

ROBUST MONETARY POLICY UNDER MODEL UNCERTAINTY IN A SMALL MODEL OF THE U.S. ECONOMY

ALEXEI ONATSKI

Harvard University

JAMES H. STOCK

Harvard University

and

NBER

This paper examines monetary policy in a two-equation macroeconomic model when the policymaker recognizes that the model is an approximation and is uncertain about the quality of that approximation. It is argued that the minimax approach of robust control provides a general and tractable alternative to the conventional Bayesian decision theoretic approach. Robust control techniques are used to construct robust monetary policies. In most (but not all) cases, these robust policies are more aggressive than the optimal policies absent model uncertainty. The specific robust policies depend strongly on the formulation of model uncertainty used, and we make some suggestions about which formulation is most relevant for monetary policy applications.

Keywords: Robust Control, Parameter Uncertainty, Brainard Uncertainty

1. INTRODUCTION

Speaking from the perspective of a regional Federal Reserve Bank president, Poole (1998) recently listed five sources of uncertainty that affect the formulation of monetary policy: future events, shocks, and disturbances; the actual workings of the economy; market reactions to current central bank policy; market expectations of future central bank policy; and limitations of the data. The first of these, which we refer to as shock uncertainty, has been studied extensively and underlies the stochastic simulations that are performed as part of state-of-the-art policy analysis exercises. The next three of these types of uncertainty can be categorized broadly as model uncertainty. Model uncertainty can arise from many sources. All models

The authors thank Ray Fair, Lars Hansen, John Leahy, Athanasios Orphanides, Glenn Rudebusch, Tom Sargent, Aaron Tornell, Mark Watson, and the participants at the Harvard Monetary and Fiscal Policy Seminar for helpful suggestions and discussions. This research was supported in part by National Science Foundation grants SBR-9409629 and SBR-9730489. Address correspondence to: James H. Stock, Kennedy School of Government, Harvard University, Cambridge, MA 02138, USA; e-mail: james.stock@harvard.edu.

are approximations, and so, any particular model has approximation error. The true economic relations might be time varying, perhaps because of structural breaks, in which case the approximate model could be time varying. Even if the economy is time invariant, the parameters of the approximate model could be time varying because of specification error. Finally, even if the model is correctly specified and is time invariant, its parameters are estimated, and so, there is econometric estimation error.

At least since Brainard's (1967) important work, economists have tried to incorporate model uncertainty into policymaking. Generally speaking, existing approaches to model uncertainty employ Bayesian decision theory; these solve for decision rules that are desirable under some set of priors on the model parameters. As is discussed below, however, this approach has important limitations, perhaps accounting for the infrequency with which formal decision theoretic methods are used in the actual conduct of monetary policy.

This paper takes a different approach to monetary policymaking under model (and shock) uncertainty. This approach draws upon advances over the past decade in the field of robust control in the engineering literature; see Dahleh and Diaz-Bobillo (1995) and Zhou et al. (1996) for recent textbook treatments.¹

This approach is fundamentally different from the Bayesian decision theoretic approach. Rather than specifying a prior and solving for the rule that minimizes the Bayes risk, the robust control approach specifies a broad, usually nonparametric, set of possible modeling errors (model perturbations) and poses a minimax problem in which the policy is chosen to minimize the maximal risk over all perturbations in this set. Typically, this problem must be solved numerically if it can be solved at all.

Our analysis focuses on a specific model, the two-equation model used by Rudebusch and Svensson (1999) (henceforth, the RS model), although some theoretical results are provided for a simplified version of this model. We consider policy rules of the generalized Taylor form, in which the monetary authority sets the interest rate based on the rate of inflation and the output gap. The RS model and generalized Taylor rules were chosen because of their simplicity so that the resulting robust rules can be given economic interpretations.

In this context, the robust control problem can be formulated at a general level. Let M denote the model at hand, and let Δ denote the unknown nonrandom deviation of the true model from the model M . In general, Δ can be time varying and state dependent, and the set of such perturbations considered is denoted by D . The model M is referred to as the base or nominal model. The perturbed nominal model is $M + \Delta$, which can be thought of as the true model that is relevant at the time the policy is placed in effect. Let $\{K\}$ be the set of policy rules under consideration, and let $R(K, M + \Delta)$ denote the risk (expected loss) of pursuing policy K when the true model is $M + \Delta$, where the expectation is taken over shock uncertainty. We consider robust control problems of the form

$$\min_{\{K\}} \sup_{\Delta \in D} R(K, M + \Delta). \quad (1)$$

In the numerical work, M is the RS model and $\{K\}$ is the set of generalized Taylor rules. Choices for the risk function R and the set of perturbations D are discussed in Sections 2 and 3, respectively.

This paper addresses two sets of questions. The first set of questions concerns the economics of the robust policies and how these policies differ from policies without parameter uncertainty. There have been three previous, limited investigations of robust rules in the RS model or variants: Stock (1999, in press) considered two restrictive choices of D , whereas Sargent (1999) studied Ball's (1999) related model using a relatively unrestricted choice of D . Intriguingly, these studies all found the robust rules to be more aggressive than the rules that would obtain without model uncertainty, in the sense that the responses to both inflation and the output gap in the modified Taylor rules were larger. This contrasts with the usual finding for Bayesian approaches to uncertainty: Model uncertainty often leads to less aggressive policies, as discussed by Brainard (1967) and, in the context of monetary policy, Rudebusch (2001) and Wieland (1998) [although, as the latter authors note, this tendency to conservatism is not generally true in the static case, cf. Chow (1973), or if there is learning]. However, how general is the proposition that robust policy rules are more aggressive than the certainty equivalence rule?

The second set of questions are methodological. Various choices for D are available, corresponding to different formulations of the robust control problem. As discussed in Section 3, there is an approximate hierarchy of generality for choice of D . This raises the questions of whether these different choices for D imply economically different policy rules and, if so, which choice of D seems most natural for the monetary policy problem.

Before proceeding to these questions, it is useful to digress by contrasting this approach to the usual decision theoretic treatment of model uncertainty. At its simplest in a static context, this approach posits a prior distribution on the model parameters and computes the policy that minimizes the Bayes risk. Rudebusch (in press, Sect. 5) provides an interesting extension of this approach in which the prior entails evolution of the parameters over time, with the stochastic nature of this evolution characterized by hyperparameters of the priors; he solves for generalized Taylor rules that minimize the Bayes risk in the RS model. More complicated dynamic treatments incorporate learning, so that posteriors are continuously updated; cf. Wieland (1998), who builds on work by Wieland (2000), Easley and Kiefer (1988), and others. Although grounded in optimal decision theory, the Bayesian approach has two related disadvantages. On the one hand, in realistic macroeconomic models, the dimensionality of the parameter space is large, and proper implementation of these methods requires elicitation of priors of policymakers to a degree that is impractical. Retreating to "uninformative" or flat priors is not a solution, both because the decision problem is compelling mainly if there are informative priors and because the construction of flat priors depends on which arbitrary nonlinear transformations of the parameters are used (dynamic responses are generally nonlinear transformations of primitive model parameters, for example). This is particularly problematic in models such as the RS model,

which has large autoregressive roots; indeed, the question of how best to specify priors in the presence of possibly large autoregressive roots in univariate autoregressive models remains unresolved (see the special issue of *Econometric Theory*, August/October 1994). On the other hand, the parametric nature of the problem limits the types of misspecifications and disturbances that can be considered: The policymaker might be worried about a possible breakdown of the model that arises from sources she cannot now identify and therefore upon which she cannot place parametric structure and prior probability. In contrast, the robust control approach specifies a neighborhood of disturbances, which can have nonparametric characteristics. The perturbations can be specified generally or can focus on the some parts of the model. Beyond specifying the general nature of her uncertainty, no further information is required from the policymaker (beyond her loss function) to construct the robust policy.

The remainder of the paper is organized as follows. The RS model, generalized Taylor rules, and policymaker loss functions are introduced in Section 2. Section 3 presents a hierarchy of different specifications of D , each of which produces robust control rules. Numerical results for these different rules are presented and discussed in Section 4. Some theoretical results for a first-order version of the RS model are presented in Section 5; these results address, from a theoretical perspective, whether the robust policy rules are generically aggressive. Section 6 concludes.

2. MODEL, LOSS FUNCTIONS, AND POLICY RULES

2.1. Model and Policy Rules

Rudebusch and Svensson's (1999) model is a two-equation model of inflation (π_t), measured at an annual percentage rate, and the output gap (y_t), measured as 100 times the log ratio of actual real output to potential output. The first equation is a Phillips curve linking the change in inflation to the past level of the output gap. The second equation links the output gap to the ex-post real interest rate. Rudebusch and Svensson (1999) estimated the model using quarterly postwar data for the United States. The estimated model, with standard errors of the coefficients in parentheses, is

$$\pi_{t+1} = \underset{(0.08)}{0.70} \pi_t - \underset{(0.10)}{0.10} \pi_{t-1} + \underset{(0.10)}{0.28} \pi_{t-2} + \underset{(0.08)}{0.12} \pi_{t-3} + \underset{(0.03)}{0.14} y_t + \varepsilon_{\pi,t+1}, \quad (2a)$$

$$y_{t+1} = \underset{(0.08)}{1.16} y_t - \underset{(0.08)}{0.25} y_{t-1} - \underset{(0.03)}{0.10} (\bar{i}_t - \bar{\pi}_t) + \varepsilon_{y,t+1}, \quad (2b)$$

where i_t is the federal funds rate at an annual rate, $\bar{i}_t = 0.25(i_t + i_{t-1} + i_{t-2} + i_{t-3})$ is the four-quarter average value of the federal funds rate, and $\bar{\pi}_t = 0.25(\pi_t + \pi_{t-1} + \pi_{t-2} + \pi_{t-3})$ is the four-quarter average value of the rate of inflation. The coefficients on lagged inflation in (2a) sum to one, so that (2a) could alternatively be expressed in terms of changes of inflation, $\pi_t - \pi_{t-1}$, rather than in levels of inflation. This accords with conventional Phillips-curve specifications for the

United States [cf. Gordon (1998)] in which the change of inflation is written as depending on a demand measure (the output gap or the unemployment rate), lagged changes in inflation, and supply shocks, although the supply shocks are not incorporated here.

The monetary authority is assumed to control the nominal interest rate using a policy rule. In general, policy rules can depend on all observable data. In this paper, however, we consider simple policy rules of the generalized Taylor form, in which interest rates are set on the basis of current values of the rate of inflation and the output gap. As emphasized by McCallum (1988, 1990) and Taylor (1993), simple rules have the advantage of being easy for policymakers to follow and also of being easy to interpret. The generalized Taylor rules considered here are

$$i_t = g_\pi \pi_t + g_y y_t, \quad (3)$$

where g_π and g_y are parameters to be chosen when designing the rule. Taylor (1993) suggested the parameters $g_\pi = 1.5$ and $g_y = 0.5$, and these will serve as a benchmark for comparing various robust policies.

It is convenient to reexpress the model and control rule in vector notation. Looking ahead, we consider loss functions that involve the change in the interest rate, $i_t - i_{t-1}$, and so, it is useful to augment the state vector to include both the change and the level of i_{t+1} . Let L denote the lag operator and $a(L)$ denote a lag polynomial. Then, (2) and (3) can be written

$$\begin{aligned} \begin{bmatrix} \pi_{t+1} \\ y_{t+1} \\ i_{t+1} - i_t \\ i_{t+1} \end{bmatrix} &= \begin{bmatrix} a_{\pi\pi}(L) & a_{\pi y}(L) & 0 & 0 \\ a_{y\pi}(L) & a_{yy}(L) & 0 & -a_{y\pi}(L) \\ a_{i\pi}(L) & a_{iy}(L) & 0 & a_{ii}(L) - 1 \\ a_{i\pi}(L) & a_{iy}(L) & 0 & a_{ii}(L) \end{bmatrix} \begin{bmatrix} \pi_t \\ y_t \\ i_t - i_{t-1} \\ i_t \end{bmatrix} \\ &+ \begin{bmatrix} 1 & 0 \\ 0 & 1 \\ g_\pi/4 & g_y \\ g_\pi/4 & g_y \end{bmatrix} \begin{bmatrix} \varepsilon_{\pi,t+1} \\ \varepsilon_{y,t+1} \end{bmatrix}, \end{aligned} \quad (4)$$

where $a_{i\pi}(L) = (g_\pi/4)(a_{\pi\pi}(L) + 1 + L + L^2) + g_y a_{y\pi}(L)$, $a_{iy}(L) = (g_\pi/4)a_{\pi y}(L) + g_y a_{yy}(L)$, $a_{ii}(L) = -g_y a_{y\pi}(L)$, $a_{\pi\pi}(L)$ is fourth order, $a_{yy}(L)$ is second order, $a_{\pi y}(L)$ is first order [so $a_{\pi y}(L) = a_{\pi y,0}$], and $a_{y\pi}(L)$ is the restricted fourth-order lag polynomial, $a_{y\pi}(L) = 0.25a_{y\pi,0}(1 + L + L^2 + L^3)$.

The combined model and control rule (4) can be written more compactly as

$$z_{t+1} = A(L)z_t + B\varepsilon_{t+1}, \quad (5)$$

where $z_t = (\pi_t, y_t, i_t - i_{t-1}, i_t)'$ and $\varepsilon_t = (\varepsilon_{\pi,t}, \varepsilon_{y,t})'$, and $A(L)$ and B are the corresponding matrices in (4). Note that $A(L)$ and B are functions of g_π , g_y , and the elements of the upper-left 2×2 block of $A(L)$.

Although no intercepts appear in any of these expressions, the analysis is invariant to reinterpreting the variables as deviations from target values. For example, if target inflation is 2% per year, then π_{t+1} would denote actual inflation minus 2%. Without loss of generality, we therefore set target values to zero.

2.2. Loss Functions

The monetary authority is taken to prefer inflation close to zero, an output gap close to zero, and a stable interest rate, that is, an interest rate with small changes. The level of the interest rate is assumed not to enter the loss function. Two loss functions are considered: quadratic (or l_2) loss and worst-case (or l_∞) loss. Let Λ denote a 4×4 diagonal matrix with diagonal elements $(1, \lambda_\pi, \lambda_{\Delta i}, 0)$.

Quadratic loss. Under quadratic loss, the loss function is

$$L_{l_2}(z_t) = z_t' \Lambda z_t = \|\Lambda^{\frac{1}{2}} z_t\|_2^2, \tag{6}$$

where $\|X\|_2$ is the l_2 norm of the $n \times 1$ vector X , $\|X\|_2 = [\text{tr}(XX')]^{\frac{1}{2}}$. The associated risk function is the expected value of this loss function,² $R_{l_2} = \text{var}(\pi_t) + \lambda_y \text{var}(y_t) + \lambda_{\Delta i} \text{var}(i_t - i_{t-1})$, where the expectation is taken over the shock uncertainty (the model perturbation is treated as nonrandom).³

Maximum value loss. Although quadratic loss is conventional because of its tractability, it might not accurately reflect the attitude of the monetary authority toward risk. For example, the monetary authority might be far more concerned about worst-case scenarios than about ensuring that the deviations from target paths are always small. One way to formalize this is to consider a loss function that focuses on worst-case outcomes for the variables of interest, where the variables are scaled by the square root of the diagonal elements of Λ . This maximum-value loss function is

$$L_{l_\infty}(z) = \sup_t \max(|\pi_t|, |\lambda_y^{\frac{1}{2}} y_t|, |\lambda_{\Delta i}^{\frac{1}{2}} (i_t - i_{t-1})|) = \|\Lambda^{\frac{1}{2}} z\|_\infty, \tag{7}$$

where $\|x\|_\infty$ is the l_∞ norm of the $n \times 1$ vector time series $x = (\dots, x_{-1}, x_0, x_1, \dots)$, $\|x\|_\infty = \sup_t \max_i |x_{it}|$. The associated risk function is the expected value, $R_{l_\infty} = E \|\Lambda^{\frac{1}{2}} z\|_\infty$. So that the risk is finite, when working with l_∞ loss, the shocks are assumed to be random variables with bounded support.

3. OPTIMAL POLICIES UNDER MODEL UNCERTAINTY

This section provides an approximate hierarchy of characterizations of model uncertainty, going from the least to the most parametric. When combined with a risk function and the specification of the policy rules (3), this completes the specification of the monetary authority’s decision problem (1), and it remains only

to solve the optimization problem. In general, these solutions are computationally difficult, and not all possible combinations of model uncertainty and risk functions constitute well-posed and computationally solvable problems. In some cases, the optimal policy can be computed, but in others only partial solutions are available, in which case typically it is feasible to compute a set that contains the optimal policy but not to compute the optimal policy itself.

In general, model uncertainty can be expressed as the true model deviating from the nominal model, so that $A(L)$ becomes $A(L) + \Delta_t$, where Δ_t is a sequence of perturbations to $A(L)$. Substituting this into (5), the perturbed model can be written as

$$z_{t+1} = A(L)z_t + \Delta_t z_t + B\varepsilon_{t+1}. \quad (8)$$

Thus, letting $\eta_{t+1} = \Delta_t z_t + B\varepsilon_{t+1}$, we have

$$z_{t+1} = A(L)z_t + \eta_{t+1} \quad (9)$$

or

$$z_{t+1} = G(L)\eta_{t+1}, \quad (10)$$

where $G(L) = [I - A(L)]^{-1}$. One interpretation of (8) is that $A(L) + \Delta_t$ represents the true model, $A(L)$ is the linear approximation, and ε_{t+1} are the shocks to the true model. The combined shock η_{t+1} includes both components.

The benchmark for the numerical analysis is the conventional linear-quadratic (LQG) control solution, for which there is no model uncertainty, and so, $\Delta_t = 0$. The general optimal policy is obtained using standard methods for linear-quadratic problems. Here, the policy is restricted to be of the generalized Taylor form (3).

The remainder of this section discusses different sets for Δ_t . Computation of the associated optimal robust policy is discussed in conjunction with the introduction of the set. The sets are approximately hierarchical and are discussed from the most general to the most specific. The most general forms of model uncertainty are essentially nonparametric and incorporate both model and shock uncertainty. The remaining forms of uncertainty focus on model uncertainty and exploit, to varying degrees, the parametric form of the nominal model.

3.1. Mixed Model and Shock Uncertainty

These formulations treat model perturbations ($\Delta_t z_t$) and shocks (ε_t) symmetrically. The robust control problem is formulated to guard against undesirable outcomes resulting from bad distributions of the combined errors η_t , which could arise from bad distributions of shocks, modeling errors, or a combination of the two. Note that, in this formulation, the worst-case combined errors against which the policymaker wants to be robust can always be constructed using shock perturbations only (for a given η_t , set $\Delta_t = 0$ and $\varepsilon_t = \eta_t$).

l_2 -Bounded perturbations, l_2 loss. In this case it is assumed that $E\|\Delta_t z_t\|_2 < \infty$ and $E\|\varepsilon_t\|_2 < \infty$, so that $E\|\eta_t\|_2 < \infty$. The monetary authority wishes to

guard against the worst-case distribution of perturbations and errors, subject to $E \|\eta_t\|_2 < \infty$. So that the problem is tractable, the monetary authority in this case is taken to have l_2 loss. This leads to the so-called H_∞ robust control problem.

Because of the linearity of (9), the scale of η_t is arbitrary, and so, without loss of generality, it is possible to restrict attention to combined shocks with $E \|\eta_t\|_2 = 1$. Thus, the optimization problem (1) becomes

$$\min_{\{g_\pi, g_y\}} \sup_{\{\eta: E \|\eta_t\|_2 = 1\}} E z_t' \Lambda z_t. \tag{11}$$

Because $E z_t' \Lambda z_t = \text{tr}(2\pi)^{-1} \int_{-\pi}^{\pi} \Lambda^{\frac{1}{2}} G(e^{i\omega}) F_\eta(\omega) G(e^{-i\omega})' \Lambda^{\frac{1}{2}} d\omega = (2\pi)^{-1} \int_{-\pi}^{\pi} \text{tr} F_\eta(\omega) G(e^{-i\omega})' \Lambda G(e^{i\omega}) d\omega$, where F_η is the spectral density matrix of η , (11) is equivalent to

$$\min_{\{g_\pi, g_y\}} \|\Lambda^{\frac{1}{2}} G\|_{H_\infty} \tag{12}$$

where $\|M\|_{H_\infty}$ is the H_∞ norm of the stable operator M , $\|M\|_{H_\infty} = \{\sup_\omega \text{maxeval}[M(e^{i\omega})' M(e^{-i\omega})]\}^{\frac{1}{2}}$, where *maxeval* denotes the maximum eigenvalue. Alternatively, (11) can be expressed as $\min_{\{g_\pi, g_y\}} \sup_{\{\eta: \|\eta\|_2 > 0\}} (\|\Lambda^{\frac{1}{2}} G \eta\|_2 / \|\eta\|_2)$. Because $\sup_{\{\eta: \|\eta\|_2 > 0\}} (\|\Lambda^{\frac{1}{2}} G \eta\|_2 / \|\eta\|_2)$ is the induced norm of $\Lambda^{\frac{1}{2}} G$, which is the H_∞ norm [e.g., Dahleh and Diaz-Bobillo (1995)], this is equivalent to (11).

l_∞ -Bounded perturbations, l_∞ loss. This case is similar to the preceding case except that the perturbations *cum* shocks η_t are assumed to be bounded in absolute value, so $\|\eta\|_\infty < \infty$, and, so that the problem is well defined and tractable, the monetary authority is assumed to have maximum value loss. Accordingly, (1) becomes

$$\min_{\{g_\pi, g_y\}} \sup_{\{\eta: \|\eta\|_\infty \leq 1\}} E \|\Lambda^{\frac{1}{2}} z\|_\infty. \tag{13}$$

Here, η_t are a series of random variables with bounded support but otherwise unrestricted. This is equivalent to the problem $\min_{\{g_\pi, g_y\}} \sup_{\{\eta: \|\eta\|_\infty \leq 1\}} \|\Lambda^{\frac{1}{2}} G \eta\|_\infty / \|\eta\|_\infty$, which is in turn equivalent to solving

$$\min_{\{g_\pi, g_y\}} \|\Lambda^{\frac{1}{2}} G\|_{l_1}, \tag{14}$$

where $\|M\|_{l_1}$ denotes the l_1 norm of the $n \times r$ operator M , $\|M\|_{l_1} = \max_{1 \leq i \leq r} \sum_{j=1}^n \sum_{k=0}^{\infty} |m_{ij,k}|$ where $M(L) = \{m_{ij}(L)\}$. Because the induced norm of the l_∞ loss/ l_∞ errors problem (13) is the l_1 norm, this is referred to as the l_1 robust control problem.

Minimum entropy control. A modification of the (deterministic) H_∞ control problem is

$$\min_{\{g_\pi, g_y\}} \sup_{\{\eta: \|\eta\|_2 \leq 1 + \eta'_0 \eta_0\}} \|\Lambda^{\frac{1}{2}} z\|_2, \tag{15}$$

where $\eta = (\eta_0, \eta_1, \eta_2, \dots)$. The solution to (15) constitutes a continuum of controllers, indexed by the Lagrange multiplier on the constraint $\|\eta\|_2 \leq 1 + \eta'_0 \eta_0$. This family of solutions is called the minimum entropy controller.

The problem (15) can be motivated in a variety of ways. It is equivalent to the risk-sensitive control problem formulated by Whittle (1990); see Hansen and Sargent (1998) for a discussion. For our purposes, the attractive feature is that if the Lagrange multiplier is zero, then the minimum entropy controller is the LQG controller, while if the Lagrange multiplier is the H_∞ norm of the nominal model (its maximum value), the controller is the H_∞ controller.

3.2. Model Uncertainty Only

In the preceding cases, the control problem guards against uncertainty both in the model and in the shocks. An alternative approach is to focus solely on model uncertainty, so that the minimax problem guards against worst-case modeling errors (misspecification, structural breaks, etc.) but not against worst-case shocks.

There are two main differences between these cases and those in the preceding subsection. First, the preceding formulations have both model and shock uncertainty, whereas these incorporate only model uncertainty. Second, the preceding formulations placed no restrictions on the perturbations Δ_t other than that $\Delta_t z_t$ have finite norm, whereas in the case of structured uncertainty these perturbations will be required to be operators and thus be defined independently of the state vector $\{z_t\}$. Here, the operator Δ denotes the deviation of the nominal model from the true model, and so, Δ captures the misspecification of the model. For example, if the true model is linear but subject to one or more structural breaks, then Δ is linear but time varying.

The various cases of structured uncertainty are differentiated by the amount of structure that is imposed on Δ . For each case, the size of the perturbations can be indexed by their norm (either H_∞ or l_1). Thus, for each case, the optimization problem becomes

$$\min_{\{g_\pi, g_y\}} \sup_{\Delta \in D_r} R, \tag{16}$$

where R denotes the risk function and D_r denotes the set of perturbations considered, indexed by r , where larger values of r denote larger sets of Δ . Calculation of r is related to the structured norm of the perturbed model, and is discussed in the Appendix.

The results that are presented for structured uncertainty are limited by the computational tools currently available. In most cases, computational methods for obtaining a unique solution to (16) (if one exists) are unavailable. However, for a given set D_r , it is generally possible to compute the set of $\{g_\pi, g_y\}$ that results in output with finite norms, that is, that result in stable controlled systems. Thus, although (16) is formulated as a robust performance problem, in fact the numerical results will mainly be for robust stability. Since stability is a property of the system and not the parameters of the loss function, these sets do not depend on Λ . For

r sufficiently large, it is generally impossible to stabilize the perturbed system, in which case no solution to (16) exists for finite risk and the set of stabilizing rules is empty. For r sufficiently small, the set of rules that stabilize the system can be large, in which case this set is relatively uninformative about the nature of the minimax rule that solves (16). A value of r that is of particular interest is the critical value, r_{\max} , such that for $r > r_{\max}$, the set of stabilizing rules is empty.

In general, the perturbation operators have the structure,

$$\Delta = \begin{bmatrix} \Delta_{\pi\pi} & \Delta_{\pi y} & 0 & 0 \\ \Delta_{y\pi} & \Delta_{yy} & 0 & -\Delta_{y\pi} \\ \Delta_{i\pi} & \Delta_{iy} & 0 & \Delta_{ii} \\ \Delta_{i\pi} & \Delta_{iy} & 0 & \Delta_{ii} \end{bmatrix}, \tag{17}$$

where [analogously to (4)] $\Delta_{i\pi} = (1/4)g_{\pi} \Delta_{\pi\pi} + g_y \Delta_{y\pi}$, $\Delta_{iy} = (1/4)g_{\pi} \Delta_{\pi y} + g_y \Delta_{yy}$, and $\Delta_{ii} = -g_y \Delta_{y\pi}$. In linear cases, the perturbations are structured so that the coefficients on the lagged rate of inflation in the first equation sum to one. Accordingly, we set $\Delta_{\pi\pi} = \tilde{\Delta}_{\pi\pi}(1 - L)$; also, for reasons of interpretation, we set $\Delta_{y\pi} = \tilde{\Delta}_{y\pi}(1 + L + L^2 + L^3)/4$. Note that all elements of Δ are functions of the elements of its first 2×2 block, and so, perturbations are henceforth specified only in terms of this block.

When entries of the Δ matrix are restricted to be constant operators (scalars), then they can be interpreted as once and for all deviations from the nominal model coefficients in (2). In this case, it is natural to measure Δ in units of the standard errors of the corresponding estimates. In the general case, we continue to measure the size of Δ in these standard error units, although strictly the standard error is the correct term only for the lead elements of Δ when Δ is a general operator.

Nonlinear time-varying bounded perturbations (general case). In this case, Δ is permitted to be a nonlinear time-varying (NTV) operator with bounded norm, and so, $D_r = \{\Delta : \max(\|\tilde{\Delta}_{\pi\pi}\|, \|\Delta_{\pi y}\|, \|\tilde{\Delta}_{y\pi}\|, \|\Delta_{yy}\|) < r\}$. For computational reasons, only results for the l_{∞} -induced norm are reported.

In principal, one could alternatively consider perturbations that are nonlinear time-invariant (NTI) or linear time-varying (LTV) operators. However, the structured norm for NTI and LTV operators is the same as for NTV operators [Dahleh and Diaz-Bobillo (1995, ch. 7)], and so, the resulting sets of stabilizing rules that are computed in lieu of unique solutions to (16) are the same for these three cases (NTV, NTI, LTV).

NTV bounded perturbations, cross effects only. In this case, $D_r = \{\Delta : \Delta_{\pi\pi} = 0, \Delta_{yy} = 0, \max(\|\Delta_{\pi y}\|, \|\tilde{\Delta}_{y\pi}\|) < r\}$. This restricts the preceding case by imposing the condition that there is no uncertainty about the ‘‘own dynamics’’ in the model, so that the only model uncertainty is associated with the cross effects (the effect of the real rate on the output gap and the effect of the output gap on inflation).

This uncertainty can include both uncertainty about the levels effect and about the cross dynamics. Results are reported for the l_∞ -induced norm.

Linear time-invariant bounded perturbations (general case). In this case, D_r is as defined for the general NTV case except that the perturbations are LTI operators. It turns out that, in general, for LTI perturbations the set of stabilizing rules is the same, for a given r , whether the norm is H_∞ or l_1 .

LTI bounded perturbations, cross effects only. In this case, Δ is LTI with $\Delta_{\pi\pi} = \Delta_{yy} = 0$. For this type of uncertainty, the same set of rules obtain for both H_∞ and l_1 norms.

LTI bounded perturbations, contemporaneous cross effects only. This case further restricts the uncertainty so that there is no uncertainty about the cross dynamics; rather, the uncertainty is only about the levels effect of the real rate on the output gap and of the slope of the Phillips curve. Thus, Δ has the form (17) with $\Delta_{\pi\pi} = \Delta_{yy} = 0$ and $\tilde{\Delta}_{y\pi}(L) = \tilde{\Delta}_{y\pi,0}$ and $\Delta_{\pi y}(L) = \Delta_{\pi y,0}$, where $\tilde{\Delta}_{y\pi,0}$ and $\Delta_{\pi y,0}$ are constants.

This formulation of uncertainty is highly restrictive but has the advantage of permitting computation of the actual optimal rule [the solution to (16)], not just a set of rules which contains the optimum. Unlike the sets computed for the other cases of structured uncertainty, this optimum policy will of course depend on Λ . The results are computed for l_2 loss. The low dimensionality of the perturbations permits direct evaluation of the loss for all perturbations of size r by grid search over D_r (to compute the worst-case loss) and over $\{g_\pi, g_y\}$ (to compute the minimax loss).

4. NUMERICAL RESULTS

4.1. Computational Issues

To simplify the calculations, the system was transformed so that the radii of allowable deviations are in standard-error units. The calculations were performed using the μ toolbox in MATLAB and programs developed for this project. The sets of stabilizing rules were computed by grid search in the domain, $g_\pi \in [1.2, 7]$ (grid of 0.2) and $g_y \in [0.1, 4.5]$ (grid of 0.1). We experimented with the size of the grid and chose the reported one because our solutions are in the interior.

4.2. Results

Mixed model and shock uncertainty. The parameters of the robust generalized Taylor rules for the cases of Section 3.1 are summarized in Figure 1. The LQG, optimal H_∞ , and optimal l_1 rules are each characterized by individual points in (g_π, g_y) space. The minimum entropy rule is indexed by a scalar parameter and the family of minimum entropy rules is a curve with endpoints equaling the LQG and

H_∞ rules. The benchmark LQG rule parameters are $(g_\pi, g_y) = (2.71, 1.58)$. The H_∞ robust control parameters are $(g_\pi, g_y) = (6.42, 2.75)$, and the l_1 robust control parameters are $(g_\pi, g_y) = (4.80, 3.33)$. The inverse of the H_∞ and l_1 norms of the controlled nominal model are shown in Figures 2 and 3.

NTV perturbations, general case. As discussed in Section 3, the results for structured perturbations generally take the form of a set in which the optimal rule parameters must fall, where the sets are indexed by the radius of the perturbations r . These sets have the property that all rules in these sets stabilize the perturbed model. These sets, for various values of r , are shown in Figure 4 for the l_1 norm. The set of rules that stabilizes the system for $r = r_{\max}$ is a singleton, and $r_{\max} = 0.805$. Thus, there is, to the limits of computational accuracy, a unique solution to the optimization problem (16) when $r = r_{\max}$, which is $(g_\pi, g_y) = (3.2, 2.3)$.

NTV perturbations, cross effects only. The sets of stabilizing policy rule parameters for this case are shown in Figure 5. The set of stabilizing rules for r_{\max} is a singleton, with $r_{\max} = 2.02$. The associated optimal robust rule for general LTI perturbations with $r = r_{\max}$ is $(g_\pi, g_y) = (3.0, 1.73)$.

LTI perturbations, general case. The sets of stabilizing policy rule parameters for general LTI perturbations are shown in Figure 6, where $r_{\max} = 1.82$.

LTI perturbations, cross effects only. The sets of stabilizing policy rule parameters are shown in Figure 7, where $r_{\max} = 3.3$. To within numerical precision, there is a plateau at $r = r_{\max}$.

LTI perturbations, contemporaneous cross effects only. Three sets of stabilizing policy rule parameters for these perturbations are shown in Figure 8: for $r_{\max} = 3.3 > r > 3.2$, for $3.2 > r > 0$, and the unstable region. The set of stabilizing rules for $r > 3.2$ is large.

As mentioned in Section 3, because of the low dimensionality of these perturbations, the optimization problem (16) can be solved numerically, given a loss function. These optimal rules were computed for l_2 loss, and the results are the curved line segment in Figure 8. The endpoints of this curve are for $r = 0$, which is the LQG rule, and for $r = 3.2 \approx r_{\max}$.

4.3. Discussion

First consider the results for the policies that are robust to mixed shock and model errors. As noted by Rudebusch and Svensson (1999) and discussed extensively by Rudebusch (2001), the LQG policy is more aggressive than the Taylor rule, with larger response coefficients on both inflation and the output gap. The policies that are robust to both shock and model uncertainty (Figure 1) are in turn more aggressive than the LQG policy.

The parameters of the optimal H_∞ and optimal l_1 robust rules appear quite different. This does not necessarily imply however that these controllers are different in terms of their objective functions. The H_∞ and l_1 objective function surfaces (inverted, so that the problem is maximization rather than minimization) are shown in Figures 2 and 3, respectively. The l_1 objective function, $\|\Lambda^{\frac{1}{2}}G\|_{l_1}$ in the notation of (14), is flat in a large region of (g_π, g_y) space. Evidently, the rule that solves the H_∞ robust control problem is nearly optimal from a l_1 perspective. The converse is not true, however, since the l_1 controller has H_∞ norm well above the minimum.

Generally speaking, as more structure is placed on the perturbations, the sets that contain the optimal policies become less aggressive, that is, they tend to exclude more aggressive policies. The optimal policy for general NTV perturbations with $r = r_{\max}$ (Figure 4) and for NTV perturbations on only the cross effects (Figure 5) are somewhat more aggressive than the LQG policy, but far less extreme than the robust policies in Figure 1. It is interesting to note that in these cases the LQG rule and even the Taylor rule have reasonably good stability properties.

The LTI perturbations are more restrictive, and the policies for the maximal perturbations become even less aggressive in the LTI case. Although the coefficient on the output gap for r_{\max} remains approximately the same as for the LQG solution, the coefficient on inflation gets smaller and approaches the limiting case of one, below which the nominal model becomes unstable. The Taylor rule and the LQG rule both have similar stability properties with LTI perturbations.

4.4. Incorporating Noisy Data

The final source of uncertainty in Poole's list in the introductory paragraph is noisy data. This source of noise has been ignored so far. Observational noise has two relevant implications. First, observational noise induced errors-in-variables bias in estimated econometric models. Second, policy rules that are optimal when there is no measurement error in general will not be so in the presence of measurement error, holding the model constant. The first of these difficulties, errors-in-variables bias, is well understood, and econometric methods are available to handle some types of measurement error. Here, we focus on the second of these implications, and examine the extent to which robust policy rules change when data noise is introduced.⁴

Orphanides (1998) argued that measurement error will temper optimal rules: A jump in inflation that in the absence of measurement error might appropriately result in an aggressive response, with measurement error, might call for a dampened response because of doubts about whether the jump is simply a consequence of mismeasurement. To examine this, we follow Orphanides (1998) and modify the nominal model to include explicitly additive noise in inflation and in the output gap. Inflation noise is white and the output gap noise is first-order autoregressive, with coefficient ρ . To simplify the computations, we considered only optimal H_∞ -robust Taylor-type rules. Three values of ρ were considered: $\rho = 0.1, 0.5,$ and 0.9 . Four sets of standard deviations of the inflation and output gap noise were

considered: (0.5, 0.5), (0.5, 1), (1, 0.5), and (2, 2). [Orphanides used (0.7, 1) for his computations, similar to our case of (0.5, 1).] Note that the standard deviation of the stochastic disturbances in the nominal model is 1.2 for the error in the inflation equation and 0.9 for the error in the output gap equation, so that the magnitude of the measurement error considered brackets the magnitude of the stochastic disturbances.

Interestingly, the H_∞ optimal rule was found to be insensitive to the introduction of this measurement error. For the nominal model, the H_∞ Taylor-rule coefficients are $(g_\pi, g_y) = (6.41, 2.75)$. When measurement error is added, the H_∞ rules become less aggressive, but only slightly so. For the noise standard deviations (0.5, 0.5), (0.5, 1), (1, 0.5), for all ρ considered, the H_∞ Taylor-rule coefficients are between $5.84 \leq g_\pi \leq 6.23$ and $2.59 \leq g_y \leq 2.76$. Even for the most extreme case, where the standard deviations of the measurement error are (2, 2), the H_∞ Taylor-rule coefficients are between $4.61 \leq g_\pi \leq 5.90$ and $2.25 \leq g_y \leq 2.80$.

Evidently, at least for these types of measurement error and H_∞ robust rules, the robust rules become less active, but only slightly so. One interpretation of this finding is that although adding measurement error can significantly change the LQG optimal rule, measurement error does not substantially change the worst possible models (or shocks) that are the focus of the H_∞ robust rules.

5. THEORETICAL RESULTS

The foregoing numerical analysis suggests that the robust rules are typically, but not necessarily, more aggressive than the LQG rule. In the case of NTV perturbations, the rule for the maximal perturbations was slightly more aggressive than the LQG rule, and the LQG rule was found to have a high degree of robust stability against these perturbations. Because the results are numerical, however, it is difficult to develop much intuition or to assess their generality.

This section undertakes a theoretical analysis of some robust rules in the simple two equations of Ball (1997) and Svensson (1997a,b). The nominal model considered is

$$\pi_{t+1} = \pi_t + \alpha y_t + \varepsilon_{\pi,t+1}, \quad (18a)$$

$$y_{t+1} = -\beta(i_t - \pi_t) + \delta y_t + \varepsilon_{y,t+1}, \quad (18b)$$

where α , β , and δ are constants. It is assumed that, for the nominal model, $\alpha > 0$ and $\beta > 0$. Note that this is the RS model in which higher-order lags are suppressed and the moving average $\bar{i}_t - \bar{\pi}_t$ is replaced by $i_t - \pi_t$. The control rules remain generalized Taylor rules. To simplify the analysis, we set $\lambda_{\Delta i} = 0$ in the loss functions (6) and (7), so that the policymaker is concerned only about inflation and the output gap.

The theoretical analysis has two parts. The first entails obtaining the LQG Taylor rule and examining the stability of the nominal controlled system. The second is an analysis of robust stability under NTV perturbations. Although we would like

to consider NTV perturbations to all three coefficients, for tractability we only consider NTV perturbations to a single key coefficient, α , the slope of the Phillips curve. The class of deviations considered thus is

$$D_r = \{\Delta : \Delta_{\pi\pi} = \Delta_{yy} = \Delta_{y\pi} = 0, \|\Delta_{\pi y}\| < r, \Delta \text{ is NTV}\}. \quad (19)$$

Two norms are considered for (19), the H_∞ norm (l_2 loss) and the l_1 norm (l_∞ loss). Under the H_∞ norm, in this simple case, the results for NTV, NTI, LTV, and LTI perturbations will coincide and, moreover, the results for the H_∞ norm are also the results for the l_1 norm with LTI perturbations [Dahleh and Diaz-Bobillo (1995, Ch. 7)]. The results for the l_1 norm with NTV perturbations are different, however, and are presented separately.

The LQG controller can be obtained analytically for this system. Write the risk function as $R = k * \text{var}(\pi) + (1 - k)\text{var}(y)$, where $0 \leq k \leq 1$. Then, the LQG optimal Taylor rule is

$$(g_\pi^{\text{LQG}}, g_y^{\text{LQG}}) = \left(1 + \frac{\zeta_1}{\beta\zeta_2}, \frac{(\alpha\zeta_1 + \delta\zeta_2)}{\beta\zeta_2}\right), \quad (20)$$

where $\zeta_1 = \frac{1}{2}\alpha k + \frac{1}{2}[\alpha^2 k^2 + 4k(1 - k)]^{\frac{1}{2}}$ and $\zeta_2 = \alpha\zeta_1 + 1 - k$. If the policymaker cares only about inflation (so $k = 1$), then

$$(g_\pi^{\text{LQG}}, g_y^{\text{LQG}}) = \left(1 + \frac{1}{\alpha\beta}, \frac{1 + \delta}{\beta}\right).$$

Alternatively, if the policymaker cares only about the output gap, then $(g_\pi^{\text{LQG}}, g_y^{\text{LQG}}) = (1, \delta/\beta)$.

It is convenient to introduce a reparameterization of the policy rule, specifically,

$$\gamma = -\delta + \beta g_y \quad \text{and} \quad \theta = \alpha\beta(g_\pi - 1).$$

This reparameterization is convenient because stability conditions depend on the parameters of the nominal model only through γ and θ . Under this transformation, the eigenvalues of the controlled system (18), x_1 and x_2 , are

$$\{x_1, x_2\} = \frac{1}{2}(1 - \gamma) \pm \frac{1}{2}[(1 + \gamma)^2 - 4\theta]^{\frac{1}{2}}. \quad (21)$$

Thus the controlled system is stable if and only if γ and θ are in the region of stability S :

$$S = \{\gamma, \theta : \theta > 2\gamma - 2, \theta < 1 + \gamma, \theta > 0\}. \quad (22)$$

We now turn to an analysis of the robust stability. The nominal model (18), controlled by the Taylor rule and subject to perturbations (19), is

$$\begin{bmatrix} \pi_{t+1} \\ y_{t+1} \end{bmatrix} = \begin{bmatrix} 1 & \alpha + \sigma \Delta_{\pi y} \\ \beta(1 - g_\pi) & \delta - \beta g_y \end{bmatrix} \begin{bmatrix} \pi_t \\ y_t \end{bmatrix} + \begin{bmatrix} \varepsilon_{\pi,t+1} \\ \varepsilon_{y,t+1} \end{bmatrix}, \quad (23)$$

where σ is a fixed constant chosen by the researcher to scale $\Delta_{\pi y}$ to be in natural units of the problem, for example, in units of standard errors of α . It is useful to rewrite this as the feedback system

$$\begin{bmatrix} \pi_{t+1} \\ y_{t+1} \\ s_t \end{bmatrix} = \begin{bmatrix} 1 & \alpha & 1 \\ \beta(1 - g_\pi) & \delta - \beta g_y & 0 \\ 0 & \sigma & 0 \end{bmatrix} \begin{bmatrix} \pi_t \\ y_t \\ h_t \end{bmatrix} + \begin{bmatrix} \varepsilon_{\pi,t+1} \\ \varepsilon_{y,t+1} \\ 0 \end{bmatrix}, \quad (24a)$$

$$h_t = \Delta_{\pi y} s_t. \quad (24b)$$

If the policy rule (g_π, g_y) stabilizes the nominal model, then introducing the additional feedback to the systems leaves it stable if and only if the feedback interconnection between h_t and s_t is stable [Dahleh and Diaz-Bobillo (1995, Ch. 5)]. In the frequency domain, the interconnection is represented by two channels: feed-forward $s = Gh$ and feedback $h = \Delta s$, where

$$\begin{aligned} G(z) &= \begin{bmatrix} 0 & \sigma \end{bmatrix} \begin{bmatrix} z^{-1} - 1 & -\alpha \\ -\beta(1 - g_\pi) & z^{-1} - \delta + \beta g_y \end{bmatrix}^{-1} \begin{bmatrix} 1 \\ 0 \end{bmatrix} \\ &= -\frac{\sigma}{\alpha} \theta z^2 [(1 - x_1 z)(1 - x_2 z)]^{-1}, \end{aligned} \quad (25)$$

where x_1 and x_2 are given in (21). Because Δ is one-dimensional, by the small-gain theorem in its necessary and sufficient form, the interconnection is stable for all Δ such that $\|\Delta\| < r$ if and only if $\|G\| \leq 1/r$, that is, $r < 1/\|G\|$, where $\|G\|$ denotes either the H_∞ or l_1 norm of G , depending on which is specified in (19).

We now are able to characterize the robust rules for the maximal perturbation, r_{\max} , for this controlled system. Results for l_2 loss [H_∞ norm in (19)] are given in Proposition 1, and results for l_∞ loss [l_1 norm in (19)] are given in Proposition 2.

PROPOSITION 1 (l_2 loss). *Let $D_r = \{\Delta : \Delta_{\pi\pi} = \Delta_{yy} = \Delta_{y\pi} = 0, \|\Delta_{\pi y}\|_{H_\infty} < r, \Delta$ is NTV}. Then the maximal radius of affordable perturbations is equal to α/σ . The set of robust rules for the maximal perturbation is $R_{H_\infty} = \{\gamma, \theta : \gamma \leq 1, 0 < \theta \leq (1 + \gamma)^2 / (3 + \gamma)\}$.*

PROPOSITION 2 (l_∞ loss). *Let $D_r = \{\Delta : \Delta_{\pi\pi} = \Delta_{yy} = \Delta_{y\pi} = 0, \|\Delta_{\pi y}\|_{l_1} < r, \Delta$ is NTV}. Then the maximal radius of affordable perturbations is equal to α/σ . The set of robust rules for the maximal perturbation is $R_{l_1} = \{\gamma, \theta : \gamma \leq 1, 0 < \theta \leq (1 + \gamma)^2 / 4\}$.*

Proofs are given in the Appendix.

The results of these propositions are summarized in Figure 9. The region of stability S is the large triangle in the figure. The area in the triangle to the right of the curved line corresponds to complex conjugate pairs of eigenvalues, whereas the area to the left of the line corresponds to real eigenvalues. The set of optimal LQG rules for different risks (i.e., different values of k , $0 \leq k \leq 1$) is given by a

straight-line segment connecting points $(0,0)$, which is optimal when all weight is on the output gap ($k = 0$), and $(1,1)$, which is optimal when all weight is on inflation ($k = 1$). The shaded area A is the region R_{H_∞} in which the radius of allowable perturbations is maximal as given in Proposition 1, and the region $A \cup B$ is the region R_{I_1} in which the radius of allowable perturbations is maximal as given in Proposition 2.

As can be seen from Figure 9, the proposition has two relevant implications. First, the LQG rules all fall in region A , and so, every LQG rule is maximally robust. Second, for any robust rule in R_{H_∞} or R_{I_1} , there exists a no less aggressive LQG rule. For these perturbations, at least, the maximally robust rules need not be more aggressive than the LQG rules.

6. DISCUSSION AND CONCLUSIONS

Several substantive conclusions emerge from this analysis. For many classes of perturbations, the robust rules are more aggressive than the LQG rules. An intuitive explanation of this finding is that a particularly undesirable situation for the monetary authorities is when the uncontrolled nominal model is unstable but monetary policy has virtually no effect; in this case, an aggressive rule is called for. This rule produces high volatility if the model is close to the nominal model, but still minimizes the worst-case risk. However, this finding that the robust rules are more aggressive than the LQG rules is not true for all types of perturbations, as shown both by the numerical analysis of Section 4 and by the theoretical results of Section 5.

The specific robust control rules vary, depending on the formulation of uncertainty. This analysis suggests some guidance about which formulation of uncertainty seems most appropriate for the monetary policy application. The formulations that involve both error and model uncertainty confound these two sources, and the worst cases can be achieved by the appropriate structure of error uncertainty alone. Thus these seem not to focus on the characterization of model uncertainty. Among the formulations that involve model uncertainty only, those with NTV perturbations allow the most general deviations of the approximate model from the true relations. The formulation of uncertainty with the greatest amount of structure, LTI perturbations to only the two leading cross terms, is closest to the structure that would arise in a Bayesian formulation (although the Bayesian formulation would in general admit uncertainty on all coefficients). This latter formulation can be thought of as a one-time shift in the model, or alternatively as the true model being of the RS form but with different coefficients. In our view, this is unduly restrictive and does not do justice to the notion that the RS model is an approximation to a much more complicated, unknown, and possibly time varying set of relations. In this light, we find the nonparametric generality of the NTV perturbations on all coefficients more appealing.

Many issues remain for future work. The computations here do not involve discounting, and one could argue that policymaker loss functions should discount

the future. We have not taken a stand on the issue of the choice of r ; although $r = r_{\max}$ is a concrete choice, this might result in unduly conservative policies. There also remains the difficult problem of computing the optimal (minimax risk) rule within the set of rules with finite maximal risk. These constitute interesting questions for future research.

NOTES

1. Applications of these concepts in economics include Tallarini (1996), Hansen and Sargent (1998), Tornell (1998), Sargent (1999), Stock (1999, in press), and Anderson et al. (2000).

2. To analyze cases in which z_t is nonstationary but ergodic, one can consider an alternative formulation of the risk function, $R_{l_2} = \lim_{N \rightarrow \infty} N^{-1} \sum_{t=1}^N \|\Lambda^{\frac{1}{2}} z_t\|_2^2$. We adopt this formulation when considering time-varying perturbations to the model.

3. The formulations of quadratic loss and risk adopted here match those in typical treatments of monetary control problems in linear-quadratic problems. They differ, however, from typical formulations in the robust-control engineering literature, in which ε and z are treated as deterministic elements of some normed space of infinite sequences. A dynamic system then becomes an operator transferring ε to z , and robust control is the control that minimizes the norm of the operator. Typically, the normed space is the space of all n -dimensional vector-valued real sequences on the integers with finite l_2 , P , or l_∞ norms, where for the sequence z , $\|z\|_2 = (\sum_{t=-\infty}^{\infty} \sum_{i=1}^n z_{it}^2)^{\frac{1}{2}}$, $\|z\|_\infty = \sup_i \max_i |z_{it}|$, and $\|z\|_P = \lim_{N \rightarrow \infty} [(1/2N) \sum_{t=-N}^N z_t' z_t]^{\frac{1}{2}}$ (power norm). In this paper, the robust control problems stated in terms of quadratic loss and risk can be reformulated in terms of norms of deterministic sequences. Perhaps the easiest way to see this is to note that the square of the power norm of a mean zero ergodic stationary process with realization $\{z_t\}$ is equal to $\text{tr } E z_t z_t'$. One also can show that, for most purposes in the paper, the formulations of problems in terms of quadratic loss and in terms of deterministic sequences with l_2 norm are interchangeable. Whether z_t is being treated as a deterministic element of a normed space or as a stochastic process with finite variance should be clear from context by the dimension of the process and by whether an expectation is taken.

4. We thank Athanasios Orphanides for suggesting this exercise to us.

REFERENCES

- Anderson, E.W., L.P. Hansen & T.J. Sargent (2000) Robustness, Detection, and the Price of Risk. Manuscript, University of Chicago.
- Ball, L. (1997) Efficient Rules for Monetary Policy. NBER working paper 5952.
- Ball, L. (1999) Policy rules for open economies. In J. Taylor (ed.), *Monetary Policy Rules*, pp. 127–144. Chicago: University of Chicago Press for the NBER.
- Brainard, W. (1967) Uncertainty and the effectiveness of policy. *American Economic Review* 57, 411–425.
- Chow, G.C. (1973) Effect of uncertainty on optimal control policies. *International Economic Review* 14, 632–645.
- Dahleh, M.A. & I.J. Diaz-Bobillo (1995) *Control of Uncertain Systems: A Linear Programming Approach*. Englewood Cliffs, NJ: Prentice Hall.
- Easley, D. & N.M. Kiefer (1988) Controlling a stochastic process with unknown parameters. *Econometrica* 56, 1045–1064.
- Gordon, R.J. (1998) Foundations of the Goldilocks economy: Supply shocks and the time-varying NAIRU. *Brookings Papers on Economic Activity* (2), 297–333.
- Hansen, L.P. & T.J. Sargent (1998) Discounted Robust Linear Decision and Estimation in the Frequency Domain. Manuscript, University of Chicago.
- McCallum, B.T. (1988) Robustness properties of a rule for monetary policy. *Carnegie-Rochester Conference Series on Public Policy* 29, 175–203.

- McCallum, B.T. (1990) Targets, indicators, and instruments of monetary policy. In W.S. Haraf & P. Cagan (eds.), *Monetary Policy for a Changing Financial Environment*, pp. 44–70. Washington DC: AEI Press.
- Orphanides, A. (1998) Monetary Policy Evaluation with Noisy Information. Manuscript, Board of Governors of the Federal Reserve System.
- Poole, W. (1998) A policymaker confronts uncertainty, pp. 3–8. In *Review*, Vol. 80, No. 5. Federal Reserve Bank of St. Louis.
- Rudebusch, G.D. (2001) Is the Fed too timid? Monetary policy in an uncertain world. *Review of Economics and Statistics* 83, 203–217.
- Rudebusch, G.D. & L.E.O. Svensson (1999) Policy rules for inflation targeting. In J. Taylor (ed.), *Monetary Policy Rules*, pp. 203–246. Chicago: University of Chicago Press for the NBER.
- Sargent, T. (1999) Comment on Ball. In J. Taylor (ed.), *Monetary Policy Rules*, pp. 144–154. Chicago: University of Chicago Press for the NBER.
- Stock, J.H. (1999) Comment on Rudebusch and Svensson. In J. Taylor (ed.), *Monetary Policy Rules*, pp. 253–259. Chicago: University of Chicago Press for the NBER.
- Stock J.H. (in press) Monetary policy in a changing economy: Indicators, rules, and the shift towards intangible output. In *Proceedings of the Eighth International Conference on Monetary Policy*, Bank of Japan.
- Svensson, L.E.O. (1997a) Inflation forecast targeting: Implementing and monitoring inflation targets. *European Economic Review* 41, 1111–1146.
- Svensson, L.E.O. (1997b) Inflation Targeting: Some Extensions. NBER working paper 5962.
- Tallarini, T.D., Jr. (1996) *Risk-Sensitive Real Business Cycles*. Ph.D. Dissertation. University of Chicago.
- Taylor, J.B. (1993) Discretion versus policy rules in practice. *Carnegie-Rochester Conference Series on Public Policy* 39, 195–214.
- Tornell, A. (1998) H_∞ Forecasting and Excess Volatility of Asset Prices. Manuscript, Harvard University.
- Whittle, P. (1990) *Risk-Sensitive Optimal Control*. New York: Wiley.
- Wieland, V. (2000) Monetary policy, parameter uncertainty and optimal learning. *Journal of Monetary Economics* 46, 199–228.
- Wieland, V. (1998) Monetary Policy and Uncertainty About the Natural Unemployment Rate. Manuscript, Federal Reserve Board.
- Zhou, K., J.C. Doyle & K. Glover (1996) *Robust and Optimal Control*. Prentice-Hall: Upper Saddle River, NJ.

APPENDIX

A.1. COMPUTATION OF RADIUS OF AFFORDABLE DEVIATIONS

As is discussed in Section 3, it is generally not computationally feasible to compute the optimum solution to (16), but it is possible to compute the set of control rules with finite risk uniformly over D_r , that is, $\{g_\pi, g_y : \sup_{\Delta \in D_r} R < \infty\}$. This set will necessarily contain the optimum. This Appendix describes the computation of this set for the structured perturbations of Section 3.2. The computation relies on the so-called structured small-gain theorem [Dahleh and Diaz-Bobillo (1995, Ch. 7)], and involves the following steps:

First, the perturbed model is rewritten as a nonperturbed model with additional input, output, and feedback connections. The general perturbed model is

$$z_{t+1} = [A(L) + \Delta]z_t + B\varepsilon_{t+1}, \quad (\text{A.1})$$

where z_t , $A(L)$, B , and ε are defined following (5) and Δ is given in (17). This system, rewritten without perturbations but with additional input, output, and feedback connections, is

$$\begin{bmatrix} z_{t+1} \\ s_{t+1} \end{bmatrix} = \begin{bmatrix} A_{11}^+ & A_{12}^+ \\ A_{21}^+ & 0 \end{bmatrix} \begin{bmatrix} z_t \\ h_t \end{bmatrix} + \begin{bmatrix} B \\ 0 \end{bmatrix} \varepsilon_{t+1} \quad (\text{A.2a})$$

$$h_t = \tilde{\Delta}s_t \quad (\text{A.2b})$$

where $\tilde{\Delta} = \text{diag}(\tilde{\Delta}_{\pi\pi}, \Delta_{\pi y}, \tilde{\Delta}_{y\pi}, \Delta_{yy})$, $A_{11}^+ = A(L)$, and

$$A_{12}^+ = \begin{bmatrix} 1 & 1 & 0 & 0 \\ 0 & 0 & 1 & 1 \\ g_\pi/4 & g_\pi & g_y & g_y \\ g_\pi/4 & g_\pi & g_y & g_y \end{bmatrix}, \quad A_{21}^+ = \begin{bmatrix} 1-L & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ \frac{1+L+L^2+L^3}{4} & 0 & 0 & -\frac{1+L+L^2+L^3}{4} \\ 0 & 1 & 0 & 0 \end{bmatrix}.$$

Next, fix g_π and g_y . If the nominal system is stable for these control parameters, then (A.2) is stable if and only if the feedback interconnection between h and s is stable. Denote the transfer function between h and s as G , so that the interconnection is $s = Gh$ and $h = \tilde{\Delta}G$. By the structured small-gain theorem, this interconnection is stable for all $\tilde{\Delta}$ such that $\|\tilde{\Delta}\| < r$ if and only if the structured norm of G , $\mu_{\Delta}(G)$, is less than or equal to the reciprocal of r , where $\mu_{\Delta}(G) = \{\inf[\|\tilde{\Delta}\| : \tilde{\Delta} \in \Delta, (I - G\tilde{\Delta})^{-1} \text{ is not stable}]\}^{-1}$, where Δ represents the structure that was assumed for $\tilde{\Delta}$. In our case, Δ contains the diagonal operators, $\text{diag}(\Delta_1, \Delta_2, \Delta_3, \Delta_4)$. The structured norm is computed as outlined by Dahleh and Diaz-Bobillo (1995, Ch. 7). The radius of allowable perturbations for this value of (g_π, g_y) is the reciprocal of the structured norm. Using a grid search over (g_π, g_y) , this radius is computed for all (g_π, g_y) for which the nominal controlled system is stable. The set $\{g_\pi, g_y : \sup_{\Delta \in D_r} R < \infty\}$ is the set of points for which the radius of allowable perturbations is at least r .

A.2. PROOFS OF PROPOSITIONS 1 AND 2

Proof of Proposition 1. Consider first the case when $x_1, x_2 \in \mathcal{R}$, $|x_1|, |x_2| < 1$. This corresponds to γ, θ such that $\theta > 0$, $\theta > 2\gamma - 2$, and $\theta \leq (1 + \gamma)^2/4$. From (25) it is clear that

$$\begin{aligned} \|G\|_{H_\infty} &= \sup_{|z|=1} \left| \frac{z^2 \sigma \theta}{\alpha(1-x_1z)(1-x_2z)} \right| = \frac{\sigma \theta}{\alpha} \inf_{|z|=1} |z-x_1||z-x_2| \\ &= \frac{\sigma \theta}{\alpha} \inf_{\omega \in [0, \pi]} |e^{i\omega} - x_1| |e^{i\omega} - x_2| \\ &= \frac{\sigma \theta}{\alpha} \inf_{\omega \in [0, \pi]} \left[(1+x_1^2 - 2x_1 \cos \omega)(1+x_2^2 - 2x_2 \cos \omega) \right]^{\frac{1}{2}}. \end{aligned}$$

It is straightforward to check that the infimum is attained at $\omega = 0$ if $\gamma \leq 1$ and at $\omega = \pi$ if $\gamma > 1$. Therefore, if $\gamma \leq 1$, $\|G\|_{H_\infty} = \sigma/\alpha$ and if $\gamma > 1$, $\|G\|_{H_\infty} = \sigma\theta/\alpha(-2\gamma + 2 + \theta) > \sigma/\alpha$.

Now consider the case in which x_1, x_2 are complex and $|x_1|, |x_2| < 1$. In terms of γ, θ , this corresponds to $\theta > (1 + \gamma)^2/4$ and $\theta < 1 + \gamma$. Substituting $z = e^{i\omega}$, $x_1 = \rho e^{i\varphi}$, and $x_2 = \rho e^{-i\varphi}$ in (25), and making some algebraic transformations, we get

$$\sup_{|z|=1} |G(z)| = \frac{\sigma\theta}{\alpha} \inf_{\omega \in [0, \pi]} (a + b \cos \omega + c \cos^2 \omega)^{\frac{1}{2}},$$

where $a = (1 - \rho^2)^2 + 4\rho^2 \cos^2 \varphi$, $b = -4\rho(1 + \rho^2) \cos \varphi$, and $c = 4\rho^2$. Optimizing with respect to ω , we find that $\sup_{|z|=1} |G(z)|$ equals σ/α if $\gamma \leq 1$ and $\theta \leq (1 + \gamma)^2/(3 + \gamma)$, and it equals $(\sigma/\alpha)\theta/(-2\gamma + 2 + \theta) > \sigma/\alpha$ if $\gamma > 1$ and $\theta < -[(1 - \gamma)^2 - 4\gamma]/(5 - \gamma)$. Finally, $\sup_{|z|=1} |G(z)| = [\sigma\theta/\alpha(1 + \gamma - \theta)] * [1 - (1 - \gamma)^2/4(\theta - \gamma)]^{-1} > \sigma/\alpha$ for the rest of θ and γ . In summary, $r_{\max} = 1/\inf_{\theta, \gamma} \sup_{|z|=1} |G(z)| = \alpha/\sigma$ is a maximal radius of affordable perturbations. The set of robust rules that correspond to it satisfy $\{\gamma, \theta : \gamma \leq 1, 0 < \theta \leq (1 + \gamma)^2/(3 + \gamma)\}$. ■

Proof of Proposition 2. From (25), it is clear that $\|G\|_{l_1} \geq \sigma\theta/\alpha$ and that equality holds for $x_1 = x_2 = 0$. This would happen if $\gamma = 1$ and $\theta = 1$. Immediately, $\min \|G\| \leq \sigma/\alpha$, $r_{\max} \geq \alpha/\sigma$, and the robust rules for the maximal perturbations have $\theta \leq 1$.

Next, it is shown that, for any rule in R , $\|G\| = \sigma/\alpha$. First write $G(z) = z^2 \sum_{n=0}^\infty G_n z^n$, where $G_n = -(\sigma\theta/\alpha)(x_1^{n+1} - x_2^{n+1})/(x_1 - x_2)$. If $\{\gamma, \theta\} \in R$, then the eigenvalues are real and, by direct calculation, $\|G\| = (\sigma\theta/\alpha) \sum_{n=0}^\infty (x_1^{n+1} - x_2^{n+1})/(x_1 - x_2) = \sigma/\alpha$; thus the maximal radius of affordable perturbations for $\{\gamma, \theta\} \in R$ is the same and is α/σ .

The rest of the rules in S are either from $R_1^c = \{\gamma, \theta : \gamma > 1, \theta \leq 1\}$, for which the eigenvalues of the nominal system are real, or from $R_2^c = \{\gamma, \theta : \theta > (1 + \gamma)^2/4, \theta \leq 1\}$, for which x_1 and x_2 are a pair of complex conjugates. First, consider the set R_1^c . For rules in this set, $\theta = (1 + \gamma)^2/4 - c$ for some $c > 0$. Along this curve, $x_1 = \frac{1}{2}(1 - \gamma) + c^{1/2}$ and $x_2 = \frac{1}{2}(1 - \gamma) - c^{1/2}$. Therefore, $|G_n| = (\sigma\theta/\alpha)|(\frac{1}{2} - \frac{1}{2}\gamma + c^{1/2})^{n+1} - (\frac{1}{2} - \frac{1}{2}\gamma - c^{1/2})^{n+1}|/2c$. It can be shown that $\partial|G_n|/\partial\gamma < 0$. It follows that, in this region, $\|G\| > \sigma/\alpha$. The eigenvalues in R_2^c are complex, and in this region it was verified numerically that $\|G\| > \sigma/\alpha$. Hence, the minimum of $\|G\|$ is σ/α , which is attained iff the rule is in R . By the small-gain theorem, the maximal radius of affordable perturbations is $r_{\max} = 1/\|G\| = \alpha/\sigma$. ■




FIGURE 1. LQG, H_∞ , minimum entropy, and l_1 , optimal rules.




FIGURE 2. Inverse of H_∞ norm of transfer function for the nominal case.




FIGURE 3. Inverse of l_1 norm of transfer function for the nominal case.




FIGURE 4. Radius of affordable l_1 NTV perturbations on four coefficients.




FIGURE 5. Radius of affordable l_1 , NTV perturbations on two coefficients.




FIGURE 6. Radius of affordable LTI perturbations on four coefficients.




FIGURE 7. Radius of affordable LTI perturbations on two coefficients.




FIGURE 8. Conservative bounds on radius and optimal rules for some radii.




FIGURE 9. Stability analysis in transformed parameter space: Large triangle = region of stability; Region A = region with maximal l_1 radius of allowable structured NTV perturbations; Region $A \cup B$ = region with maximal H_∞ radius of allowable structured NTV perturbations; LQG rules = straight-line segment connecting $(0,0)$ and $(1,1)$.