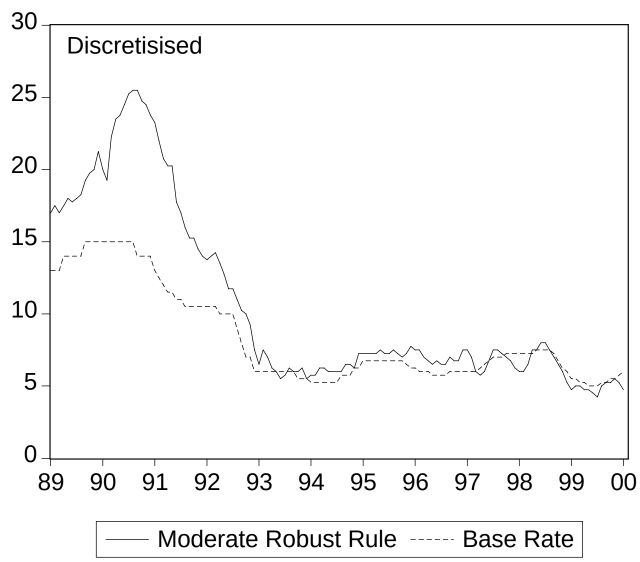

Figure 9: The moderately robust rule and the base rate.

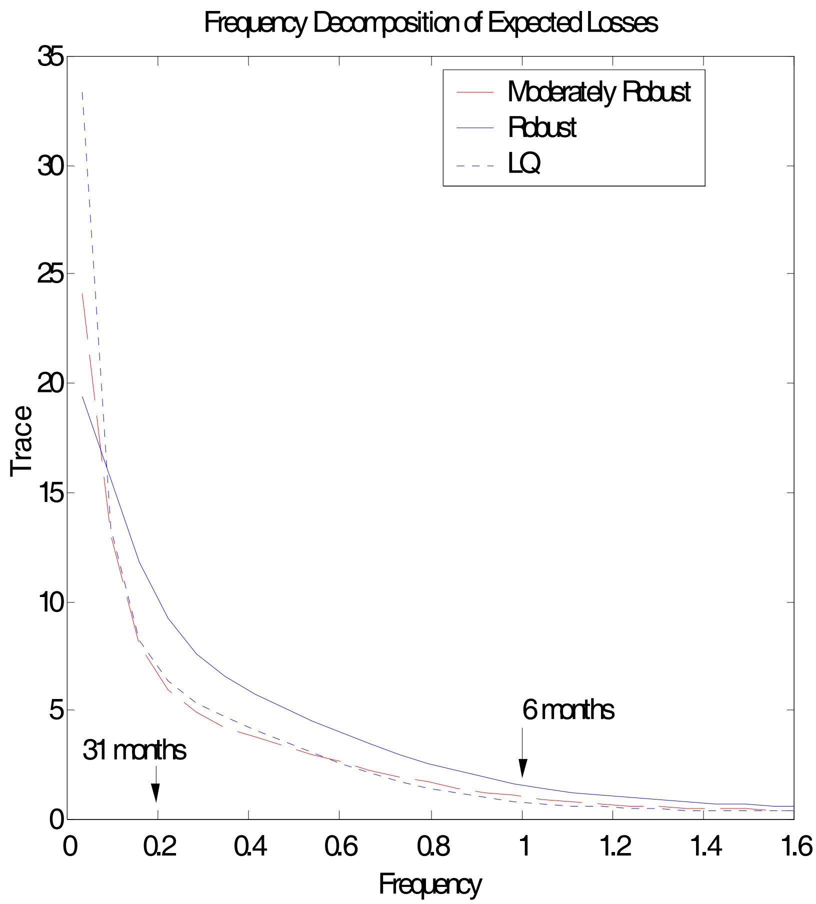

Figure 10: Performance in the Frequency Domain