

Robustifying Learnability

Robert J. Tetlow* Peter von zur Muehlen†

November 2005.

Abstract

In recent years, the learnability of rational expectations equilibria (REE) and determinacy of economic structures have rightfully joined the usual performance criteria among the sought-after goals of policy design. Some contributions to the literature, including Bullard and Mitra (2001) and Evans and Honkapohja (2002), have made significant headway in establishing certain features of monetary policy rules that facilitate learning. However a treatment of policy design for learnability in worlds where agents have potentially misspecified their learning models has yet to surface. This paper provides such a treatment. We begin with the notion that because the profession has yet to settle on a consensus model of the economy, it is unreasonable to expect private agents to have collective rational expectations. We assume that agents have only an approximate understanding of the workings of the economy and that their learning the reduced forms of the economy is subject to potentially destabilizing perturbations. The issue is then whether a central bank can design policy to account for perturbations and still assure the learnability of the model. Our test case is the standard New Keynesian business cycle model. For different parameterizations of a given policy rule, we use structured singular value analysis (from robust control theory) to find the largest ranges of misspecifications that can be tolerated in a learning model without compromising convergence to an REE.

In addition, we study the cost, in terms of performance in the steady state of a central bank that acts to robustify learnability on the transition path to REE. (Note: This paper contains full-color graphics)

- **JEL Classifications:** C6, E5.
- **Keywords:** monetary policy, learning, E-stability, learnability, robust control.

* *Contact author:* Robert Tetlow, Federal Reserve Board, 20th and C Streets, NW, Washington, D.C. 20551. *Email:* rtetlow@frb.gov. We thank George Evans, Steve Durlauf, John C. Williams and participants at a conference on learning and monetary economics at the UC-Santa Cruz for helpful comments, and Brian Ironside and Sean Taylor for help with the charts. The views expressed in this paper are those of the authors alone and do not represent those of the Board of Governors of the Federal Reserve System or other members of its staff. This paper and others may be found at <http://www.roberttetlow.com>

† von zur Muehlen & Associates, Vienna, VA 22181. *E-mail:* pvzmuehlen@cox.net

1 Introduction

It is now widely accepted that policy rules—and in particular, monetary policy rules—should not be chosen solely on the basis of their performance in a given model of the economy. There is simply too much uncertainty about the true structure of the economy to warrant taking the risk of so narrow a criterion for selection. Rather, policy should be designed to operate "well" in a wide range of models. There has been substantial progress in a relatively short period of time in the literature on robustifying policy. The first strand of the literature examines the performance of rules given the presence of measurement errors in either model parameters or unobserved state variables.¹ The second strand focuses on comparing rules in rival models to see if their performance spanned reasonable sets of alternative worlds.² The third considers robustifying policy against unknown alternative worlds, usually by invoking robust control methods.³

At roughly the same time, another literature was developing on the learnability (or E-stability) of models.⁴ The learnability literature takes a step back from rational expectations and asks whether the choices of uninformed private agents could be expected to converge on a rational expectations equilibrium (REE) as the outcome of a process of learning. Important early papers in this literature include Bray [5], Bray and Savin [6] and Marcet and Sargent [33]. Evans and Honkapohja summarize some of their many contributions to this literature in their book [17].

The question arises: could monetary policy help or hurt private agents learn the REE? The common features of the robust policy literature include, first, that it is the government that does not understand the true structure of the economy, and second, that the government's ignorance will not vanish simply with the collection of more data.⁵ By contrast, in the

¹ Brainard [4] is the seminal reference. Among the many, more recent references in this large literature are Sack [40], Orphanides et al. [38], Soderstrom [42] and Ehrmann and Smets [16].

² See, e.g., Levin *et al.* [28] and [29].

³ Hansen and Sargent [26] and [27], Tetlow and von zur Muehlen [46] and Coenen [13]. These strands of the robustness literature are named in the text in chronological order but the three methods should be seen as complementary rather than substitutes.

⁴ In this paper, as in most of the rest of the literature, the terms learnability, E-stability and stability under learning will all be used interchangeably. These terms are distinct from stable—without the "E-" or "under learning" added—which should be taken to mean saddle-point stable. The term saddle-point stable, determinate and regular are taken as equivalent adjectives describing equilibria.

⁵ The concept of "truth" is a slippery one in learning models. In some sense, the truth is jointly determined by the deep structural parameters of the economy and what people believe them to be. Only in steady state, and then only under some conditions, will this be solely a function of deep parameters and not of beliefs.

learning literature it is usually the private sector that is assumed not to have the information necessary to form rational expectations, but this situation has at least the prospect of being alleviated with the passage of time and the collection of more data. In this paper, we take the robust policy rules literature and marry it with the learnability literature.

Since the profession has been unable to agree on a generally acceptable workhorse model of the economy, it is unreasonable to expect private agents to have rational expectations at all points in time. The most that one can expect is that agents have an approximate understanding of the workings of the economy and that they are on a transition path toward learning the true structure. And as Evans and McGough [21] show, designing policy as if the process of learning has already been completed can result in indeterminacy or unstable equilibria. So we retain the assumption of adaptive learning that is common with most contributions to this literature.

If the presumption of rational expectations is questionable, and the hazards of learning cannot be ignored, then from a policy perspective, it follows that the job of facilitating the transition to REE is logically prior to the job of maximizing the performance of the economy once the transition is complete. In this paper, we consider two issues. The first is how a policy maker might choose policy to maximize the set of worlds inhabited by private agents that are able to learn the REE. The second is an assessment of the welfare cost of assuring learnability in terms of forgone stability in equilibrium. Or, put differently, we measure the welfare cost of learnability insurance. Each of these questions is important. In worlds of model uncertainty, an ill-chosen policy rule—or policy maker—could lead to explosiveness or indeterminacy. At the same time, excessive concern for learnability will imply costs in terms of forgone welfare.

Ours is not the first paper to consider the issue of choosing monetary policies for their ability to deliver determinacy and learnability. Bernanke and Woodford [2] argue that inflation-forecast-based (IFB) policy rules—that is, rules that feed back on forecasts of future output or inflation—can lead to indeterminacy in linear rational expectations (LRE) models. Clarida, Gali and Gertler [11] show that violation of the so-called Taylor principle in the

Nevertheless, in this paper, when we refer to a "true model" or "truth" we mean the REE upon which successful learning eventually converges.

context of an IFB rule may have been the source of the inflation of the 1970s.⁶ Bullard and Mitra [8] in an important paper show that higher persistence in instrument setting—meaning a large coefficient on the lagged instrument in a Taylor-type rule—can facilitate determinacy in the same class of models. Evans and Honkapohja [19] note similar problems in a wider class of rules and argue for feedback on structural shocks, although questions regarding the observability of such shocks leave open the issue of whether such a policy is implementable. Evans and McGough [21] compute optimal simple rules conditional on their being determinate in rival models. Each of these papers makes an important contribution to the literature, but all are special cases within broader sets of policy choices. In this paper, we follow a somewhat different approach and consider the design of policies to maximize learnability of the economy.

The remainder of the paper is organized as follows. The second section lays out the theory, beginning with a review of the literature on least-squares learnability and determinacy, and following with methods from the robust control literature. Section 3 introduces the models with which we work, beginning with the case of the very simple Cagan model of money demand in hyperinflations and then moving on to the New Keynesian business cycle (NKB) model. For the NKB model, we study the design of time-invariant simple monetary policy rules to robustify learnability of three types: a lagged-information rule, a contemporaneous information rule and a forecast-based policy rule. We close the section by covering the insurance cost of robustifying learnability. A fourth and final section sums up and concludes.

2 Theoretical overview

2.1 Expectational equilibrium under adaptive learning

The theory of *E-stability* or *learnability* in linear rational expectations models dates back more than 20 years to Bray [5] who showed that agents using recursive least squares would, if the arguments to their regressions were properly specified, eventually converge on the correct REE. This convergence property gave a considerable shot in the arm to rational

⁶ Levin *et al.* [29] and Batini and Pearlman [1] study the robustness properties of different types of inflation-forecast based rules for their stability and determinacy properties.

expectations applications since proponents had an answer to the question "how could people come to have rational expectations?" The theory has been advanced by the work of Marcat and Sargent [33] and Evans and Honkapohja [various]. Our rendition follows Evans and Honkapohja [17], chapters 8-10.

Begin with the following linear rational expectations model:

$$y_t = A + ME_t y_{t+1} + N y_{t-1} + P v_t, \quad (1)$$

where y_t is a vector of n endogenous variables, including, possibly, policy instruments, and v_t comprises all m exogenous variables. Equation (1) is general in that both non-predetermined variables, $E_t y_{t+1}$, and predetermined variables, y_{t-1} , are represented. By defining auxiliary variables, e.g., $y_t^j = y_{t+j}$, $j \neq 0$, arbitrarily long (finite) lead or lag lengths can also be accommodated. Finally, extensions to allow lagged expectations formation; e.g., $E_{t-1} y_t$, and exogenous variables are straightforward to incorporate with no significant changes in results. Next, define the prediction error for y_{t+1} , to be $\xi_{t+1} = y_{t+1} - E_t y_{t+1}$. Under rational expectations, $E_t \xi_{t+1} = 0$, a martingale difference sequence. Evans and Honkapohja [17] show that for at least one rational expectations equilibrium to exist, the stochastic process, y_t , that solves (1), must also satisfy:

$$y_{t+1} = -M^{-1}A + M^{-1}y_t - M^{-1}N y_{t-1} - M^{-1}P v_t + \xi_{t+1} \quad (2)$$

We can express (2) as a first-order system:

$$\begin{bmatrix} y_{t+1} \\ y_t \end{bmatrix} = \begin{bmatrix} -M^{-1}A \\ 0 \end{bmatrix} + \begin{bmatrix} M^{-1} & -M^{-1}N \\ I_n & 0 \end{bmatrix} \begin{bmatrix} y_t \\ y_{t-1} \end{bmatrix} + \begin{bmatrix} -M^{-1} \\ 0 \end{bmatrix} P v_t + \begin{bmatrix} 1 \\ 0 \end{bmatrix} \xi_{t+1}$$

or, rewriting:

$$Y_{t+1} = F + B Y_t + C v_t + D \xi_{t+1}, \quad (3)$$

where $Y = [y_t, y_{t-1}]'$. Then we can easily show when (3) satisfies the Blanchard-Kahn [3] conditions for stability, namely, that the number of characteristic roots of the matrix B of norm less than unity equal the number of predetermined variables (taking y_t to be scalar, this is one), then the model is *determinate*, and there is just one martingale difference sequence, ϵ_{t+1} , that will render (2) stationary; if there are fewer roots inside the unit circle than there are predetermined variables, the model is *explosive* meaning that there is no

martingale difference sequence that will satisfy the system; and if there are more roots inside the unit circle than there are predetermined variables, the model is said to be *indeterminate*, and there are infinite numbers of martingale difference sequences that make (2) *saddle-point stable*. The roots of B are determined by the solution to the characteristic equation: $\lambda^2 - M^{-1}\lambda + M^{-1}N = 0$.

Determinacy is one thing, learnability is quite another. As Bullard and Mitra [8] have emphasized, determinacy does not imply learnability, and indeterminacy does not imply a lack of learnability. We can address this question by postulating a representation for the REE that a learning agent might use. For the moment, we consider the *minimum state variable* (MSV) representation, advanced by McCallum [34]. Let us assume that v_t is observable and follows a first-order stochastic process,

$$v_t = \rho v_{t-1} + \epsilon_t,$$

where ϵ_t is an iid white noise process. The ρ matrix is assumed to be diagonal.

Under these assumptions, we can write the following *perceived law of motion* (PLM):

$$y_t = a + by_{t-1} + cv_t. \quad (4)$$

Rewrite equation (1) slightly, and designate expectations formed using adaptive learning with a superscripted asterisk on the expectations operator, E_t^* :

$$y_t = A + ME_t^* y_{t+1} + Ny_{t-1} + Pv_t. \quad (5)$$

Then, leading (4) one period, taking expectations, substituting (4) into the result, and finally into (5), we obtain the *actual law of motion*, (ALM), the model under the influence of the PLM :

$$y_t = A + M(I + b)a + (N + Mb^2)y_{t-1} + (M(bc + c\rho) + P)\epsilon_t. \quad (6)$$

So the MSV solution will satisfy the mapping from PLM to ALM:

$$\begin{aligned} A + M(I + b)a &= a, \\ N + Mb^2 &= b, \\ M(bc + c\rho) + P &= c. \end{aligned}$$

Learnability depends then on the mapping of the PLM on to the ALM, defined from (6):

$$T(a, b, c) = [A + M(I + b)a, N + Mb^2, M(bc + c\rho) + P] \quad (7)$$

The fixed point of this mapping is a MSV representation of a REE, and its convergence is given by the matrix differential equation:

$$\frac{d}{d\tau}(a, b, c) = T(a, b, c) - (a, b, c). \quad (8)$$

Convergence is assured if certain eigenvalue conditions for the following matrix differential equations are satisfied.

$$\begin{aligned} \frac{da}{d\tau} &= [A + M(I + b)]a - a, \\ \frac{db}{d\tau} &= Mb^2 + N - b, \\ \frac{dc}{d\tau} &= M(bc + c\rho) + P - c. \end{aligned} \quad (9)$$

As shown by Evans and Honkapohja (2001), the necessary and sufficient conditions for E-stability are that the eigenvalues of the following matrices have negative real parts:

$$\begin{aligned} DT_a - I &= M(I + b) - I, \\ DT_b - I &= b' \otimes M + I \otimes Mb - I, \\ DT_c - I &= \rho' \otimes M + I \otimes Mb - I. \end{aligned}$$

The important points to take from equations (9) are that the conditions are generally multivariate in nature—meaning that the coefficients constraining the intercept term, a , can be conflated with those of the slope term, b , and that the coefficients of both the PLM and the ALM come into play. Learnability applications in the literature to date have been confined to very simple, small-scale models where these problems rarely come into play.⁷ In the kind of medium- to large-scale models that policy institutions use, these issues cannot be safely ignored.⁸ Without taking away anything from the important contributions of Bullard and

⁷ A notable exception is Garratt and Hall [22], but even then the learning problem was constrained to exchange rate determination. The rest of the London Business School model that they used was taken as known.

⁸ At the Federal Reserve Board, for example, the staff use a wide range of models to analyze monetary policy issues, including a variety of reduced-form forecasting models, a calibrated multi-country DSGE model called SIGMA, a medium-scale DSGE U.S. model, and the FRB/US model, a larger-scale, partly micro-founded estimated model.

Mitra [8] and Evans and Honkapohja [19], the choice of monetary policy rules must not only consider how they foster learnability in a given model but whether they do so for the broader class of models within which the true learning model might be found. Similarly, taking as given the true model, the initial beliefs of private agents can affect learnability both through the inclusion and exclusion of states to the PLM and through the initial values attached to parameters. In the context of the above example, values of a , b , and c that are initially "too far" from equilibrium can block convergence. The choice of a particular policy can shrink or expand the range of values for a , b , and c that is consistent with E-stability.⁹ This is our concern in this paper: how can a policy maker deal with uncertainty in agents' learning mechanisms in the choice of his or her policy rule and thereby maximize the prospect that the economy will converge successfully on a rational expectations equilibrium? For this, we work with perturbations to the T-mapping described by equations (8) or systems like it. We take this up in the next subsection.

2.2 Structured robust control

In the preceding subsection, we outlined the theory of least-squares learning in a relatively general setting. In this subsection we review some useful methods from robust control theory. Recall that our objective is to uncover the conditions under which monetary policy can maximize the prospect that the process of learning will converge on a REE—that is, to robustify learnability—so the integration of the theories of these two subsections is what will provide us with the tools we seek.

The argument that private agents might have to learn the true structure of the economy takes a useful step back from the assumption of known and certain linear rational expectations models. However, what the literature to date has usually taken as given is, first, that agents use least-squares learning to adapt their perceptions of the true economic structure, and second, that they know the correct linear or linearized form of the REE solution.¹⁰ Both of these assumptions can be questioned. It is a common-place convenience of macroecono-

⁹ In fact, in this example, the intercept coefficient, a , turns out to be irrelevant for the determination of learnability, although this result is not general.

¹⁰ Evans and Honkapohja [17] survey variations on least-squares learning, including under- and over-parameterized learning models and discounted (or constant-gain) least squares learning. Still, in general, either least-squares learning or constant gain learning is assumed. An exception is Marcet and Nicolini [32].

mists to formulate a dynamic stochastic general equilibrium model and then linearize that model. It is certainly possible that ill-informed agents use only linear approximations of their true decision rules. But it is hard to argue that the linearized decision rule is any more valid than some other approximation. Similarly, least-squares learning is the subject of research more because of its analytical tractability than its empirical plausibility. The utility of tractable, linear formulations of economic forms is undeniable; at the same time, however, the risk in over reliance on such forms should be just as apparent. There would appear to be a least a *prima facie* case for retaining the simplicity of linear state-space representations and linear rules, while taking seriously the consequences of such approximations.

With this in mind, we retain the assumption of a linear reference model, and least-squares learning on the part of agents, but assume that the process of learning is subject to uncertainty. Such uncertainty may arise because agents take their decision rules as simplifications of truly optimal decision rules due to the complexity of such rules. Attributing model uncertainty to agents can also be justified if we assume that least-squares learning is untenable in worlds where agents pick and choose the information to which they respond in forming and updating beliefs. The point is that from the perspective of the monetary authority, there are good reasons to be wary of both the assumed learning rules and the underlying models, and yet there is very little guidance on how to model those doubts. We motivate the present approach by positing a central bank that is concerned about ensuring learnability of the true model when learning is to subject errors, the distributions of which are unknown. Accordingly, we analyze these doubts using only a minimum amount of structure, drawing on the literature on *structured model uncertainty* and robust control.

For the most part, treatments of robust control have regarded model uncertainty as unstructured, that is, as uncertainty not ascribed to particular features of the model but instead represented by one or more additional shock variables wielded by some "evil agent" bent on causing harm.¹¹ The approach taken here differs in that we consider a central bank worried about how large agents' learning errors can be before learning fails to drive the economy towards eventual equilibrium. The central bank will employ techniques of robust

¹¹ See, in particular, Sargent [41], Giannoni [24], Hansen and Sargent ([26], [27]), Tetlow and von zur Muehlen ([45], [46]), Onatski and Stock [36], and Onatski and Williams [37].

control to set the parameters of her policy rule in a way that makes room for the largest degree of estimation errors by agents without rendering the true model unlearnable. To develop strategies for setting policies that determine maximum allowable misspecifications in agents' learning models while keeping the economy just shy of becoming unlearnable, we need to consider *structured model uncertainty*. Structured model uncertainty shares with its unstructured sibling a concern for uncertainty in the sense of Knight—meaning that the uncertainty is assumed to be nonparametric. But structured robust control differs in that it associates the uncertainty with perturbations to particular parts of the model. More importantly, to such perturbations we can assign a variety of structures, including time-series specifications. Work in this field was initiated by Doyle [15] and further developed in Dahleh and Diaz-Bobillo [52] and Zhou *et al.* [51], among others. Recent applications of this strand of robust control to monetary policy can be found in Onatski and Stock [36], Onatski [35], and Tetlow and von zur Muehlen [45].

While most contributions to the literature on monetary policy have been concerned with maximizing economic performance, our concern is maximizing the prospects for learnability. Thus our metric of success is not the usual one of maximized utility or a quadratic approximation thereof, although we will have a look at the "insurance premium" for robustness.

Boiled down to its essence, the five steps to designing policies subject to the constraint that agents must adapt to those policies and a learning model that may be misspecified are:

1. Write down a structural model of the economy and compute the conditions necessary for the model to attain a unique saddle-point stationary equilibrium.
2. Given the structural model, formulate the central bank's depiction of the perceived law of motion used by agents to learn the structural model. Substitute this into the structural model to arrive at the actual law of motion. We refer to this as the reference model. While the reference model is the central bank's best guess of the ALM, the bank understands that the reference model is only an approximation of the true ALM, and doubts remain about its local accuracy.¹²

¹² It is sometimes argued that robust control—by which people mean minmax approaches to model uncertainty—is unreasonable on the face of things. The argument is that the worst-case assumption is too extreme, that to quote a common phrase, "if I worried about the worst case outcome every day, I

3. Specify a set of perturbations to the reference model structured in such a way as to isolate the possible misspecifications to which the reference model is regarded to be most vulnerable.
4. For a given policy, use structured singular value analysis to determine the maximum allowable perturbations to the ALM that will bring the economy up to, but not beyond, the point of E-instability.
5. Finally, compute the policy for which the allowable range of misspecifications is the largest.

When, in the agents' learning model—the MSV-based PLM described in (4)—the parameters a , b , and c or Π , have been correctly estimated by agents, this model should be considered to be the true reduced form of the structural model in (1). Note, however, that even if individuals manage to specify their learning model correctly in terms of included variables and lag structures, the expectations of future output and inflation they base on these estimates are (at best) in a period of transition towards being truly rational. The model that agents actually estimate may differ from (1) in various ways that may be persistent. We want to determine how far off the true model agents' learning model can become before it becomes in principle unlearnable.

To begin, we rewrite the ALM from (6) and vectorize the disturbance, ϵ_t , to emphasize the stochastic nature of the estimating problem faced by agents,

$$Y_{t+1} = \Pi Y_t + \tilde{\epsilon}_t,$$

where $Y_t = [1, y_{t-1}, v_t]'$ is of dimension $n + 1$, $\tilde{\epsilon}_t = [0 \ 0 \ \epsilon_t]'$ and

$$\Pi = \begin{bmatrix} 1 & 0 & 0 \\ A + M(I + b)a & (N + Mb^2) & M(bc + c\rho) + P \\ 0 & 0 & \rho \end{bmatrix}.$$

Notice that by using the ALM, we are modeling the problem from the policy authority's point of view. The authority is taken as knowing, up to the perturbations we are about to

wouldn't get out of bed in the morning". Such remarks miss the point that the worst-case outcome should be thought of as local in nature. Decision makers are envisioned as wanting to protect against uncertainties that are empirically indistinguishable from the data generating process underlying their reference models.

add to the model, the structure of private agents' learning problems. As a consequence, the authority is in a position to influence the resolution of that problem.

Potential errors in parameter estimation are then represented by a *perturbation block*, Δ . In principle, the Δ operator can be structured to implement a variety of misspecifications, including alternative dynamic features. Robust control theory is remarkably rich in how it allows one to consider omitted lag dynamics, inappropriate exogeneity restrictions, missing nonlinearities, and time variation. This being the first paper of its kind, we keep our goals modest: in the language of linear operator theory, we will confine our analysis to *linear time-invariant scalar* (LTI-scalar) perturbations. LTI-scalar perturbations represent such events as one-time shifts and structural breaks in model parameters, as agents perceive them. Such perturbations have been the subject of study of parametric model uncertainty; see, e.g., Bullard and Euseppi [7].¹³ With this restriction, the perturbed model becomes:¹⁴

$$\begin{aligned}\tilde{\epsilon}_t &= Y_{t+1} - [\Pi + W_1\Delta W_2]Y_t, \\ &= [Q - W_1\Delta W_2]Y_t,\end{aligned}\tag{10}$$

where $Q = I_n L^{-1} - \Pi$, L is the lag operator, Δ is a $k \times k$ linear, time-invariant block-diagonal operator representing potentially destabilizing learning errors, and W_1 and W_2 are, respectively, $(n+1) \times k$ and $k \times (n+1)$ selector matrices of zeros and ones that select which parameters in which equations are deemed to be subject to such errors. Either W_1 or W_2 can, in addition, be chosen to attach scalar weights to the individual perturbations so as to reflect relative uncertainties with which model estimates are to be regarded. The second line is convenient for analyzing stability of the perturbed model under potentially destabilizing learning errors. Using this construction, the perturbation operator, Δ , and the weighting matrices can be structured so that misspecifications are focused on particular features of the model deemed especially susceptible to learning errors involving the model's variables for any chosen lag or lags.

The essence of this paper is to find out how large, in a sense to be defined presently, the misspecifications represented by the perturbations in (10)—called the *radius of allowable*

¹³ See Evans and Honkapohja [17] for some treatment of learning with an over-parameterized PLM.

¹⁴ Multiplicative errors in specification would be modeled in a manner analogous to (10): $\epsilon_t = [A(1 - W_1\Delta W_2)]Y_t$.

perturbations—can become without eliciting a failure of convergence to rational expectations equilibrium. Any policy that expands the set of affordable perturbations is one that allows the widest room for misspecifications committed by agents and thus offers an improved chance that policy will not be destabilizing. To do this we bring the tools of *structured robust control* analysis mentioned earlier.

Let \mathcal{D} denote the class of allowable perturbations to the set of parameters of a model defined as those that carry with them the structure information of the perturbations. Let $r > 0$ be some finite scalar and define D_r as the set of perturbations in (10) that obey $\|\Delta\| < r$, where $\|\Delta\|$ is the *induced norm* of Δ considered as an operator acting in a normed space of random processes¹⁵. The scalar, r , can be considered a single measure of the maximum size of errors in estimation. A policy authority wishing to operate with as much room to maneuver as possible will act to maximize this range. For the tools to be employed here, norms will be defined in complex space. In what follows, much use is made of the concept of *maximum singular value*, conventionally denoted by $\bar{\sigma}$ ¹⁶. For reasons that will become clearer below, the norm of Δ that we shall use will be the L_∞ norm of the function $\Delta(e^{i\omega})$, defined as the largest singular value of $\Delta(e^{i\omega})$ on the frequency range $\omega \in [-\pi, \pi]$:

$$\|\Delta\|_\infty = \left\{ \sup_{\omega} \max \text{eig}[\Delta'(e^{-i\omega})\Delta(e^{i\omega})] \right\}^{1/2}, \quad (11)$$

where $\max \cdot \text{eig}$ denotes the maximum eigenvalue. The choice of $\|\Delta\|_\infty$ as a measure of the size of perturbations conveys a sense that the authority is concerned with worst-case outcomes.

Imagine two artificial vectors, $h_t = [h_{1t}, h_{2t}, \dots, h_{kt}]'$ and $p_t = [p_{1t}, p_{2t}, \dots, p_{kt}]'$, connected

¹⁵ Induced norms are defined as follows. Let X be a vector space. A real-valued function $\|\cdot\|$ defined on X is said to be a norm on X if it satisfies: (i) $\|x\| \geq 0$, (ii) $\|x\| = 0$ only if $x = 0$, (iii) $\|\alpha x\| = |\alpha| \|x\|$ for any scalar α , (iv) $\|x + y\| \leq \|x\| + \|y\|$ for any $x \in X$ and $y \in X$. For $x \in C^n$, the \mathcal{L}_p vector p-norm on x is defined as $\|x\|_p = (\sum_{i=1}^n |x_i|^p)^{1/p}$, where $1 \leq p \leq \infty$. For $p = 2$, $\mathcal{L}_2 = \|x\|_2 = \sqrt{\sum_{i=1}^n |x_i|^2}$, that is, the quadratic problem. Corresponding to \mathcal{L}_2 , we also have $\mathcal{L}_1 = \sum_{i=1}^n |x_i|$, and $\mathcal{L}_\infty = \max_{1 \leq i \leq n} |x_i|$. Finally, let $A = [a_{ij}] \in C^{m \times n}$ in an equation $y_t = A_t x_t$, where x_t may some random vector. The matrix norm *induced* by a vector p-norm, $\|x\|_p$, is $\|A\|_p \equiv \sup_{x \neq 0} \frac{\|Ax\|_p}{\|x\|_p}$. Note that for $p = 2$, $\|A\|_2 = \sqrt{\max \text{eig}(A \cdot A)}$ and $\|A\|_1 = \max_{1 \leq j \leq n} \sum_{i=1}^m |a_{ij}|$. More details are given in Tetlow and von zur Muehlen [45].

¹⁶ As is apparent from the expression in (11), the largest singular value, $\bar{\sigma}(X)$, of a matrix, X , is the largest eigenvalue of $X'X$.

to each other and to Y_t via¹⁷

$$\begin{aligned} p_t &= W_2 Y_t \\ h_t &= \Delta \cdot p_t. \end{aligned}$$

Then we may recast the perturbed system (10) as the *augmented feedback loop*¹⁸

$$\begin{bmatrix} Y_{t+1} \\ p_t \end{bmatrix} = \begin{bmatrix} \Pi & W_1 \\ W_2 & 0 \end{bmatrix} \begin{bmatrix} Y_t \\ h_t \end{bmatrix}, \quad (12)$$

$$h_t = \Delta \cdot p_t. \quad (13)$$

A reduced-form representation of this loop (from h_t to Y_t and p_t) is the transfer function

$$\begin{bmatrix} Y_t \\ p_t \end{bmatrix} = \begin{bmatrix} G_1 \\ G_2 \end{bmatrix} h_t, \quad (14)$$

where $G_1 = (I_n L^{-1} - \Pi)^{-1} W_1$, and $G_2 = W_2 (I_n L^{-1} - \Pi)^{-1} W_1$ is a $n \times k$ matrix, where k is the number of diagonal elements in Δ . As we shall see, the stability of the interconnection between h_t and p_t , representing a feedforward $p_t = G_2 h_t$ and a feedback $h_t = \Delta \cdot p_t$, is critical. Note first that, together, these two relationships imply the homogenous matrix equation

$$0 = (I_k - G_2 \Delta) p_t. \quad (15)$$

An E-stable ALM is also dynamically stable, meaning that Π has all its eigenvalues inside the unit circle. To make this link, the following theorem is critical.

Theorem 1 *The Small Gain Theorem.*

Let $Re(s)$ denote the real part of $s \in \mathcal{C}$, where \mathcal{C} is the field of complex numbers, and let \mathcal{H}_∞ denote the set of \mathcal{L}_∞ functions analytic in $Re(s) > 0$. Furthermore, let \mathcal{RH}_∞ designate the set of real, rational values in the \mathcal{H}_∞ -normed space. Suppose $G_2 \in \mathcal{RH}_\infty$ and $r > 0$. Then the interconnected system in (12)-(13) is well posed and internally stable for all $\Delta(s) \in \mathcal{RH}_\infty$ with

(a) $\|\Delta\| \leq 1/r$, if and only if $\|G_2(s)\|_\infty < r$,

(b) $\|\Delta\| < 1/r$, if and only if $\|G_2(s)\|_\infty \leq r$.

¹⁷ See Dahleh and Bobillo [14], chapter 10.

¹⁸ Because the random errors in this model play no role in what follows, we leave out the ϵ vector.

Proof: See Zhou and Doyle,[52] p. 137. \square

By assumption, Q , defined in (10), is invertible on the unit circle, allowing us to write¹⁹

$$\begin{aligned} \det(Q)\det(I_k - G_2\Delta) &= \det(Q)\det(I_k - W_2Q^{-1}W_1\Delta) \\ &= \det(Q)\det(I_k - Q^{-1}W_1\Delta W_2) \\ &= \det(Q - W_1\Delta W_2). \end{aligned}$$

The preceding expressions establish the link between stability of the interconnection, G_2 , and stability of the perturbed model: if $\det(I_k - G_2\Delta) = 0$, then the perturbed model (10) is no longer invertible on the unit circle, hence unstable, and vice versa. Thus, any policy rule that stabilizes the G_2 also stabilizes the augmented system (12)-(13). The question to be asked then is how large, in the sense $\|\cdot\|_\infty$, can Δ become without destabilizing the feedback system (12)-(13).

The settings we consider involve linear time-invariant perturbations, where the object is to find the minimum of the largest singular value of the matrix, Δ , from the class of D_r such that $I - G_2\Delta$ is not invertible. The inverse of this minimum, expressed in the frequency domain, is the *structured singular value*²⁰ of G_2 with respect to D_r , defined at each frequency, $\omega \in [-\pi, \pi]$,

$$\mu[G_2(e^{i\omega})] = \frac{1}{\min\{\bar{\sigma}[\Delta(e^{i\omega})] : \Delta \in D_r, \det(I - G_2\Delta)(e^{i\omega}) = 0\}}, \quad (16)$$

with the provision that if there is no Δ such that $\det(I - G_2\Delta)(e^{i\omega}) = 0$, then $\mu[G_2(e^{i\omega})] = 0$.

The small gain theorem then tells us that, for some $r > 0$, the loop (12)-(13) is well posed and internally stable for all $\Delta(\cdot) \in \mathcal{D}_r$ with $\|\Delta\| < r$, if and only if $\sup_{\omega \in \mathcal{R}} \mu[G_2(e^{i\omega})] \leq 1/r$.

Let ϕ denote a vector of policy parameters.

Thus we can now formally state the problem that interests us as seeking a best $\phi = \phi^*$ by finding a maximum value of $\mu = \bar{\mu}$, satisfying

$$\bar{\mu}(\phi^*) = \inf_{\phi} \sup_{\omega \in \mathcal{R}} \mu[G_2(e^{i\omega})]$$

¹⁹ The small gain theorem links the stability of the loop between p and h under perturbations to the full system subject to model uncertainty. For some sufficiently large number r , such that $\|\Delta\|_\infty < r$, the determinant $\det(I_k - G_2\Delta) \neq 0$. Now raise r to some value \bar{r} such that $\det(I_k - G_2\Delta) = \det(I_k - W_2Q^{-1}W_1\Delta) = 0$.

²⁰ The singular value is said to be "structured" in recognition of structure built into the perturbation matrix, Δ . Through selection of the structure, Δ can encompass uncertainties about one-time shifts in selected parameters, unmodeled dynamics in parts of the model, and nonlinearities, among other things.

subject to the satisfaction of the saddle-point stability condition for the relevant model.

The solution to this problem is not amenable to analytical methods, except in special cases, an example of which we explore in the next section. Instead, we will employ efficient numerical techniques to find the lower bound on the structured singular value. The minimum of $\mu^{-1}(G_2)$ over $\omega \in [0, \pi]$ is exactly the maximal allowable range of misspecification for a given policy. A monetary authority wishing to give agents the widest latitude for learning errors that nevertheless allow the system to converge on REE selects those parameters in its policy rule that yield largest value of r .

Figure 1 below provides a schematic representation of what is done. Given a policy rule, the ALM (and hence the reference model) is represented by the transition matrix, Π . By assumption it is in the stable and determinate region of the space. The central bank chooses perturbation, Δ , to Π . The largest feasible perturbation, Π_{Δ}^* —the one that renders the largest radius, r , defines a region shown by the ellipse within which any ALM, including but not restricted to the reference model ALM, will converge on a rational expectations equilibrium. The weights, $W1$ and $W2$, determine the shape of the ellipse. Perturbations larger (in norm) than r will push the ellipse over the line into indeterminate or explosive regions; smaller perturbations provide less robustness than the optimal one.

3 Two examples

We study two sample economies, one the very simple model of money demand in hyperinflations of Cagan [10], the other the linearized neo-Keynesian model originated by Woodford ([48], [49]), Rotemberg and Woodford [39] and Goodfriend and King [25]. Closed-form solutions for μ , being non-linear functions of the eigenvalues of models, are not generally feasible. However, some insight is possible through considering simple scalar example economies like the Cagan model. The second has the virtue of having been studied extensively in the literature on monetary policy design. It thus provides some solid benchmarks for comparison.




Figure 1: Schematic representation of robust learnability

3.1 A simple univariate example

Consider a version of Cagan's monetary model, cited in Evans and Honkapohja [17], although our rendition differs slightly. The model has two equations, one determining (the log of) the price level, p_t , and the other a simple monetary feedback rule determining the (log of the) money supply, m_t :

$$m_t - p_t = -\kappa(E_t p_{t+1} - p_t)$$

$$m_t = \chi - \phi p_{t-1}.$$

Normally, all parameters should be greater than zero; for κ this means that money demand is inversely related to expected inflation. We will relax this assumption a bit later. Combining the two equations leads to:

$$p_t = \alpha + \beta E_t p_{t+1} - \gamma p_{t-1}, \tag{17}$$

where $\alpha = \chi/(1 + \kappa)$, $\beta = \kappa/(1 + \kappa)$, and $\gamma = \phi/(1 + \kappa)$. To set the stage for what follows, let us consider the conditions for a unique rational expectations equilibrium. First off, we will clearly want to assume that $\beta, \gamma \neq 0$ to avoid degenerate models. We will go a bit further and assume that $\beta - \gamma \neq 1$, which is a mild invertibility assumption. We will also

assume that $\kappa \neq 0$ and $\kappa \neq -1$. Finally, for simplicity we will focus on the case where solutions are real. Standard methods then reveal:

$$\beta\lambda_i^2 - \lambda_i - \gamma = \frac{\kappa}{1 + \kappa}\lambda_i^2 - \lambda_i - \frac{\phi}{1 + \kappa} = 0 \quad (18)$$

where $\lambda_i, i = 1, 2$ are the eigenvalues of the quadratic equations associated with (17). As is well known from Blanchard and Kahn [3], among other sources, Equation (18) is the key to establishing existence and uniqueness of saddle-point equilibrium. Letting $\lambda_1 \geq \lambda_2$, without loss of generality, if $1 < \lambda_2 \leq \lambda_1$, then the model is *explosive*, meaning that there are no initial conditions which can establish a saddle-point equilibrium; if $\lambda_2 \leq \lambda_1 < 1$, then the model is said to be *irregular*, meaning that every set of initial conditions leads to a different equilibrium. Putting the same thing in different words, the equilibrium is said to be *indeterminate* (or non-unique). Finally, the condition $\lambda_2 < 1 < \lambda_1$ is the condition for saddle-point stability; that is, the condition by which all (local) sets of initial conditions converge, in expectation, on the same equilibrium. In this instance, the equilibrium is said to be *determinate* (or regular, or unique). Inspection of the equation to the right of the first equality in equation (18) shows that the determinacy of the model is governed by the interaction between the structural money-demand parameter, κ , and the policy feedback parameter, ϕ .

Proposition 1 *Assume that $\kappa \neq -1$. For $\kappa > -1$, determinacy requires: $\phi > -1$ and $\phi < 1 + 2\kappa$. For $\kappa < -1$, determinacy requires $\phi < -1$ and $\phi > 1 + 2\kappa$.*

Proof: Equation (18) means that $|\beta - \gamma| = |(\kappa - \phi)/(1 + \kappa)| < 1$ is the condition for exactly one eigenvalue to be above unity. There are two cases. Assume first that $\kappa > -1$, which means that the denominator is always positive. Then simple arithmetic shows that $\{\phi, \kappa\} \in \{\mathcal{R} : |\frac{\kappa - \phi}{1 + \kappa}| < 1\} = \{\phi > -1\} \cup \{\phi < 1 + 2\kappa\}$ is implied. Now consider $\kappa < -1$. In this instance, $\{\phi, \kappa\} \in \{\mathcal{R} : |\frac{\kappa - \phi}{1 + \kappa}| < 1\} = \{\phi < -1\} \cup \{\phi > 1 + 2\kappa\}$. \square

Now let us assume that agents form expectations employing adaptive learning, and designate expectations formation in this way with the operator, E_t^* . The perceived law of motion for this model is assumed to be $p_t = a + bp_{t-1}$ —the minimum state variable (MSV) solution,

implying $E_t^* p_{t+1} = (1+b)a + b^2 p_{t-1}$. The actual law of motion is found by substituting the PLM into the structural model:

$$p_t = [\alpha + \beta a(1+b)] + (\beta b^2 - \gamma)p_{t-1}. \quad (19)$$

Following the steps outlined earlier, the ALM is $p_t = T_a(a, b) + T_b(a, b)p_{t-1}$ where T defines the mapping $\begin{pmatrix} a \\ b \end{pmatrix} = T\left(\begin{pmatrix} a \\ b \end{pmatrix}\right)$. This is

$$T_a: a = \alpha + \beta a(1+b) \quad (20)$$

$$T_b: b = \beta b^2 - \gamma. \quad (21)$$

It is equation (21) that is key for the learnability of the model. But notice that (21) is identical to equation (18) with $b = \lambda_i$. This means that the conditions for learnability and determinacy are tightly connected in this particular model, as we shall discuss in detail below. The solutions to equations (20) and (21) are:

$$a = \alpha/[1 - \beta(1+b)] \quad (22)$$

$$b = .5[1 \pm \sqrt{1 + 4\beta\gamma}]/\beta. \quad (23)$$

Equation (23) is quadratic with one root greater than or equal to unity, and the other less than unity. Designate the larger of the two values for b as b^+ and the smaller as b^- . Existence of the REE requires us to choose the smaller root; otherwise, $b^+ \geq b^- > 1$. The ordinary differential equation system implied by this mapping is

$$\frac{d\begin{pmatrix} a \\ b \end{pmatrix}}{d\tau} = T\left(\begin{pmatrix} a \\ b \end{pmatrix}\right) - \begin{pmatrix} a \\ b \end{pmatrix},$$

for which the associated DT matrix is derived by differentiating $[T_a \ T_b]'$ with respect to a and b :

$$DT = \begin{bmatrix} \beta(1+b) & a\beta \\ 0 & 2\beta b \end{bmatrix}.$$

The eigenvalues of $DT - I$ are,

$$\psi_1 = 2\beta b - 1 \quad (24)$$

$$\psi_2 = \beta(1+b) - 1. \quad (25)$$

Satisfaction of the weak E-stability condition requires that both eigenvalues be negative.

Proposition 2 *Determinate solutions of the Cagan model that are real are also E-stable.*

Proof: Substitute the expression for b^- into (24) to get $\psi_1 = -[1 + 4\kappa\phi/(1 + \kappa)^2]^{1/2}$ and do the same for (25) to arrive at $\psi_2 = \kappa/(1 + \kappa) - \frac{1}{2} + \psi_1$. Substitute ψ_1 into b^- to arrive at: $b^- = \frac{1+\kappa}{2\kappa}[1 + \psi_1]$. A necessary and sufficient condition for $\psi_1 < 0$, $\psi_1 \in \Re$ is (P1): $\phi > (1 + \kappa)^2/(4\kappa) \equiv S$. Proposition 1 imposes restrictions on ϕ as a function of κ to ensure determinacy. For $\kappa > -1$, these are $\phi > -1$ and $\phi < 1 + 2\kappa$. Substituting these into the expression for ψ_1 gives $\psi_1 < -[(1 + 3\kappa)/(1 + \kappa)]^2$ which readily yields $\psi_1 < 0$ and is real whenever $|b^-| < 1$ provided the solution is real. Simple substitution of ψ_1 into ψ_2 shows that ψ_2 is also negative. For $\kappa < -1$, a similar proof applies. \square

Having outlined the connection between b and β (or κ) for E-stability, let us now consider unstructured perturbations to the ALM. Let $X_t = [1 \ p_t]'$. The reference ALM model is then written as $X_t = \Pi X_{t-1}$, where

$$\Pi = \begin{bmatrix} 1 & 0 \\ \alpha + \beta a(1 + b) & \beta b^2 - \gamma \end{bmatrix}$$

is the model's transition matrix. For simplicity, let us focus on b as the object of concern to policy makers, and let the policy maker apply structured perturbations to Π , scaled by the parameter, σ_b . The scaling parameter can be thought of as a standard deviation, but need not be. Letting $W_1 = [0 \ \sigma_b]'$ and $W_2 = [0 \ 1]$, write the perturbed matrix Π as:

$$\Pi_\Delta = \begin{bmatrix} 1 & 0 \\ \alpha + \beta a(1 + b) & \beta b^2 - \gamma + \Delta \end{bmatrix}.$$

As in (14), the relevant matrices are defined in complex space. Accordingly, let $z = e^{i\omega}$, $\omega \in [-\pi, \pi]$. To find the maximal allowable perturbation, write

$$G = \begin{bmatrix} z^{-1} - 1 & 0 \\ -\alpha - \beta a(1 + b) & z^{-1} - \beta b^2 + \gamma \end{bmatrix} = I \cdot z^{-1} - \Pi,$$

which, defining $W_1 = [0 \ \sigma_b]'$ and $W_2 = [0 \ 1]$, is used to form G_2 :

$$\begin{aligned} G_2 &= W_2 G^{-1} W_1 \\ &= [0 \ 1] \begin{bmatrix} \frac{z}{1-z} & 0 \\ \frac{(\alpha - \beta a(1+b))z}{(1-z)(1 - (\beta b^2 - \gamma)z)} & \frac{z}{1 - (\beta b^2 - \gamma)z} \end{bmatrix} \begin{bmatrix} 0 \\ \sigma_b \end{bmatrix} \\ &= \frac{\sigma_b z}{1 - (\beta b^2 - \gamma)z}. \end{aligned}$$

It is for this expression that we seek the smallest structured singular value, μ , as indicated by (16).

In the multivariate case, the scaling parameter σ_b , can be parameterized as the standard deviation of b relative to a , although other methods of parameterization can be entertained. Doing so would reflect a concern for robustness of the decision maker and thus could also be thought of as a taste parameter. Since it is a *relative* term, it will turn out to be irrelevant in this scalar case, and so from here we set it to unity without loss of generality. The structured norm of G_2 —equal to the absolute value of this last expression (see footnote 15)—is μ . It is also easily established that the maximum of μ over the frequency range $\{-\pi, \pi\}$, let us call it $\bar{\mu}$ is $\bar{\mu} = \|\Delta\|_\infty = |G_2|$ arises at frequency π . Also, since at frequency π , $z = -1$, it follows that $|G_2| = \bar{\mu} = \frac{1}{1+b}$, or equivalently, the *allowable perturbation* is:

$$\begin{aligned}\Delta &= \frac{1}{\bar{\mu}} = 1 + b \\ &= 1 + \beta b^2 - \gamma \\ &= 1 + \frac{1}{2\beta} [1 - \sqrt{1 + 4\beta\phi/(1 + \kappa)}],\end{aligned}\tag{26}$$

which depends inversely on the policy parameter, ϕ .²¹ Note also that while we have derived this expression for Δ by applying perturbations to the ALM, we would have obtained exactly the same result by working with the PLM.

If equation Δ is the allowable perturbation, conditional on a given ϕ , then we can define a ϕ^* as the policy maker's optimal choice of ϕ , where optimality is defined in the sense of choosing the largest possible perturbation to b —call it Δ^* —such that the model will retain the property of E-stability. Let us call this the *maximum allowable perturbation*. It is the Δ^* and the associated ϕ^* that is at a boundary where Δ is just above -1 :

$$\phi^* = 1 + 2\kappa - \epsilon,\tag{27}$$

where $\phi < 1 + 2\kappa$ maintains stable convergence toward a REE and ϵ is an arbitrarily small positive constant necessary to keep $b + \Delta$ off the unit circle. Note that this expression for ϕ^* indicates that the monetary authority will always respond more than one-for-one to

²¹ Note that at frequency π , $1 - G_2\Delta = 1 - \frac{\sigma_b}{(1+b)} \frac{(1+b)}{\sigma_b} = 0$ as required by the definition of μ .

deviations in lagged prices from steady state, with the extent of that over-response being a positive function of the slope of the money demand function. Substituting these expressions back into our perturbed transition matrix,

$$\begin{aligned}
\Pi_{\Delta}^* &= \begin{bmatrix} 1 & 0 \\ \alpha - \beta a(1+b) & \beta b^2 - \gamma + \Delta \end{bmatrix} \\
&= \begin{bmatrix} 1 & 0 \\ \alpha - \beta a(1+b) & 1 + 2(\beta b^2 - \gamma) \end{bmatrix} \\
&= \begin{bmatrix} 1 & 0 \\ \alpha - \beta a(1+b) & -1 + \eta \end{bmatrix}, \tag{28}
\end{aligned}$$

where η is an arbitrarily small number, as determined by ϵ in (27). The preceding confirms that the authority's policy is resilient to a perturbation in the learning model that pushes the transition matrix is to the borderline of instability. In other words, setting a ϕ that allows for the maximal stable misspecification of the learning model is one that permits convergence to the REE.

Figure 2 shows the regions of dynamic stability and learnability for the Cagan model as functions of the structural parameters: the absolute interest elasticity of money demand, κ , and the monetary policy feedback parameter, ϕ . As noted in the legend to the right, the determinate regions of the structural model are the blue areas, to the northeast and southwest of the figure. Proposition 1 above, warns the monetary authority to stay out of the indeterminate regions, the sliver of purple toward the southeast of the chart, that is possible for some $\kappa > 0$ when $\phi < -1$, and the larger region of purple to the west. Also to be avoided are the orange regions of explosive solutions to the north and south.

The E-stable region is the large area between the two dashed lines. The first thing to note is that E-stability does not imply determinacy: convergence in learning on indeterminate equilibria in the area where both $-1 < \kappa < -1/2$ and $-1 < \phi < 0$, is possible, corroborating a point made by Evans and McGough [20] in a different context. In addition, learnability of unstable equilibria is also possible as shown by the orange regions between the two dashed lines. Indeed, even if one were to accept *a priori* that $\kappa > 0$, as Cagan assumed, there are unstable equilibria that are learnable. At the same time, the figure clearly shows what Proposition 2 noted: for this model, determinate models are always E-stable; the blue region is entirely within the area bordered by the two dashed lines. It follows that in the special




Figure 2: Regions of determinacy and E-stability in the Cagan model

case of the Cagan model, robustifying learnability is equivalent to maximizing the basin of attraction for the rational expectations equilibrium of the model. The loci of robust policies, ϕ^* , conditional of values of κ , is shown by the thick diagonal line running from the south west to north east of the chart, and marked $\phi^* = 1 + 2\kappa$. The line shows that contrary to what unguided intuition might suggest, the robust policy does not choose a rule that is in the middle of the blue determinate and E-stable region, but rather chooses a policy that might be quite close to the boundary of indeterminacy for the REE. Doing so increases the region of E-stable ALMs—something that cannot be seen in the chart—and thereby enhances the prospects for convergence on an REE.

3.2 The canonical New Keynesian model

We now turn to an analysis of the canonical New Keynesian business cycle model of Rotemberg and Woodford [39], Goodfriend and King [39] and others. Clarida, Gali, and Gertler [12] used this model to derive optimal discretionary as well as optimal commitment rules. Their

version includes a specified process for exogenous natural output. Evans and Honkapohja [18] study this model to explore issues of determinacy and learnability for several optimal commitment rules. Bullard and Mitra [9] likewise use the Woodford model to examine determinacy and learnability of variants of the Taylor rule.

The behavior of the private sector is described by two equations. The aggregate demand (IS) equation is a log-linearized Euler equation derived from optimal consumer behavior,

$$x_t = E_t^* x_{t+1} - \sigma[r_t - E_t^* \pi_{t+1} - r_t^n], \quad (29)$$

and the aggregate supply (AS) equation—indeed, the price setting rule for monopolistically competitive firms is,

$$\pi_t = \kappa x_t + \beta E_t^* \pi_{t+1}, \quad (30)$$

where x is the log deviation of output from potential output, π is inflation, r is a short-term interest rate controlled by the central bank, and r^n is the natural interest rate. For the application of Bullard and Mitra’s [8] (BM) example, we assume that r_t^n is driven by a first-order autoregressive process,

$$r_t^n = \rho_r r_{t-1}^n + \epsilon_{r,t}, \quad (31)$$

$0 \leq |\rho_r| < 1$, and $\epsilon_{r,t} \sim iid(0, \sigma_r^2)$. This is essentially Woodford’s [48] version of this model, which specifies that aggregate demand responds to the deviation of the real rate, $r_t - E_t \pi_{t+1}$ from the natural rate, r_t^n .

We need to close the model with an interest-rate feedback rule. We study three types of policy rules. In the first set of experiments described in Section 3.3, a central bank chooses an interest rate setting in each period as a reaction to observed events, such as inflation and the output gap, without explicitly attempting to improve some measure of welfare. Instead, the policy authority is mindful of the effect its policy has on the prospect of the economy reaching REE and designs its rule accordingly. Bullard and Mitra [8] study such rules for their properties in promoting learnable equilibria and consider that effort as prior to one of finding optimal policy rules consistent with REE. We take this analysis further by seeking to find policy rules that maximize learnability of agents’ models when policy influences the outcome.

The information protocol in these experiments is as follows. The central bank knows the structural model and has access to the data. Economic agents see the data, which change over time, and formulate the perceived law of motion. Agents form expectations based on recursive (least-squares) estimation of a reduced form. The data are regenerated each period, subject to the authority having implemented its policy and agents' having made investment and consumption decisions based on their newly formed expectations.

We assume that agents mistakenly specify a vector-autoregressive model in the endogenous and exogenous variables of the model. That means we assume the learning model to be overparameterized in comparison with the model implied by the MSV solution. The scaling factors used in W_1 to scale the perturbations to the PLM are the standard errors of the coefficients obtained from an initial run of a recursive least squares regression of such a VAR with data being updated by the true model, given an arbitrary but determinate parameterization of the policy rule being studied. As noted earlier, an alternative approach would be to revise the scalings with each trial policy, given that the VAR would likely change with each parameterization of policy. We leave this for a revision.

3.3 Simple interest-rate feedback rules

This section describes two versions of the Taylor rule analyzed by Bullard and Mitra [8]. The complete system comprises equations (29)-(32), and the exogenous variable, r_t^n . The policy instrument is the nominal interest rate, r_t . The first policy rule specifies that the interest rate responds to lagged inflation and the lagged output gap. In their paper, BM study the role of interest-rate inertia and so include a lagged interest rate term.

$$r_t = \phi_\pi \pi_{t-1} + \phi_x x_{t-1} + \phi_r r_{t-1} \quad (32)$$

McCallum has advocated such a *lagged data* rule because of its implementability, given that contemporaneous data are generally not available in real time to policy makers.

Some research suggests that forward-looking rules perform well in theory (see, e.g., Evans and Honkapohja [18]) as well as in actual economies, such as Germany, Japan, and the US (see Clarida, Gali, and Gertler [11]). Accordingly, BM propose the rule

$$r_t = \phi_\pi E_t^* \pi_{t+1} + \phi_x E_t^* x_{t+1} + \phi_r r_{t-1}. \quad (33)$$

The expectations operator E^* has an asterisk to indicate that expectations need not be rational.

Finally, the most popular rules of this class are contemporaneous data rules, of which the following is our choice:

$$r_t = \phi_\pi \pi_t + \phi_x x_t + \phi_r r_{t-1} \quad (34)$$

where as before, we allow the lagged federal funds rate to appear to capture instrument-smoothing behavior by uncertainty averse decision makers.

3.4 Results

We adopt BM's calibration for the New Keynesian model's parameters, $\sigma = 1/.157$, $\kappa = .024$, $\beta = .99$, and $\rho = .35$, the same calibration as in Woodford [49]. We also set $\sigma_r = 0.01$. For reference purposes, it is useful to compare our results against those of rules that are not parameterized with robust learnability in mind. To facilitate this, we employ a standard quadratic loss function:

$$L_t = \frac{1000}{2} \sum_{j=0}^{\infty} \beta^j [(\pi_{t+j} - \pi^*)^2 + \lambda_x x_{t+j}^2 + \lambda_r (r_{t+j} - r^*)^2]. \quad (35)$$

Walsh [47] shows that with the values $\lambda_x = .077$ and $\lambda_i = .027$, equation (35) is the quadratic approximation to the social welfare function of the model. Rules that are computed to maximize the prospect of convergence to REE under the greatest possible misspecification of the ALM model in the manner described above will be referred to as "robust" or "robust learnable" rules. A credible benchmark against which to compare these robust rules, are what we shall refer to as *optimized rules*. These are rules that minimize (35) subject to (29), (30), (31) and one of either (32), (34) or (33). Such rules can be optimized using a standard hill-climbing algorithm using methods well described in the appendix to Tetlow and von zur Muehlen [44] among other sources.

Let us consider the lagged-data rule first. BM find that the determinacy of a unique rational expectations equilibrium, as well as convergence toward that equilibrium when agents learn adaptively, is extremely sensitive to the policy parameters, ϕ_r , ϕ_x , and ϕ_π . Without some degree of monetary policy inertia, ($\phi_r > 0$), this model is determinate and learnable,

with the above calibrations, only if the Taylor principle holds, ($\phi_\pi > 1$), **and** the response to the output gap is modest, ($\phi_x \leq 0.5$). Insufficient or excessive responsiveness to either inflation or the output gap can in some instances lead to explosive instability or indeterminacy. Through simulation, BM establish the regions for the parameters that lead to determinacy as well as E-stability.

Table 1 shows our results. The table is broken into three panels. The upper panel—the rows marked (1) to (3)—shows optimized rules. The second panel, contains some results for the generic Taylor rule. Finally, the third panel shows our robust learnable rules. The next-to-last column of the table gives a measure of the total uncertainty that the PLM can tolerate under the cited policy. It is a measure of the maximal allowable deviation embodied in $1/\mu$.²² The last column shows the loss as measured by (35).

Table 1 : Standard and robust learnable rules

	row	ϕ_x	ϕ_π	ϕ_r	radius ¹	L^2
<i>optimized rules:</i>						
lagged data rule	(1)	0.052	0.993	1.13	1.07	3.679
contemporaneous data rule	(2)	0.053	0.995	1.12	1.06	3.626
forecast-based rule	(3)	0.286	0.999	1.32	0.88	3.628
<i>standard rules:</i>						
Taylor rule	(4)	0.500	1.500	0	0.85	5.690
<i>robust learnable rules:</i>						
lagged data rule	(5)	0.065	0.40	1.10	1.16	3.712
contemporaneous data rule	(6)	0.052	1.21	1.41	1.13	3.701
forecast-based rule	(7)	0.040	2.80	0.10	2.32	4.434

1. Magnitude of the largest allowable perturbation. $r = \|W_1 \Delta W_2\|_\infty$

2. Asymptotic loss, calculated according to eq. (35) in REE under the reference model.

Let us concentrate initially on our optimized rules along with the Taylor rule to provide some context for the robust learnable rules. The lagged data rule, shown in row (1), and the contemporaneous data rule, (2), are essentially the same. They both feature very small feedback on the output gap, and strong responses to inflation. Moreover, they also feature funds rate persistence that amounts to a first-difference rule; that is, a rule where the dependent variable is Δr rather than r . The forecast-based rule, in line (3), has much stronger feedback on the output gap, although proper interpretation of this requires noting that in equilibrium the expectation of future output gaps will always be smaller than actual

²² For comparison of the trials with each other and also to give a sense of natural units related to the scalings we employed, the radius is calculated as the H_∞ norm of the scaled perturbations to the PLM model: radius = $\|W_1 \Delta W_2\|_\infty$.

gaps because of the absence of expected future shocks and the internalization of future policy in the formulation of that expectation. Thus, the response of the funds rate to the expected future gap will not be as large as the feedback coefficient alone might lead one to believe.

These three rules confirm the received wisdom of monetary control in New Keynesian models, to wit: strong feedback on inflation, comparatively little on output, and strong persistence in funds rate setting. These rules are chosen to minimize the loss shown in the right-hand column of the table; the losses for all three are very similar, at a little over 3.6.

The results for the Taylor rule demonstrate, indirectly, the oft-discussed advantages of persistence in funds rate setting for monetary control. Without such persistence, the Taylor rule produces losses that are substantially higher than those of the optimized rules.

Now let us turn to the robust learnable rules in the bottom panel of the table, concentrating for the moment on the lagged data and contemporaneous data rules shown in lines (5) and (6). The first thing to note is that the results confirm the efficacy of persistence in instrument setting. The robust learnable rules are at least as persistent—if persistence greater than unity is a meaningful concept—as the optimized rules. At the same time, while persistence is evidently useful for learnability, our results do not point to the hyper-persistence result, ($\phi_r \gg 1$), that BM hint at. To understand this outcome, it is important to realize that while our results are related to the BM results, there are conceptual differences. BM describe the range of policy-rule coefficients for which the model is learnable, *taking as given the model*. We are describing the range of policy coefficients that maximizes the range of models that are still learnable. So while large values for ϕ_r are beneficial to learnability *holding constant the model and its associated ALM*, at some point, they come at a cost in terms of the perturbations that can be withstood in other dimensions.

Now let us look at the costs and benefits of these two rules in comparison with their optimized counterparts. We measure the benefits by comparing the radii of robustness from the column second from the right, for various rules. For the optimized, outcome-based rules, shown in the first two rows of the table, the radii are about 1.06 or so, while those of their robustified counterparts range from 1.13 to 1.16. Thus the improvement in robustness of learnability would appear to be moderate. Costs are inferred by comparing the losses shown in the right-hand column of the table. The results show that the cost of

maximizing learnability measured in terms of foregone performance in the REE is very small. Evidently, learnability can be robustified, to some degree, without much of any concomitant loss in economic performance, at least in the canonical NKB model.

Before moving on to forecast-based rules, let us consider the classic Taylor rule shown in the fourth row. Recall that the Taylor rule has been advocated as a policy that is at least reasonably robust across a fairly wide range of models. Here, however, the radius associated with the Taylor rule is shown to be quite small at 0.85. At the same time, the performance of the rule in terms of loss is relatively weak. Thus, to the extent that we can take claims of the robustness of the Taylor rule with its original parameterization as applying to the issue of learnability, the rule would appear to come up a bit short.

Now let us examine the results for the forecast-based policy shown in the seventh row. Here the prescribed robust learnable policy is much different from the optimized rule shown in line (3). The robust rule essentially removes the policy persistence that the optimized policy calls for. The policy performance in the rational expectations equilibrium of the forecast-based robustly learnable rule is somewhat worse than its optimized counterpart, but notice that the radius of learnability is nearly triple that of the optimized rule.

While the superiority in terms of robustness of an (almost) non-intertial forecast-based rule is superficially at odds with Bullard and Mitra, the result really should not be all that surprising. Forecast-based rules leverage heavily the rational expectations aspects of the model—even more so than the contemporaneous and lagged data rules since there are rational expectations in the model itself and in the policy rule—and there is risk in leverage. The learnability of the economy is highly susceptible to misspecification in this area. This is, of course, just a manifestation of the problem that Bernanke and Woodford [2] and others have warned about.

We can obtain a deeper understanding of the effects of a concern for robust learnability on policy design by examining the properties of different calibrations of policy rules for their effects on the allowable perturbations. The magnitude of perturbations that a given model can tolerate, conditional on a policy rule, is given by the radius. The radii for the rules shown in Table 1 are in the column second from the right. We can, however, provide a visualization of radii mapped against policy-rule coefficients and judge how policy affects

robust learnability.

Figure 3 provides one such visualization: contour maps of radii against the output-gap feedback coefficient, ϕ_x , and inflation feedback coefficient, ϕ_π , in this case for the contemporaneous data rule. The third dimension of policy, the feedback on the lagged fed funds rate, ϕ_r , is being held constant in these charts, at zero in the upper panel and at unity in the lower. The colors of the chart index the radii of allowable perturbations for each rule, with the bar at the right-hand side showing the tolerance for misspecification. The area in deep blue, for example, represents policies with no tolerance for misspecification of the model or learning whatsoever, either because the rule fails to deliver E-stability in the first place, or because it is very fragile. The sizable region of deep blue in the upper panel shows the area that violates the Taylor principle. The right of the deep blue region—where $\phi_\pi > 1$ —we enter regions of green, where there is modest tolerance for misspecification that allows learnability. In general, with no interest-rate smoothing, there is little scope for misspecification.

Now let us look at the case where $\phi_r = 1$ in the bottom panel. Now the region of deep blue is relegated to the very south-west of the chart, as is the region of green. To the north-east of those are expansive areas of higher tolerance for misspecification. Evidently, at least some measure of persistence in policy is useful for robustifying learnability. Notice how there is a deep burgundy sliver of fairly strong robustness in the north-east part of the panel.

Figure 4 continues the analysis for the contemporaneous data rule by showing contour charts for two more levels of ϕ_r . The upper panel shows the value for the rule that allows the maximum allowable perturbation as shown in line (6) of the table. With $\phi_r = 1.41$ the burgundy region of highest robustness is at its largest and the policy rule shown in line (6) of the table is within that region. More generally, the area of significant robustness—the redder regions—are collectively quite large. Finally, we go to the bottom panel of the figure which shows the results for a relatively high level of ϕ_r . What has happened is that the regions shown in the top panel have rotated down and to the right as ϕ_r has risen. The burgundy region is now gone, and the red regions command much less space. Thus, while policy persistence is good for learnability, in terms of robustness of that result to misspecification, one can go too far.





Figure 3: Contours of radii for the NKB model, contemporaneous data rule, selected ϕ_r





Figure 4: Contours of radii for the NKB model, contemporaneous data rule, selected ϕ_r

Figures 3 and 4 cover the case of the contemporaneous data rule. We turn now to forecast-based rules. The results here look quite different, but the underlying message is very much the same. As before, Figure 5 shows the results for low levels of persistence in policy setting. The upper panel shows the static forecast-based rule. The deep blue areas to the left of $\phi_\pi = 1$ are areas of indeterminacy, as they were in Figure 3. There are, however, numerous blue "potholes" elsewhere in the panel. These are areas where the learnable equilibrium is feasible, but fragile.²³ Notice, however, that these blue regions border very closely to burgundy regions where the allowable perturbations exceed 2; that is, the allowable perturbations are very large. The bottom panel shows contours covering the policy persistence level that is optimal, as shown in line (7) of table 1. There are fewer potholes. The optimally robust policy is toward the top of this chart.

Finally, let us examine Figure 6. The top panel shows that a small increase in ϕ_r from 0.10 to 0.12, reduces the number of potholes to nearly zero. The radii shown in the rest of the chart remain high, but the optimal policy is not in this region.²⁴

The bottom panel of the chart shows the contours for a modest and conventional value of funds rate persistence, $\phi_r = 0.50$. The potholes have now completely disappeared, but the large red region is less robust than the burgundy regions in the previous charts. Not shown in these charts are still higher levels of persistence. These involve still lower levels of robustness, with radii for $\phi_r > 1$ associated with radii that are less than half the magnitude of the maximum allowable perturbation for this rule. Higher levels of persistence in policy setting are deleterious for robustification of model learnability in inflation-forecast based policy rules.²⁵

Of course these particular results are contingent on the relative weightings for pertur-

²³ Since degree of robustness is a function of the model's eigenvalues and those are non-linear functions of the parameters of the model and of the learning mechanism, it is not possible to identify the source of these potholes. That said, it isn't necessary either. The idea behind the methods described in this paper is to avoid the pitfalls of *nonparametric* errors.

²⁴ The presence of the "potholes" in the chart for $\phi_r = 0.10$, wherein the optimally robust rule is found, and their near-absence for the chart $\phi_r = 0.12$ points to another concept of robustness. We assume the monetary authority knows the structural model. As a result, the economy cannot accidentally fall into one of the potholes shown in the figure. A worthwhile extension would be to allow the authority to have doubts about the structural parameters of the model, in addition to the learning mechanism. However the current paper is—the first in this area—is already ambitious enough and so we leave the issue for future research.

²⁵ We tested ϕ_r up to nearly 20. What we found is that the radii fell as ϕ_r rose for intermediate levels, and then rose slowly again for $\phi_r \gg 1$. However, for no level of ϕ_r could we find radii that came anywhere close to the maximum allowable perturbation shown in row (7) of the table.





Figure 5: Contours of radii for the NKB model, forecast-based rule, selected ϕ_r





Figure 6: Contours of radii for the NKB model, forecast-based rule, selected ϕ_r

bations, captured in W_1 , and our selection is just one of many that could have been made. For the numerical experiments, the weightings were set equal to the standard deviations of the coefficients of a first-order VAR for the output gap, inflation, the interest rate, and the natural rate, estimated at the beginning of each experiment using recursive least squares. This approximates the private sector's problem, and should give a rough idea of the *relative* uncertainties associated with the coefficients of the PLM. Whether using estimated standard deviations to scale the relative impact of Knightian model uncertainty on the learning mechanism is proper or desirable can be debated, of course. For now the salient point is that robustness of learning in the presence of model uncertainty is not the same thing as choosing the rule parameters for which the E-stable region of a given model is largest.

4 Concluding remarks

We have argued that model uncertainty is a serious issue in the design of monetary policy. On this score we are in good company. Many authors have advanced that minimizing a loss function subject to a given model presumed to be known with certainty is no longer best practice for monetary authorities. Central bankers must also take model uncertainty and learning into account. Where this paper differs from its predecessors is that we unify uncertainty about the learning mechanism used by private agents with the steps the monetary authority can take to address the problem. In particular, we examine a central bank that designs monetary policy to maximize the possible worlds in which ill-informed private agents need to learn about their particular world and still allow convergence on the rational expectations equilibrium.

The motivation for this approach is straightforward: if economics as a profession cannot agree on what the true model of the economy is, it is a leap of faith to expect private agents to agree, coordinate, and find the REE by themselves. Policy makers can play a role in facilitating (or frustrating) the process of learning the REE through the design of policy. This paper begins from the premise that best practice for a monetary authority using simple instrument rules is to parameterize the rule that provides good performance

not just in the steady state—that is, when convergence to REE has been achieved—but also in out-of-equilibrium behavior of the system. We further argue that foremost among the considerations of what constitutes good out-of-equilibrium behavior of an uncertain system under learning should be the prospects for converging on an REE.

In pursuit of this goal, this paper has married the literature on adaptive learning to that of structured robust control to examine what policy makers can do to facilitate learning. We have introduced some tools with which the questions that Bullard and Mitra [8] are asking can be broadened and generalized.

More narrowly we have also found that the warnings of Bernanke and Woodford [2] are well placed; inflation-forecast-based monetary policy rules do present dangers. We have also shown that the conclusion that Bullard and Mitra [8] point to is not as general as one might initially suppose.

Looking ahead, we see new research directions that broaden the scope of robustness by addressing a wider range of uncertainties against which a policy maker may wish to protect. Evans and McGough [21], for example, show how to compute Taylor-type rules that will converge on an REE in a variety of models, taking the learning mechanism as given, whereas we take the structural model as given and investigate the implications of misspecification of learning rules. A fusion of the two approaches would seem to be worth investigating.

References

- [1] Batini, N. and Pearlman, J. (2002) "Too much too soon: instability and indeterminacy with forward-looking rules" unpublished manuscript, Bank of England.
- [2] Bernanke, B. and Woodford, M. (1997) "Inflation forecasts and monetary policy" *Journal of Money, Credit and Banking* 24: 653-684.
- [3] Blanchard, O. and Kahn, C. (1980) "The solution of linear difference equations under rational expectations" *Econometrica*, 48: 1305-1311.
- [4] Brainard, W. (1967) "Uncertainty and the effectiveness of monetary policy" *American Economic Review* 57(2): 411-425.
- [5] Bray, M. (1982) "Learning, estimation and the stability of rational expectations equilibria" *Journal of Economic Theory*, 26: 318-339.
- [6] Bray, M. and Savin, N. (1986) "Rational expectations equilibria, learning and model specification" *Econometrica*, 54: 1129-1160.
- [7] Bullard, J. and Eusepi, S. (2003) "Did the great inflation occur despite policymaker commitment to a Taylor rule?" Federal Reserve Bank of St. Louis working paper no. 2003-13.
- [8] Bullard, J. and Mitra, K. (2003) "Determinacy, learnability and monetary policy inertia" Federal Reserve Bank of St. Louis working paper 2000-030A (revised version: March 2003)
- [9] Bullard, J. and Mitra, K. (2002), "Learning about monetary policy rules", *Journal of Monetary Economics*, 49: 1105-1139.
- [10] Cagan, P (1956) "The monetary dynamics of hyperinflations" in M. Friedman (ed.) *Studies in the Quantity Theory of Money* (Chicago: University of Chicago Press).
- [11] Clarida, Richard, Jordi Gali, and Mark Gertler (1998) "Monetary Policy Rules in Practice: Some International Evidence" *European Economic Review*, 42: 1033-1067.

- [12] Clarida, Richard, Jordi Gali, and Mark Gertler (1999) "The Science of Monetary Policy: A New Keynesian Perspective" *Journal of Economic Literature*,70: 807-824.
- [13] Coenen, G. (2003) "Inflation persistence and robust monetary policy design" European Central Bank working paper no. 290 (November).
- [14] Dahleh, M and Diaz-Bobillo, I. (1995) *Control of Uncertain Systems: A Linear Programming Approach* (Englewood Hills, NJ: Prentice Hall).
- [15] Doyle, J. "Analysis of feedback systems with structured uncertainties" *IEEE Proceedings*,133 (part D, no. 2): 45-56.
- [16] Ehrmann, M. and Smets, F. (2003) "Uncertain potential output: implications for monetary policy" *Journal of Economic Dynamics and Control*,27: 1611-1638.-
- [17] Evans, G. and Honkapohja, S. (2001) *Learning and Expectations in Macroeconomics* (Princeton: Princeton University Press).
- [18] Evans, G.W. and Honkapohja, S.(2002) "Monetary Policy, Expectations, and Commitment" unpublished manuscript, University of Oregon and Oregon State University (May).
- [19] Evans, G. and Honkapohja, S. (2003) "Expectations and the stability for optimal monetary policies" *Review of Economic Studies*,70: 807-824.
- [20] Evans, G. and McGough, B. (2005a) "Monetary policy, indeterminacy and learning" *Journal of Economic Dynamics & Control*,29: 1809-1840.
- [21] Evans, G. and McGough, B. (2005b) "Optimal constrained monetary policy rules" unpublished manuscript, Department of Economics, University of Oregon http://economics.uoregon.edu/papers/UO-2005-9_Evans_Optimal_Constrained.pdf.
- [22] Garratt, A. and Hall, S.(1997) "E-equilibria and adaptive expectations: output and inflation in the LBS model" *Journal of Economic Dynamics and Control*,21: 87-96.

- [23] Giordani, P. and Soderlind, P. (2004) "Solution of macromodels with Hansen-Sargent robust policies: some extensions" *Journal of Economic Dynamics & Control*,28: 2367-2397.
- [24] Giannoni, M.P. (2002) "Does model uncertainty justify caution?: model uncertainty in a forward-looking model" *Macroeconomic Dynamics*,6(1): 111-144.
- [25] Goodfriend, M. and King, R. (1997) "The new neo-classical synthesis and the role of monetary policy"
- [26] Hansen, L and Sargent, T. (2003) *Misspecification in Recursive Macroeconomic Theory* (unpublished monograph, November 2003)
- [27] Hansen, L.P. and Sargent, T.J. (2003) "Robust control of forward-looking models" *Journal of Monetary Economics* 50(3): 581-604.
- [28] Levin, A.T., Wieland, V. and Williams, J.C. (1999) "Monetary policy rules under model uncertainty" in J.B. Taylor (ed.) *Monetary Policy Rules* (Chicago: University of Chicago Press).
- [29] Levin, A.T., Wieland, V. and Williams, J.C. (1999) "The performance of forecast-based monetary policy rules under model uncertainty" *American Economic Review*,93(2): 622-645.
- [30] Lubik, T.A. and Schorfheide, F. (2003) "Computing sunspot equilibria in linear rational expectations models" *Journal of Economic Dynamics and Control*,28(2): 273-285.
- [31] Lubik, T.A. and Schorfheide, F. (2004) "Testing for indeterminacy: an application to U.S. monetary policy" *American Economic Review*,94(1): 190-217.
- [32] Marcet, A and Nicolini, J.P. (2003) "Recurrent hyperinflations and learning" *American Economic Review*,96:1476-1498.
- [33] Marcet, A. and Sargent, T .J.(1989) "Convergence of least squares learning mechanisms in self-referential linear stochastic models" *Journal of Economic Theory*,48(2): 337-368

- [34] McCallum, B. (1983) "On non-uniqueness in rational expectations models: an attempt at perspective" *Journal of Monetary Economics*,11: 134:168.
- [35] Onatski, A (2003) "Robust monetary policy under model uncertainty: incorporating rational expectations" unpublished manuscript, Columbia University.
- [36] Onatski, A. and Stock, J. (2002) "Robust monetary policy under model uncertainty in a small model of the U.S. economy" *Macroeconomic Dynamics*
- [37] Onatski, A and N. Williams (2003) "Modeling model uncertainty" *Journal of the European Economic Association*,1: 1087-1122.
- [38] Orphanides, A., Porter, R., Reifschneider, D., Tetlow, R., and Finan, F. (2000) "Errors in the measurement of the output gap and the design of monetary policy" *Journal of Economics and Business*,52(1/2): 117-141.
- [39] Rotemberg, J. and Woodford, M. (1997) "An optimization-based econometric framework for the evaluation of monetary policy" in B. Bernanke and J. Rotemberg (eds.) *NBER Macroeconomics Annual* (Cambridge, MA: MIT Press): 297-345.
- [40] Sack, B. (1999) "Does the fed act gradually: a VAR analysis" *Journal of Monetary Economics*,46(1): 229-256.
- [41] Sargent, T. (1999) "Comment" in J.B. Taylor (ed.) *Monetary Policy Rules* (Chicago: University of Chicago Press): 144-154
- [42] Soderstrom, U. (2002) "Monetary policy with uncertain parameters" *Scandinavian Journal of Economics*,104(1): 125-145.
- [43] Taylor, J.B.(1993) "Discretion versus policy rules in practice" *Carnegie-Rochester Conference Series on Public Policy*,39: 195-214.
- [44] Tetlow, R. and von zur Muehlen, P. (2001) "Simplicity versus optimality: the choice of monetary policy rules when agents must learn" *Journal of Economic Dynamics and Control*,25(1/2): 245-279.

- [45] Tetlow, R. and von zur Muehlen, P. (2001) "Robust monetary policy with misspecified models: does model uncertainty always call for attenuated policy?" *Journal of Economic Dynamics and Control* 25(6/7): 911-949.
- [46] Tetlow, R. and von zur Muehlen, P. (2004) "Avoiding Nash Inflation: Bayesian and robust responses to model uncertainty" 7,*Review of Economic Dynamics*, 4 (October): 869-899.
- [47] Walsh, C. E (2004) "Parametric misspecification and robust monetary policy rules," unpublished manuscript, University of California, Santa Cruz.
- [48] Woodford, M.(1999) "Optimal Monetary Policy Inertia" *NBER working paper no. 7261* (July 1999)
- [49] Woodford, M. (2003) *Interest and Prices: Foundations of a theory of monetary policy* (Princeton: Princeton University Press).
- [50] Zames, G. (1966) "On the input-output stability of nonlinear time-varying feedback systems, parts I and II" *IEEE Transactions on Automatic Control*,AC-11: 228, 465.
- [51] Zhou, K., Doyle,J.C., and Glover, K. (1996) *Robust and Optimal Control* (Englewood Cliffs, NJ.: Prentice-Hall).
- [52] Zhou, K., with Doyle,J.C. (1998) *Essentials of Robust Control* (Englewood Cliffs, NJ.: Prentice-Hall).