Macroeconomic Dynamics, **13** (2009), 421–449. Printed in the United States of America.
doi:10.1017/S136510050908016X

ARTICLES

LEARNING, COMMITMENT, AND MONETARY POLICY

THE CASE FOR PARTIAL COMMITMENT

GEORGE A. WATERS
Illinois State University

This paper examines a class of interest rate rules that respond to public expectations and to lagged variables and considers varying levels of commitment that correspond to varying degrees of response to lagged output. Under commitment the policymaker adjusts the nominal rate with lagged output to impact public expectations. Within this class of rules, I provide a condition for nonexplosive and determinate solutions. Expectational stability obtains for any nonnegative response to lagged output. Simulation results show that modified commitment is best under least-squares learning. However, in the presence of parameter uncertainty and/or measurement error in the policymaker's data on public expectations, the best policy is one of partial commitment, where the response to lagged output is less than under modified commitment. The case for partial commitment is strengthened if the gain parameter in the learning mechanism is high, which can be interpreted as the use of few lags by public agents in the formation of expectations or as an indication of low credibility of the policymaker. The appointment of a conservative central banker ameliorates these concerns about modified commitment.

Keywords: Learning, Monetary Policy, Interest Rate Rules, Determinacy, Expectational Stability

1. INTRODUCTION

The gains from commitment for a monetary policymaker have been established both theoretically [Barro and Gordon (1983) and Clarida et al. (1999)] and empirically [McCallum and Nelson (2004)][1]. Another strand of work on monetary policy [Howitt (1992) and Bullard and Mitra (2002)][2] advocates the view that monetary policy rules should be learnable, meaning they should lead to dynamically stable systems when public expectations deviate from full rationality. Setting policy under commitment implies that the policymaker acts to impact public expectations, so the way the public forms expectations could be critical to the outcome of such

The author thanks the editor, two anonymous referees, and participants at Learning Week at the Federal Reserve Bank of St. Louis for helpful comments. All mistakes are my own. Address correspondence to: George A. Waters, Department of Economics, Campus Box 4200, Illinois State University, Normal, IL 61790-4200, USA; e-mail: gawater@ilstu.edu.

© 2009 Cambridge University Press 1365-1005/09

https://doi.org/10.1017/S136510050908016X Published online by Cambridge University Press

an approach. The present work addresses the question of whether there are gains from commitment under learning.

Interest rate rules that respond explicitly to public expectations, studied by Evans and Honkapohja (2003, 2006), have the desirable properties of determinacy and expectational stability under both commitment and discretionary policies[3]. Within this class, the optimal policy rule under rational expectations is identical in both cases with one exception. Under commitment, policy responds to lagged values of output while under discretion it does not. As Woodford (1999b) emphasizes, if public agents are forward-looking, commitment requires that policy be *history-dependent*. The more the policymaker wants to impact public expectations, the greater the response to lagged variables.

Since the existence of gains to commitment under learning is an open question, I study a continuum of responses to lagged output that includes the optimal rational expectations discretionary and commitment policies as special cases. There are other reasons to consider such a range of policies. First, if the policymaker has biased estimates of the model parameters, he or she may overreact or underreact to lagged output under commitment. I also consider the possibility that the policymaker's knowledge of public expectations is imprecise. For the expanded set of rules, I provide a condition under which there are non-explosive and determinate rational expectations equilibria. These equilibria are expectationally stable for any nonnegative response to lagged output.

Simulations of the model under learning allow a comparison of policy outcomes for various specifications of policy and the learning mechanism. Public expectations are formed with recursive least squares using constant gain. The gain parameter is a key component, indicating agents' sensitivity to new information, and has natural interpretations in terms of the number of lags used in the estimation, the rationality of the agents, and the credibility of the policymaker.

One can regard any policy response to lagged output as a type of commitment in that it is an attempt to influence expectations. The commitment optimum under rational expectations, as derived[4] in Evans and Honkapohja (2006), is referred to as *full commitment*. Blake (2001) argues for a particular[5] intermediate response to lagged output called *modified commitment* for the present study. The results from the simulations demonstrate that, under least-squares learning, the policy outcome under modified commitment minimizes the policymaker's loss in comparison to full commitment, discretion, and all other settings for the response to lagged output. However, the results also indicate that if there is uncertainty about the parameters of the model, the policymaker should not respond to lagged output as strongly as the full or modified commitment settings suggest. I refer to such a policy prescription as *partial commitment*.

Implementing expectations-based interest rate rules raises questions about the policymaker's knowledge of public expectations. Evans and Honkapohja (2006) note that the introduction of white noise errors into the policymaker's estimates of public expectations does not alter their results concerning expectational stability.

https://doi.org/10.1017/S136510050908016X Published online by Cambridge University Press

However, whether such errors affect the policy outcomes under learning remains an open question.

The results presented here show that the introduction of such errors substantiate the case for partial commitment. In the presence of errors in the policy rule, the response to lagged output should be weaker than the full or modified commitment recommendations. Moreover, for higher levels of the gain parameter, the optimal response to lagged output falls further. If agents are placing heavy emphasis on recent information, possibly due to a lack of credibility of the policymaker, partial commitment is the best approach under learning.

The results are verified for different parameter values found in the literature. The model studied here is the standard New Keynesian approach where the policymaker faces a trade-off between stabilizing output and inflation, described in Clarida et al. (1999). Results are very similar for different values of the model parameters, but they are sensitive to changes in the policymaker's preference parameter measuring his or her relative interest in these two goals.

A significant factor underlying the optimality of partial commitment is the asymmetry of the losses across the magnitude of the response to lagged output, in that an excessive response can lead to very poor outcomes from the policymaker's perspective. This paper provides support for a conservative central banker (Rogoff 1985) since a lower emphasis on output stabilization diminishes the impact of this asymmetry and increases the range for nonexplosive and determinate policies.

The learning rule in the simulated model uses constant gain, meaning the emphasis agents place on new information does not change over time. Therefore, learning has a significant impact on the dynamics of the model at all times. Orphanides and Williams (2002, 2006) study an alternative model of monetary policy under constant-gain learning and show that the nature of optimal policy changes when expectations deviate from full rationality. Constant gain[6] is an important aspect of the dynamics in Bullard and Cho (2005), who show how the economy can shift to a deflationary equilibrium when agents continually update their estimate of the policymaker's inflation target.

The approach to expectation formation employed here follows a large literature on learning and convergence to rational expectations that relaxes the extreme informational requirements on agents required by rational expectations. Bray (1982) is one of the earliest efforts, while Marcet and Sargent (1989) give convergence criteria for least squares learning. Evans and Honkapohja (2001) provide an overview of this literature and define expectational stability for a wide variety of learning mechanisms and models. Howitt (1992) contains one of the earliest suggestions that monetary policy rules should be stable under learning. More recently, Bullard and Mitra (2002) and Honkapohja and Mitra (2004) have studied the expectational stability of various interest rate rules for monetary policy.

The paper is organized as follows. Section 2 describes the model and derives conditions for optimal policy for varying levels of commitment. Section 3 shows the computation of the rational expectations equilibrium, and Section 4 introduces the expectations based interest rate rule and demonstrates the condition for

https://doi.org/10.1017/S136510050908016X Published online by Cambridge University Press

non-explosiveness and determinacy. Section 5 describes the learning mechanism and the proposition concerning expectational stability. Section 6 reports the initial simulation results, whereas Section 7 discusses those for the model with errors in the policy rule and examines the simulation results for alternative parameters values. Section 8 concludes.

2. THE MODEL

The following New Keynesian model has become standard for the study of interest rate rules in monetary policy. It has micro foundations described in Woodford (2003) including price stickiness that allows the policymaker to play a stabilizing role in the economy:

$$x_t = -\varphi(i_t - E_t^*\pi_{t+1}) + E_t^* x_{t+1} + g_t \tag{1}$$

$$\pi_t = \lambda x_t + \beta E_t^* \pi_{t+1} + u_t. \tag{2}$$

These equations are expectations-augmented IS and Phillip's curve relationships. The variables x_t and π_t are the deviations of output and inflation from their target values. The notation E_t^* indicates private sector expectations formed in time t where the $(*)$ is used to show that expectations might not be rational. The policymaker controls the nominal interest rate i_t. The parameters φ, λ, and β are all positive and the discount rate $\beta < 1$. The stochastic terms g_t and u_t both have autoregressive structure:

$$\begin{pmatrix} g_t \\ u_t \end{pmatrix} = F \begin{pmatrix} \widehat{g}_{t-1} \\ u_{t-1} \end{pmatrix} + \begin{pmatrix} \widetilde{g}_t + \widetilde{w}_t \\ \widetilde{u}_t \end{pmatrix} \quad \text{for} \quad F = \begin{pmatrix} \mu & 0 \\ 0 & \rho \end{pmatrix}. \tag{3}$$

The parameters μ and ρ lie in the interval $(-1, 1)$, and the shocks $\widetilde{u}_t$, $\widetilde{w}_t$, and $\widetilde{g}_t$ and are iid with standard deviations σ_u, σ_w and σ_g, respectively. The structure of g_t follows McCallum and Nelson (2004), who decompose this term into a preference shock $\widetilde{g}_t$ and an AR(1) process $\widehat{g}_t = \mu\widehat{g}_{t-1} + \widetilde{w}_t$ that accounts for uncertainty about the evolution of the natural rate that enters into the forecast error for output in the IS equation (1) and represents the portion of the demand shocks about which the public has some information to use for forecasting.

The policymaker sets the nominal interest rate to stabilize the endogenous variables. Formally the task is to set i_t to minimize the loss function

$$L = E_t \sum_{s=0}^{\infty} \beta^s \left(\pi_{t+s}^2 + \alpha x_{t+s}^2\right), \tag{4}$$

assuming rational expectations for the following derivation. The parameter α measures the relative importance of output and inflation stabilization for the policymaker, the value $\alpha = 0$ corresponding to pure inflation targeting. Minimizing the loss function[7] over x_{t+s} and π_{t+s} under the constraint of the Phillip's curve (2)

https://doi.org/10.1017/S136510050908016X Published online by Cambridge University Press

yields the following first-order conditions:

$$E_t(2\alpha x_{t+s} + \lambda\omega_{t+s}) = 0 \quad \text{for} \quad s = 0, 1, 2, 3 \ldots \tag{5}$$

$$E_t(2\pi_{t+s} + \omega_{t+s-1} - \omega_{t+s}) = 0 \quad \text{for} \quad s = 1, 2, 3 \ldots \tag{6}$$

$$2\pi_t - \omega_t = 0. \tag{7}$$

The variable ω_{t+s} is the Lagrange multiplier on the constraint for each $s = 0, 1, 2, 3, \ldots$. Optimal policy in time t is governed by (7) but policy in succeeding periods is determined by condition (6). The time consistency problem is evident because, when the policymaker solves the problem in the period $t+1$, the policy given by (7) will be different from the policy prescribed by (6) in the period t solution.

To find the condition for policy under discretion, the policymaker adopts the optimal policy each period, solving out ω_t for $s = 0$ from the first-order conditions (5) and (7):

$$\lambda\pi_t + \alpha x_t = 0. \tag{8}$$

Under discretion, the policymaker takes expectations to be fixed and ignores (6); see Clarida et al. (1999) for a detailed discussion. However, if the policymaker has a "timeless perspective" (Woodford 1999a), he or she uses the policy that would have been prescribed in past periods, ignoring condition (7), using (5) and (6) for $s = 0$. In this case, which I refer to as *full commitment*, the policymaker acts to influence future expectations, yielding the condition

$$\lambda\pi_t + \alpha\,(x_t - x_{t-1}) = 0. \tag{9}$$

Policy under commitment differs from discretion in its response to the previous period's output gap. The reason can be seen in (6), where the lagged term appears because the policymaker reacts to the expected inflation in the Phillip's curve (2). Under discretion, however, the policymaker ignores the change in expected inflation, and lagged output does not enter the policy condition (8).

There are several reasons to question full commitment, meaning the policymaker should not be restricted to these extreme cases. First, some agents may not believe the policymaker will commit to a time inconsistent policy. A related issue is that gains to commitment depend on a forward looking public, so if the public is using a learning mechanism to form expectations, the policymaker may not want to fully commit. Waters (2006) shows that discretion is superior to commitment under adaptive expectations, when the public forecast of a variable is simply a weighted of the previous period's forecast and realized value. Although the least-squares mechanism studied in the present paper is a more sophisticated method of forming expectations, it remains an open question whether full commitment is optimal. There are other potential problems with commitment if the policymaker lacks full information. If estimates of model parameters or public expectations were biased, the policymaker would over- or underreact to lagged variables under commitment.

https://doi.org/10.1017/S136510050908016X Published online by Cambridge University Press

Furthermore, Blake (2001) argues that the first-order condition (6) improperly assumes certainty equivalence, and he derives an optimal policy that differs from full commitment. Using a version of the loss function (4) without discounting over a finite number of periods, he gives an intuitive argument[8] for the policymaker to respond to lagged output, but not to the degree indicated by the full commitment condition (9). This approach captures the spirit of the timeless perspective, because including discounting overemphasizes the loss in initial periods. For these reasons, we consider the general condition

$$\lambda\pi_t + \alpha\,(x_t - \kappa x_{t-1}) = 0. \tag{10}$$

The condition under discretion is a special case of (10) where $\kappa = 0$, and full commitment is equivalent to setting $\kappa = 1$.

The effect of Blake's (2001) approach without discounting is that the second term in (6) is now $\beta\omega_{t+s-1}$. In this case, optimal policy corresponds to the condition (10), where κ equals the discount factor $\beta < 1$, which I call *modified commitment*. Woodford (2003) argues in favor of full commitment by using a slightly criterion different from the loss function (4) for evaluating policy. Under rational expectations the optimal level of κ is either full commitment at $\kappa = 1$ or modified commitment at $\kappa = \beta$ depending on the approach to the derivation. I refer to any setting for κ such that $0 < \kappa < \beta$, when the policymaker's response to lagged inflation is less than for modified commitment, as *partial commitment*.

3. RE EQUILIBRIA

One can now solve for the unique minimum state variables[9] rational expectations equilibrium (REE) for a given κ under the general condition (10). Postulating solutions of the form

$$\begin{aligned} x_t &= b_x x_{t-1} + c_x u_t \\ \pi_t &= b_\pi x_{t-1} + c_\pi u_t \end{aligned} \tag{11}$$

implies that $E_t\pi_{t+1} = b_\pi(b_x x_{t-1} + c_x u_t) + c_\pi \rho u_t$. Using the method of undetermined coefficients with (10) and (2) yields

$$\begin{aligned} b_\pi &= \lambda b_x + \beta b_\pi b_x \\ \lambda b_\pi &= \alpha(\kappa - b_x) \\ c_\pi &= \lambda c_x + \beta\,(b_\pi c_x + c_\pi \rho) + 1 \\ \lambda c_\pi &= -\alpha c_x. \end{aligned} \tag{12}$$

Combining the first two equations from (12) gives us

$$\beta b_x^2 - \left(\frac{\lambda^2}{\alpha} + \beta\kappa + 1\right) b_x + \kappa = 0. \tag{13}$$

https://doi.org/10.1017/S136510050908016X Published online by Cambridge University Press

The larger root of the quadratic equation (13) is always such that $b_x > 1$, so the smaller root is the potentially nonexplosive solution for b_x in (11) and is the focus throughout the paper. Given b_x, the solution for b_π is determined by the first and third equations in (12). Note that the discretionary value $\kappa = 0$ admits the solution $b_x, b_\pi = 0$, which corresponds to the fact that the minimum state variables solution in this case depends only on the supply shock u_t, as in Evans and Honkapohja (2003). Further, for $\kappa > 0$, the relevant solutions for b_x and b_π are real and positive (see Appendix A.2). Hence, public expectations respond to lagged output only when the policymaker does, demonstrating the policymaker's influence on expectations under commitment. The solutions for the coefficients under full commitment for $\kappa = 1$ correspond to those in Evans and Honkapohja (2006). Similarly, the third and fourth equations from (12) may be combined to give

$$c_x = \lambda(\alpha\beta\rho - \alpha - \lambda^2 - \beta\lambda b_\pi)^{-1}$$

and then the fourth equation in (12) determines c_π, completing the solution.

4. EXPECTATIONS-BASED INTEREST RATE RULES

Evans and Honkapohja (2003, 2006) advocate for monetary policy to be conducted with interest rate rules that respond explicitly to public expectations. One can compute the optimal form under rational expectations for the expectations-based interest rate rule associated with the policy condition (10). This section also provides a condition for the nonexplosive and determinate solutions, whereas the following section on learning addresses expectational stability.

Using (10) to substitute for inflation in the Phillip's curve (2) gives

$$x_t = \lambda(\lambda^2 + \alpha)^{-1}(\alpha\kappa\lambda^{-1}x_{t-1} - \beta E_t\pi_{t+1} - u_t).$$

Substituting out x_t in the IS equation with the above expression yields the interest rate rule

$$i_t = \delta_L x_{t-1} + \delta_\pi E_t^*\pi_{t+1} + \delta_x E_t^* x_{t+1} + \delta_g g_t + \delta_u u_t, \tag{14}$$

where the parameters are

$$\delta_L = -\varphi^{-1}(\lambda^2 + \alpha)^{-1}\alpha\kappa, \quad \delta_\pi = 1 + \varphi^{-1}(\lambda^2 + \alpha)^{-1}\lambda\beta,$$
$$\delta_x = \varphi^{-1}, \quad \delta_g = \varphi^{-1}, \quad \delta_u = \varphi^{-1}(\lambda^2 + \alpha)^{-1}\lambda.$$

Note that the extent to which the interest rate responds to lagged output, shown by δ_L, depends on κ, but the other parameters in the rule do not. Under discretion, κ is zero and i_t is unaffected by x_{t-1}, but for any other value, including the full, modified, and partial commitment settings, the interest rate responds directly to x_{t-1}.

The interest rate rule (14) demonstrates the potential impact of parameter uncertainty on the conduct of policy. For example, if the policymaker's estimate

https://doi.org/10.1017/S136510050908016X Published online by Cambridge University Press

φ_p is less than the actual φ, the magnitude of the response of i_t to x_{t-1} will be too large, which is equivalent to setting κ higher than intended, and simulation results indicate that such a policy could lead to a large loss.[10] For this reason, it is important to study the behavior of the model when $\kappa > 1$, even though the policymaker would never knowingly make such a choice. To guard against such a possibility, the policymaker may not fully commit and may choose a smaller response to lagged output than would be indicated by (9).

The next tasks are to determine conditions under which interest rate rules of the form (14) are nonexplosive, determinate, and expectationally stable. A determinate equilibrium is unique in that there are no other similar equilibria based on extraneous variables. Expectational stability implies that the expectations of agents using a reasonable learning rule converge to the REE. All of these features of interest rate rules are desirable for a policymaker trying to stabilize endogenous variables.

For analysis of determinacy and expectational stability and for simulations, it is convenient to write the New Keynesian model with (1), (2) and (14) in matrix form:

$$\begin{pmatrix} x_t \\ \pi_t \end{pmatrix} = \begin{pmatrix} (\lambda^2+\alpha)^{-1}\alpha\kappa & 0 \\ (\lambda^2+\alpha)^{-1}\lambda\alpha\kappa & 0 \end{pmatrix} \begin{pmatrix} x_{t-1} \\ \pi_{t-1} \end{pmatrix} + \begin{pmatrix} 0 & -(\lambda^2+\alpha)^{-1}\beta\lambda \\ 0 & (\lambda^2+\alpha)^{-1}\beta\alpha \end{pmatrix} \begin{pmatrix} E_t^* x_{t+1} \\ E_t^* \pi_{t+1} \end{pmatrix}$$
$$+ \begin{pmatrix} 0 & -(\lambda^2+\alpha)^{-1}\lambda \\ 0 & (\lambda^2+\alpha)^{-1}\alpha \end{pmatrix} \begin{pmatrix} g_t \\ u_t \end{pmatrix}. \quad \textbf{(15)}$$

From the above, one can see that the policymaker acts so that the shock to the IS equation g_t does not affect output and inflation, but the supply shock u_t does. Note that if $\alpha = 0$, meaning the policymaker targets only inflation, the supply shock does not affect inflation. These observations reflect the trade-off faced by the policymaker between output and inflation stabilization that arises from the presence of the supply shock.

The next proposition gives a single condition on κ guaranteeing nonexplosiveness and determinacy.

PROPOSITION 1. *Under rational expectations, for κ such that*

$$0 \leq \kappa < 1 + \frac{\lambda^2}{\alpha(1-\beta)},$$

the model defined by the policy condition (10) and the Phillip's curve (2) has a nonexplosive solution such that $b_x < 1$, and the model given by (15) is determinate.

Proof. See Appendix A.1. ■

The condition for nonexplosiveness applies to the matrix form (15), as the policy rule (14) is derived from the general condition (10). The discretionary and full commitment cases of $\kappa = 0$ and 1 yield a determinate model, as shown

https://doi.org/10.1017/S136510050908016X Published online by Cambridge University Press

in Evans and Honkapohja (2003, 2006), as do any intermediate values of κ, including modified and partial commitment, as well as some values of $\kappa > 1$. In the absence of extraneous variables in the simulated model, indeterminacy should not be a crucial factor in the policy outcomes for the simulations reported here, but given the number of candidate variables (stock market indices, real estate market conditions) that agents might consider for use in forming expectations in practice, determinacy remains a desirable property for a policy rule.

One might argue that values of κ such that $\kappa > 1$ are not interesting, as there is no justification for such policy. However, as noted above, the policymaker with mistaken estimates of the model parameters could easily overreact to lagged output. This study reports simulation results using parameters from three papers in the literature, McCallum and Nelson (2004), Clarida et al. (2000), and Woodford (1999b), referred to as MN04, CGG00, and W99, respectively. The following table gives their estimates for the parameters in equations (1) and (2):

	β	φ	λ
MN04	0.99	0.4	0.05
CGG00	0.99	4.0	0.075
W99	0.99	$(0.157)^{-1}$	0.024

(**16**)

There is quite a difference of opinion regarding φ and λ that could affect the response of the interest rate to lagged output, represented by δ_L in (14). If the policymaker fully commits ($\kappa = 1.0$), places relatively low emphasis on output stabilization with $\alpha = 0.1$, and takes the parameter values from MN04 to be correct, the value for δ_L is $\delta_{L,MN} = -2.439$, but if the true values are those in CGG00, this parameter should be $\delta_{L,CGG} \approx -0.237$. Here, the policymaker would mistakenly set policy corresponding to a value of $\kappa \approx 10.3$, which exceeds the bound of 5.625 for determinacy and nonexplosiveness given in Proposition 1. So parameter uncertainty could lead the policymaker to set policy in an explosive and indeterminate range. Evans and McGough (2005) make a similar argument in the context of Taylor rules in the New Keynesian model.

Proposition 1 provides an additional rationale for the appointment of a conservative central banker meaning one with a low weight on output stabilization α. Rogoff (1985) shows a conservative central banker can diminish the problem of inflation bias[11] in models where the policymaker wants to raise output above the natural rate. In the present framework, the policymaker cares only about deviations from the natural rate, but a low α may still be desirable because it increases the bound for determinacy and nonexplosiveness, making such issues less problematic.

https://doi.org/10.1017/S136510050908016X Published online by Cambridge University Press

5. LEARNING

The next task is to specify the learning mechanism agents use to form expectations. As is common in the learning literature, I use recursive least squares.[12] This mechanism has a natural relationship with expectational stability and can be used to simulate the model, which also allows for a comparison of policy outcomes for varying parameter values. Further goals are to find a condition for expectational stability and to give interpretations for the gain, a key parameter in the learning mechanism.

Assume agents update expectations using a model, called the perceived law of motion (PLM), that corresponds in structure to the minimum state variables solutions (11). Let the vectors y_t and v_t be such that $y_t = (x_t, \pi_t)'$ and $v_t = (g_t, u_t)'$. Now, the PLM can be written as

$$y_t = a + by_{t-1} + cv_t. \tag{17}$$

The REE corresponds to the case where

$$a = \begin{pmatrix} 0 \\ 0 \end{pmatrix}, \quad b = \begin{pmatrix} b_x & 0 \\ b_\pi & 0 \end{pmatrix}, \quad \text{and} \quad c = \begin{pmatrix} 0 & c_x \\ 0 & c_\pi \end{pmatrix},$$

but, under learning, agents do not know these values and must update their estimate of $\xi_t = (a_t, b_t, c_t)$ over time. The PLM (17) with time subscripts for a_t, b_t, and c_t denotes' agents beliefs at time t about parameter values and implies that $E_t^* y_{t+1} = a_t + b_t E_t^* y_t + c_t E_t^* v_{t+1}$, or

$$E_t^* y_{t+1} = a_t + b_t (a_t + b_t y_{t-1} + c_t v_t) + c_t F v_t. \tag{18}$$

This equation shows how agents use past information about the endogenous variables, estimates about parameter values in the PLM, and current values of the shocks to form expectations.

Given such a description of expectations formation and the model (15), one can describe the actual law of motion (ALM) governing the endogenous variables. Agents update the parameters to find ξ_t and use the new shocks v_t to form expectations $E_t^* y_{t+1}$ with (18), which then determines the contemporaneous value of y_t via the model (15).[13] One key question is whether the parameters ξ_t will converge over time to their values in the REE, given by (12).

Evans and Honkapohja (2001) define expectational stability, which implies the convergence of many learning mechanisms including least-squares learning (see Appendix A.2). For the present model, one can show the following.

PROPOSITION 2. *The REE given by* (11, 12) *for model* (15) *under the PLM* (17) *is expectationally stable for any* $\kappa \geq 0$.

Proof. See Appendix A.2. ■

Expectational stability obtains for any level of response to lagged output, providing further support for interest rate rules of the form (14) and extending the

https://doi.org/10.1017/S136510050908016X Published online by Cambridge University Press

expectational stability results of Evans and Honkapohja (2003, 2006). The condition in Proposition 1 thereby guarantees nonexplosiveness, determinacy, and expectational stability.

To formally define the least-squares learning mechanism, let z_t be such that

$$z_t = (1, y'_{t-1}, v'_t).$$

The recursive least-squares algorithm[14] is given by

$$\begin{aligned} R_t &= R_{t-1} + \tau_t(z_{t-1}z'_{t-1} - R_{t-1}) \\ \xi_t &= \xi_{t-1} + \tau_t R_t z_{t-1}\,(y_{t-1} - \xi_{t-1}z_{t-1})', \end{aligned} \tag{19}$$

where R_t acts as an updated covariance matrix. The realization for ξ_t then determines the expectations in (15), which then determine x_t and π_t. So the dynamics of the model are defined by the recursive least-squares updating equation (19), the expectations (18) derived from the PLM, the model (15), and the stochastic shocks.

A key parameter in the learning mechanism is the gain parameter τ_t, which indicates the emphasis that agents place on recent information. To form an expectation, agents update their forecasting model (the PLM) from the previous period using the new realizations of the endogenous variables. A high gain indicates that agents' estimates of (a_t, b_t, c_t) are more sensitive to recent information.

If the gain parameter is decreasing over time so that $\tau_t = 1/t$, then the updating equations (19) are equivalent to recursive least squares using all lags, given an appropriate initial condition. The expectational stability result in Proposition 2 implies that recursive least squares will converge to the REE from (11) and (12), given that the PLM (17) is of the same form as the minimum state variables solutions (11) (see Evans and Honkapohja (2001), Sec. 2.6 and Ch. 10).

To study the impact of learning on policy, decreasing gain is not an appropriate assumption, because the impact of learning on the dynamics is decreasing over time. Following much of the learning literature, I adopt a constant gain $\tau_t = \bar{\tau}$, the value of which has a number of interrelated interpretations. Orphanides and Williams (2002) explain that the learning mechanism (19) with constant gain updates expectations using a rolling window of past data. A lower gain parameter means that agents are using more lags to make their estimates.[15] Furthermore, because least-squares learning with decreasing gain approaches the rational expectation equilibrium, one can think of the gain parameter as a measure of the rationality of agents, lower gain indicating a higher degree of rationality. A third interpretation can be found in Waters (2007), who argues that low gain implies high credibility of the policymaker. Given that the policymaker is trying to keep output and inflation near certain targets, if agents have faith in the abilities of the policymaker, they will not be excessively swayed by recent data and will tend to keep their beliefs about the target values that they have formed over time. Conversely, if agents do not trust the policymaker to keep the economy near the

https://doi.org/10.1017/S136510050908016X Published online by Cambridge University Press

targets or even to keep fixed targets, they will place greater emphasis on recent realizations when forming expectations, corresponding to high gain.

6. SIMULATIONS

Simulating the New Keynesian model with expectations-based interest rate rules under constant-gain learning allows a comparison of policy outcomes for varying degrees of commitment. The primary goals are to determine whether there are gains to commitment and to find the optimal magnitude of response to lagged output for the interest rate, which is parameterized by κ in the policy rule (14). Because the full commitment value of $\kappa = 1.0$ is close to the bound for nonexplosive and determinate solutions for some reasonable parameter values, it is questionable whether this policy remains optimal under learning. The simulation results show that the modified commitment setting $\kappa = \beta$ minimizes the loss under least-squares learning with constant gain. However, this and the following section show that consideration of parameter uncertainty and the introduction of measurement error in the policymaker's observations of public expectations complicates these results.

The choice of parameter values for this section follows McCallum and Nelson (2004), MN04 in (16), who study full commitment versus discretionary outcomes for a number of different policy rules under full rationality. They report gains from commitment over discretion for the policymaker for a number of policy rules under different degrees of output stabilization. The parameters are calibrated so that the time periods correspond to quarterly data. The parameter values that determine the behavior of the stochastic variables g_t and u_t in (1) and (2) are as follows:

μ	ρ	σ_w	σ_g	σ_u
0.95	0.8	0.005	0.02	0.005

Further, let α in the loss function (4) be $\alpha = 0.25$, which is an intermediate value of those studied by McCallum and Nelson (2004) and indicates equal weight on output and inflation stabilization. I report results for values of the gain parameters from $\bar{\tau} = 0.01$ to $\bar{\tau} = 0.15$, which correspond to agents using a rolling window of 200 and $3.\bar{3}$ years of past data, respectively. The smaller value 0.01 is the lowest considered by Orphanides and Williams (2006), whereas $\bar{\tau} = 0.15$ is much larger than the values they consider, but in line with the values used in other studies.[16]

Figures 1 and 2 give sample paths for the output and inflation gaps for simulations under recursive least squares with $\bar{\tau} = 0.1$. Agents begin the simulations with beliefs about the parameter values corresponding to the REE, but with constant gain their estimates may continue to fluctuate. For this simulation in Figure 1, the parameter κ is set to $\kappa = 1.3$, which is in the range where the REE is determinate and nonexplosive. The RE value for $b_x \approx 0.98$ is close to one, leading to the persistence in the series in Figure 1. In contrast, for Figure 2, $\kappa = 0.9$, the

https://doi.org/10.1017/S136510050908016X Published online by Cambridge University Press

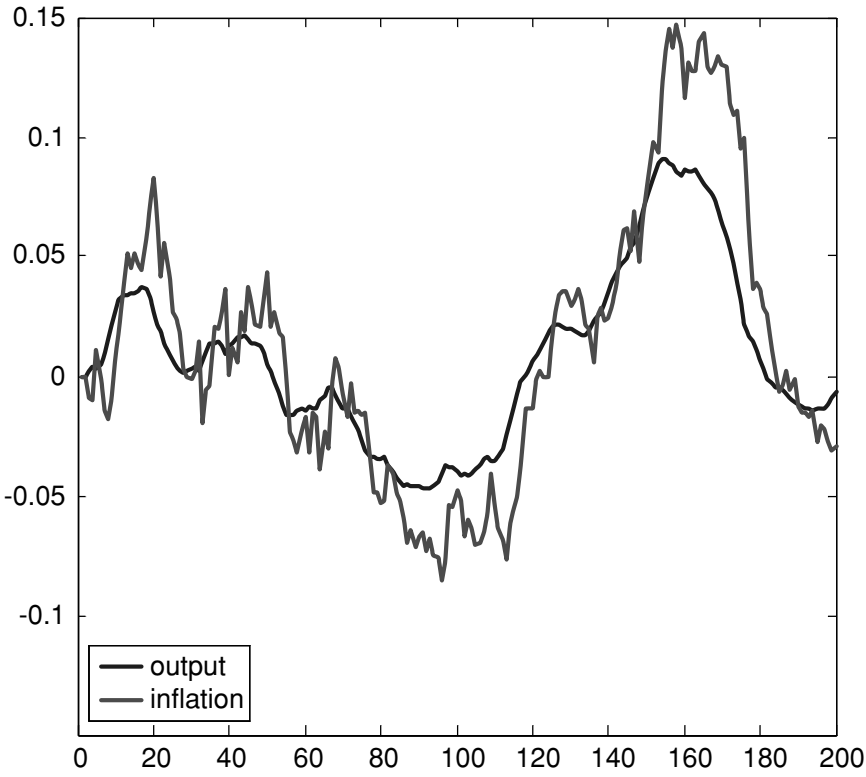


FIGURE 1. The value of κ is $\kappa = 1.3$, which is in the range for determinate and nonexplosive equilibria. The value of the gain is $\bar{\tau} = 0.1$.

corresponding value is $b_x \approx 0.85$, and the policymaker is able to keep output and inflation closer to their target values.

To gain intuition about the problem facing the policymaker, Figure 3 graphs the losses for $\kappa \in [0.0, 1.1]$ and $\alpha \in [0.1, 1.0]$, taking, an average of the losses from 10,000 runs.[17] A number of observations can be made from Figure 3. First, the loss is minimized in the neighborhood of $\kappa = 1.0$ for all levels of α, though there is an asymmetric change in the loss for κ above and below the minimizing value. Here is evidence in favor of partial commitment if there is uncertainty about the true parameter values, because accidentally using a policy equivalent to an excessively large κ could lead to a very bad outcome. It is notable, however, that the asymmetry is less severe for lower levels of α, in line with Proposition 1, which shows that a lower alpha raises the bound on κ for determinacy and nonexplosiveness. Again, a conservative central banker reduces the risks associated with an excessive response to lagged output. Preliminary simulations indicate that this asymmetry is present under rational expectations as well, though I leave a full investigation of this issue for future work.

https://doi.org/10.1017/S136510050908016X Published online by Cambridge University Press

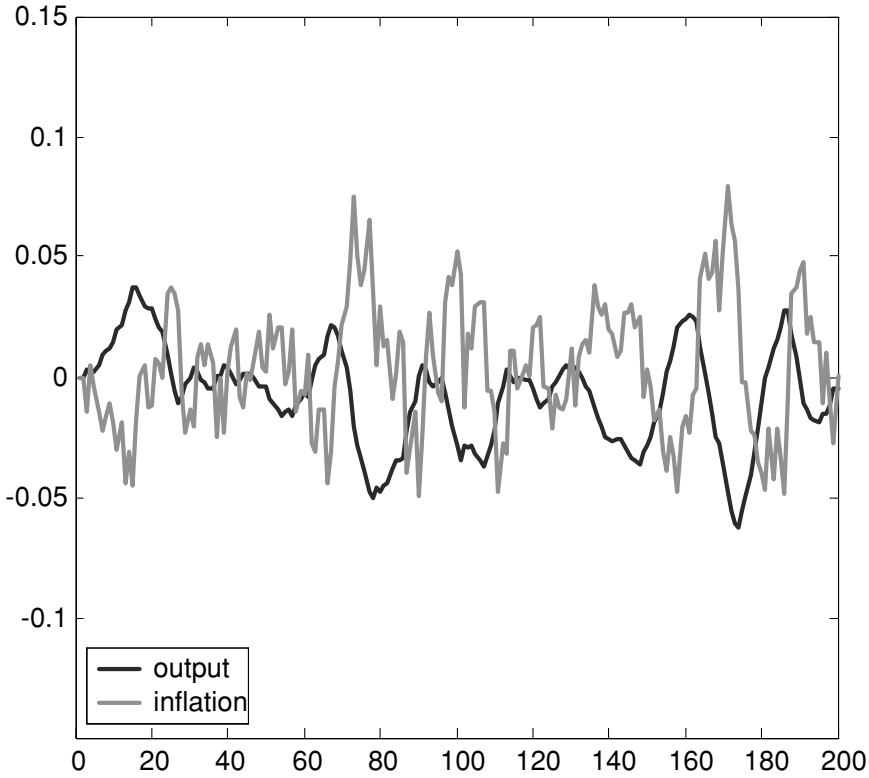


FIGURE 2. The value of κ is $\kappa = 0.9$, which is in the range for determinate and nonexplosive equilibria. The value of the gain is $\bar{\tau} = 0.1$.

To give some precision to these observations, Table 1 gives numerical values of losses for different values of κ, including those for discretion, partial commitment, and full commitment, with $\alpha = 0.25$. For this section the reader should focus on the first column of losses. Using the parameters of MN04, as in Figure 3, the minimum loss is achieved at the partial commitment setting of $\kappa = 0.99$. Hence, modified commitment is optimal under least-squares learning.

7. IMPERFECT INFORMATION FOR THE POLICYMAKER

To this point I have made the assumption that the policymaker has perfect knowledge of public expectations. Evans and Honkapohja (2006) state that the results concerning expectational stability hold with white noise measurement error in the expectations in the policy rule (14). Here, I study whether the introduction of such errors would impact the policy outcomes under learning. Simulation results reported in this section show that such errors can have a definite effect. In the

https://doi.org/10.1017/S136510050908016X Published online by Cambridge University Press

TABLE 1. Each entry shows the mean loss over 10,000 runs for varying κ across different parameterizations of the model. The cells in bold are the minimum losses for each column.

	MN04		CGG00		W99	
κ		$\sigma_x = 0.01$		$\sigma_x = 0.01$		$\sigma_x = 0.01$
0.0	0.12792	0.13042	0.11547	0.11809	0.13677	0.13857
0.90	0.08064	0.08950	0.05368	0.06729	0.11492	0.12571
0.91	0.07920	0.08950	0.05776	0.06686	0.11343	0.12525
0.92	0.07780	0.08887	0.05692	0.06653	0.11182	0.12490
0.93	0.07648	0.08847	0.05615	0.06633	0.11013	**0.12474**
0.94	0.07525	**0.08828**	0.05546	**0.06627**	0.10836	0.12487
0.95	0.07416	0.08836	0.05488	0.06635	0.10657	0.12547
0.96	0.07324	0.08876	0.05440	0.06661	0.10485	0.12677
0.97	0.07252	0.08955	0.05405	0.06704	0.10332	0.12913
0.98	0.07204	0.09078	0.05382	0.06768	0.10218	0.13306
0.99	**0.07186**	0.09254	**0.05374**	0.06854	**0.10171**	0.13924
1.00	0.07202	0.09490	0.05382	0.06963	0.10231	0.14863
1.01	0.07257	0.09794	0.05407	0.07098	0.10449	0.16251
1.02	0.07357	0.10176	0.05449	0.07259	0.10890	0.18256

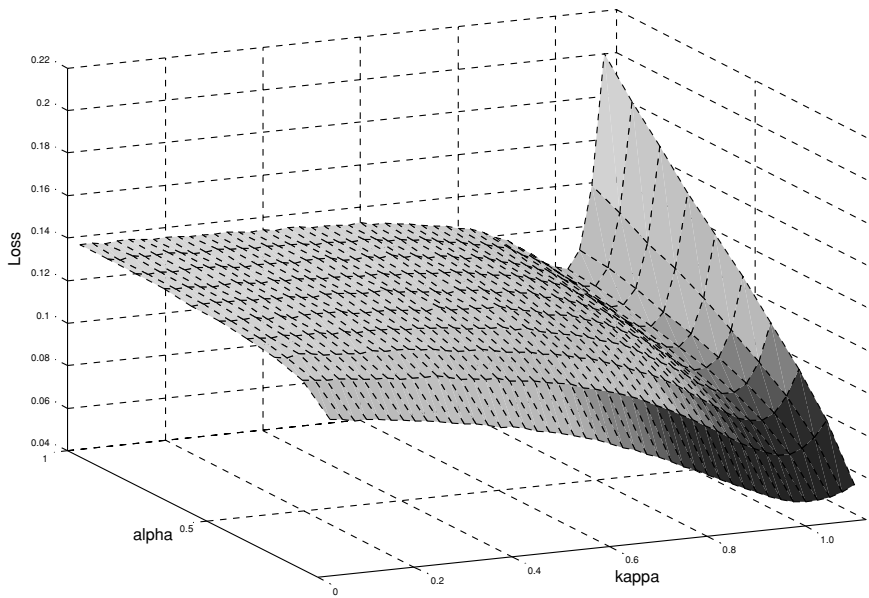


FIGURE 3. Mean losses over 10,000 runs for varying α and κ.

https://doi.org/10.1017/S136510050908016X Published online by Cambridge University Press

presence of such uncertainty, the policymaker should not respond to lagged output as strongly as the full or modified commitment settings indicate.

To examine this issue, I include the error terms $\varepsilon_{x,t}$ and $\varepsilon_{\pi,t}$, both white noise with variances σ_x and σ_π, in the policy rule (14) to form[18]

$$i_t = \delta_L x_{t-1} + \delta_\pi (E_t^* \pi_{t+1} + \varepsilon_{\pi,t}) + \delta_x (E_t^* x_{t+1} + \varepsilon_{x,t}) + \delta_g g_t + \delta_u u_t. \quad \textbf{(20)}$$

Substituting this equation into (1) shows that the model is identical to (15), but adds the additional noise

$$-\varepsilon_{x,t} - \left(1 + \frac{\lambda\beta}{\alpha + \lambda^2}\right) \varepsilon_{\pi,t}$$

to the preference shock $\widetilde{g}_t$. Introducing measurement error has the effect of adding unobservable shocks to the demand side of the model. Besides measurement error, the extra noise could be interpreted as any unseen factors affecting the connection between the policymaker's instrument i_t and the endogenous variables, decisions within financial institutions being one of many examples. Because the source of the errors is unimportant to the simulation results, I make the simplifying assumption that $\varepsilon_{\pi,t} = 0$ and run simulations for different levels of σ_x, adjusting σ_g so that the magnitude of the sum of the shocks does not change.[19]

Introduction of the policy rule errors has an obvious effect on the public's forecasting ability under learning. Figure 4 shows agents' parameter estimates over time for b_x, b_π, c_x, and c_π and demonstrates this conclusion clearly. In the case of small or zero errors in the interest rate rules, estimates of these parameters appear to be constant[20] near their RE values, given by (12). However, for larger $\sigma_x = 0.01$, as in Figure 4, the variation of these parameters is quite visible. Greater variation in the parameter values leads to higher volatility in expectations, see (18), and in output and inflation, as well. Larger errors in the interest rate rule (20) lead to policy mistakes, making the task of estimating the values ξ_t of the perceived model (17) more difficult. The problem is increasingly severe at higher levels of commitment, when the policymaker is influencing public expectations, and for higher values of the gain parameter $\bar{\tau}$, when agents are putting a great deal of emphasis on recent information when estimating ξ_t.

In the cases with significant variation in ξ_t, the estimated values in the PLM (17) can easily slide into regions when the system appears to be explosive. To ensure that the policy rules are operational, I impose a bound or projection facility[21] on the elements of ξ_t. Each element of ξ_t is bounded by a unit interval centered at its RE value. The size of the interval is set to be large enough so that the variation in the elements of ξ_t can have an impact on the endogenous variables but also to be strict enough so that the simulated data does not exhibit unreasonably wild swings. In Table 2, the estimates of c_x hit the bound a number of times. Grandmont (1998) has questioned whether imposing a projection facility excessively restricts agents' behavior, and our results are sensitive to changes in the bound. Imposing a projection facility implies that agents have a certain

https://doi.org/10.1017/S136510050908016X Published online by Cambridge University Press

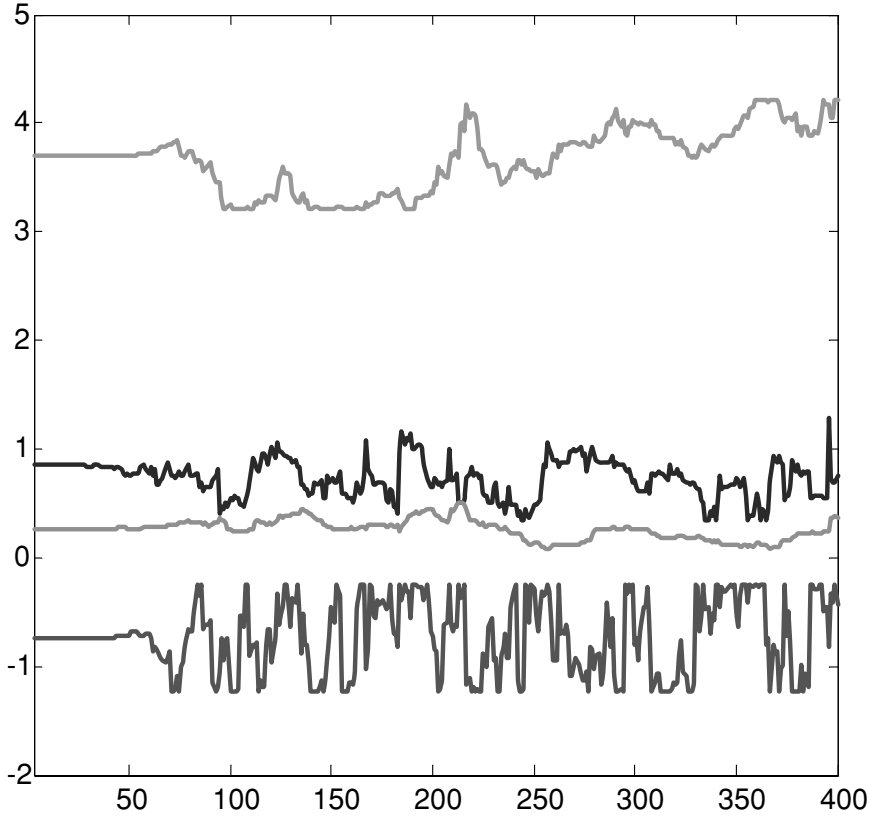


FIGURE 4. The above graph shows the evolution of agents' estimates of c_π, b_x, b_π, and c_x, from top to bottom. The value of the gain is $\bar{\tau} = 0.15$, and the standard deviation of the policy rule errors is $\sigma_x = 0.01$.

amount of faith that the economy will not fluctuate excessively. Investigating whether and how such bounds should be imposed and interpreted is a large area for future research. Here, I simply note that expanding or removing the projection facility makes the errors in the policy rule (20) more of a concern for the policymaker.

For each of the parameterizations in (16), Table 1 reports losses for varying κ for the case of no policy errors and for errors with $\sigma_x = 0.01$, where their impact is equal to that of the preference shocks. For all three parameterizations under learning, the modified commitment setting $\kappa = \beta = 0.99$ remains optimal within the class of rules (14), but within the class of rules (20) with errors in the policy rule, a partial commitment setting of $\kappa < \beta$ is optimal, $\kappa = 0.94$ in the case of MN04.

Table 1 also reports losses in the discretionary case where $\kappa = 0.0$ to ascertain the magnitude of the gains to commitment. Assuming rational expectations,

https://doi.org/10.1017/S136510050908016X Published online by Cambridge University Press

TABLE 2. Each entry shows the loss-minimizing value of κ for different values of σ_x and α. These are determined by calculating the mean loss across 10,000 runs of 200 periods each for $\kappa = 0.00, 0.01, 0.02, \ldots, 1.02$.

a) $\tau = 0.025$					
$\sigma_x \backslash \alpha$	0.01	0.1	0.25	0.5	1
0	0.99	0.99	0.99	0.99	0.99
0.001	0.99	0.99	0.99	0.99	0.99
0.01	0.98	0.97	0.94	0.90	0.42
0.02	0.98	0.92	0.84	0.30	0.05
b) $\tau = 0.15$					
$\sigma_x \backslash \alpha$	0.01	0.1	0.25	0.5	1
0	0.99	0.99	0.99	0.99	0.99
0.001	0.99	0.99	0.99	0.96	0.93
0.01	0.99	0.95	0.91	0.84	0.26
0.02	0.98	0.9	0.68	0.03	−0.04

McCallum and Nelson (2004) report a ratio of losses between the discretionary and fully rational cases of 1.29. Under learning, the ratios between discretion and both full and modified commitment are larger at approximately 1.78. Using the alternative parameterizations, CGG00 and W99, the qualitative conclusions are unchanged.

Tables 2(a) and 2(b) report the loss-minimizing κ's under learning for constant gain $\bar{\tau} = 0.025$ and $\bar{\tau} = 0.15$, respectively, for varying levels[22] of the policymaker preference parameter α and magnitude of policy rule shocks σ_x. The surface of losses in Figure 3 shows that the loss-minimizing setting for κ does not depend on α, and, indeed, modified commitment where $\kappa = 0.99$ is best for all choices of α in the cases where there are no or minimal errors ($\sigma_x = 0.001$) in the policy rule, according to Table 2(a). However, for larger magnitudes of policy rule errors, partial commitment values of $\kappa < 0.99$ minimize the policymaker's loss and vary inversely with α. The more significant the measurement error in the policy rule and the greater the policymaker's emphasis on stabilizing output, the less the interest rate rule should respond to lagged inflation. Results using constant gain of $\bar{\tau} = 0.01$ and $\bar{\tau} = 0.05$ were nearly identical to Table 2(a).

The case for partial commitment is even stronger for the higher constant-gain parameter $\bar{\tau} = 0.15$, as in Table 2(a). Modified commitment is still best for any α in the absence of policy rule errors, but partial commitment with $\kappa < 0.99$ is optimal for minimal errors ($\sigma_x = 0.001$) if the the policymaker is sufficiently concerned about output stabilization. With one exception,[23] higher gain leads to a lower optimal level of commitment. The finding of a negative κ that minimizes the loss in the bottom right-hand corner is counterintuitive, but the loss surface in the neighborhood of $\kappa = 0$ is quite flat, and this outcome should be interpreted

https://doi.org/10.1017/S136510050908016X Published online by Cambridge University Press

TABLE 3. Each column reports summary statistics over 10,000 runs for the chosen parameter values.

	$\sigma_x = 0.001$ $\tau = 0.025$	$\sigma_x = 0.001$ $\tau = 0.15$	$\sigma_x = 0.01$ $\tau = 0.025$	$\sigma_x = 0.01$ $\tau = 0.15$
b_x				
Mean deviation from RE	−3.36E-04	−2.96E-03	−1.84E-02	−1.17E-01
Mean squared deviation	1.52E-05	5.28E-07	1.14E-03	2.53E-02
b_p				
Mean deviation from RE	1.98E-04	−1.29E-03	1.02E-02	6.67E-02
Mean squared deviation	5.39E-06	2.59E-06	2.78E-04	2.66E-02
c_x				
Mean deviation from RE	1.48E-04	−9.67E-04	1.09E-02	2.12E-02
Mean squared deviation	3.40E-04	4.26E-03	3.02E-02	1.13E-01
c_p				
Mean deviation from RE	−5.99E-05	9.14E-03	−1.19E.02	−4.11E-02
Mean squared deviation	5.52E-05	1.47E-02	4.89E-03	8.95E-02
x forecast				
Mean squared error	1.33E-05	1.50E-05	1.29E-04	2.71E-04
π forecast				
Mean squared error	5.35E-05	5.31E-05	5.46E-05	5.58E-05
Skewness of losses	0.748	1.819	0.838	39.282

as a recommendation in favor of discretionary policy for the indicated parameter values.

A significant deterioration of the performance of modified commitment requires some combination of policymaker concern about output stabilization, policy rule errors, and/or high gain. Introduction of any one of these factors by itself does not alter greatly the conclusion that modified commitment is preferable. For example, if policy rule errors are at their maximum $\sigma_x = 0.02$ when the policymaker is primarily focused on inflation stabilization $\alpha = 0.01$ and the gain is small $\bar{\tau} = 0.025$, the loss-minimizing κ is 0.98 [Table 2(a)] and the gain over modified commitment $\kappa = 0.99$ is less than 10^{-5}. However, increasing α with the higher policy rule errors changes this conclusion, making partial commitment better, and raising the gain parameter in addition lowers the loss-minimizing κ even more.

To clarify the results concerning the optimality of partial commitment for this class of interest rate rules (20), I closely examine four cases under modified commitment where the standard deviation of the policy rule errors takes the value of $\sigma_x = 0.001$ or 0.01, and the gain parameter takes the values $\bar{\tau} = 0.025$ and 0.15. Table 3 reports the mean deviation of the estimates of b_x, b_π, c_x, and c_π from their RE values, their squared deviation from the mean of the simulated values, the mean squared error of the forecasts for output and inflation, and the skewness of the losses over the 10,000 runs. Orphanides and Williams (2002) report systematic

https://doi.org/10.1017/S136510050908016X Published online by Cambridge University Press

bias in agents' estimates even in the absence of stochastic shocks in the model. Here, bias and fluctation in the estimates of the parameters are evident and coincide with larger forecast errors.

The first two columns are cases with minimal policy rule errors when modified commitment $\kappa = \beta = 0.99$ is best. In both cases, the bias and fluctuation in the parameter estimates and the forecast errors are smaller than the cases in the third and fourth columns. The presence of significant policy rule errors has a detrimental effect on agents forecasting ability, most notably for output. The losses in the first two columns show positive skewness, as expected given the asymmetry in Figure 3, though it is not sufficient to affect the optimality of modified commitment. The most notable entry in Table 3 is the high positive skewness in the case with large policy rule errors and gain, $(\sigma_x, \bar{\tau}) = (0.01, 0.15)$. In this case, the prevalence of some losses far above the mean makes a setting of $\kappa = 0.91$ [see Table 2(b)] better than modified commitment. However, comparing the second and third columns, the case where $(\sigma_x, \bar{\tau}) = (0.001, 0.15)$ corresponds to a loss-minimizing $\kappa = 0.99$ whereas the case where $(\sigma_x, \bar{\tau}) = (0.01, 0.025)$ has a lower loss-minimizing $\kappa = 0.94$, even though the skewness is less than in the second column. Hence, larger bias and/or variation in the parameter estimates or the larger policy rule errors in the third column must also be a factor behind the optimal κ being below modified commitment.

Simulations of the model under partial commitment in the case where $(\sigma_x, \bar{\tau}) = (0.01, 0.025)$ for a fixed ξ, further disentangle the influence of these factors. Fixing the agents' estimate of ξ at the biased value from Table 3 yielded a loss of 0.09120, which is a slight improvement over the 0.0925 shown in Table 1 but not nearly as good as the loss achieved at the optimal $\kappa = 0.94$. Fixing the estimate of ξ at the rational expectations value yielded a loss of 0.09219, which is worse than the outcome under the biased ξ, so the bias actually improved the policy outcome in the presence of policy rule errors. Hence, the fluctuations in the estimates of ξ play a role in the deterioration of the performance of modified commitment, but the presence of the policy rule errors is a major factor.

The deterioration of policy outcomes at modified commitment for larger policy rules errors and higher gain is evident from the distribution of the losses shown in Figure 5. These curves track the height of the midpoints of a histogram over 100 intervals of width 0.003. For the case where $(\sigma_x, \bar{\tau}) = (0.01, 0.15)$ the mean is to the right, and the tail is fatter compared to the other distributions. The jump at 0.3 for this case indicates there are a number of cases with losses above the maximum of 0.3 as well.

Close examination of these four cases shows that larger policy rule errors, greater variation in the estimates of ξ, and higher gain lead to increased incidence of large losses and worse policy outcomes under modified commitment, necessitating a partial commitment setting for $\kappa < \beta$. Conversely, for smaller policy rule shocks and low gain, the fluctuations in ξ are minimal, the asymmetry of the losses seen in Figure 3 does not play a significant role, and the modified commitment setting $\kappa = \beta$ is optimal under learning.

https://doi.org/10.1017/S136510050908016X Published online by Cambridge University Press

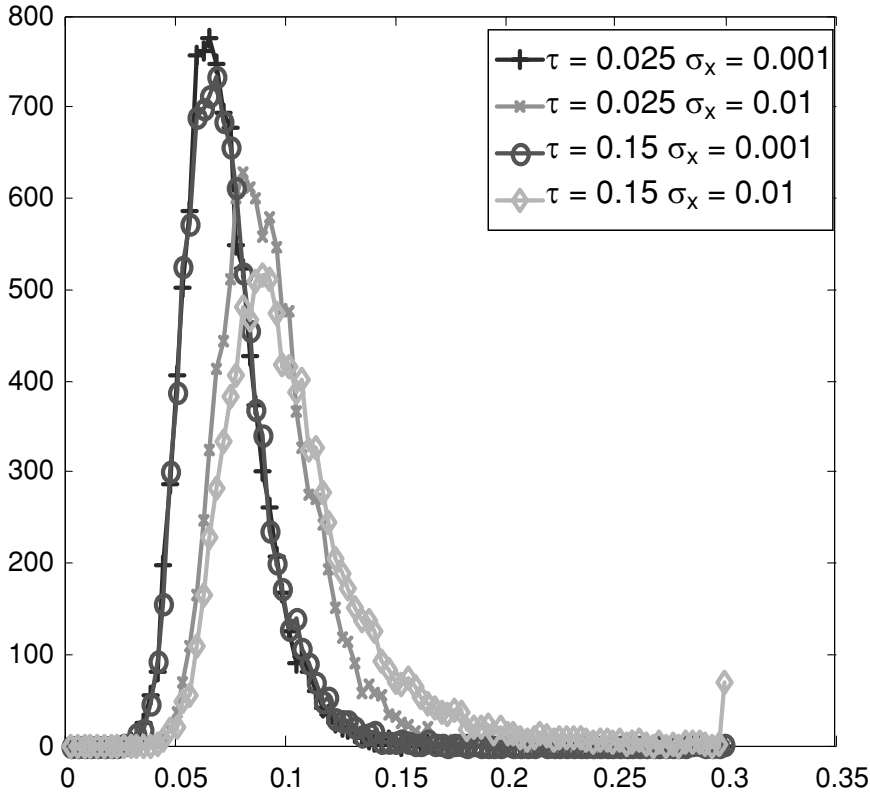


FIGURE 5. Each curve tracks the distribution of the mean losses of 10,000 runs of 200 periods each for the chosen parameters. The histograms are generated across 100 intervals of width 0.003.

The interpretation of the results with respect to σ_x is straightforward. Larger errors in the policy rule indicate a lower level of knowledge on the part of the policymaker and a public that is less able to correctly learn the values of the forecasting model. Policymakers cannot set interest rates as accurately and expectations show more instability, leading to worse outcomes from a policy perspective.

The impact of a higher $\bar{\tau}$ is more open to interpretation. The gain parameter can represent the lags of data used by agents, in their estimation, the rationality of agents, and/or the credibility of the policymaker. A marked policy deterioration only appears when $\bar{\tau}$ is raised to 0.15, a relatively high value indicating that agents are using only 3.3 years of data to make their estimates. This higher gain could be reasonable if there were a perceived shift of regime in the economy or if the policymaker had very low credibility. The policymaker is less able to commit when he or she has low credibility or poor information about the connection between the policy instrument and the other endogenous variables in the economy.

https://doi.org/10.1017/S136510050908016X Published online by Cambridge University Press

The conclusion that a policymaker with weak credibility should not aggressively commit rests on the assumption of constant gain, however. If the model is expanded to allow the gain parameter to fall as the policymaker gains credibility, then the policymaker may have an incentive to stick to a higher level of commitment. Such a model is studied in Waters (2007), who uses an endogenous gain mechanism introduced by Marcet and Nicolini (2004).

8. CONCLUSION

Relaxing the assumption of full rationality for the formation of expectations raises a number of related issues for the design of monetary policy rules. One concern is whether a rule is stable under learning, whereas another is whether a rule minimizes the policymaker's loss under learning. Evans and Honkapohja (2003, 2006) introduce a class of expectations-based interest rate rules and study the former issue. The present work extends their analysis to a broader class of rules allowing for varying degrees of response to lagged variables and goes on to study the latter question for such rules.

Within the present framework, one can meaningfully discuss varying levels of commitment. The greater the policymaker's degree of commitment, the more he or she will adjust the interest rate in response to lagged output to affect public expectations. The commitment optimum under rational expectations or full commitment is a special case, as is the modified commitment value advocated by Blake (2001). I provide a condition for nonexplosive and determinate equilibria for this class of rules and show that expectational stability holds for any nonnegative response to lagged output.

Simulation results show that modified commitment minimizes the policymaker's loss under least-squares learning in the baseline case. However, parameter uncertainty, errors in the policy rule, and high gain are all factors that could, in combination, make a policy of partial commitment best, implying a lower response to lagged output than modified commitment. An excessive response to lagged output can lead to very bad outcomes for the policymaker, so he or she should act to minimize that possibility. The presence of a conservative central banker makes such issues less problematic, however.

The gain parameter has important economic interpretations in terms of the lags used to form expectations, the rationality of the agents, and the credibility of the policymaker. When agents are using only recent data (high gain), indicating low credibility, the case for partial commitment becomes stronger.

Expectations-based interest rate rules continue to demonstrate desirable characteristics for monetary policy, but the present work does offer a few caveats for practical implementation. For an extended range of responses to lagged output they exhibit nonexplosive, determinate, and expectationally stable solutions. Some degree of commitment minimizes the loss for the policymaker when agents use a learning mechanism to form expectations. However, it is often the case that the policymaker should be cautious about the commitment optimum under

https://doi.org/10.1017/S136510050908016X Published online by Cambridge University Press

rational expectations and engage in partial commitment. The prerequisites for partial commitment, parameter uncertainty and policy rule errors, clearly exist to some degree, so the policymaker must be vigilant.

These results point to some important areas for future work, namely determining the extent to which parameter uncertainty and policy rule errors are present in estimated models of monetary policy. Furthermore, credibility is an important aspect of monetary policy but has proved to be an elusive feature to model formally. Interpreting the gain parameter as an indication of credibility may open new avenues for research on monetary policy. The present paper gives practical guidance on the design of rules and areas that need attention for deeper understanding of monetary policy.

NOTES

1. Time consistency of monetary policy has been a relevant issue going back to Kydland and Prescott (1977) and continuing with Woodford's (1999a, 2003) discussion of a *timeless perspective.*

2. Other examples include Honkapohja and Mitra (2003) and Beechey (2004).

3. Evans and McGough (2005) study determinancy and expectational stability for different policy rules and specifications of the Phillip's curve.

4. Woodford (1999b) and Blake (2001) refer to this as the timeless perspective policy.

5. Jensen and McCallum (2002) also discuss this point.

6. Honkapohja and Mitra (2003) give theoretical stability results under bonded memory learning. Carceles-Poveda and Giannitsarou (2006) use constant gain to study asset price dynamics.

7. The present approach follows Woodford (1999a), Clarida et al. (1999), and Evans and Honkapohja (2006).

8. He also checks the result with a formal derivation. Jensen and McCallum (2002) also discuss this issue. Evans and McGough (2006) refer to this approach as the MJB-alternative in their study of the New Keynsian model with inertia.

9. See McCallum (1983, 1997).

10. If the policymaker has a biased estimate of λ, the change in policy would be equivalent to the change caused by varying the policymaker preference parameter α.

11. See Walsh (2003) for a full discussion of inflation bias.

12. As noted, Orphanides and Williams (2002) is closely related. Fuhrer and Hooker (1993) is another example of a monetary policy analysis using least-squares learning.

13. Note that exogenous variables and parameter estimates dated at time t are used to form expectations of future variables, but endogenous variables dated at time t are not.

14. See Evans and Honkapohja (2001) and Marcet and Sargent (1989) for details.

15. A constant gain parameter $\bar{\tau}$ corresponds to $2/\bar{\tau}$ lags of data used.

16. See Evans and Honkapohja (1993) for an example where the gain parameter with the optimal performance was 0.15. Orphanides and Williams (2002) use semiannual data, so a gain of 0.05 is equivalent to a gain of 0.025 here.

17. For all reported results, the loss for each run is calculated over 200 periods using (4) after 600 periods for initialization. The initial values ξ_0 are set to their RE values. Because losses are computed with discounting, longer runs do not provide much extra information.

18. The form of this rule is identical to that found in Evans and Honkapohja (2003).

19. Formally, $\sigma_g^2 + \sigma_x^2 = (0.02)^2$.

20. For example, of the four paramters, c_x always has the largest mean squared deviation (MSD), and for $\sigma_x = 0.001$ its MSD is 0.0005, but for $\sigma_x = 0.01$, as shown in Figure 3, the MSD is 300 times higher at 0.15.

21. This term originally comes from Marcet and Sargent's (1989) study of least squares learning.

https://doi.org/10.1017/S136510050908016X Published online by Cambridge University Press

22. The range for α corresponds to McCallum and Nelson (2004).

23. At $\alpha = 0.01$, and $\sigma_x = 0.01$, the loss-minimizing κ is 0.98 for $\bar{\tau} = 0.025$ but 0.99 for $\bar{\tau} = 0.15$. This result was verified for 20,000 runs, but the difference in loss between $\kappa = 0.99$ and $\kappa = 0.98$ for the lower gain was less than 10^{-6}.

REFERENCES

Barro, Robert J. and David B. Gordon (1983) A positive theory of monetary policy in a natural rate model. *Journal of Political Economy* 91, 589–610.

Beechey, Meredith (2004) Excess Sensitivity and Volatility of Long Interest Rates: The Role of Limited Information in Bond Markets. Working Paper, University of California Berkeley.

Blake, A.P. (2001) A "Timeless Perspective" on Optimality in Forward-Looking Rational Expectations Models. Working Paper, National Institute for Economic and Social Research.

Blanchard, Oliver and Charles Kahn (1980) The solution of linear difference models under rational expectations. *Econometrica* 48(5), 1305–1311.

Bray, Margaret (1982) Learning, estimation, and the stability of rational expectations. *Journal of Economic Theory* 26(2), 312–339.

Bullard, James and In-Koo Cho (2005) Escapist policy rules. *Journal of Economic Dynamics and Control*, 29(11), 1841–65.

Bullard, James and Kaushik Mitra (2002) Learning about monetary policy rules. *Journal of Monetary Economics* 49, 1105–1129.

Carceles-Poveda, Eva and Chryssi Giannitsarou (in press) Asset pricing with adaptive learning. *Review of Economic Dynamics*.

Clarida, Richard, Jordi Gali, and Mark Gertler (1999) The science of monetary policy: A New Keynesian perspective. *Journal of Economic Literature* 37, 1661–1707.

Clarida, Richard, Jordi Gali, and Mark Gertler (2000) Monetary policy rules and macroeconomic stability: Evidence and some theory. *Quarterly Journal of Economics* 115, 147–180.

Evans, George and Seppo Honkapohja (1993) Adaptive forecasts, hysteresis and endogenous fluctuations. *Federal Reserve Bank of San Francisco Economic Review* 1993 (1), 3–13.

Evans, George and Seppo Honkapohja (2001) *Learning and Expectations in Macroeconomics*. Princeton, NJ: Princeton University Press.

Evans, George and Seppo Honkapohja (2003) Expectations and the stability problem for optimal monetary policies. *Review of Economic Studies* 70, 807–824.

Evans, George and Seppo Honkapohja (2006) Monetary policy, expectations and commitment. *Scandinavian Journal of Economics* 108, 15–38.

Evans, George and Bruce McGough (2005) Monetary policy, indeterminacy and learning. *Journal of Economic Dynamics and Control* 29, 1809–1840.

Fuhrer, Jeffrey and Mark Hooker (1993) Learning about monetary regime shifts in an overlapping wage contract model. *Journal of Economic Dynamics and Control* 17(4), 531–53.

Grandmont, Jean-Michel (1998) Expectations formation and stability of large socioeconomic systems. *Econometrica* 66, 741–781.

Honkapohja, Seppo and Kaushik Mitra (2003) Learning with bounded memory in stochastic models. *Journal of Economic Dynamics and Control* 27, 1437–1457.

Honkapohja, Seppo and Kaushik Mitra (2004) Are non-fundamental equilibria learnable in models of monetary policy? *Journal of Monetary Economics* 51(8), 1743–1770.

Howitt, Peter (1992) Interest rate control and nonconvergence to rational expectations. *Journal of Political Economy* 100, 776–800.

Jensen, Christian and Bennet McCallum (2002) The non-optimality of proposed monetary policy rules under timeless-perspective commitment. *Economics Letters*, 77 (2), 163–168.

Kydland, Finn E. and Edward C. Prescott (1977) Rules rather than discretion: The inconsistency of optimal plans. *Journal of Political Economy* 85, 473–493.

https://doi.org/10.1017/S136510050908016X Published online by Cambridge University Press

Marcet, Albert and Juan Nicolini (2004) Recurrent hyperinflations and learning. *American Economic Review* 93 (5), 1476–1498.
Marcet, A., and T. Sargent (1989) Convergence of least squares learning mechanisms in self-referential linear stochastic models. *Journal of Economic Theory* 48, 337–368.
McCallum, Bennet T. (1983) On non-uniqueness in linear rational expectations models: An attempt at perspective. *Journal of Monetary Economics* 11, 139–168.
McCallum, Bennet T. (1997) The role of the minimum state variables criterion in rational expectations models. *Journal of International Tax and Finance* 6(4), 621–639.
McCallum, Bennet T. and Edward Nelson (2004) Timeless perspective vs. discretionary monetary policy in forward looking models. *Federal Reserve Bank of St. Louis Review* 86(2), 43–56.
Orphanides, Athanios and John C. Williams (2002) Imperfect Knowledge, Inflation Expectations and Monetary Policy. Working Paper 9884, National Bureau of Economic Research.
Orphanides, Athanios and John C. Williams (2006) Monetary policy with imperfect knowledge. *Journal of the European Economic Association* 4 (2–3), 366–375.
Rogoff, Kenneth (1985) The optimal degree of commitment to an intermediate monetary target. *Quarterly Journal of Economics* 100(4), 1169–89.
Walsh, Carl E. (2003) *Monetary Theory and Policy*. Cambridge, MA: MIT Press.
Waters, George (2006) Dangers of commitment: Monetary policy with adaptive learning. *Journal of Economics and Finance* 30(1), 93–104.
Waters, George (2007) Regime changes, learning and monetary policy. *Journal of Macroeconomics* 29(2), 255–282.
Woodford, Michael (1999a) Commentary: How should monetary policy be conducted in an era of price stability? In *New Challenges for Monetary Policy*, pp. 277–316. Federal Reserve Bank of Kansas City.
Woodford, Michael (1999b) Optimal monetary policy inertia. *The Manchester School, Supplement* 67, 1–35.
Woodford, Michael (2003) *Interest and Prices*. Princeton, NJ: Princeton University Press.

APPENDIX

A.1. PROOF OF PROPOSITION 1

To demonstrate determinacy, first define a new variable x_t^L for the lag of the output gap, such that $x_{t+1}^L = x_t$. Under rational expectations $E_t\pi_{t+1} = \pi_{t+1} + \varepsilon_{t+1}$ for some iid, mean-0-shock ε_t. Then the model (15) can be written in the form

$$\begin{pmatrix} 1 & 0 & \dfrac{-\alpha\kappa}{\alpha+\lambda^2} \\ 0 & 1 & \dfrac{-\lambda\alpha\kappa}{\alpha+\lambda^2} \\ 1 & 0 & 0 \end{pmatrix} \begin{pmatrix} x_t \\ \pi_t \\ x_t^L \end{pmatrix} = \begin{pmatrix} 0 & \dfrac{-\lambda\beta}{\alpha+\lambda^2} & 0 \\ 0 & \dfrac{\alpha\beta}{\alpha+\lambda^2} & 0 \\ 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} x_{t+1} \\ \pi_{t+1} \\ x_{t+1}^L \end{pmatrix} + K \begin{pmatrix} u_t \\ g_t \\ \varepsilon_t \end{pmatrix}.$$

Multiplying both sides by the inverse of the 3×3 matrix on the left-hand side gives the model the form

$$\begin{pmatrix} x_t \\ \pi_t \\ x_t^L \end{pmatrix} = J \begin{pmatrix} x_{t+1} \\ \pi_{t+1} \\ x_{t+1}^L \end{pmatrix} + \tilde{K} \begin{pmatrix} u_t \\ g_t \\ \varepsilon_t \end{pmatrix},$$

https://doi.org/10.1017/S136510050908016X Published online by Cambridge University Press

where $\tilde{K}$ is the appropriate transformation of K. The condition for determinacy [see Blanchard and Kahn (1980) and Evans and Honkapohja (2006)] is that the matrix J must have two eigenvalues with modulus less than one and one eigenvalue with modulus greater than one.

A straightforward computation shows that one eigenvalue e of J is 0, whereas the other two must satisfy

$$e^2 - \left(\beta + \frac{\alpha + \lambda^2}{\alpha\kappa}\right) e + \frac{\beta}{\kappa} = 0.$$

For one root to be greater than one and one root to be less than one, the right-hand side of the equation above must be negative for $e = 1$, which is equivalent to the condition

$$\kappa < 1 + \frac{\lambda^2}{\alpha\,(1-\beta)}.$$

Nonexplosiveness requires that the negative solution of b_x in the quadratic condition (13) be less than one. Because the positive root is always greater than one, the nonexplosiveness condition requires that the left-hand side of (13) be less than zero at $b_x = 1$, which is true for

$$\kappa < 1 + \frac{\lambda^2}{\alpha\,(1-\beta)}$$

as well. ■

A.2. PROOF OF PROPOSITION 2

To prove Proposition 2, I first prove the following lemma.

LEMMA 3. *Given $\kappa \geq 0$ and the associated REE value of b_x,*

(i) *b_x is real,*
(ii) *$b_x - \kappa \leq 0$,*
(iii) *$b_x - \kappa > -1$.*

Proof. The relevant root of (13) is

$$b_x = (2\beta)^{-1}\left[\frac{\lambda^2 + \alpha}{\alpha} + \beta\kappa - \sqrt{\left(\frac{\lambda^2 + \alpha}{\alpha} + \beta\kappa\right)^2 - 4\beta\kappa}\right],$$

which is real if the radicand rad is positive. The radicand can be written

$$\text{rad} = \beta^2\kappa + 2\beta\kappa\left(\frac{\lambda^2 - \alpha}{\alpha}\right) + \left(\frac{\lambda^2 + \alpha}{\alpha}\right)^2$$

and is minimized at

$$\kappa = -\frac{1}{\beta}\left(\frac{\lambda^2 - \alpha}{\alpha}\right),$$

so, substituting, the minimum value attained by rad is $4\lambda^2/\alpha > 0$. Hence, the radicand is always positive for positive parameters, and b_x is real.

https://doi.org/10.1017/S136510050908016X Published online by Cambridge University Press

The difference $b_x - \kappa$ can be written

$$b_x - \kappa = (2\beta)^{-1}\left[\frac{\lambda^2+\alpha}{\alpha} - \beta\kappa - \sqrt{\left(\frac{\lambda^2+\alpha}{\alpha} - \beta\kappa\right)^2 - 4\beta\kappa}\right].$$

Therefore $b_x - \kappa \leq 0$ if $4\beta\kappa > 0$ which is true if $\kappa \geq 0$.

The condition in (iii), $b_x - \kappa > -1$, is equivalent to

$$2\beta + \frac{\lambda^2+\alpha}{\alpha} - \beta\kappa > \sqrt{\left(\frac{\lambda^2+\alpha}{\alpha} - \beta\kappa\right)^2 - 4\beta\kappa}.$$

Squaring both sides and simplifying yields another equivalent condition,

$$\beta + \frac{\lambda^2+\alpha}{\alpha} > -\kappa(1-\beta),$$

which is true for $\kappa \geq 0$, $\beta < 1$, and positive parameters. ■

To study expectational stability of (15), consider the general model

$$y_t = ME_t^* y_{t+1} + Ny_{t-1} + Pv_t$$

with $v_t = Fv_{t-1}$ from (3). Expectations $E_t^* y_{t+1}$ are formed according to (18) and substituting them into equation above gives the ALM in the form $y_t = T(a_{t-1}, b_{t-1}, c_{t-1})z_{t-1}$ for the map

$$T(a,b,c) = [M(a+ba), Mb^2 + N, M(bc+F) + P],$$

suppressing the time subscripts in the $T-$map. For the model (15) with REE (11) the matrices are such that

$$M = \begin{pmatrix} 0 & \dfrac{-\lambda\beta}{\alpha+\lambda^2} \\ 0 & \dfrac{\alpha\beta}{\alpha+\lambda^2} \end{pmatrix} \quad \text{and} \quad B = \begin{pmatrix} b_x & 0 \\ b_\pi & 0 \end{pmatrix}.$$

Expectational stability is defined according to the matrix differential equation

$$\frac{d}{ds}(a,b,c) = T(a,b,c) - (a,b,c),$$

where s is notional time, distinct from the periods denoted by t. The REE solutions are fixed points of the differential equation where $T(a,b,c) = (a,b,c)$, and expectational stability is defined in terms of stability of the differential equation.

The associated REE from (11) has the form

$$y_t = By_{t-1} + Cv_t.$$

The conditions for expectational stability [see Evans and Honkapohja (2001), Ch. 10)] are that the matrices

$$DT_a = M(I+B)$$
$$DT_b = B' \otimes M + I \otimes MB$$

https://doi.org/10.1017/S136510050908016X Published online by Cambridge University Press

$$DT_c = F' \otimes M + I \otimes MB$$

must have eigenvalues with modulus less than one.

The proof of Proposition 2 follows.

Proof. The matrices for the condition for expectational stability are

$$DT_a = \begin{pmatrix} \dfrac{-b_\pi \lambda\beta}{\alpha+\lambda^2} & \dfrac{-\lambda\beta}{\alpha+\lambda^2} \\ \dfrac{b_\pi \alpha\beta}{\alpha+\lambda^2} & \dfrac{\alpha\beta}{\alpha+\lambda^2} \end{pmatrix}$$

$$DT_b = \begin{pmatrix} \dfrac{-b_\pi \lambda\beta}{\alpha+\lambda^2} & \dfrac{-b_x \lambda\beta}{\alpha+\lambda^2} & 0 & \dfrac{-b_\pi \lambda\beta}{\alpha+\lambda^2} \\ \dfrac{b_\pi \alpha\beta}{\alpha+\lambda^2} & \dfrac{b_x \alpha\beta}{\alpha+\lambda^2} & 0 & \dfrac{b_x \alpha\beta}{\alpha+\lambda^2} \\ 0 & 0 & \dfrac{-b_\pi \lambda\beta}{\alpha+\lambda^2} & 0 \\ 0 & 0 & \dfrac{b_\pi \alpha\beta}{\alpha+\lambda^2} & 0 \end{pmatrix}$$

$$DT_c = \begin{pmatrix} \dfrac{-b_\pi \lambda\beta}{\alpha+\lambda^2} & \dfrac{-\mu\lambda\beta}{\alpha+\lambda^2} & 0 & 0 \\ \dfrac{b_\pi \alpha\beta}{\alpha+\lambda^2} & \dfrac{\mu\alpha\beta}{\alpha+\lambda^2} & 0 & 0 \\ 0 & 0 & \dfrac{-b_\pi \lambda\beta}{\alpha+\lambda^2} & \dfrac{-\rho\lambda\beta}{\alpha+\lambda^2} \\ 0 & 0 & \dfrac{b_\pi \alpha\beta}{\alpha+\lambda^2} & \dfrac{\rho\alpha\beta}{\alpha+\lambda^2} \end{pmatrix}.$$

The eigenvalues of DT_a are 0 and

$$\frac{\alpha\beta}{\alpha+\lambda^2}\left(1 - \frac{\lambda}{\alpha} b_\pi\right).$$

Using the second equation from (12), the nonzero eigenvalue is

$$\frac{\alpha\beta}{\alpha+\lambda^2}\left(1 + b_x - \kappa\right),$$

which is less than one because $\alpha/(\alpha+\lambda^2)$, $\beta < 1$, and $b_x - \kappa \leq 0$ from Lemma 3.

The the matrix DT_b has two eigenvalues equal to 0. It also has eigenvalues,

$$\frac{\alpha\beta}{\alpha+\lambda^2}\left(b_x - \frac{\lambda}{\alpha} b_\pi\right) \quad \text{and} \quad \frac{-b_\pi \lambda\beta}{\alpha+\lambda^2}.$$

Using the second equation from (12), the two nonzero eigenvalues are

$$\frac{\alpha\beta}{\alpha+\lambda^2}\left(b_x + b_x - \kappa\right) \quad \text{and} \quad \frac{\alpha\beta}{\alpha+\lambda^2}\left(b_x - \kappa\right),$$

https://doi.org/10.1017/S136510050908016X Published online by Cambridge University Press

the second of which is less than one, since $\alpha/(\alpha+\lambda^2)$, $\beta < 1$, and $b_x - \kappa \leq 0$ from Lemma 3. The condition on the first nonzero eigenvalue of DT_b,

$$\frac{\alpha\beta}{\alpha+\lambda^2}(b_x + b_x - \kappa) < 1,$$

is equivalent to

$$1 - \frac{1}{2}\left(\frac{\alpha}{\alpha+\lambda^2}\right)\left[\sqrt{\left(\frac{\lambda^2+\alpha}{\alpha}+\beta\kappa\right)^2 - 4\beta\kappa} + \sqrt{\left(\frac{\lambda^2+\alpha}{\alpha}-\beta\kappa\right)^2 - 4\beta\kappa}\right] < 1,$$

using the expressions for b_x and $b_x - \kappa$ in Lemma 3. The condition is satisfied as long as both radicands from the roots of b_x and $b_x - \kappa$ are positive. Lemma 3 proves that b_x is real, and hence $b_x - \kappa$ is as well, both radicands are positive, and the eigenvalue is less than one.

The matrix DT_c has two repeated eigenvalues of 0 and

$$\frac{\alpha\beta}{\alpha+\lambda^2}\left(\mu - \frac{\lambda}{\alpha}b_\pi\right) = \frac{\alpha\beta}{\alpha+\lambda^2}(\mu + b_x - \kappa),$$

using the second equation from (12) again. The nonzero eigenvalue is less than one because $\frac{\alpha}{\alpha+\lambda^2}$, β, $\mu < 1$ and $b_x - \kappa < 0$ from Lemma 3.

The lemma also guarantees that $b_x - \kappa > -1$, so all the nonzero eigenvalues for the three matrices are greater than -1 as well.

Therefore, the moduli of all the eigenvalues of DT_a, DT_b, and DT_c are all less than one and the model (15) is expectationally stable for any REE. ■

https://doi.org/10.1017/S136510050908016X Published online by Cambridge University Press