

Optimal Constrained Interest-rate Rules

George W. Evans Bruce McGough
University of Oregon Oregon State University

May 19, 2005

Abstract

We show that if policy-makers compute the optimal unconstrained interest-rate rule within a Taylor-type class, they may be led to rules that generate indeterminacy and/or instability under learning. This problem is compounded by uncertainty about structural parameters since an optimal rule that is determinate and stable under learning for one calibration may be indeterminate or unstable under learning under a different calibration. We advocate a procedure in which policy-makers restrict attention to rules constrained to lie in the determinate learnable region for all plausible calibrations, and that minimize the expected loss, computed using structural parameter priors, subject to this constraint.

JEL classification: E52, E32, D83, D84.

Keywords: Monetary policy, Taylor rules, indeterminacy, learning, E-stability, parameter uncertainty, robust rules.

1 Introduction

The development of tractable forward looking models of monetary policy, together with the influential work of [38], has led to considerable interest in the performance of Taylor-type interest rate rules.¹ These rules take the nominal interest rate as the policy instrument and direct the central bank to set this rate according to a simple dependence on current, lagged or expected

¹For a recent survey and extended analysis, see [35].

inflation and output gap, and possibly on a term generating interest rate smoothing. Extended Taylor-type rules would allow for a dependence on all these variables and also on observable exogenous shocks.

While these simple policy rules have clear advantages, it has been noted by a number of authors, e.g. [4], [43], [36] and [10], that the corresponding models exhibit indeterminate steady-states for large regions of the reasonable parameter space. This is undesirable because associated with each indeterminate steady-state is a continuum of equilibria depending on extraneous variables known as “sunspots”, and the particular equilibrium on which agents ultimately coordinate may be suboptimal.

A distinct and equally troubling concern, which can arise whether the steady state is determinate or indeterminate, is the possibility that the choice of interest rate rule results in the intended equilibrium being unstable under least-squares learning by private agents; see [6], [7], [15], [16] and [17]. In such cases, the economy under learning would fail to deliver the intended equilibrium and might either follow a divergent path or converge on a sunspot equilibrium. The issue of whether sunspot solutions can be stable under learning, in New Keynesian models, is also of interest and has recently received some attention.²

The current paper examines these issues for the New Keynesian model in greater detail. [6] and [17] show that indeterminacy and instability under learning can arise for plausible but *ad hoc* Taylor rules. Here we assess the hazards of indeterminacy and instability under learning when policy-makers choose a Taylor-type rule optimally in the sense that they minimize their loss function. One might anticipate that optimized Taylor rules would generate determinate steady-states and learnable equilibria, but we show that this is not the case: unstable indeterminacy, stable indeterminacy and unstable determinacy are all possibilities.³

²[42] demonstrated the possibility of sunspot solutions being stable under learning in simple overlapping generations models. For local stability conditions in purely forward-looking models see [11] and [14]. [19] provide corresponding stability conditions for models with a predetermined variable. For the possibility of stable sunspots in the New Keynesian Monetary model, see [25] and [17]. Results on PPP rules in a small open-economy set-up are given in [45].

³Some partial results along these lines have already been obtained. [32] evaluated optimal policy for several classes of Taylor-type rules and found cases in which the optimal rule lay within the indeterminate region. However, they did not investigate stability under learning. [15], [16] showed that certain types of rules designed to be fully optimal can be unstable under learning, but their examples are restricted to “fundamentals-based” rules

Similar analysis can be conducted on extended Taylor-type rules that depend on exogenous shocks as well as endogenous variables. Such rules are sufficiently general that they are in principle able to implement unconstrained optimal policy, i.e. minimize the government loss function. However, there is a multiplicity of “optimal” rules of this type, and we find that optimal rules consistent with both stable and unstable indeterminacy abound. A numerical search algorithm, if left unconstrained, would not distinguish between a rule leading to stable determinacy and a rule subject to instability and/or indeterminacy problems. The clear conclusion from the first part of our analysis is therefore that, whether one is considering plausible “ad hoc” rules or rules that aim to implement optimal policy, unconstrained policy can yield undesirable outcomes and it is crucial to explicitly impose both determinacy and stability under learning constraints.

Having established this point, we turn to the issue of parameter uncertainty. Alternative calibrations of the New Keynesian model differ greatly both in terms of the structural parameters controlling the interest elasticity of demand and output elasticity of inflation, and in terms of the degree of inertia. Are optimal constrained policies obtained under one specification robust to alternative calibrations? Our second main conclusion is that parameter uncertainty greatly compounds instability and indeterminacy concerns: a rule that is constrained optimal with respect to one calibration may result in stable sunspots, instability or even explosiveness under another calibration. This point holds regardless of the true degree of inertia in the structural model.

To address the additional problems arising from parameter uncertainty, we consider the existence of “robust” policy. At issue here is whether there exist policy rules that yield stable determinacy across calibrations. Strikingly we find that such rules do exist, and thus we can search for the optimal rule within this class. That is, given a *prior* distribution over structural parameters, we can use a “Bayesian” approach to compute the optimal interest-rate rule subject to the requirement that it generate determinate stable solutions across all calibrations which obtain with positive probability.⁴

that depend only on exogenous shocks or lagged endogenous variables. Our analysis also goes beyond these papers in investigating policy across calibrations and in the presence of structural parameter uncertainty.

⁴Our approach is similar in spirit to the one advocated by [5], but our sense of robustness is different: they use a Bayesian approach that incorporates a minimax element placing additional weight on worst-case outcomes, whereas we guard against indeterminacy and

Our concrete numerical results are easily summarized. We use a *prior* with subjective weights that reflects the wide diversity of opinion on appropriate parameter values. The robust optimal constrained Taylor-type rule has heavy interest-rate smoothing and places substantial weight on expected future, expected current, and lagged inflation and output, and on observed exogenous shocks. However, we also provide a quite simple robust Taylor-type rule, which has only a 1.3% deterioration in performance, and we further find that fairly satisfactory results can even be obtained using simple Taylor-rules with interest-rate smoothing and well-chosen coefficients.

Our results are derived within a standard but very simple linearized New Keynesian model. However, our approach has general applicability and could in principle be applied to more detailed and realistic models. The key steps are, first, to obtain the set of policies that are robust in the specific sense that they yield a model that is determinate and an equilibrium that is stable under learning, and then to obtain the policy that maximizes the expected value of the policy-maker objective subject to this constraint.

The paper is organized as follows. Section 2 summarizes the monetary policy framework and the tools for analyzing determinacy and stability under learning. Section 3 presents the basic issues in the context Taylor-type rules and summarizes the analogous results for extended Taylor rules. Section 4 extends our analysis of optimal constrained policy to allow for parameter uncertainty, and Section 5 concludes.

2 Framework

We study optimal policy using several variants of the New Keynesian Monetary model that have in common the following forward-looking linearized IS-AS curves:

$$IS : x_t = -\phi(i_t - E_t\pi_{t+1}) + \delta E_t x_{t+1} + (1 - \delta)x_{t-1} + g_t \quad (1)$$

$$AS : \pi_t = \beta(\gamma E_t\pi_{t+1} + (1 - \gamma)\pi_{t-1}) + \lambda x_t + u_t. \quad (2)$$

Here x_t is the proportional output gap, π_t is the inflation rate, i_t is the nominal interest rate, and g_t and u_t are independent, exogenous, stationary, zero mean AR(1) shocks with damping parameters $0 \leq \rho_g < 1$ and $0 \leq \rho_u < 1$ respectively.

instability under learning. [5] also address estimation issues.

The first equation is a formulation of the forward-looking IS curve amended to include inertia. This functional form may be obtained from a linearized model of optimizing behavior on the part of consumers. In some cases we also allow for an inertial term x_{t-1} , which is present due to habit formation: see for example [34]. The second equation is the forward-looking Phillips curve. When $\gamma = 1$, equation (2) is the pure forward-looking New Keynesian “AS” relationship based on “Calvo pricing,” and employed in [9] and Ch. 3 of [44].⁵ Here $0 < \beta < 1$ is the discount factor. Again, this equation is obtained as the linearization around a steady state.⁶ The specification of the AS curve in the case $0 < \gamma < 1$ incorporates an inertial term and is similar in spirit to [20], the Section 4 model of [21], and the Ch. 3, Section 3.2 model of [44], each of which allows for some backward looking elements. Models with $0 < \gamma < 1$ are often called “hybrid” models, and we remark that in some versions, such as [20], $\beta = 1$, so that the sum of the forward and backward looking components is one, while in other versions $\beta < 1$ is possible.⁷

The structural model is closed by specifying a policy rule describing how interest rates are set. The indeterminacy and instability regions depend critically on the specification of this policy rule. We analyze a number of policy rules, which we parameterize as follows:

$$PR_1 : i_t = \alpha_\pi E_t \pi_t + \alpha_x E_t x_t \tag{3}$$

$$PR_2 : i_t = \alpha_\pi \pi_{t-1} + \alpha_x x_{t-1} \tag{4}$$

$$PR_3 : i_t = \alpha_\pi E_t \pi_{t+1} + \alpha_x E_t x_{t+1}. \tag{5}$$

We previously studied the determinacy and stability properties of this set of rules (as well as others) in [17], but here we consider the issue of optimal policy. PR_1 , PR_2 , and PR_3 are the rules examined by [6]. We have omitted the intercepts for convenience, and in each policy rule π_t can be interpreted as the deviation of inflation from its target. These are all Taylor-type rules in the spirit of [38]. We assume throughout that $\alpha_\pi, \alpha_x \geq 0$ and thus the $\alpha_\pi E_t \pi_t$ term in PR_1 indicates the degree to which monetary policy authorities raise i_t in response to an upward deviation of $E_t \pi_t$ from its target.

⁵For the version with mark-up shocks see [44] Chapter 6, Section 4.6.

⁶One limitation of our approach is that our set-up employs a linearized structural model. When non-linearities are present other REE may exist even if the linearized model is locally determinate, a point emphasized in [2]. When nonlinearities are important it may be desirable to strengthen the constraints on policy rules that we advocate.

⁷To remain consistent with the work of [20], when we include inertia in our analysis, we set $\beta = 1$. For cases without inertia we use the standard calibration value $\beta = 0.99$.

Taylor’s original formulation specified dependence on current values of endogenous variables, but the assumption that current data on inflation and the output gap are available to policymakers when interest rates are set has been met with criticism: see for example [28]. [6] look at three natural alternatives: a slight modification of Taylor’s formulation yields PR_1 in which policy makers condition their instrument on expected values of current inflation and the output gap⁸; in PR_2 policy makers respond to the most recent observed values of these variables; and in PR_3 they respond instead to forecasts of future inflation and the output gap.

This list of rules is far from exhaustive. In particular, it is quite natural to include dependence on fundamental shocks and to analyze more general rules that nest $PR_1 - PR_3$ as special cases; in fact, some forms of these more general rules allow for the implementation of the best possible equilibrium as measured by the government’s objective. We consider some of these more general rules below, but to provide better context, we put off their discussion until Section 3.3.

The model is said to be determinate if there is a unique nonexplosive REE and indeterminate if there are multiple nonexplosive solutions.⁹ The determinacy of a model can be analyzed by writing the reduced form equation as a difference equation with the associated extraneous noise terms capturing the errors in the agents’ forecasts of the free variables. If the nonexplosive requirement of a rational expectations equilibrium pins down the forecast errors, that is, if the dimension of the unstable manifold is equal to the number of free variables, then the model is determinate. The methodology for assessing determinacy in linearized models is well known, and we refrain from presenting the details. For the monetary models and interest rate rules considered in this paper the specifics are given in our earlier paper [17].

If the model is indeterminate, there exist multiple stationary solutions, including those that depend on extrinsic fluctuations or “sunspots” as well as solutions that “over-react” to intrinsic shocks. Furthermore, a given sunspot solution has alternative representations, a point that is important if one is interested in whether sunspot solutions are stable under learning. Again, these issues are discussed at length in [17].

⁸The time t expectations of $t + s$ -dated endogenous variables ($s = 0, 1$) are assumed conditioned on t -dated exogenous variables, and all $t - 1$ -dated information.

⁹By “nonexplosive” we mean that for each t the conditional expectations $E_t |x_{t+s}|$ and $E_t |\pi_{t+s}|$ are uniformly bounded over s . For a detailed discussion of this and related concepts see [19].

We now discuss the issue of stability under learning. If the model is determinate, so that there is a unique non-explosive REE (rational expectations equilibrium), it is desirable that the solution be stable under learning. By this we mean that there is convergence to the solution if private agents in the economy estimate and update the coefficients of their forecast functions using least squares regressions. Because the models are self-referential, i.e. the evolution of the economy depends on how agents form expectations, the stability of an REE under least squares learning cannot be taken for granted.

More specifically, the structural model combined with the interest rate rule can be written in reduced form as follows:

$$y_t = AE_t^* y_{t+1} + BE_t^* y_t + Cy_{t-1} + D\hat{g}_t, \quad (6)$$

where $y_t' = (x_t, \pi_t)$ and $\hat{g}_t' = (g_t, u_t)$. We now write E_t^* to indicate that we no longer impose rational expectations. Thus we are treating the IS and AS equations (1)-(2) as arising from the aggregation of individual decisions, which depend on expected and lagged output and inflation.¹⁰ At issue is how agents form their time t expectations E_t^* . In the determinate case the unique nonexplosive solution takes the form

$$y_t = \bar{a} + \bar{b}y_{t-1} + \bar{c}\hat{g}_t, \quad (7)$$

for particular values \bar{a} , \bar{b} and \bar{c} .

Under least squares learning (7) is treated as the econometric specification of a forecasting rule, the parameters of which are estimated by the private agents. The specification is often referred to as a Perceived Law of Motion (PLM). Combining the regressors into the vector $X_t' = (1, y_{t-1}', \hat{g}_t')$ and writing the parameters as $\Theta' = (a, b, c)$, the PLM can be written as $y_t = \Theta' X_t$. Under learning agents obtain least squares estimates $\Theta_t' = (a_t, b_t, c_t)$ using data through time t and then use the estimated PLM to form their forecasts $E_t^* y_{t+1}$ and $E_t^* y_t$, which in turn influence the path of y_t . The question is then whether or not $(a_t, b_t, c_t) \rightarrow (\bar{a}, \bar{b}, \bar{c})$ as $t \rightarrow \infty$. If so, we say that the solution is stable under learning.

We use expectational stability as our criterion for judging whether agents may be able to coordinate on specific solutions, including in particular sunspot equilibria. This is because, for a wide range of models and solutions, E-stability has been shown to govern the local stability of REE under least

¹⁰For detailed discussion on this point for the New Keynesian model, see [16].

squares learning. In many cases this correspondence can be proved, and in cases where this cannot be formally demonstrated the “E-stability principle” has been validated through simulations. For a thorough discussion of E-stability see [12].

The E-stability technique is based on a mapping from the PLM to the corresponding Actual Law of Motion (ALM) parameters. For the case at hand, if agents believed in the PLM (a, b, c) then their corresponding forecasts would be given by $E_t^*y_{t+1} = a + bE_t^*y_t + cE_t^*\hat{g}_{t+1}$. Let ρ denote the 2×2 diagonal matrix with elements ρ_g, ρ_u . Using $E_t^*y_t = a + by_{t-1} + c\hat{g}_t$, and assuming for convenience that ρ is known, so that $E_t^*\hat{g}_{t+1} = \rho\hat{g}_t$, yields

$$E_t^*y_{t+1} = (I_2 + b)a + b^2y_{t-1} + (bc + c\rho)\hat{g}_t.$$

Inserting $E_t^*y_t$ and $E_t^*y_{t+1}$ into (6) and solving for y_t as a linear function of an intercept, y_{t-1} and \hat{g}_t yields the corresponding ALM parameters induced by the PLM.

$$a \rightarrow A(I_2 + b)a + Ba \tag{8}$$

$$b \rightarrow Ab^2 + Bb + C \tag{9}$$

$$c \rightarrow A(bc + c\rho) + Bc + D. \tag{10}$$

Equations (8)-(10) define a mapping from PLM parameters Θ to the ALM parameters $T(\Theta)$. The REE $\bar{\Theta} = (\bar{a}, \bar{b}, \bar{c})$ is a fixed point of this map and is said to be E-stable if it is locally asymptotically stable under the differential equation

$$\frac{d\Theta}{d\tau} = T(\Theta) - \Theta. \tag{11}$$

The E-stability principle tells us that E-stable representations are locally learnable for Least Squares and closely related algorithms. That is, if Θ_t is the time t estimate of the coefficient vector Θ , and if Θ_t is updated over time using recursive least squares, then $\bar{\Theta}$ is a possible convergence point, i.e. locally $\Theta_t \rightarrow \bar{\Theta}$, if and only if $\bar{\Theta}$ is E-stable. Computing E-stability conditions is often straightforward, involving computation of eigenvalues of the Jacobian matrices of (11).

Determinacy and stability under learning are both clearly desirable properties for a policy rule. If a policy rule yields indeterminacy then in addition to the intended REE there exist other solutions depending on sunspot variables that may be substantially inferior, in terms of the policy-maker objective function. If the policy rule yields determinacy but is unstable under

learning, then the economy will fail to converge to the intended solution. The earlier literature has shown that these are independent properties and so both must be checked.

A further issue of interest is whether, in the case of indeterminacy, sunspot solutions are stable under learning. Recent research has found that sunspot solutions can in some cases be stable under learning in monetary models of the type considered here, and that stability can depend on the particular representation of the solution that forms the basis of the agents' PLM. We examine different representations when considering the stability properties of sunspot equilibria, and we will say that an equilibrium is learnable (or stable) if it is E-stable for at least one representation. The representations used for our stability analysis are discussed in detail in [17].

Above we stated that the model is closed with an instrument rule for the interest rate. However, we can treat the parameters of the policy rule as assessed, or pinned down, by imposing optimizing behavior on the part of the government. As is common in the literature, we assume, for the government's criterion, a loss function that is quadratic in π and x .¹¹ The government chooses its policy parameters to minimize this criterion subject to the structural model of the economy. For example, suppose the government faces the problem

$$\min_{\alpha_x, \alpha_\pi} \psi \text{Var}(x|\alpha) + \text{Var}(\pi|\alpha) \quad (12)$$

where equations (1) and (2) hold and the interest rate is determined by, say, PR_1 . ψ is the relative weight assigned to the variance of the output gap, and $\text{Var}(\cdot|\alpha)$ is the unconditional variance of “.” given the policy parameters. Here the government has only two choice variables, α_x and α_π , but for the more general rules considered below this number may increase to nine. An interest rate variance term is sometimes included in the loss function (12). This is omitted in our formulation in part for simplicity, but also because neither its inclusion nor its specific form are widely agreed upon.

The value of the government's objective can be computed by determining the rational expectations equilibrium associated with the relevant policy parameters α . In the indeterminate case, this value is not well defined due to the presence of multiple equilibria. For the analysis in this paper we examine the “minimal state variable” (MSV) solution, i.e. a solution of the

¹¹This is consistent with a second order approximation to expected average utility: see [44]. There are, however, deeper issues involved when comparisons are made across model type, as discussed below.

form (7), and if there are multiple such solutions, we take the loss to equal the minimum across solutions of this form.¹² We make this choice because MSV solutions are the ones normally computed in optimal policy exercises and because policy-makers are aiming to minimize expected loss.¹³

We will refer to the government’s problem just described as being “unconstrained.” It is unconstrained in the sense that the government, when choosing its optimal policy, does not restrict attention to the region corresponding to stable determinacy. Our first contention is that when unconstrained the solution may be a policy that yields instability under learning, indeterminacy, or both.

3 Results for Taylor-type Rules

Our central concern in this section is to investigate the possibility that unconstrained optimization may result in undesirable outcomes. More specifically, we examine whether choosing the policy that yields an MSV solution minimizing the government’s loss function can lead to indeterminacy, instability, or even the presence of stable sunspots. Here, and throughout the remainder of the paper, we use the terms “stability” and “instability” to refer to stability under least squares learning, as discussed above.¹⁴ We will find that for forward-looking specifications of the model, all of these outcomes can indeed arise. We conclude that it is essential that the optimal policy problem be explicitly constrained to deliver both stability under learning and determinacy.

Analytic results are not tractable and so we proceed numerically. We begin with Taylor rules of the form (3)-(5) and analyze each policy with respect to three different calibrations of the parameters in the IS-AS curves, due to [43], [10], and [29], as well as a fourth *Variant* calibration given by the largest values of ϕ and λ across the three calibrations: the relevant parameter values are given in Table 1.¹⁵ All calibrations have in common $\rho_g = \rho_u = .9$.

¹²If no lagged variables are present in the structural model or in the policy rule then the (unique) minimal state variable solution takes the form $y_t = a + c\hat{g}_t$.

¹³[37] actually suggested that, in the presence of indeterminacy, a plausible selection criterion might be the solution delivering the minimum variance.

¹⁴All solutions examined in this Section are stochastically stationary, since otherwise the variances being minimized by policy would be infinite. Explosive nonstationary solutions arise in Section 4 when policy-makers face structural parameter uncertainty.

¹⁵The calibrations are for quarterly data with π_t and i_t measured as quarterly changes. The CGG parameters have been adjusted accordingly.

Table 1: Calibrations

Name	ϕ	λ
W	1/.157	.024
CGG	4	.075
MN	.164	.3
V	1/.157	.3

In this Section we mainly focus on purely forward-looking specifications in which $\beta = .99$ and $\delta = \gamma = 0$. Finally, we set the conditional variance of g and u each at 0.1.¹⁶

For each policy rule and calibration, and for policy objective function weights $\psi \in \{.1, 1, 10\}$, a lattice was analyzed in the region of policy space given by $0 \leq \alpha_x, \alpha_\pi \leq \bar{\alpha}$. Unless otherwise stated we set $\bar{\alpha} = 5$. The stability and determinacy properties of the model corresponding to each lattice point were computed, and the value of the government's objective function (expected loss) was determined. These values were then used to numerically compute contours, and hence a graphical representation of the government's indifference curves was obtained. Finally, a numerical optimization algorithm was used to compute the optimal policy parameters within the specified $\bar{\alpha} \times \bar{\alpha}$ benchmark policy space.

3.1 Results for Forward-looking Model

Table 2 presents a complete summary of the results obtained in our numerical analysis for the non-inertial specification. This table gives the value of the government's loss function (12) assuming the numerically computed optimal policy is used, together with the associated stability and determinacy properties of the equilibrium, for all permutations of policy rules, calibrations, and policy weights ψ . To identify the stability and determinacy properties, we use the notation SD (stable determinacy), UD (unstable determinacy), SI (stable indeterminacy), and UI (unstable indeterminacy).

For example, under the W calibration, using PR_1 , and assuming $\psi = 1$,

¹⁶Altering the value of the conditional variance appears only to change the value of the government's objective at the optimum, and not the parameter values corresponding to optimal policy or the stability and determinacy properties of the associated economy.

Table 2: Forward-Looking Model

Calibration	PR	$\psi = .1$		$\psi = 1$		$\psi = 10$	
W	1	29.84*	SD	42.25*	SD	44.09*	SD
	2	31.39	UI	134.22	UI	1094.03	UI
	3	29.84*	SI	42.25*	SI	44.09*	SI
CGG	1	7.73	SD	30.07*	SD	42.31*	SD
	2	6.22*	SD	40.73	SD	295.25	UI
	3	7.73	SI	30.07*	SI	42.31*	SI
MN	1	1.70 (52%)	SD	6.17 (13%)	SD	28.40* (8%)	SD
	2	1.52* (49%)	SD	5.99* (12%)	SD	28.54 (4%)	SD
	3	2.03 (57%)	SD	6.46 (16%)	SD	28.99 (9%)	SD
V	1	.58*	SD	5.17	SD	25.21*	SD
	2	.59	SD [†]	4.59*	UI	38.02	UD
	3	.58*	SI	5.17	SI	25.21*	SI

the optimal policy yields a stable determinate equilibrium and results in an objective value of 42.25. An objective value marked with an asterisk indicates that across rules it is the smallest value associated with that calibration and objective weight: see Section 3.2.3 below. The symbol [†] indicates a solution that is very near the origin. In this case there is a solution not near the origin and within the 5×5 space that yields a value for the objective function close to the optimum. In many cases, the artificial constraint that α_x and α_π be in the 5×5 benchmark space was binding. When this occurred, the boundary was extended to 10×10 and the optimal policy again computed. If the new optimal policy resulted in an objective value that differed by more than 0.1% from the value originally obtained, the percentage difference was recorded in parentheses.¹⁷

This table illustrates the first result of this paper: unconstrained optimal policy can yield SD, UD, SI, or UI. Thus, not only do regions of UI, UD, and SI exist, but optimization algorithms may seek them out; against this possibility policy makers must stand guard.

¹⁷In Table 1 a significant difference arises only in the MN calibration. The contours have the same general shape as in other calibrations (i.e. appear to form a valley: see Figure 1) but for MN the flat part of the objective is not near the origin.

It is useful to discuss briefly some of the specific results. The stability and determinacy properties of the purely forward-looking New Keynesian model (1) and (2) closed with PR_1 have been characterized analytically by [6]. They found that the regions in policy space corresponding to determinacy and stability coincide – in particular, there are no stable sunspots – and this desirable feature inclined the authors to recommend this rule.¹⁸ Their argument is strengthened by the results in the table. Under PR_1 , for all calibrations and governmental objectives concerned, the resulting optimal policy is stable and determinate; and, this is the only rule for which SD always obtains. Figure 1 presents our numerical analysis under the W calibration with $\psi = .1$. For Figures 1 – 3, the contours represent policy indifference curves and regions corresponding to SD, SI, UI and UD are separated by bold curves. The large black dot represents the location of the optimal policy parameters as determined by the search algorithm.

Figure 1 Here

Optimal policy chosen by the search algorithm (and consistent with the contours) lies on the eastern boundary of our artificially constrained 5×5 policy space, and continues to be near the boundary (and again, is SD) for a 100×100 grid with a difference of only 0.001% in the objective function.

In practice many Central Banks are believed to use forward-looking rules of the form PR_3 . Since under learning expectations are not fully rational, we interpret this rule as $i_t = \alpha_x E_t^* x_{t+1} + \alpha_\pi E_t^* \pi_{t+1}$ and note that one can in principle distinguish between the forecasts of the private sector, which enter the IS and AS curves, and the forecasts of the Central Bank, which enter policy rule PR_3 . We adopt the simplest assumption for studying stability under learning, which is that the forecasts for the private sector and the Central Bank are identical, either because private agents and the Central Bank use the same least squares learning scheme or because one group relies on the others' forecasts. In the latter case, for example, the Central Bank might be setting interest rates as a reaction to private sector forecasts, as in [4] or [13]. For further discussion, see [6], [26], [17] and [13].

From Table 2 we see that optimizing policy makers may choose forward-looking rules that result in SI, which implies the existence of sunspot equilibria that are stable under learning. Indeed, only the MN calibration is

¹⁸[17] extended the result of Bullard and Mitra to include inertia in the Phillips Curve, thus further strengthening the argument for rules of PR_1 form.

free of this possibility. Figure 2 presents the results for PR_3 under the same calibration, etc. as was used in Figure 1. All of the features, including the shape and values of the level curves, and the location of and value at the optimum, are essentially identical. However, much of the region which, in Figure 1 corresponded to SD, here corresponds to SI.

Figure 2 Here

Our finding that unconstrained optimizing policy choice can lead to stable indeterminacy, under interest-rate rules of the form PR_3 , is particularly striking and argues forcefully against the view that stable sunspots are not a serious concern.

Finally, consider the lagged Taylor rule, PR_2 , which gives interesting results not seen with PR_1 or PR_3 . As noted by Bullard and Mitra, under PR_2 there exist determinate cases for which the REE is not stable under learning. We find that such cases might even be selected by optimizing policy makers. As an example, Figure 3 illustrates an optimum within the region of unstable determinacy for the V calibration with $\psi = 10$.

Figure 3 Here

While this calibration and choice of ψ are perhaps relatively implausible, the example does forcefully illustrate our point that policy optimization does not automatically deliver a solution that is stable under learning, even if the solution is determinate.

3.2 Further Discussion

3.2.1 Flatness of the objective

As mentioned above in our discussion of Figure 1, there is a tendency for the government's objective function to be very flat near the optimum. This has the potential benefit of rendering precision irrelevant when attempting to determine the optimal policy, but flatness can also make it difficult to pin down an appropriate optimal rule in the presence of multiple nearly optimal rules.

As an example, consider the W calibration with either PR_1 or PR_3 and with $\psi = 0.1$. Figure 2 suggests that the objective is nearly flat for a non-empty sub-region of the benchmark space, and perhaps even constant along a

positively sloped line. To analyze this more closely, we had the optimization algorithm solve the policy problem twice, yielding two different optimal policies (both with essentially the same value for the objective). These two points were used to construct a line with the specification $\alpha_x = 0.433\alpha_\pi - 0.365$. We then allowed α_π to vary from 1 to 5, used our constructed line to choose α_x , and computed the value of the government’s objective.

The result is plotted in Figure 4, assuming PR_3 is used. Here a dashed line indicates the corresponding model is SD and a solid line indicates SI. The government’s objective is almost constant across these parameters.¹⁹

Figure 4 here

One implication of these results is that imposing determinacy and stability requirements in the policy optimization problem can in some cases have negligible cost in terms of the policy objectives.

3.2.2 Inertial specifications

We also considered inertial specifications with $\beta = 1$ and $\delta = \gamma = .5$, which is broadly consistent with [34]. For this specification the results from the forward-looking model are greatly mitigated: optimal policy results in SD for all calibrations except V. For the V calibration, however, PR_2 results in UI for $\psi = 10$ and PR_3 is SI for $\psi = 0.1$ or $\psi = 1$. Furthermore, low levels of inertia yield results similar to those obtained in the purely forward-looking case, e.g. optimal policy may be consistent with stable sunspots under W and CGG calibrations. We consider inertial specifications more thoroughly below, when we also allow for structural parameter uncertainty

3.2.3 Optimum across rules

Thus far we have considered the implications of policy making via unconstrained optimization for a given form of the policy rule. Can indeterminacy or instability arise if policy-makers can choose across the rules? This issue is also addressed in Table 2: for fixed calibration and objective weight, the rule(s) yielding the lowest loss value are marked with an asterisk.

We see that even optimizing across rules does not avoid the difficulties we have identified. For example, for the variant calibration V, the optimum

¹⁹Extending the line to $\alpha_\pi = 100$ does not alter this finding.

across rules yields UI for $\psi = 0.1$. Also, for both the W and CGG specifications, PR₁ and PR₃ can produce the same or nearly the same minimum value of the government’s loss function, but with PR₁ yielding SD while PR₃ results in SI, i.e. stable sunspots.²⁰

3.3 Some Results on Extended Taylor-type Rules

So far, we have restricted attention to policy rules that depend only on the model’s endogenous variables; however, such a restriction is not necessary. If policy makers can observe fundamental shocks, and if these shocks contain information orthogonal to that provided by the endogenous variables, then policy makers would do well to condition their policy accordingly. Furthermore, as shown by [16] (EH), it may be possible for the overall policy optimum – that is, the REE yielding the minimum value of the loss function over all possible policy rules: see equation (15) below – to be implemented by a rule (which we call an “Extended Taylor-type rule”) that depends also on fundamental shocks. In our companion paper [18], we consider in detail the stability and determinacy issues surrounding optimal extended Taylor-type rules; because of the close connection between those issues and the subject of this paper, we briefly summarize some of the results here.

For ease of presentation, consider the non-inertial specification of the New-Keynesian model. Following EH and the related literature, we vary slightly from the government’s criterion presented in Section 2, and assume that monetary policy minimizes²¹:

$$E_t \sum_{s=0}^{\infty} \beta^s (\pi_{t+s}^2 + \psi x_{t+s}^2). \quad (13)$$

As shown by Woodford (and others), the “timeless perspective” optimal policy with commitment must satisfy

$$\lambda \pi_t = -\psi(x_t - x_{t-1}). \quad (14)$$

²⁰Table 2 never indicates an absolute preference for PR₃, but when inertial specifications are considered, any of the rules PR₁ – PR₃ can dominate, depending on the calibration and objective weight (and, this holds regardless of whether the MN calibration is considered).

²¹Our original objective is equivalent to the expected value of (13) across initial conditions, provided the asymptotic distribution of the state variables is used to compute the expected value.

This dynamic equation may be combined with the AS curve to obtain a representation of the unique optimal REE taking the form

$$y_t = Ay_{t-1} + B\hat{g}_t, \quad (15)$$

where both the second column of A and the first column of B are zero (see EH for details).

EH note that the optimal REE is consistent with two different interest rate rules: one, which they call the “fundamentals based rule,” specifies a dependence on lagged output gap and current exogenous shocks; the other, which we will call the EH-rule, includes an additional dependence on expectations of future π and x . EH find that the fundamentals based rule may lead to indeterminacy and is always unstable under learning, while the EH-rule always generates stable determinacy.²²

That there are two possible policy rules consistent with the optimal REE raises the question, “what does the collection of all policy rules consistent with the optimal REE look like, and what are the associated stability and determinacy properties?” In the companion paper, we address this question by characterizing the collection of all policy rules (restricted within a certain class) that are capable of implementing the optimal REE (15). We postulate a general policy rule of the form

$$i_t = \alpha^f E_t y_{t+1} + \alpha^L y_{t-1} + \alpha^{\hat{g}} \hat{g}_t, \quad (16)$$

where $\alpha^f = (\alpha_x^f, \alpha_\pi^f)$, $\alpha^L = (\alpha_x^L, \alpha_\pi^L)$, and $\alpha^{\hat{g}} = (\alpha_g^{\hat{g}}, \alpha_u^{\hat{g}})$. We define the “optimal policy manifold” to be the subset of \mathbb{R}^6 corresponding to the collection of all policy parameters α yielding rules capable of implementing the optimal REE. We show that the optimal policy manifold is two dimensional, that is, it is a 2-manifold in 6-space.

We here consider briefly the implications of choosing different rules from the optimal policy manifold. By finding two policy rules with distinct determinacy and stability properties, EH have already shown that the stability and determinacy properties vary across the manifold. To study the impact on stability and determinacy of using alternate optimal policy rules we specify a calibration and project the optimal policy manifold onto the $(\alpha_\pi^f, \alpha_x^f)$ -plane,

²²The approach followed in EH and here is to look for rules that are optimal under rational expectations and that are stable under learning. [22] consider the implications of a Central Bank solving its dynamic optimization problem while explicitly taking into account the least squares learning rule of private agents.

and restrict attention to the 5×5 benchmark square. We then compute the stability and determinacy properties of the model closed with the corresponding optimal policy: see Figure 5. For this figure, the V parameterization was used, with the modification that $\lambda = 1$; admittedly this is a value of λ that is larger than estimates found in the literature, but it is consistent with certain theoretical models: see for example [8]. We use this calibration for emphasis and note that while less dramatic, the same conclusions apply to other calibrations. The large dot is the location of the EH-rule.²³

Figure 5 Here

Figure 5 demonstrates that while rules capable of implementing the optimal REE abound, a large proportion of these rules have associated with them either stable or unstable indeterminacy. This result is punctuated by the location of the EH-rule: while it does lie in the region of stable determinacy, this region is a small oasis surrounded by a sea of trouble. Because all policies represented in this figure are consistent with the optimal REE, an unconstrained optimizer cannot distinguish between the oasis and the sea.

4 Robust Optimal Constrained Policy

4.1 The Dangers of Model Uncertainty

The results of the previous section warn of the need to impose both stability under learning and determinacy. When we allow for model uncertainty, these potential hazards become even more acute. Our second main result is that a policy that is optimal, stable and determinate for a given calibration, can be UI, SI or explosive (E), under alternate calibrations. These possibilities are shown in Tables 3.1 and 3.2, which give the outcomes under Taylor-rules PR_2 or PR_3 for various alternative calibrations (labeled “Truth”) given policy rule parameters chosen as constrained optimal (i.e., constrained to yield stable determinacy) for the perceived calibration. These tables are computed for $\psi = 1$ and both purely forward-looking ($\delta = \gamma = 0$ and $\beta = 0.99$) and “lag” specifications ($\delta = \gamma = 0.5$ and $\beta = 1$) are considered.

²³Note that in contrast to the exercises producing the figures in the previous section, when producing Figure 5 the policy parameters not referenced in the figure (such as α_x^L , etc) are allowed to vary; in fact, they are required to vary in order to maintain a rule consistent with the optimal REE.

Although policy rule PR_1 does always yield SD, both PR_2 or PR_3 encounter problems under various alternative calibrations. Specifically, PR_3 in some cases leads to stable indeterminacy and PR_2 can generate either unstable indeterminacy or explosive outcomes.

Similar analysis can be conducted using the extended Taylor-type rules discussed in section 3.3, and we find that similar results obtain. For example, the EH-rule, specified under a perceived calibration, which necessarily yields stable determinacy with respect to that perception, can produce UI, SI or E outcomes under alternate calibrations.

Table 3.1: Model Uncertainty: Policy Rule PR_2

<i>Perception</i>	<i>Truth</i>							
	W	CGG	MN	V	W-Lag	CGG-Lag	MN-Lag	V-Lag
W	–	UI	UI	UI	UI	UI	UI	UI
CGG	SD	–	SD	SD	SD	SD	SD	SD
MN	E	E	–	E	E	E	SD	E
V	SD	SD	SD	–	SD	SD	SD	SD
W-Lag	SD	SD	SD	SD	–	SD	SD	SD
CGG-Lag	SD	SD	SD	SD	SD	–	SD	SD
MN-Lag	E	E	SD	E	E	E	–	E
V-Lag	SD	SD	SD	SD	SD	SD	SD	–

Table 3.2: Model Uncertainty: Policy Rule PR_3

<i>Perception</i>	<i>Truth</i>							
	W	CGG	MN	V	W-Lag	CGG-Lag	MN-Lag	V-Lag
W	–	SD	SD	SI	SD	SD	SD	SI
CGG	SI	–	SD	SI	SI	SD	SD	SI
MN	SI	SI	–	SI	SI	SI	SD	SI
V	SD	SD	SD	–	SD	SD	SD	SD
W-Lag	SD	SD	SD	SD	–	SD	SD	SD
CGG-Lag	SI	SD	SD	SI	SI	–	SD	SI
MN-Lag	SD	SD	SD	SI	SD	SD	–	SI
V-Lag	SD	SD	SD	SD	SD	SD	SD	–

4.2 Robust Optimal Constrained Policy

The results in Tables 3.1 and 3.2, together with the range of estimates of structural parameters available in the literature, indicate the need for a policy that is robust to model uncertainty. The method of robust analysis we describe is a Bayesian technique, similar in spirit, but more elementary than the method prescribed by [5]. Put simply, we have policy makers assign a distribution over possible models and the policy maker’s objective may then be computed as the expected value of the objective conditional on this *prior* distribution.²⁴ However, we now add the constraint that policy be determinate and stable under learning for all parameter values that have positive probability.²⁵

Put more formally, we consider a set of structural model parameters $\omega \in S$ and assume that we have a *prior* probability distribution over their possible values with support $\bar{S} \subset S$. We examine a specified class of interest-rate rules parameterized by $\xi \in X$ and determine the set $P \subset X$ defined by

$$P = \{\xi \in X : \text{determinacy and learning stability hold for all } \omega \in \bar{S}\}.$$

Given a policy loss function $\mathcal{L}(\xi, \omega)$ that is well-defined for all $\xi \in P$ and $\omega \in \bar{S}$, and assuming that P is non-empty, the robust optimal constrained policy is defined by

$$\bar{\xi} = \arg \min_{\xi \in P} E\mathcal{L}(\xi, \omega),$$

where the expectation is taken over the *prior* distribution on \bar{S} .

We illustrate how this method can be implemented by assigning probabilities to the various calibrations, parameterized by the key parameters $\omega = (\phi, \lambda, \gamma, \delta)$. Thus we allow for uncertainty with respect to both the structural parameters and the degree of inertia. For the structural parameters ϕ, λ we set the conditional probability weights at:

Table 4.1: Structural Parameter Specifications

Calibration	W	CGG	MN	V
Probability	0.3	0.3	0.3	0.1

²⁴[5] further advise incorporating a minimax technique to guard against very bad outcomes.

²⁵[1] also impose determinacy requirements in their study of inflation-forecast-based rules. However, they do not address the issue of stability under learning.

For the inertial parameters γ, δ we set the conditional probability weights at:

Table 4.2: Inertial Specifications

$\gamma = \delta$	1	0.75	0.5	0.25	0.01
Probability	0.1	0.2	0.4	0.2	0.1

The two sets of probability weights are treated as independently distributed. We refer to this *prior* as the “full *prior*”. For comparison we also look at the case of “non-inertial *prior*” in which the structural parameter calibrations receive the above probability weights but $\gamma = \delta = 0$ is assumed.²⁶ We consider both simple and extended Taylor-type interest rate rules.

4.3 Robust Taylor Rules

We begin by considering robust rules that depend only on endogenous variables. Specifically, we consider PR₁ – PR₃, with the additional possibility of including an interest rate smoothing term parameterized by θ . Thus PR₁ with smoothing takes the form

$$i_t = \alpha_\pi E_t \pi_t + \alpha_x E_t x_t + \theta i_{t-1}, \quad (17)$$

and PR₂ – PR₃ are modified accordingly. In terms of the definition of robust optimal constrained policy just given, for these policy rule classes we are setting $\xi = (\alpha_\pi, \alpha_x, \theta)$. It is straightforward to verify numerically that for both policy rules PR₁ and PR₃ there is a policy ξ such that the associated model is stable determinate for all calibrations and all inertial specifications on which the *prior* places positive weight. Thus the constraint set P of our optimization problem is non-empty for the policy rules classes PR₁ and PR₃. On the other hand, PR₂ does not deliver feasible robust policies, i.e. for this class of rules P is empty.

Using Matlab to perform the optimization subject to stability, determinacy, and $\alpha_x, \alpha_\pi \in [0, 5]$, $\theta \in [0, 2]$, we obtain the results presented in Table 5, corresponding to $\psi = 1$. We report the results for PR₁ and PR₃ both for the case in which the interest-rate smoothing term is restricted to $\theta = 0$ and

²⁶There is sharp disagreement on the extent of inertia in the economy. For example, [33] employ a set-up that is fully inertial, while [34] estimates the degree of inertia at one-half and [30] provides evidence that apparent persistence may be due to adaptive learning in a fully forward-looking model.

the case in which it is allowed to vary freely. Results are given both for the full *prior* described above and for the non-inertial *prior*.²⁷

Table 5.1: Taylor-type Rules, Non-inertial Prior

Policy Rule	Restriction	α_x	α_π	θ	Value
PR ₁	$\theta = 0$	5	3.78	–	27.05 [‡]
	$\theta \in [0, 2]$	1.54	.82	1.09	25.42
PR ₃	$\theta = 0$.27	1.31	–	52.34
	$\theta \in [0, 2]$.58	.92	.80	32.74

Table 5.2: Taylor-type Rules, Full Prior

Policy Rule	Restriction	α_x	α_π	θ	Value
PR ₁	$\theta = 0$	2.43	5	–	141.82 [‡]
	$\theta \in [0, 2]$	1.76	3.02	.88	140.83
PR ₃	$\theta = 0$.24	1.47	–	166.52
	$\theta \in [0, 2]$.55	1.34	.83	143.36

Here ‡ indicates that the optimum lies on the boundary of the policy region considered.

Notice that, when the smoothing term θ is set equal to zero then PR₁ performs much better than PR₃. This reflects the prevalence of indeterminacy when policy rules depend on forward expectations (e.g. compare Figures 1 and 2). However, the inclusion of the interest-rate smoothing term $\theta > 0$ significantly mitigates this effect, and in the case of PR₃ the benefits of a large θ are considerable.

Note that the expected loss is much higher when full *priors* are considered. This reflects the much smaller extent of control available to policy-makers when economic outcomes are more heavily determined by lagged inflation and output than by expected future macroeconomic conditions. We remark that in evaluating policy we continue to use the policy loss function (12). We think this is realistic, but we should point out that in New Keynesian models

²⁷With the non-inertial prior PR₂ was also feasible but its performance was substantially worse than the other two rules.

with inertia, the expected utility of the representative agent would dictate different forms of objective function, and that the precise form would in fact depend on the specific source of inertia. This raises deep issues, since in principle the appropriate policy objective would depend on whether inertia is present and on the reason for inertia. By using (12) to evaluate results across models we side-step this issue, but in a way that seems realistic from the viewpoint of practical policy-making.

4.4 Robust Extended Taylor-type Rules

Precisely the same method can be employed to obtain robust optimal constrained extended Taylor-type rules. We now consider the class of interest-rate rules

$$i_t = \theta i_{t-1} + \alpha^f E_t y_{t+1} + \alpha^c E_t y_t + \alpha^L y_{t-1} + \alpha^{\hat{g}} \hat{g}_t. \quad (18)$$

For this exercise, we again set $\psi = 1$.

Table 6: Extended Taylor-type Rules, Full Prior

α_x^f	α_π^f	θ	α_x^c	α_π^c	α_x^L	α_π^L	$\alpha_g^{\hat{g}}$	$\alpha_u^{\hat{g}}$	Value
0	2.98	1.07	5	0	-4.31	-1.89	-.04	3.91	132.03
2.31	12.62	1.08	5.15	-12.74	-5.95	2.61	.35	9.8	130.37

Table 6 reports two policy rules for the full *prior*. The first row imposes that policy parameters lie in the range $\alpha_x, \alpha_\pi \in [-5, 5]$, $\theta \in [0, 2]$, which turns out to be a constraint for α_x^c . In the second row we give the results for expanding the box for α_x, α_π to $[-100, 100]$. This provides some additional though limited improvement in the expected loss function, consistent with our previous findings that the objective function is quite flat near the optimal policy.

As with the robust Taylor rules, the smoothing parameter θ is chosen to be large. Here θ is actually larger than one (giving “superinertial” rules), though the value is close to one (this is also the case when the non-inertial prior is used). This suggests that while fully optimal policy may be implemented without an interest rate smoothing term when parameter values are known with certainty (as we saw in Section 3.3), a strong smoothing term appears indicated when parameter uncertainty is present.

4.5 A Simple Nearly Optimal Rule

Given the flatness of the objective function, it is of interest to ask whether simpler rules can be devised that meet our requirements that the rule be determinate and stable across all specifications that receive positive weight in the *prior*, and which come close, in terms of minimizing the expected loss function, to the performance of the optimal rule just computed. A search across simpler rules yields the following specification:

$$\Delta i_t = E_t \pi_{t+1} + 2(E_t \pi_{t+1} - \pi_{t-1}) + 0.6E_t x_t + 4.5(E_t x_t - x_{t-1}) + 4u_t. \quad (19)$$

This interest-rate rule is stable and determinate for all calibrations and inertial specifications with positive prior probability, and yields an objective value of 132.22, which represents slightly less than a 1.3% increase over the expected loss achieved by the optimal rule.

The rule (19) is easily interpretable, taking an “integral” form in which interest rates are increased in response to higher levels of the expected future inflation rate and its increase relative to last period’s inflation rate, to the expected current output gap and its increase relative to last period’s output gap, and to observed exogenous inflation shocks.

Remarkably, even simpler rules lead to only modest reductions in performance. Returning to simple Taylor rules with inertia of the form (17), given earlier, we see that specification $\alpha_x = 1.76$, $\alpha_\pi = 3.02$, and $\theta = 0.88$ leads to an objective value of 140.83, which is an additional 6.5% loss over that achieved by (19).

4.6 Discussion

Given the wide range in the literature for structural parameter and inertial calibrations of New Keynesian models, a variety of calibrations should be considered when searching for a good policy rule. To simultaneously consider multiple calibrations, we employed a simple “Bayesian” procedure, weighting each calibration according to a specified *prior* probability. We then computed, within a given class of rules, the optimal policy parameters subject to the constraint that the economy be stable under learning and determinate for all calibrations with positive probability.

An important finding of this exercise is that the constraint set is not empty, despite allowing for a very wide range of calibrations that received positive weight. This suggests that our procedure may provide a feasible and

fruitful way to think about the optimal choice of interest-rate rules in serious applied models, while imposing determinacy and stability constraints and allowing for structural parameter uncertainty.

We also found that a relatively simple rule, specifically (19), is capable of simultaneously yielding stability and determinacy, as well as producing an objective value within 1.3% of the optimum. Admittedly, we have only considered a range of simple, stylized models. These results, however, suggest that the technique advanced above may provide similarly simple nearly optimal robust rules in more general and more realistic modeling environments.²⁸

We conclude this Section by noting that the existence of a non-empty constraint set is a requirement that may not be met in future applications of the technique and thus warrants further discussion. Although various modifications would be possible in this case, our preference would be to extend the procedure in the following way. Given $0 \leq \eta \leq 1$, let $P(\eta)$ denote the set

$$P(\eta) = \left\{ \xi \in X : \Pr_{\omega \in \bar{S}}(\text{determinacy and stability under learning}) \geq 1 - \eta \right\},$$

and let $\bar{\eta}$ be the smallest value for which $P(\eta)$ is nonempty. For any given $\eta \geq \bar{\eta}$ the robust optimal constrained policy $\bar{\xi}(\eta)$ would then be defined by

$$\bar{\xi}(\eta) = \arg \min_{\xi \in P(\eta)} \tilde{E}\mathcal{L}(\xi, \omega),$$

where \tilde{E} denotes expectation now computed conditional on parameters for which the equilibrium is determinate and stable. We think of η as a parameter chosen by policy makers and representing the largest acceptable probability that the assessed constrained optimal policy will yield an outcome that is not SD.

Policy-makers could then either specify a sufficiently low value for $\eta \geq \bar{\eta}$ that they consider acceptable, or be presented with the trade-off between η

²⁸We emphasize that our policies are robust in a very specific sense. We do not consider unstructured model uncertainty in the sense of [24], robust optimality to exogenous disturbances in the sense of [23], or robustness to qualitatively different structural models, as in [27]. Furthermore, our learnability constraint concerns local stability under learning under a fixed structure. Other relevant concerns are maximizing the basin of attraction for least-squares learning, one of the issues addressed in [41], and policy choice with structural shifts and perpetual learning, which is the focus of [31]. Clearly a number of these approaches are complementary in the sense that they could be combined with our analysis.

and $\min_{\xi \in P(\eta)} \tilde{E}\mathcal{L}(\xi, \omega)$. Indeed, even in the case found here in which $\bar{\eta} = 0$, one could examine the potential gains, in terms reduced expected loss, from permitting $\eta > 0$. We refrain from presenting these results on the assumption that policy-makers would prefer to be maximally robust.

5 Conclusion

We have demonstrated the potential for policy makers, attempting to choose the optimal rule within a class, to be directed towards rules that lie in the indeterminacy and/or learning instability regions. Taylor-type rules have an appealing simplicity, with the key inflation and output coefficients traditionally chosen based on plausible rules of thumb. One might expect that improved performance would be obtained by choosing these policy parameters optimally for a given calibrated model. Paradoxically this may not be the case, because searching for the optimal policy rule within a given class may fail to deliver a rule that produces a determinate equilibrium that is stable under learning. It is therefore imperative that the search for optimal policy rules be constrained to the determinate stable region. Considering a class of extended Taylor rules sufficiently general that it includes fully optimal solutions does not avoid this problem since it will contain some “fully optimal” rules that are subject to indeterminacy and/or instability problems.

These problems are compounded by structural parameter uncertainty. Policy rules that lead to determinacy and stability under learning and that are fully optimal for one set of structural parameters can lead to indeterminacy, instability or even explosiveness, for another set of parameters. We therefore advocate a “robust” optimization procedure, in which policy makers select the optimal constrained rule. Such a rule is computed as the one which minimizes the policy-makers expected loss, based on *prior* probabilities for the structural parameters, but which is constrained to satisfy the condition that it lies within the region of stable determinacy for every calibration that has positive probability.

Our key points about the hazards of an unconstrained search for optimal policy have been developed in the context of the basic New Keynesian framework and hybrid versions incorporating inertia, and we have used this framework to show how to compute the robust optimal constrained interest-rate rule. However, the approach is quite general and could be applied to more elaborate structural macroeconomic models.

References

- [1] Nicoletta Batini, Alejandro Justiniano, Paul Levine, and Joseph Pearlman. Robust inflation-forecast-based rules to shield against indeterminacy. Working paper, IMF, 2004.
- [2] Jess Benhabib and Stefano Eusepi. The design of monetary and fiscal policies: A global perspective. mimeo., New York University, 2005.
- [3] Ben S. Bernanke and Michael Woodford, editors. *Inflation Targeting*. University of Chicago Press, Chicago, 2005.
- [4] B.S. Bernanke and M. Woodford. Inflation forecasts and monetary policy. *Journal of Money, Credit, and Banking*, 24:653–684, 1997.
- [5] William A. Brock, Steven N. Durlauf, and Kenneth D. West. Policy evaluation in uncertain environments. Working Paper SSRI 2003-15, University of Wisconsin - Madison, 2003.
- [6] James Bullard and Kaushik Mitra. Learning about monetary policy rules. *Journal of Monetary Economics*, 49:1105–1129, 2002.
- [7] James Bullard and Kaushik Mitra. Determinacy, learnability and monetary policy inertia. Working paper, Federal Reserve Bank of St. Louis, 2003.
- [8] V. V. Chari, Patrick J. Kehoe, and Ellen R. McGrattan. Sticky price models of the business cycle: Can the contract multiplier solve the persistence problem. *Econometrica*, 68:1151–1179, 2000.
- [9] Richard Clarida, Jordi Gali, and Mark Gertler. The science of monetary policy: A new Keynesian perspective. *Journal of Economic Literature*, 37:1661–1707, 1999.
- [10] Richard Clarida, Jordi Gali, and Mark Gertler. Monetary policy rules and macroeconomic stability: Evidence and some theory. *Quarterly Journal of Economics*, 115:147–180, 2000.
- [11] George W. Evans and Seppo Honkapohja. On the local stability of sunspot equilibria under adaptive learning rules. *Journal of Economic Theory*, 64:142–161, 1994.

- [12] George W. Evans and Seppo Honkapohja. *Learning and Expectations in Macroeconomics*. Princeton University Press, Princeton, New Jersey, 2001.
- [13] George W. Evans and Seppo Honkapohja. Adaptive learning and monetary policy design. *Journal of Money, Credit and Banking*, 35:1045–1072, 2003.
- [14] George W. Evans and Seppo Honkapohja. Existence of adaptively stable sunspot equilibria near an indeterminate steady state. *Journal of Economic Theory*, 111:125–134, 2003.
- [15] George W. Evans and Seppo Honkapohja. Expectations and the stability problem for optimal monetary policies. *Review of Economic Studies*, 70:807–824, 2003.
- [16] George W. Evans and Seppo Honkapohja. Monetary policy, expectations and commitment. *Scandinavian Journal of Economics*, forthcoming, 2004.
- [17] George W. Evans and Bruce McGough. Monetary policy, indeterminacy and learning. *Journal of Economic Dynamics and Control*, forthcoming, 2004.
- [18] George W. Evans and Bruce McGough. On the potential instability of optimal monetary policy. mimeo, 2005.
- [19] George W. Evans and Bruce McGough. Stable sunspot solutions in models with predetermined variables. *Journal of Economic Dynamics and Control*, 29:601–625, 2005.
- [20] Jeff Fuhrer and George Moore. Inflation persistence. *Quarterly Journal of Economics*, 110:127–159, 1995.
- [21] Jordi Gali and Mark Gertler. Inflation dynamics: A structural econometric approach. *Journal of Monetary Economics*, 44:195–222, 1999.
- [22] Vitor Gaspar, Frank Smets, and David Vestin. Optimal monetary policy under adaptive learning. mimeo., European Central Bank, 2005.
- [23] M. Giannoni and M. Woodford. Optimal interest-rate rules: II. Applications. NBER working paper series, no. 9420, 2002.

- [24] Lars Hansen and Thomas J. Sargent. Robust control of forward-looking models. *Journal of Monetary Economics*, 50:581–604, 2003.
- [25] Seppo Honkapohja and Kaushik Mitra. Are non-fundamental equilibria learnable in models of monetary policy? *Journal of Monetary Economics*, 51:1743–1770, 2004.
- [26] Seppo Honkapohja and Kaushik Mitra. Performance of monetary policy with internal central bank forecasting. *Journal of Economic Dynamics and Control*, forthcoming, 2005.
- [27] Andrew Levin, Volker Wieland, and John Williams. The performance of forecast-based monetary policy rules under model uncertainty. *American Economic Review*, 93:622–645, 2003.
- [28] Bennett T. McCallum. Issues in the design of monetary policy rules. In Taylor and Woodford [40], chapter 23, pages 1483–1530.
- [29] Bennett T. McCallum and Edward Nelson. Performance of operational policy rules in an estimated semi-classical model. In Taylor [39], pages 15–45.
- [30] Fabio Milani. Expectations, learning and macroeconomic persistence. mimeo., 2004.
- [31] Athanasios Orphanides and John C. Williams. Robust monetary policy with imperfect knowledge. *Journal of Economic Dynamics and Control*, forthcoming, 2004.
- [32] Julio J. Rotemberg and Michael Woodford. Interest-rate rules in an estimated sticky price model. In Taylor [39], pages 55–119.
- [33] Glenn D. Rudebusch and Lars E. O. Svensson. Policy rules for inflation targeting. In Taylor [39].
- [34] Frank Smets. Maintaining price stability: How long is the medium term? *Journal of Monetary Economics*, 50:1293–1309, 2003.
- [35] Lars E. O. Svensson. What is wrong with Taylor rules? Using judgment in monetary policy through targeting rules. *Journal of Economic Literature*, 41:426–477, 2003.

- [36] Lars E. O. Svensson and Michael Woodford. Implementing optimal policy through inflation-forecast targeting. In Bernanke and Woodford [3].
- [37] John B. Taylor. Conditions for unique solutions in stochastic macroeconomic models with rational expectations. *Econometrica*, 45:1377–1386, 1977.
- [38] John B. Taylor. Discretion versus policy rules in practice. *Carnegie-Rochester Conference Series in Public Policy*, 39:195–214, 1993.
- [39] John B. Taylor, editor. *Monetary Policy Rules*. University of Chicago Press, Chicago, 1999.
- [40] John B. Taylor and Michael Woodford, editors. *Handbook of Macroeconomics, Volume 1*. Elsevier, Amsterdam, 1999.
- [41] Robert J. Tetlow and Peter von zur Muehlen. Robustifying learnability. mimeo., Federal Reserve Board, 2005.
- [42] Michael Woodford. Learning to believe in sunspots. *Econometrica*, 58:277–307, 1990.
- [43] Michael Woodford. Optimal monetary policy inertia. *The Manchester School, Supplement*, 67:1–35, 1999.
- [44] Michael Woodford. *Interest and Prices: Foundations of a Theory of Monetary Policy*. Princeton University Press, Princeton, NJ, 2003.
- [45] Luis-Felipe Zanna. PPP rules, macroeconomic (in)stability and learning. mimeo., Board of Governors of the Federal Reserve System, 2004.

Figure 1: Woodford Calibration, PR_1 , No lag, $\psi = .1$




Figure 2: Woodford Calibration, PR_3 , No Lag, $\psi=.1$




Figure 3: V Calibration, PR_2 , $\psi=10$




Figure 4: W Cal., No Inertia, $\alpha_x = .433 \times \alpha_\pi - .365$




Figure 5: Optimal Manifold, $\lambda = 1$, $\phi = 6.37$, $\beta = .99$, $\psi = 1$

