

Monetary and Fiscal Policy under Learning in the Presence of a Liquidity Trap

George W. Evans

This paper reports on the findings of Evans, Guse, and Honkapohja (2007) concerning the global economic dynamics under learning in a New Keynesian model in which the interest rate rule is subject to the zero lower bound. Under normal monetary and fiscal policy, the intended steady state is locally but not globally stable. Large pessimistic shocks to expectations can lead to deflationary spirals with falling prices and falling output. To avoid this outcome, we recommend augmenting normal policies with inflation threshold policies: if under normal policies inflation would fall below a suitably chosen threshold, these policies should be replaced by aggressive monetary and fiscal policies that guarantee this lower bound on inflation.

Keywords: Adaptive learning; Monetary policy; Fiscal policy; Zero interest rate lower bound; Indeterminacy

JEL Classification: E52, E58, E63

Department of Economics, University of Oregon (E-mail: gevans@uoregon.edu)

This paper was prepared in part while the author was a visiting scholar at the Institute for Monetary and Economic Studies (IMES), Bank of Japan (BOJ). This report is based on Evans, Guse, and Honkapohja (2007) and other work in progress with Professor Seppo Honkapohja and Dr. Eran Guse of the University of Cambridge. Support for this research was provided by National Science Foundation Grant No. SES-0617859. I am grateful for comments received during seminar presentations at the University of Tokyo and IMES, and for the comments and suggestions of the referees. Views expressed in this paper are those of the author and do not necessarily reflect the official views of the BOJ.

I. Introduction

A standard characterization of monetary policy is that the central bank follows a rule in which the interest rate R_t responds more than one-for-one to deviations of the inflation rate π_t from its target π^* . This is sometimes called the “Taylor principle,” and it implies that when inflation is too high the monetary authorities raise real interest rates.¹ This is a natural monetary policy principle, since higher real interest rates discourage consumption and investment and this puts downward pressure on the inflation rate.

As is now widely recognized, the zero lower bound on nominal net interest rates has the potential to generate a “liquidity trap” with possibly major implications for economic performance. One way to view the problem is that under weak additional assumptions such a policy rule will entail a second steady state at a lower inflation rate. The reason for this can be seen in Figure 1.

The straight line represents the Fisher equation $R = \beta^{-1}\pi$, since in many baseline theories with a subjective discount factor $0 < \beta < 1$ the real interest factor is β^{-1} . The curved line $R = 1 + f(\pi)$ represents a “global” interest rate rule, giving the response of the gross nominal interest rate R to the inflation factor π_t . At the target inflation factor π^* , the slope is assumed to be larger than β^{-1} . This is the discrete time analog of the Taylor principle that net interest rates respond more than one-for-one to changes in net inflation rates (Taylor [1993] recommended a coefficient of 1.5). The dashed horizontal line indicates the zero lower bound, which constrains R to be at least one (i.e., the net interest rate $R - 1$ to be nonnegative). Continuity of the Taylor rule implies another, unintended, steady state $\pi_L < \pi^*$. Here π is the inflation factor, in other words, $\pi_t = P_t/P_{t-1}$. Thus, constant prices over time correspond to $\pi = 1$ and values $0 < \pi < 1$ represent deflation. Depending on the value of β and the specific policy rule $f(\pi)$, the value of π_L may be at positive or negative net rates of inflation.²

The multiple equilibrium issue was emphasized, in particular, by Benhabib, Schmitt-Grohé, and Uribe (2001b), who showed that under perfect foresight (or rational expectations [RE]) this second, unintended, low-inflation steady state not only necessarily exists but also would possess a multiplicity of paths converging to it. This has been interpreted as implying a significant risk that the economy might follow one of these “liquidity trap” paths. However, the analysis of convergence of paths to the low-inflation steady state π_L relies heavily on perfect foresight (or, in stochastic versions, on RE).


There is a substantial literature that has discussed the plausibility of the economy becoming trapped in a deflationary state, and what macroeconomic policies would be able to avoid or extricate the economy from a liquidity trap.³ The view in Evans, Guse, and Honkapohja (2007), as well as in the earlier paper Evans and Honkapohja (2005), is that the evolution of expectations plays a key role in the dynamics of the economy and

1. Taylor (1993) suggested that an appropriate coefficient was 1.5. He also suggested that R_t should respond to the output gap.

2. In our numerical illustrations, in Section VII, π_L corresponds to a deflation rate of about 2.5 percent per annum.

3. See Krugman (1998) for a recent seminal discussion and Adam and Billi (2007), Coenen, Orphanides, and Wieland (2004), and Eggertsson and Woodford (2003, 2004) for representative recent analyses and further references. Braun and Waki (2006) provide a calibrated model for Japan that incorporates the zero interest rate lower bound.

Figure 1 Multiple Steady States with the Global Taylor Rule



that the tools from learning theory are needed for a comprehensive analysis of these issues. In this report, I outline the reasons for this perspective, the main theoretical results that emerge and the implications for monetary and fiscal policy.

The importance of expectations in the liquidity trap is now widely accepted. It is implicit in the perfect foresight analysis of Benhabib, Schmitt-Grohé, and Uribe (2001a, b), and in the Eggertsson and Woodford (2003) emphasis on the importance of policy commitment for influencing expectations under the RE assumption. The learning perspective alters both the assessment of the plausibility of particular dynamics and the impact of policy.

Under learning, private agents are assumed to form expectations using an adaptive forecasting rule, which they update over time in accordance with standard statistical procedures. In many standard setups, least-squares learning is known to converge asymptotically to RE, but cases of instability can also arise. In Evans and Honkapohja (2005), we examined a flexible price model with a global Taylor rule. We found that while the intended steady state π^* was locally stable under learning, the low-inflation steady state π_L was not locally stable,⁴ and there was also the possibility under learning of inflation slipping below π_L . In that paper, we showed that switching to a sufficiently aggressive monetary policy at an appropriate inflation threshold could avoid these unstable trajectories. In contrast, fiscal policy in these circumstances was ineffective.

4. See also McCallum (2002) for an argument that the low-inflation steady state is not stable under learning. In contrast to these findings, Bullard and Cho (2005) show, within a (linearized) New Keynesian model, that there are “escape paths” toward a low nominal interest rate, low inflation rate outcome. Thus, the targeted equilibrium of the monetary authority is locally stable under least-squares learning, but escapes can occur under constant gain learning.

The analysis of Evans and Honkapohja (2005), however, was conducted in a flexible-price model with exogenous output. In Evans, Guse, and Honkapohja (2007), and in other work in progress, we employ a New Keynesian model to reexamine these issues in an economy in which output can deviate from the flexible-price equilibrium.⁵

This approach leads to a number of striking results. The possibility of liquidity traps, that is, net interest rates near zero, combined with deflation and falling output, emerges as a serious concern. Although the targeted steady state is locally stable under learning, a large pessimistic shock to expectations can result, under learning, in a self-reinforcing process in which inflation falls over time, eventually leading to deflation and a declining consumption path. Unstable paths of this type will be referred to as a “deflationary spiral.” We consider a number of policies to insulate the economy from this outcome. Each of these policies maintains the Taylor rule over most of the range but switches to aggressive policies if inflation or output falls below some threshold.

We first consider an inflation threshold policy in which aggressive monetary policy is used whenever inflation falls below, or threatens to fall below, some specified threshold. It turns out that this policy, although it does offer some protection, is not sufficient if the negative expectations shock is very large. Next, we augment the preceding policy by adding aggressive fiscal policy if monetary policy alone is inadequate to keep inflation at or above the threshold. We demonstrate that this combination of aggressive policies with a threshold chosen at a suitable level can always eliminate the possibility of deflationary spirals and ensure global stability of the targeted steady state. This is the central policy finding that emerges from the adaptive learning perspective.

Our central policy result leads to several further questions. One natural question is whether an output threshold could be substituted for an inflation threshold. Surprisingly, the answer is no: using an output threshold to trigger aggressive monetary and fiscal policies will not necessarily avoid deflationary spirals. Another question concerns the timing for implementing our recommended policy in which normal monetary and fiscal policy is augmented by inflation threshold policies. Using simulations, we show that it is better to adopt inflation threshold policies earlier rather than later. Ideally, our inflation threshold policy is in place before substantial negative expectation shocks impact the economy.

II. The Model

We adopt a fairly standard representative agent model along the lines of Benhabib, Schmitt-Grohé, and Uribe (2001b, section 3), except that we allow for stochastic shocks and conduct the analysis in discrete time.⁶ There is a continuum of household-firms units, which produce a differentiated consumption good under conditions of

5. Our analysis provides a theoretical framework for the potential role of fiscal policy in combating liquidity traps, which has been a controversial topic in the empirical literature on Japan's slump. See Ball (2005), Kuttner and Posen (2002), and Perri (2001).

6. We develop our analysis within a closed-economy model. For discussions of liquidity traps in open economies, see, for example, McCallum (2000) and Svensson (2003).

monopolistic competition and price adjustment costs. We allow for both fiscal and monetary policy and for the government to issue debt.

A. Private Sector

The objective for agent j is to maximize expected, discounted utility subject to a standard flow budget constraint:

$$\text{Max } E_0 \sum_{t=0}^{\infty} \beta^t U_{t,j} \left(c_{t,j}, \frac{M_{t-1,j}}{P_t}, h_{t,j}, \frac{P_{t,j}}{P_{t-1,j}} - 1 \right), \quad (1)$$

$$\text{s.t. } c_{t,j} + m_{t,j} + b_{t,j} + \tau_{t,j} = m_{t-1,j} \pi_t^{-1} + R_{t-1} \pi_t^{-1} b_{t-1,j} + \frac{P_{t,j}}{P_t} y_{t,j}, \quad (2)$$

where $c_{t,j}$ is the Dixit-Stiglitz consumption aggregator, $M_{t,j}$ and $m_{t,j}$ denote nominal and real money balances, $h_{t,j}$ is the labor input into production, $b_{t,j}$ denotes real bonds held by the agent at the end of period t , $\tau_{t,j}$ is the lump-sum tax collected by the government, R_{t-1} is the nominal interest rate factor, $P_{t,j}$ is the price of consumption good j , $y_{t,j}$ is output of good j , P_t is the aggregate price level, and the inflation rate is $\pi_t = P_t/P_{t-1}$. The utility function has the parametric form

$$U_{t,j} = \frac{c_{t,j}^{1-\sigma_1}}{1-\sigma_1} + \frac{\chi}{1-\sigma_2} \left(\frac{M_{t-1,j}}{P_t} \right)^{1-\sigma_2} - \frac{h_{t,j}^{1+\varepsilon}}{1+\varepsilon} - \frac{\gamma}{2} \left(\frac{P_{t,j}}{P_{t-1,j}} - 1 \right)^2.$$

The final term parameterizes the cost of adjusting prices in the spirit of Rotemberg (1982).

The production function for good j is

$$y_{t,j} = h_{t,j}^\alpha,$$

where $0 < \alpha < 1$. Output is differentiated, and firms operate under monopolistic competition. Each firm faces a downward-sloping demand curve given by

$$P_{t,j} = \left(\frac{y_{t,j}}{Y_t} \right)^{-1/\nu} P_t.$$

Here $P_{t,j}$ is the profit-maximizing price set by firm j consistent with its production $y_{t,j}$. The parameter ν is the elasticity of substitution between two goods and is assumed to be greater than one.

B. Fiscal and Monetary Policy

The government's budget constraint is

$$b_t + m_t + \tau_t = g_t + m_{t-1} \pi_t^{-1} + R_{t-1} \pi_t^{-1} b_{t-1}, \quad (3)$$

where g_t denotes government consumption of the aggregate good and τ_t is the lump-sum tax collected. We assume that fiscal policy will follow a linear tax rule as in Leeper (1991):

$$\tau_t = \kappa_0 + \kappa b_{t-1} + \eta_t, \quad (4)$$

where η_t is a white noise shock. Provided $\beta^{-1} - 1 < \kappa < 1$, fiscal policy is “passive” in the sense that increases in debt lead to an increase in taxes that is at least sufficient to pay the steady-state interest on the extra debt.⁷ We also assume that g_t is stochastic, with

$$g_t = \bar{g} + u_t, \quad (5)$$

where u_t is an observable exogenous stationary AR(1) mean zero shock. From market clearing, we have

$$c_t = h_t^\alpha - g_t. \quad (6)$$

Monetary policy is assumed to follow a global interest rate rule:

$$R_t - 1 = \theta_t f(\pi_t). \quad (7)$$

The function $f(\pi)$ is taken to be non-negative and non-decreasing, while θ_t is an exogenous, stationary positive AR(1) shock with mean 1, representing random shifts in the behavior of the monetary policymaker.⁸ We assume the existence of π^* , R^* such that $R^* = \beta^{-1} \pi^*$ and $f(\pi^*) = R^* - 1$. π^* can be viewed as the inflation target of the central bank. In the numerical analysis, we will use the functional form

$$f(\pi) = (R^* - 1) \left(\frac{\pi}{\pi^*} \right)^{AR^*/(R^*-1)},$$

which implies the existence of a nonstochastic steady state at π^* . Note that $f'(\pi^*) = AR^*$, which we assume is bigger than β^{-1} . Thus, the “Taylor principle” holds locally at the target steady state π^* and Figure 1 depicts the global monetary policy rule.

Equations (4), (5), and (7), with fiscal policy passive and the Taylor principle satisfied at π^* , constitute what we call “normal policy.” We focus on this policy benchmark because interest rate rules satisfying the Taylor principle at the target inflation rate appear to be a good description of actual monetary policy in many countries, and because this form of monetary policy in combination with passive fiscal policy is widely believed to have desirable properties. As we will see, normal policy does lead to a locally unique solution that is stable under learning.⁹ Our concern is with the global stability properties of the “normal” or “benchmark” policies under learning.

7. If monetary policy obeys the Taylor principle, and fiscal policy is “active” in the terminology of Leeper (1991), that is, $|\beta^{-1} - \kappa| > 1$, then under RE there are no non-explosive solutions.

8. For simplicity, we only include monetary and fiscal random shocks. However, it would be straightforward also to allow, for example, for productivity and taste shocks.

9. Bullard and Mitra (2002) study determinacy and stability under learning in linearized New Keynesian models, for different forms of interest rate rules. Evans and Honkapohja (2007) study determinacy and stability under learning in linearized flexible-price models, for different monetary and fiscal regimes.

In the first part of the paper, we examine the system under normal policy. Later we consider modifications to normal policy when inflation or output falls below some stated threshold.

C. Key Equations

Private-sector optimization yields the key equations

$$-h_t^{1+\varepsilon} + \frac{\alpha\gamma}{\nu}(\pi_t - 1)\pi_t + \alpha\left(1 - \frac{1}{\nu}\right)h_t^\alpha c_t^{-\sigma_1} = \beta\frac{\alpha\gamma}{\nu}E_t[(\pi_{t+1} - 1)\pi_{t+1}], \quad (8)$$

$$c_t^{-\sigma_1} = \beta R_t E_t(\pi_{t+1}^{-1} c_{t+1}^{-\sigma_1}), \quad (9)$$

$$m_t = (\chi\beta)^{1/\sigma_2} \left(\frac{(1 - R_t^{-1})c_t^{-\sigma_1}}{E_t\pi_{t+1}^{\sigma_2-1}} \right)^{-1/\sigma_2}, \quad (10)$$

to which we add the equations (3)–(7).

Consider first the nonstochastic steady states in the absence of random shocks. For any steady state π , equation (9) implies that the nominal interest rate factor satisfies the Fisher equation

$$R = \beta^{-1}\pi. \quad (11)$$

As emphasized by Benhabib, Schmitt-Grohé, and Uribe (2001b), and as discussed above, because $f(\cdot)$ is non-negative, continuous (and differentiable), and has a steady state π^* with $f'(\pi^*) > \beta^{-1}$, there must be a second steady state $\pi_L < \pi^*$ with $f'(\pi_L) < \beta^{-1}$. For our parameterization of $f(\cdot)$, there are no steady states other than the intended steady state π^* and the unintended low-inflation steady state π_L .

The other steady-state equations are given by

$$c = h^\alpha - \bar{g}, \quad (12)$$

$$-h^{1+\varepsilon} + \frac{\alpha\gamma}{\nu}(1 - \beta)(\pi - 1)\pi + \alpha\left(1 - \frac{1}{\nu}\right)h^\alpha c^{-\alpha_1} = 0, \quad (13)$$

and a steady-state version of (10). For a given steady state $\pi \geq 1$, there is a corresponding unique interior steady state $c > 0$ and $h > 0$. For steady states $\pi < 1$, there continue to be unique values for c and h provided π is close to one and $\bar{g} > 0$.

Near each nonstochastic steady state, a corresponding stochastic steady state can be shown to exist provided the support of the exogenous shocks is sufficiently small. Furthermore, each steady state is locally determinate (locally unique), provided the steady state of the corresponding linearized system is determinate. Numerically, these results appear to carry over to the case of large shocks. Based on the linearization there

is an approximate solution, near each steady state, taking the form

$$\begin{pmatrix} c_t \\ \pi_t \end{pmatrix} = \begin{pmatrix} c \\ \pi \end{pmatrix} + \begin{pmatrix} G_{cu} & G_{c\theta} \\ G_{\pi u} & G_{\pi\theta} \end{pmatrix} \begin{pmatrix} u_t \\ \theta_t - 1 \end{pmatrix}, \quad (14)$$

where the coefficients depend on the steady state in question. In Evans, Guse, and Honkapohja (2007), we show the following (for sufficiently small shocks):¹⁰

Under normal monetary and fiscal policy, there are two steady states $\pi^ > \pi_L$. Provided the degree of price stickiness $\gamma > 0$ is sufficiently small, the steady state $\pi = \pi^*$ is locally determinate and the steady state $\pi = \pi_L$ is locally indeterminate.*

Thus, while the π^* steady state is locally unique, the π_L steady state is not. In addition to the stochastic steady state near π_L , there are stochastic paths that converge toward π_L .

III. Learning and Expectational Stability

We now formally introduce learning to the model in place of the hypothesis that RE prevails in all periods. In the current section, we study the system under learning when normal monetary and fiscal policy are in place. We will see that normal policy usually performs well in the sense that the targeted steady state π^* is locally stable under learning: small or even moderate deviations of expectations from the RE values at the intended steady state return over time, under learning, to the rational expectations equilibrium (REE). However, the stability is not global: certain large shocks to expectations lead to unstable trajectories.

In the modeling of learning, it is assumed that private agents make forecasts using a reduced-form econometric model of the relevant variables and that the parameters of this model are estimated using past data. The forecasts are the input to the agents' decision rules and in each period the economy attains a temporary equilibrium, that is, an equilibrium for the current-period variables given the forecasts of the agents. The temporary equilibrium provides a new data point, which in the next period leads to reestimation of the parameters and updating of the forecasts and, in turn, to a new temporary equilibrium. The sequence of temporary equilibria may generate parameter estimates that converge to a fixed point corresponding to an REE for the economy, provided the form of the econometric model that agents use for forecasts is consistent with the REE. When convergence takes place, we say that the REE is stable under learning.

This particular version of bounded rationality, that private agents, when making forecasts, are modeled as econometricians, satisfies the "cognitive consistency" principle that we should model our economic agents as being about as smart as economists. Economists do not know the exact stochastic process followed by the economy. When we need to make forecasts, we do so using estimated models. When new forecasts are required, we update our coefficient estimates and use them to make the forecasts based on our current information set. That is how the adaptive learning approach models how economic agents make forecasts.

10. For proofs of propositions and other derivations, see Evans, Guse, and Honkapohja (2007).

The literature on adaptive learning has shown that there is a close connection between the possible convergence of least-squares learning to an REE and a stability condition, known as E-stability, based on a mapping from the Perceived Law of Motion (which private agents are estimating) to the implied Actual Law of Motion generating the data under these perceptions. E-stability is defined in terms of local stability, at an REE, of a differential equation based on this map. For a general discussion of adaptive learning and the E-stability principle, see Evans and Honkapohja (2001).

For the case at hand, when the exogenous shocks u_t and g_t are stationary AR(1) processes, the appropriate forecast rule based on (14) is for private agents to estimate the linear projections of c_{t+1} and π_{t+1} onto an intercept and the exogenous shocks u_t and θ_t . That is, agents use a version of least squares to estimate

$$c_{t+1} = a_c + du_t + e\theta_t + \varepsilon_{c,t+1},$$

$$\pi_{t+1} = a_\pi + fu_t + g\theta_t + \varepsilon_{\pi,t+1}.$$

The usual timing assumption made in the learning literature is that at the end of period $t-1$, agents estimate the parameters using data on all variables through time $t-1$. This yields estimates $a_{c,t-1}$, d_{t-1} , e_{t-1} , $a_{\pi,t-1}$, f_{t-1} , g_{t-1} . Then, at the start of time t agents form forecasts using these estimates and exogenous data at t ,

$$c_{t+1}^e = a_{c,t-1} + d_{t-1}u_t + e_{t-1}\theta_t,$$

$$\pi_{t+1}^e = a_{\pi,t-1} + f_{t-1}u_t + g_{t-1}\theta_t.$$

Based on these expectations, households and firms determine actual current period values of c_t , π_t . Then, at the end of period t the parameters are updated using the extra data point, and the process continues.

It is now convenient to make a simplification, which does not in any way affect our key theoretical results. It turns out that the stability under learning of the two different steady states is governed by stability of the intercepts, not by the coefficients of the exogenous shocks. We will therefore focus on the case in which the exogenous shocks u_t and θ_t are *i.i.d.* processes. In this case, the RE solutions for π_t and c_t described above are simply noisy steady states, that is, *i.i.d.* processes, and forecasts π_{t+1}^e and c_{t+1}^e will not depend on current values of the exogenous variables u_t and θ_t . This simplifies the presentation of the analysis of learning since it is now natural for private agents to omit these variables from their regressions and forecast by simply estimating the mean values of π_t and c_t . In the learning literature, this is often called “steady-state learning.”

Under steady-state learning, agents treat (14) as a Perceived Law of Motion and for each variable they estimate simply the intercept or mean. We can thus identify expectations of the variables with the estimates of their means, and this has a simple formulation as recursive algorithms:

$$\pi_{t+1}^e = \pi_t^e + \phi_t(\pi_{t-1} - \pi_t^e), \quad (15)$$

$$c_{t+1}^e = c_t^e + \phi_t(c_{t-1} - c_t^e), \quad (16)$$

where ϕ_t is known as the gain sequence. Under least-squares learning, the gain sequence is usually taken to be $\phi_t = t^{-1}$, often termed a “decreasing-gain” sequence, whereas under “discounted least-squares” or “constant-gain” learning it is set to $\phi_t = \phi$, where $0 < \phi < 1$ is a small positive constant. Decreasing gains have the advantage that they can asymptotically converge to RE, while constant-gain learning rules are more robust to structural change.

In what follows, we analyze both theoretically and numerically the model under various specifications of monetary and fiscal policy. The theoretical results for learning are based on E-stability analysis of the system under the learning rules (15)–(16). When we say that an equilibrium π^* or π_L is stable (or unstable) under learning, this implies that it is stable (or not) under these learning rules with decreasing gain, so that $\pi_{t+1}^e \rightarrow \pi^*$ or $\pi_{t+1}^e \rightarrow \pi_L$ (or not) as $t \rightarrow \infty$. In the simulations, we instead use a small constant gain. Under small constant gain, when an equilibrium is E-stable there is local convergence of learning in a weaker sense to a random variable that is centered near and tightly distributed around the equilibrium.¹¹

In studying the economy under learning, we return to the nonlinear model so that we can examine the global dynamics of the system. In doing so, it is convenient to make the additional assumption of point expectations, for example, replacing the expectation of $\pi_{t+1}^{-1} c_{t+1}^{-\sigma_1}$ by $(\pi_{t+1}^e)^{-1} (c_{t+1}^e)^{-\sigma_1}$. This allows us to deal directly with expectations of future consumption and inflation rather than with expectations of nonlinear functions of these quantities, and this is anyway a plausible assumption under bounded rationality. Using also the production function to substitute out h_t leads to the system

$$\begin{aligned} \beta \frac{\alpha \gamma}{\nu} (\pi_{t+1}^e - 1) \pi_{t+1}^e &= -(c_t + g_t)^{(1+\varepsilon)/\alpha} + \frac{\alpha \gamma}{\nu} (\pi_t - 1) \pi_t \\ &\quad + \alpha \left(1 - \frac{1}{\nu}\right) (c_t + g_t) c_t^{-\sigma_1}, \end{aligned} \quad (17)$$

$$c_t = c_{t+1}^e \left(\frac{\pi_{t+1}^e}{\beta R_t} \right)^{\sigma_1}, \quad (18)$$

where $g_t = \bar{g} + u_t$. These equations, together with the interest rate rule (7), implicitly define the temporary equilibrium values for c_t and π_t given values for expectations c_{t+1}^e , π_{t+1}^e and given the exogenous shocks u_t , θ_t . Formally, we write the temporary equilibrium map as

$$\pi_t = F_\pi(\pi_{t+1}^e, c_{t+1}^e, u_t, \theta_t),$$

$$c_t = F_c(\pi_{t+1}^e, c_{t+1}^e, u_t, \theta_t),$$

where it follows from the implicit function theorem that such a map exists in a neighborhood of each steady state (the linearization was given above as [14]).

11. For formal details, see Evans and Honkapohja (2001, section 7.4).

The dynamic system for c_t and π_t under learning is then given by (17)–(18) and (7) together with (15)–(16). The full dynamic system under learning augments these equations with the money equation

$$m_t = (\chi\beta)^{1/\sigma_2} \left(\frac{(1 - R_t^{-1})c_t^{-\sigma_1}}{(\pi_{t+1}^e)^{\sigma_2-1}} \right)^{-1/\sigma_2},$$

and the bond equation (3).

The stability of a steady-state REE under learning is determined by E-stability. The REE is said to be E-stable if the differential equation (in notional time τ)

$$\begin{pmatrix} d\pi^e/d\tau \\ dc^e/d\tau \end{pmatrix} = \begin{pmatrix} T_\pi(\pi^e, c^e) \\ T_c(\pi^e, c^e) \end{pmatrix} - \begin{pmatrix} \pi^e \\ c^e \end{pmatrix}$$

is locally asymptotically stable at a steady state (π, c) , where here

$$T_\pi(\pi^e, c^e) = EF_\pi(\pi^e, c^e, u_t, \theta_t)$$

$$T_c(\pi^e, c^e) = EF_c(\pi^e, c^e, u_t, \theta_t)$$

is the mapping from the Perceived Law of Motion to the corresponding Actual Law of Motion. $T(\cdot)$ gives the actual means for π_t and c_t when private agents have expectations (π^e, c^e) . E-stability is determined by the Jacobian matrix DT of $T = (T_\pi, T_c)'$ at the steady state. We have the following result for low levels of price stickiness (small $\gamma > 0$): *Under normal policy, the steady state $\pi = \pi^*$ is locally stable under learning and the steady state $\pi = \pi_L$ is locally unstable under learning, taking the form of a saddle point.*

The saddle point property of π_L creates a region in which there can be deflationary spirals. We illustrate this using a numerically constructed phase diagram. This also allows us to examine larger $\gamma > 0$ and conduct a global analysis. Parameters are set at $A = 2.5$, $\pi^* = 1.05$, $\beta = 0.96$, $\sigma_1 = 0.95$, $\alpha = 0.75$, $\gamma = 5$, $\nu = 1.5$, $\varepsilon = 1$, and $\bar{g} = 0.1$. Figure 2 shows the E-stability dynamics under normal monetary and fiscal policy. These indicate how, under learning, expectations will on average adjust over time when the economy is perturbed from its steady-state equilibrium. It can be seen that while the π^* REE, indicated by a star, is locally stable, the low steady state $\pi_L \approx 0.969$ is a saddle. π_L is therefore locally unstable under learning.

What might cause deviations of expectations from equilibrium values? Under constant-gain learning, although expectations remain centered in mean around REE values, there are continual deviations as coefficients are updated to recent data that remain subject to random shocks. Provided the system remains in the stability region, learning dynamics will tend to return the system to the stable REE π^* . However, there is always the possibility that a particular, relatively unlikely, sequence of shocks pushes expectations over time sufficiently far from equilibrium to “escape” from the region of local stability. Furthermore, in actual economies, there is also the possibility of a major shock to expectations, arising from unexpected and possibly unmodeled events, that are rightly or wrongly perceived by economic agents to require a substantial revision in


their expectations of the future course of the economy. We have in mind, for example, a large decline in optimism following a substantial decline in equity or other asset prices, which in turn could have been triggered by various precipitating events.

The dashed line in Figure 2 shows the dividing line between the regions of stability and instability. Under learning, normal policy works satisfactorily for moderate-sized perturbations from the targeted steady state π^* : any initial position above and to the right of the dashed line leads to expectation paths, under the mean learning dynamics, that return to the π^* REE. However, there are also starting points that lead to instability. In particular, if an exogenous shock leads to a strong downward revision of expectations, relative to the normal steady state, these pessimistic expectations can generate paths leading to a deflationary spiral. This is illustrated by point A, which is inside the unstable region. At A, both c^e and π^e are lower than the π^* REE values. π^e declines steadily over time, and although c^e initially rises, eventually it too declines over time along the deflationary spiral path indicated in the figure.

Although at point A both c^e and π^e are below the values corresponding to the targeted steady state π^* , it can be seen that a sufficiently large negative shock either to inflation expectations, or to consumer expectations, can put the economy into the unstable region.

The intuition for the instability of the low steady state π_L is as follows. Near π_L we are close enough to the zero net interest rate lower bound, so that a reduction in

Figure 2 π^e and c^e Dynamics under Normal Policy



π_t can only result in a small lowering of R_t . If π_{t+1}^e is slightly below π_L , this must therefore lead to an increase in the real interest rate, to lower c_t through the household Euler equation and to lower π_t through the New Keynesian Phillips curve. This sets in motion downward movements in both c_t and π_t , which are reinforced as they feed into expectations. Of course, along these paths it is likely that something would eventually change, in other words, private agents or policymakers would alter their reactions. We think the most plausible scenario is that policymakers would respond to the deteriorating situation with major changes in policy. The goals of this paper are, first, to show that normal policies, while locally stable, have the potential for instability after major expectational shocks and, second, to propose policies that move the economy out of a deflationary spiral as well as to insulate the economy against these unstable outcomes.

The results of this section indicate the need for more aggressive policies when expectations are overly pessimistic. We begin by considering changing to an aggressive monetary policy when inflation threatens to become too low. As we will see, it may be important also to alter fiscal policy in certain circumstances.

IV. Adding Aggressive Monetary Policy

We first consider modifying monetary policy so that it follows the normal interest rate rule as long as $\pi_t \geq \tilde{\pi}$, but cuts interest rates to a low-level floor \hat{R} if inflation threatens to get below a threshold $\tilde{\pi}$, which we set so that $\pi_L < \tilde{\pi} < \pi^*$. Thus,

$$R_t = \begin{cases} 1 + \theta_t f(\pi_t) & \text{if } \pi_t > \tilde{\pi} \\ \hat{R} & \text{if } \pi_t < \tilde{\pi}, \end{cases}$$

and

$$\hat{R} \leq R_t \leq 1 + \theta_t f(\pi_t) \quad \text{if } \pi_t = \tilde{\pi},$$

where we will think of \hat{R} as very slightly more than 1.¹² The modified interest rate rule is shown in Figure 3.

A policy question of major importance is whether an aggressive monetary policy of this form is sufficient to eliminate the possibility of deflationary spirals arising when expectations are pessimistic. It can be shown that aggressive monetary policy will not always be adequate to avoid these outcomes. We have the following result for policy in which an inflation threshold $\tilde{\pi}$ triggers aggressive monetary policy.¹³


Incorporating aggressive monetary policy triggered by an inflation threshold $\tilde{\pi}$ leads to existence of an additional steady state at $\hat{\pi} = \beta \hat{R}$, which is a saddle point under learning.

We illustrate this point numerically using a phase diagram showing expectational dynamics. We here set $\hat{R} = 1.0001$, so that net nominal interest rates are cut almost

12. In our numerical examples, we set $\hat{R} = 1.0001$. We set \hat{R} above one to keep money demand finite under our parameterization.

13. In Evans, Guse, and Honkapohja (2007) we formally demonstrate this and subsequent results for low degrees of price stickiness, that is, small $\gamma > 0$.

Figure 3 Inflation Threshold $\tilde{\pi}$, $\pi_L < \tilde{\pi} < \pi^*$, for Aggressive Monetary Policy



to zero when aggressive monetary policy is triggered. Figure 4 shows the impact of setting a value $\tilde{\pi} = 1.01 > \pi_L$. Other parameter values are as in the “normal policy” case. Although the threshold policy eliminates the unstable steady state at π_L , the deflationary spiral still exists for sufficiently pessimistic expectations. There are now two steady states: the targeted steady state at π^* , which is locally stable, and a low-level steady state at $\hat{\pi} = \beta \hat{R} < \pi_L$, which is a saddle with nearby deflationary paths.

The conclusion from this analysis is that aggressive monetary policy will not always be sufficient to eliminate deflationary spirals and stagnation.¹⁴ We therefore now take up fiscal policy as a possible additional measure.

14. This contrasts with the findings of Evans and Honkapohja (2005). There we found for the flexible-price case that switching to a sufficiently aggressive money growth rule, at the threshold $\pi_t = \tilde{\pi}$, would render the π^* steady state globally stable.


policy inadequate to the task, then aggressive fiscal policy is deployed. For the policy to successfully eliminate the possibility of deflationary spirals, we need to choose $\tilde{\pi}$ so that $\pi_L < \tilde{\pi} < \pi^*$. We have the following result:

Consider policy that incorporates aggressive monetary and fiscal policy triggered by an inflation threshold $\tilde{\pi}$ with $\pi_L < \tilde{\pi} < \pi^$. Then π^* is the unique steady state.*

We have thus eliminated all steady states other than the one intended by policymakers. How is this possible, in view of Figures 1 and 3, which seem to render inevitable the existence of a second low-inflation steady state? Under the policy recommended in the present section, the interest factor R_t continues to be set according to the rule shown in Figure 3. However, the only part of Figure 3 that is reached are inflation factors $\pi_t \geq \tilde{\pi}$. Because $\tilde{\pi}$ triggers aggressive monetary and fiscal policies, and because these policies can always ensure $\pi_t \geq \tilde{\pi}$, however pessimistic expectations may be, actual inflation rates $\pi_t < \tilde{\pi}$ are no longer realizable. Consequently, low-inflation steady states no longer exist.

Based on earlier results, we know that the stochastic steady state at π^* is locally determinate and locally stable under learning. In fact, numerical computations indicate that the π^* equilibrium is now globally unique and globally stable under learning. Figure 5 illustrates the result. We set $\hat{R} = 1.0001$, so that when aggressive monetary policy is triggered the nominal interest rate is cut almost all the way to the zero lower bound, if required, as discussed in the previous section. (Other parameters are as

Figure 5 Inflation Threshold $\tilde{\pi}$, $\pi_L < \tilde{\pi} < \pi^*$, for Aggressive Monetary Policy and, If Needed, Aggressive Fiscal Policy



before.) In Figure 5, we set $\pi_L < \tilde{\pi} = 1 < \pi^*$. There is now a unique steady state at π^* and it is evident from the figure that it is globally stable.

Our main finding is that a combination of aggressive monetary and fiscal policy to maintain a sufficiently high lower bound on inflation will eliminate the possibility of a deflationary spiral. Choosing $\pi_L < \tilde{\pi} < \pi^*$ eliminates the π_L steady state, but does not create any new ones. The key reason for this is that the inflation threshold $\pi_t \geq \tilde{\pi}$ is achievable by bringing in aggressive fiscal policy, if necessary, to supplement aggressive monetary policy. Having set the policy to ensure this inflation threshold, we simultaneously ensure that the system is restricted to a region in which there are stable learning dynamics.

VI. An Output Threshold for Policy?


The preceding discussion naturally raises the question whether another type of threshold might be used for triggering aggressive policies. Consider in particular the possibility that the policy authorities choose a minimum output threshold, so that policies ensure $c_t + g_t \geq \tilde{y}$ by first dropping interest rates as needed to ensure the threshold, subject to their not falling below the floor \hat{R} . If setting $R = \hat{R}$ is not sufficient to meet the output threshold, then also g_t is raised as required to ensure $y_t = \tilde{y}$. Thus, this policy is analogous to the one recommended in Section V, except that we now have a minimum output threshold instead of an inflation threshold.

Surprisingly, it turns out that this form of policy does not always eliminate deflationary spirals. There is again the possibility of an unintended steady state, which is a saddle under learning. The theoretical details are somewhat complicated, so I will just give a numerical example and some intuition.

Suppose we set the output threshold so that $y_L < \tilde{y} < y^*$. In particular, we set \tilde{y} at 99.5 percent of the high steady-state output (the other parameters are unchanged). In this case, a constrained steady state at $\hat{\pi} = \beta \hat{R}$ exists, which again is locally a saddle under learning. Figure 6 shows that deflationary spirals exist at the bottom-left corner of the phase diagram.

On these deflationary spiral paths, consumption falls steadily after a certain point. Output is then sustained by ever-increasing government spending. The intuition is that in a deflationary spiral, even at a near-zero net nominal interest rate $R_t = \hat{R}$, the *ex ante* real interest rate increases, which depresses private consumption. Simply maintaining output is not enough. To put a floor on consumption, it is critical to put an upper bound on real interest rates, and this can only be done by stabilizing inflation. One might think that stabilizing output at a high-enough level is enough to stabilize π_t , but this is not the case. In the temporary equilibrium Phillips curve (17), π_t depends separately on output $y_t = c_t + g_t$ and on consumption, c_t . In particular, π_t depends negatively on the marginal utility of consumption. Consequently, if $y_t = \tilde{y}$ is maintained by increasing g_t in the face of falling c_t , inflation will continue to fall because households/firms become more willing to reduce prices as the marginal utility of consumption rises.

Figure 6 Learning Dynamics under an Output Threshold with $\tilde{y} > y_L$




VII. Stochastic Simulations

We now illustrate our recommended policy using real-time stochastic simulations. We here assume a constant-gain form of the learning rule with a small gain. Simulations confirm local convergence to the stable targeted steady state under normal policy and global convergence under our recommended policy, in which normal policy is augmented by aggressive monetary and fiscal policy if π_t threatens to fall below a threshold $\tilde{\pi} < \pi_L$.

It is beneficial to have our recommended policies in place before a collapse in expectations. We illustrate how our policies work, in the face of pessimistic expectations, if initially normal policies are used, and then our recommended policies are implemented after some point t_1 . For the simulations, we have chosen $\pi^* = 1.02$, corresponding to an inflation target of 2 percent per year. With an interest rate rule parameter of $A = 1.8$, the low-inflation steady state π_L is approximately $\pi_L = 0.975$, a deflation rate of 2.5 percent per year. Other parameters are close to those used earlier.¹⁵ For the inflation threshold that triggers aggressive monetary and fiscal policy we choose $\tilde{\pi} = 1$, that is, zero net inflation or price stability.

15. Parameters are $A = 1.8$, $\pi^* = 1.02$, $\beta = 0.96$, $\sigma_1 = 0.95$, $\alpha = 0.75$, $\gamma = 5$, $v = 1.5$, $\varepsilon = 1$, $g = 0.1$, $\hat{R} = 1.002$. Other parameters are $\phi = 1/30$, $\sigma_\theta = 0.02$, $\sigma_u = 0.000001$, $\sigma_\psi = 0.000001$, $\sigma_\eta = 0.001$, $\kappa_0 = -0.005$, $\kappa = \beta^{-1} - 1 + 0.15$, and $\chi = 0.0005$.

Figure 7 Dynamics of π and π^e after Pessimistic Expectations Shock

We consider the impact, under real-time learning, of a negative expectations shock. We start in the targeted steady state, with $\pi^* = 1.02$ and $c^* = 0.52864$. Then, at $t = 1$ there is a negative shock to expectations, in which π^e falls to 1.01 and c^e falls to 0.486. This is a substantial fall in consumption expectations, of just over 8 percent, combined with a drop in inflation expectations. The magnitude of these expectation shocks, which we treat as an exogenous pessimistic shift that is not rooted in fundamentals, turns out to be just sufficient to put the economy on a path toward a deflationary spiral under normal policy. We consider the impact if our recommended policy is not implemented until $t_1 = 150$ versus implementation at $t_1 = 80$, and we compare both to the outcomes if recommended policy is initially in place. Figures 7–9 give the results in the form of time paths of π , π^e , c , c^e , R , g , and b .

For $t_1 = 150$, the figures show consumption diverging to low values before the augmented policies are introduced. Inflation is on a steady downward trajectory when only normal policy rules are in place. Introduction of the aggressive policies at t_1 leads to a recovery of inflation and consumption to the targeted steady-state values. It is seen that interest rates fall to the floor level \hat{R} and debt gradually rises under the normal policy regime in which government spending is constant. At time 150, when the augmented policies are introduced, this leads to an increase in government spending and consequently a further substantial increase in debt in a short interval in time. With the new policy, government spending is gradually reduced as expectations of inflation and consumption recover. This also allows debt to return gradually to the steady state. Interest rates also return to normal levels and inflation converges toward π^* .

Figure 8 Dynamics of c and c^e after Pessimistic Expectations Shock





Figure 9 Dynamics of R , g , and b after Pessimistic Expectations Shock



The results for $t_1 = 80$ show that introduction of our policies at an earlier time avoids the worst part of stagnation. Consumption does not fall as much and returns to normal levels much earlier, and the debt level does not rise nearly as much. Finally, if our policies are in place at the time of the expectations shocks, the impact of the shocks is much less severe. In fact, in this case aggressive monetary policy is enough to maintain inflation at $\tilde{\pi} = 1$ in the face of the shocks, and consequently aggressive fiscal policy is never required. These results clearly show that incorporation of an inflation threshold policy can prevent the economy from sliding into a deflationary spiral and can then greatly attenuate the impact of pessimistic expectations shocks.

However, monetary policy alone is not *always* sufficient. Consider the economy with everything the same except that the initial drop in c^e is larger, to 0.47. Simulations show that these shocks are sufficiently large that they cannot be offset by monetary policy even if interest rates are dropped immediately to the floor. Some use of fiscal policy is needed to stabilize prices and achieve $\pi_t \geq \tilde{\pi}$. However, only a modest use of fiscal policy is needed if the threshold policy is in place when the shocks occur. In contrast, waiting to implement our recommended policies leads to lower consumption, and greater use of fiscal policy with a larger (though temporary) buildup of debt.

Since the impact of aggressive monetary policy is limited by the zero lower bound, one might expect that a higher inflation target π^* would lead to a lower likelihood of needing countercyclical fiscal policy. This turns out to be the case. Consider changing the inflation target to $\pi^* = 1.05$ and suppose that the random shocks and other parameters are the same except that π^e falls from 1.05 to 1.03 instead of from 1.03 to 1.01. Suppose the initial drop in c^e is to the same level as before, and we keep the inflation threshold at $\tilde{\pi} = 1$. Simulations show that in this case there is now no need for fiscal policy because there is greater room for aggressive monetary policy. Of course, although a higher π^* provides additional flexibility for monetary policy, this must be set against the greater inefficiency of having a higher steady-state inflation rate.

VIII. Further Discussion and Extensions

Our analysis raises a number of questions, some of which may lead to fruitful extensions. I will briefly discuss several of these points, and also return to the issue of how the learning approach differs from an approach that simply assumes RE.

First, noting the critical role of fiscal policy in stabilizing inflation, one might ask whether we could dispense entirely with aggressive monetary policy and simply resort to aggressive fiscal policy whenever π_t threatens to fall below $\tilde{\pi}$. While the answer is yes, we think our recommended policy is clearly preferable, because there are good reasons to treat monetary policy as the primary tool for countercyclical macroeconomic policy. While we have not included an analysis of the benefits of government spending, it is reasonable to assume that its mean levels have been set to balance costs and benefits. If extensive government spending is used to guarantee the inflation threshold, then it is likely that much of the spending will be wasteful in the sense that private consumption would be more highly valued. We therefore prefer to use fiscal policy as a policy of last resort to ensure the inflation threshold.

Second, does fiscal policy need to take the form of changes in government spending? In our setup we have lump-sum taxes, and under RE Ricardian Equivalence holds. Consequently, it is variations in g_t , not in taxes, that must be used. Continuing with this point, if the variations in g_t are balanced by equal changes in lump-sum taxes, then the temporary debt buildup, which sometimes accompanied our recommended policy, could be avoided. Our tax rule was merely set to ensure some target level of real debt asymptotically. Of course, lump-sum taxes are unrealistic and a useful extension would be to look at a model that includes tax distortions, to make sure that our recommended policy continues to guarantee global stability in this setup. With distortionary taxes, there is an efficiency advantage to tax rate smoothing, so one would again expect a temporary buildup of debt whenever aggressive fiscal policy is required.

How does our approach compare with the policies for avoiding the liquidity trap recommended by Benhabib, Schmitt-Grohé, and Uribe (2002), based on a purely RE approach?¹⁶ Under RE/perfect foresight, the issue of concern is the existence of paths converging to π_L . Benhabib, Schmitt-Grohé, and Uribe (2002) argue that these paths can be eliminated by altering tax policy, if inflation falls to a neighborhood of π_L , so that real lump-sum taxes plus seigniorage are set equal to a negative coefficient times total real government liabilities $b + m$. Along an inflation path converging to π_L , there would then be an explosive increase in total government liabilities, which would lead to a violation of the household transversality condition. Because satisfaction of the transversality condition is a necessary condition for a perfect foresight equilibrium, under this fiscal policy neither the π_L steady state nor paths to π_L are possible equilibrium outcomes. In effect, the government eliminates the liquidity trap paths by threatening to implement unsustainable tax cuts at low inflation rates. Private agents are assumed to recognize that the resulting time path could not be an equilibrium and consequently coordinate instead on the intended equilibrium at π^* . This argument relies heavily, and in my view implausibly, on the perfect foresight assumption.

From the adaptive learning perspective of the current paper, the results of such a policy would be very different. Under the learning dynamics of Section III, the dynamic system in π_t , c_t , π_{t+1}^e , and c_{t+1}^e is autonomous with respect to m_t and b_t . A switch to the tax reduction policy, described in the previous paragraph, when inflation falls below some threshold level, would not avert deflationary spiral paths of the type shown in Figure 2, but would lead to these paths being accompanied by even larger increases in debt.

The aggressive fiscal policy we propose relies on a different fiscal instrument, namely, increases in government spending, and a different and more direct economic mechanism. Government purchases directly affect the demand for goods and hence raise the rate of inflation. Using aggressive fiscal policy, if necessary, to supplement aggressive monetary policy ensures that inflation will not fall below $\tilde{\pi} > \pi_L$. Because, under learning, expectations are grounded in the data, this must eventually lead to $\pi^e \geq \tilde{\pi}$. With interest rates at or near the zero lower bound, the resulting low *ex ante* real interest rates must then lead to $c > c^e$ and hence to a recovery of consumer spending. Eventually

16. Benhabib, Schmitt-Grohé, and Uribe (2002) study nonstochastic continuous-time flexible-price economies under perfect foresight, but the same points could be made in our setup.

this process lifts inflation above $\tilde{\pi}$ and the inherent stable learning dynamics return the economy to the intended steady state at π^* .

Under our proposed policy, the economy also behaves differently under learning than it would under fully RE. With our policy in place, there is a unique steady-state REE. Under RE, autonomous adverse shocks to expectations cannot occur. If they did arise due to some unmodeled disturbance, expectations would return immediately to their RE values. Under adaptive learning, however, this process plays out over time, with expectation coefficients updated in response to actual economic data.

Another issue concerns the potential role of commitment or announcements of future policy changes. The RE literature, on the benefits of commitment to a monetary policy rule, stresses the impact of this commitment on expectations. One might wonder whether this expectational channel is absent when RE is replaced by learning. This is not the case. Commitment to an optimal policy rule can readily be handled within an adaptive learning approach, as discussed by Evans and Honkapohja (2003, 2006) for monetary policy in normal times. Since optimal policy with commitment includes history dependence, this alters the form of the REE. The history dependence introduced into the economy by the central bank will be reflected in the form of the Perceived Law of Motion used by private agents: the list of explanatory variables they use for forecasting would be augmented to include, for example, lagged GDP. If the REE is E-stable, then it is learnable using this augmented forecasting model. The general orientation of the adaptive learning approach is that commitment to a specific policy rule will affect private-agent expectations over time, as the variables and the parameters of the forecasting model adapt statistically to observed outcomes. All the expectational channels that are present under RE are available also under adaptive learning. However, the learning approach extends and qualifies the RE analysis, by examining local and global stability of the REE under learning and by studying the potential for additional learning dynamics.

Similar issues arise in connection with monetary policy when interest rates are at or near the zero lower bound. In the RE analysis of Eggertsson and Woodford (2003) and Braun and Waki (2006), commitment to a future period of zero net interest rates, after the zero lower bound constraint ceases to bind, plays an important role. This introduces a specific form of nonlinear history dependence through which monetary policy affects expectations. In our analysis, we have purposely kept very simple the class of policy rules studied, to provide a systematic analysis of global learning dynamics within the nonlinear structural model. Clearly, it would be of interest to extend our learning analysis to examine more general interest rate rules incorporating various forms of history dependence.

A related issue concerns the planning horizon assumed for our boundedly rational agents. The approach we have adopted here is based on “Euler equation learning,” in which we treat the Euler equations (17)–(18) as the behavioral equations that determine π_t and c_t . This is a valid and convenient approach to modeling bounded rationality, since the Euler equations express necessary first-order conditions for optimum decision-making. As Seppo Honkapohja and I have stressed elsewhere, Euler-equation learning converges to RE equilibria in a variety of contexts, including real business cycle models,

simple overlapping generations models, and New Keynesian models with appropriate interest rate rules. An alternative approach, stressed in Preston (2005), retains adaptive learning, but asks agents each to forecast infinitely far into the future and to re-solve their dynamic optimization problem each period. Frequently, these approaches do not come to significantly different qualitative conclusions concerning stability.¹⁷ Again, it would be of interest to know whether any of our results are affected by the planning horizon of private agents.

A final issue worth pursuing concerns whether the inclusion of assets other than money and bonds would affect the possibility of deflationary spirals or alter our main policy conclusions. Extending our approach to include models with capital would certainly be desirable. In current work in progress, Eran Guse, Seppo Honkapohja, and I have extended our analysis to an open-economy setting in which foreign assets can be accumulated. This provides an additional exchange rate channel for monetary policy, and we are studying the implications of this for alternative policy rules under private-agent learning.

IX. Conclusions

The recent theoretical literature on the zero lower bound to nominal interest rates has emphasized the possibility of multiple equilibria and liquidity traps when monetary policy is conducted using a global Taylor rule. Most of this literature has focused on models with perfect foresight or fully RE. We take these issues very seriously, but our findings for these models under adaptive learning are quite different and in some ways much more alarming than suggested by the RE viewpoint. We have shown that under standard monetary and fiscal policy, the steady-state equilibrium targeted by policymakers is locally stable. In normal times, these policies will appropriately stabilize inflation, consumption, and output. However, the desired steady state is not globally stable under normal policies. A sufficiently large pessimistic shock to expectations can send the economy along an unstable deflationary spiral.

To avoid the possibility of deflation and stagnation, we recommend a combination of aggressive monetary and fiscal policy triggered whenever inflation threatens to fall below an appropriate threshold. Monetary policy should immediately reduce nominal interest rates, as required, even (almost) to the zero net interest rate floor if needed, and this should be augmented by fiscal policy, if necessary, in the form of increased government purchases. Intriguingly, using an aggregate output threshold in the same way will not always successfully reverse a deflationary spiral.

When aggressive fiscal policy is necessary, this will lead to a temporary buildup of government debt. However, government spending and debt will gradually return to their steady-state values. An earlier implementation of the recommended policies will mitigate the use of government spending, and if our recommended policy is already in place at the time of the shocks, the immediate use of aggressive monetary policy can in

17. One situation where the planning horizon is important is when private agents confidently anticipate unique future structural or policy changes that have not yet been implemented. How to treat this within an adaptive learning framework is analyzed in Evans, Honkapohja, and Mitra (2007).

some (but not all) cases entirely avoid the need to use fiscal policy. Raising the inflation target π^* is an alternative way of reducing the likelihood of needing to employ fiscal policy, but this may be undesirable for other reasons.

References

- Adam, K., and R. M. Billi, "Discretionary Monetary Policy and the Zero Lower Bound on Nominal Interest Rates," *Journal of Monetary Economics*, 54, 2007, pp. 728–752.
- Ball, L., "Fiscal Remedies for Japan's Slump," NBER Working Paper No. 11374, National Bureau of Economic Research, 2005.
- Benhabib, J., S. Schmitt-Grohé, and M. Uribe, "Monetary Policy and Multiple Equilibria," *American Economic Review*, 91, 2001a, pp. 167–186.
- , ———, and ———, "The Perils of Taylor Rules," *Journal of Economic Theory*, 96, 2001b, pp. 40–69.
- , ———, and ———, "Avoiding Liquidity Traps," *Journal of Political Economy*, 110, 2002, pp. 535–563.
- Braun, R. A., and Y. Waki, "Monetary Policy during Japan's Lost Decade," *Japanese Economic Review*, 57, 2006, pp. 324–344.
- Bullard, J., and I.-K. Cho, "Escapist Policy Rules," *Journal of Economic Dynamics and Control*, 29, 2005, pp. 1841–1866.
- , and K. Mitra, "Learning about Monetary Policy Rules," *Journal of Monetary Economics*, 49, 2002, pp. 1105–1129.
- Coenen, G., A. Orphanides, and V. Wieland, "Price Stability and Monetary Policy Effectiveness When Nominal Interest Rates Are Bounded at Zero," *Advances in Macroeconomics*, 4 (1), Article 1, 2004.
- Eggertsson, G. B., and M. Woodford, "The Zero Bound on Interest Rates and Optimal Monetary Policy," *Brookings Papers on Economic Activity*, 1, 2003, pp. 139–233.
- , and ———, "Optimal Monetary and Fiscal Policy in a Liquidity Trap," working paper, 2004.
- Evans, G. W., E. Guse, and S. Honkapohja, "Liquidity Traps, Learning and Stagnation," mimeo, 2007 (forthcoming in *European Economic Review*).
- , and S. Honkapohja, *Learning and Expectations in Macroeconomics*, Princeton, New Jersey: Princeton University Press, 2001.
- , and ———, "Adaptive Learning and Monetary Policy Design," *Journal of Money, Credit and Banking*, 35, 2003, pp. 1045–1072.
- , and ———, "Policy Interaction, Expectations and the Liquidity Trap," *Review of Economic Dynamics*, 8, 2005, pp. 303–323.
- , and ———, "Monetary Policy, Expectations and Commitment," *Scandinavian Journal of Economics*, 108, 2006, pp. 15–38.
- , and ———, "Policy Interaction, Learning and the Fiscal Theory of Prices," *Macroeconomic Dynamics*, 11, 2007, pp. 665–690.
- , ———, and K. Mitra, "Anticipated Fiscal Policy and Adaptive Learning," mimeo, 2007.
- Krugman, P. R., "It's Baaack: Japan's Slump and the Return of the Liquidity Trap," *Brookings Papers on Economic Activity*, 2, 1998, pp. 137–205.
- Kuttner, K. N., and A. S. Posen, "Fiscal Policy Effectiveness in Japan," *Journal of the Japanese and International Economies*, 16, 2002, pp. 536–558.
- Leeper, E. M., "Equilibria under 'Active' and 'Passive' Monetary and Fiscal Policies," *Journal of Monetary Economics*, 27, 1991, pp. 129–147.
- McCallum, B. T., "Theoretical Analysis Regarding a Zero Lower Bound on Nominal Interest Rates," *Journal of Money, Credit and Banking*, 32, 2000, pp. 870–904.

- , “Inflation Targeting and the Liquidity Trap,” in N. Loayza and R. Soto, eds. *Inflation Targeting: Design, Performance, Challenges*, Santiago, Chile: Central Bank of Chile, 2002, pp. 395–437.
- Perri, F., “The Role of Fiscal Policy in Japan: A Quantitative Study,” *Japan and the World Economy*, 13, 2001, pp. 387–404.
- Preston, B., “Learning about Monetary Policy Rules When Long-Horizon Expectations Matter,” *International Journal of Central Banking*, 1, 2005, pp. 81–126.
- Rotemberg, J. J., “Sticky Prices in the United States,” *Journal of Political Economy*, 90, 1982, pp. 1187–1211.
- Svensson, L. E., “Escaping from a Liquidity Trap and Deflation: The Foolproof Way and Others,” *Journal of Economic Perspectives*, 17, 2003, pp. 145–166.
- Taylor, J., “Discretion versus Policy Rules in Practice,” *Carnegie-Rochester Conference Series in Public Policy*, 39, 1993, pp. 195–214.

