Policy Interaction, Expectations and the Liquidity Trap*

George W. Evans
University of Oregon

Seppo Honkapohja
University of Helsinki and the Bank of Finland

This version: April 30, 2003

Abstract

We consider inflation and government debt dynamics when monetary policy employs a global interest rate rule and private agents forecast using adaptive learning. Because of the zero lower bound on interest rates, active interest rate rules are known to imply the existence of a second, low inflation steady state, below the target inflation rate. Under adaptive learning dynamics we find the additional possibility of a liquidity trap, in which the economy slips below this low inflation steady state and is driven to an even lower inflation floor that is supported by a switch to an aggressive money supply rule. Fiscal policy alone cannot push the economy out of the liquidity trap. However, raising the threshold at which the money supply rule is employed can dislodge the economy from the liquidity trap and ensure a return to the target equilibrium.

JEL Classification: E63, E52, E58.

*Financial support from the US National Science Foundation, the Academy of Finland, Yrjö Jahnsson Foundation, Bank of Finland and Nokia Group is gratefully acknowledged. The views expressed are those of the authors and do not necessarily reflect the views of the Bank of Finland. We thank In-Koo Cho, Noah Williams and the participants in the Federal Reserve Bank of Atlanta Conference on Monetary Policy and Learning for their comments.

1 Introduction

The possibility of a liquidity trap equilibrium has recently received considerable attention as a possible explanation for recent episodes of low inflation or deflation, close to zero nominal interest rates and low growth, such as seen in Japan over the last ten years. Global analysis of some standard macroeconomic models in which monetary policy is conducted using a nonlinear instrument rule, often called a Taylor rule, to set the nominal interest rate, has shown the possibility of two steady states, including an indeterminate low inflation "liquidity trap" equilibrium π_L in addition to the desired target equilibrium π_H. Moreover, the liquidity trap is not only a theoretical curiosity. (Benhabib, Schmitt-Grohe, and Uribe 2001) and (Benhabib, Schmitt-Grohe, and Uribe 2002) have shown that there are a "large number" of perfect foresight paths that start from initial values near π_H and converge to π_L. These results give a clear warning that a well-meaning regime of monetary policy may lead to undesirable outcomes. The possibility of convergence of the economy to π_L under perfect foresight raises several issues worthy of further study.

The demonstration of the existence of convergent paths to π_L relies heavily on the assumption of perfect foresight in a context involving strongly nonlinear global dynamics, see Section 6 of (Benhabib, Schmitt-Grohe, and Uribe 2001). In such settings the hypothesis of rational expectations is worked very hard as agents must be able to compute these nonlinear convergent paths exactly correctly, i.e. they must have perfect foresight over these paths. It is important to raise the question of whether the conclusion about the possibility of convergence to a "liquidity trap" is robust to a natural weakening of the perfect foresight/rational expectations hypothesis to the alternative assumption that agents have much less information and try instead to learn adaptively the equilibria of the system. In other words, are either π_L, or paths converging to π_L, stable under adaptive learning? In this paper we take up these issues and their ramifications. Because of the complexity of the economy under learning we conduct the analysis using the simplest possible framework, i.e. a flexible price endowment economy.

There have been a couple of previous studies that have analyzed liquidity traps in the context of adaptive learning. Using a linearized model (McCallum 2001b) suggests that the low inflation, low interest rate equilibria are not stable under adaptive learning and thus are not very probable outcomes. (Bullard and Cho 2002) instead view the liquidity trap as a tem-

porary "escape path" from the usual steady state within a linearized model. The deviation is caused by the interaction of large shocks, agents' use of constant gain learning rules and the policy maker's inflation target adapting to inflation expectations.

Our approach differs from these studies in that we examine the global dynamics of learning within a nonlinear model. Agents are assumed to use linear forecasts functions that provide good approximations to the dynamics locally near any steady state. Naturally, these linear forecast functions differ greatly between the different types of steady states. For monetary policy we follow (Benhabib, Schmitt-Grohe, and Uribe 2001) in assuming that it is conducted using a global nonlinear Taylor rule. However, we introduce two major modifications. First, we assume that monetary policy has a "second pillar" taking the form of a money supply rule that supersedes the interest rate rule if inflation reaches a specified floor $\tilde{\pi}$.[1] Second, we explicitly consider the interaction of monetary and fiscal policies. This interaction turns out to be crucial for the stability of the different solutions under learning and for the design of appropriate policies to avoid a liquidity trap.

We will see that while there is the theoretical possibility of paths converging to π_L under learning, the liquidity trap primarily takes the form of inflation slipping below π_L and converging to the floor $\tilde{\pi}$. Appropriate specifications of monetary and fiscal policy can eliminate this threat and even dislodge the economy from the liquidity trap if this has arisen from inappropriate past policies. The appropriate policies will ensure convergence of the economy to the target inflation rate and to stable levels of public debt.

2 The Model

We conduct the analysis in a stochastic representative agent model with perfect competition. For simplicity, we will also postulate an endowment economy in which output is constant and thus liquidity traps are equilibria with very low inflation or even deflation. This model was introduced in (Evans and Honkapohja 2002) and it is closely related to (Leeper 1991) and

[1]Though we use the phrase "two pillars of monetary policy", our model should not be viewed as an attempt to formalise the monetary policy strategy of the European Central Bank (ECB). See Chapter 3 of (European Central Bank 2001) for a description of the monetary policy strategy of the ECB, which emphasizes both analysis based on money and analysis of a broad set of indicators.

(McCallum 2001a).[2]

Households are assumed to maximize the utility function

$$\max E_t \left\{ \sum_{s=t}^{\infty} \beta^{s-t} \left[(1-\sigma_1)^{-1} c_s^{1-\sigma_1} + \mathcal{A}(1-\sigma_2)^{-1} (m_{s-1}\pi_s^{-1})^{1-\sigma_2} \right] \right\}.$$

Here c_s denotes consumption in period s and $m_s = M_s/P_s$, where M_s is the money supply and P_s is the price level at s. Note that real money balances enter utility as $m_{s-1}\pi_s^{-1} = (M_{s-1}/P_{s-1})(P_{s-1}/P_s) = M_{s-1}/P_s$. The household's flow budget constraint is

$$c_s + m_s + b_s + \tau_s = y + m_{s-1}\pi_s^{-1} + R_{s-1}\pi_s^{-1} b_{s-1}, \tag{1}$$

where $b_s = B_s/P_s$, $\pi_s = P_s/P_{s-1}$ is the gross inflation rate, and τ_s is a real lump-sum tax. Note that B_s is the end of period s nominal stock of bonds. R_{s-1} is the gross nominal interest rate on bonds, set at time $s-1$ but paid in the beginning of period s. The household has a constant endowment y of consumer goods each period.

We assume that there is a constant flow of government purchases $g \geq 0$. As shown in (Evans and Honkapohja 2002), household optimality and market clearing conditions imply the Fisher equation

$$R_t^{-1} = \beta E_t \pi_{t+1}^{-1} \tag{2}$$

and the equation for money market equilibrium, in period t,

$$\mathcal{A}\beta m_t^{-\sigma_2} E_t \pi_{t+1}^{\sigma_2 - 1} = (y-g)^{-\sigma_1}(1 - \beta E_t \pi_{t+1}^{-1}). \tag{3}$$

In addition, the equilibrium must satisfy the transversality conditions

$$\lim_{t\to\infty} \beta^t m_{t+1} = 0 \text{ and } \lim_{t\to\infty} \beta^t b_{t+1} = 0. \tag{4}$$

The above equations (2) and (3) are usually derived under rational expectations (RE), but, as discussed below, we can also treat them as holding in a temporary equilibrium for given subjective expectations.[3] To simplify the

[2]For the basic model and specification of fiscal policy rules we follow Leeper, but we use McCallum's more general class of utility functions and also his timing in which utility depends on beginning of period money balances.

[3]This assumption means that the Euler equations for household optimality are taken to describe the behavior rules of the household. We have argued elsewhere that this is one reasonable way to model bounded rationality. An alternative would be to use optimality conditions over an infinite horizon as the behavioral rule, see (Sargent 1993), pp.122-125 and (Preston 2003) for the latter approach.

analysis we assume point expectations π^e_{t+1}, so that the Fisher equation (2) and the demand for money (3) can be written as a function of the nominal (gross) interest rate:

$$R_t = \beta^{-1}\pi^e_{t+1}, \tag{5}$$

$$m_t = m(R_t) \equiv (\mathcal{A}\beta)^{1/\sigma_2}(y-g)^{\sigma_1/\sigma_2}[(1-R_t^{-1})(\beta R_t)^{1-\sigma_2}]^{-1/\sigma_2}. \tag{6}$$

The specification of the model is completed by giving the government budget constraint and policy rules. The government budget constraint, written in real terms, is

$$b_t + m_t + \tau_t = g + m_{t-1}\pi_t^{-1} + R_{t-1}\pi_t^{-1}b_{t-1}. \tag{7}$$

For fiscal policy we use the linear tax rule as in (Leeper 1991):

$$\tau_t = \gamma_0 + \gamma b_{t-1} + \psi_t + \varepsilon_t, \tag{8}$$

where ψ_t is an observed exogenous random shock and we have also introduced an unobserved shock ε_t. For simplicity both shocks are assumed *iid* with mean zero. We will make the natural assumption that $0 \leq \gamma \leq \beta^{-1}$ and introduce the definition of active and passive fiscal policy.

Definition 1 *Fiscal policy is said to be "active" (AF) if* $\beta^{-1} - 1 > \gamma$ *and "passive" (PF) if* $\beta^{-1} - 1 < \gamma$.

This follows the terminology of (Leeper 1991), and was also adopted in (Evans and Honkapohja 2002).

The key novel feature in the model is the specification of monetary policy in terms of a global interest rate rule

$$R_t - 1 = \theta_t f(\pi_t). \tag{9}$$

Here $f(\pi)$ is assumed to be a non-negative and non-decreasing function, while θ_t is an exogenous, *iid* and positive random shock with mean 1. 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. As first noted by (Benhabib, Schmitt-Grohe, and Uribe 2001), whenever $f(.)$ is continuous (and differentiable) and has a steady state π_H with $f'(\pi_H) > 1$, in accordance with the Taylor principle, non-negativity of the (net) nominal interest rate

implies the existence of a second low inflation steady state π_L with $f'(\pi_L) < 1$. 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 $\pi_H = \pi^*$. Note that $f'(\pi^*) = AR^*$, which we assume is bigger than 1.

We also assume that the interest rate rule (9) is applied only as long as inflation remains within some specified upper and lower bounds, denoted by $\hat{\pi}$ and $\tilde{\pi}$, respectively. Such bounds can be imposed if the Central Bank switches to a money supply rule if inflation becomes too low or too high. If inflation is at the floor level $\pi_t = \tilde{\pi}$, then using the money demand (6) we get

$$M_t = P_{t-1}\tilde{\pi}m(\beta^{-1}\pi^e_{t+1}). \tag{10}$$

If inflation expectations are observable, (10) can be used to ensure that inflation does not get below the lower bound $\tilde{\pi}$.[4] For any given P_{t-1} and π^e_{t+1} the floor inflation rate $\tilde{\pi}$ can be attained by expanding money supply to the level given by (10). Incorporating this "second pillar" of monetary policy, we obtain the policy relationship

$$\pi_t = \min[\max(f^{-1}((R_t - 1)/\theta_t), \tilde{\pi}), \hat{\pi}]. \tag{11}$$

Figure 1 illustrates the interest rate rule, in the absence of the random shock θ_t, together with the Fisher equation (5). When $\tilde{\pi} < \pi_L$, the lower bound $\tilde{\pi}$ constitutes a new boundary steady state for inflation and real balances. This can be seen as follows. From Figure 1 we see that $\pi^e_{t+1} = \tilde{\pi}$ implies $f^{-1}(\beta^{-1}\tilde{\pi} - 1) < \tilde{\pi}$, which would lead to further reduction in π if the constant inflation floor $\tilde{\pi}$ were not imposed.[5] There are thus three steady states in the model, provided fiscal policy is set so that the process for real bonds b_t is stationary at these points.

FIGURE 1 HERE

Near an interior steady state we can derive a linear approximation of (11), which can be written as $R_t = \alpha_{0,i} + \alpha_i\pi_t + \delta_i\theta_t$, $i = L, H$, with $\alpha_i = f'(\pi_i)$

[4] A similar argument can be applied at the upper bound $\hat{\pi}$ (not shown in Figure 1).

[5] At the upper boundary $\hat{\pi}$ we have $f^{-1}(\beta^{-1}\hat{\pi} - 1) < \hat{\pi}$, which implies a "permissible" reduction in inflation. Thus $\hat{\pi}$ does not constitute a boundary steady state.

and $\delta_i = f(\pi_i)$. Locally near a steady state we thus have the linearization of the model

$$\begin{aligned} \pi_t &= (\alpha_i\beta)^{-1}\pi^e_{t+1} - \alpha_i^{-1}\theta_t + k_{i,1} && (12)\\ 0 &= b_t + \varphi_{i,1}\pi_t + \varphi_{i,2}\pi_{t-1} - (\beta^{-1} - \gamma)b_{t-1} && (13)\\ &\quad +\psi_t + \varepsilon_t + \varphi_{i,3}\theta_t + \varphi_{i,4}\theta_{t-1} + k_{i,2}, \end{aligned}$$

where the coefficients α_i, $\varphi_{i,1}, \ldots, \varphi_{i,4}$ and the intercepts $k_{i,1}, k_{i,2}$ are specific to the steady state i. The formal details, including the expressions for the coefficients $\varphi_{i,j}$, are given in the Appendix of (Evans and Honkapohja 2002). We can then introduce a modified form of the terminology suggested by (Leeper 1991).

Definition 2 *Monetary policy is locally active (LAM) at steady state $i = L, H$ if $|\alpha_i\beta| > 1$ and locally passive (LPM) if $|\alpha_i\beta| < 1$.*

At the boundary steady state $\tilde{\pi}$ the linear approximation (12)-(13) does not exist, but it can be thought of as a limiting case where $f'(\tilde{\pi}) \approx +\infty$ (i.e. $\alpha^{-1} \approx 0$). By this criterion monetary policy is locally active at π_H and $\tilde{\pi}$ and policy is locally passive at π_L in Figure 1.

Using the linearization (12)-(13) and the definitions of active and passive fiscal and monetary policy, we have the following results on local uniqueness of stationary rational expectations equilibria (REE):

Proposition 3 *(i) The linearization (12)-(13) has a locally unique stationary REE near the high steady state π_H when fiscal policy is passive, i.e. PF prevails.*
(ii) The linearization (12)-(13) has a locally unique stationary REE near the low interior steady state π_L when fiscal policy is active, i.e. AF prevails.
(iii) The low boundary steady state $\tilde{\pi}$ is a locally unique stationary REE under PF provided the support of θ_t is sufficiently small.

Parts (i) and (ii) are a consequence of the results proved in (Evans and Honkapohja 2002). In that paper it is shown that the linearization yields a locally unique stationary REE if monetary policy is (locally) active and fiscal policy is passive, i.e. under LAM/PF, or if monetary policy is (locally) passive and fiscal policy is active, i.e. under LPM/AF.

Part (iii) is established as follows. For π^e_{t+1} close to $\tilde{\pi}$ and θ_t close to one we have $f^{-1}((\beta^{-1}\pi^e_{t+1} - 1)/\theta_t) < \tilde{\pi}$ and hence $\pi_t = \tilde{\pi}$ for all t. Thus under RE $\pi^e_{t+1} = \tilde{\pi}$ and the unique REE is

$$\begin{aligned} \pi_t &= \tilde{\pi},\ R_t = \beta^{-1}\tilde{\pi} \\ 0 &= b_t - (\beta^{-1} - \gamma)b_{t-1} + \psi_t + \varepsilon_t + \text{constant}, \end{aligned}$$

which is stationary under PF.

It can also be shown that with (LPM/PF) there is local indeterminacy of REE, i.e. locally there are multiple stationary equilibria, and thus π_L is an indeterminate steady state. However, it can be shown that this case involves instability under learning and is thus not of interest under our approach, see (Evans and Honkapohja 2002).

With (LAM/AF) the system is locally explosive. In this case it is possible that inflation and real balances remain stationary while the stock of real bonds grows in an explosive fashion. These rational "Euler paths" satisfy all of the conditions for an equilibrium, except for the transversality condition on bonds. To side-step this issue we will assume that the government sets upper and lower bounds on government debt, achieved by changing the tax rule at these debt thresholds. Effectively, this would convert fiscal policy to become passive at these bounds.[6] For simplicity, we do not explicitly incorporate this feature into the analysis. However, we will comment later on cases in which these thresholds might be reached.

3 Learning: Introduction

We now formally introduce learning to the model of Section 2 in place of the hypothesis that RE prevails in all periods. 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 input to agent's decision rules and in each period the economy attains a temporary equilibrium, i.e. an equilibrium for the current period variables given the forecasts of the

[6] Thus, strictly speaking, the tax rules (8) should be classified as "locally passive" or "locally active" fiscal policy. We are, of course, assuming that the bounds are set at levels that do not constrain the steady state levels of debt implied by (8). In particular, the upper debt threshold can be set arbitrarily high.

agents. The temporary equilibrium provides a data point, which in the next period leads to re-estimation 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 the convergence takes place, we say that the REE is stable under learning.[7]

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 (that 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), and for detailed theoretical analysis of the linearized model studied in this paper, see (Evans and Honkapohja 2002).

If there are multiple REE, the nature of the perceived law of motion used by the agents in forecasting, i.e. their econometric model, can in some cases determine which types of REE can be outcomes of the learning process. The simplest case is learning of stochastic steady states. In this case agents think that the economy is near a steady state, and they try to estimate the (constant) mean value of inflation, which they use to forecast future inflation. Another possibility is that agents believe that the process for the endogenous variables takes a more complex form, for example a VAR process. We next discuss these two formulations.

3.1 Steady State Learning

Formally, the temporary equilibrium of the economy is given by the following equations: the demand for real balances (6), the Fisher equation (5), the interest rate rule incorporating the inflation bounds (11), and the flow budget constraint for the government defined by (7) and (8). Given inflation expectations, and the exogenous and predetermined variables, these equa-

[7]Stability under learning is often used as a selection criterion between possible REE. In other words, an REE is "reasonable" if it is a stable outcome of a learning process just outlined.

tions jointly determine the endogenous variables.[8] In particular, by (5), the interest rate R_t depends on inflation expectations π^e_{t+1}. Our next step is to formulate how these expectations are formed. We begin by examining steady state learning.

For the linearized system (12)-(13) there are REE near both π_H and π_L in which π_t is *iid*. These REE take the form

$$\begin{aligned} \pi_t &= \pi_i - \alpha_i^{-1}\theta_t,\ R_t = \beta^{-1}\pi_i \\ 0 &= b_t - (\beta^{-1} - \gamma)b_{t-1} + \psi_t + \varepsilon_t + \text{constant}, \end{aligned}$$

for $i = L, H$. In addition there is the low inflation boundary steady state REE $\pi_t = \tilde{\pi}$ described in the previous section. Note that in each case the bond process is nonexplosive if and only if fiscal policy is passive.

For steady state learning the agents are assumed to treat inflation as an *iid* process with an unknown mean, which they try to estimate by least squares, i.e. by computing the (possibly weighted) sample mean from past data. Agents then forecast that inflation in the next period is the estimated value of the steady state. The evolution of forecasts π^f_{t+1} is formally determined by the recursive algorithm

$$\pi^f_{t+1} = \pi^f_t + \phi_t(\pi_{t-1} - \pi^f_t), \tag{14}$$

where ϕ_t is known as the gain parameter. Two possibilities for ϕ_t are commonly used in the literature and we discuss them below. We will assume that the forecasts determined by (14) are subject to additional (white noise) expectation shocks and account is also taken of the bounds $\hat{\pi}$ and $\tilde{\pi}$, so that actual expectations are determined by

$$\pi^e_{t+1} = \min(\max(\pi^f_{t+1} + \eta_t, \tilde{\pi}), \hat{\pi}).$$

In the theoretical analysis we ignore the expectation shock η_t, as it can be shown that, if sufficiently small, such shocks do not affect the local stability properties of the equilibria under learning.

3.1.1 Learning under Decreasing Gain

The first case we consider is that agents might be computing (possibly weighted) averages of the past data in which case ϕ_t would be a decreasing

[8] See (Evans and Honkapohja 2002) for further discussion.

sequence of positive weights such that $\sum_{t=1}^{\infty} \phi_t = \infty$. Under this assumption the dynamics of inflation and inflation expectations can be written formally as a stochastic recursive algorithm (SRA) and, under certain conditions, inflation expectations will converge to a constant value $\bar{\pi}$ with actual inflation following the stochastic steady state process

$$\pi_t = \min[\max(f^{-1}(\beta^{-1}\bar{\pi} - 1)/\theta_t), \tilde{\pi}), \hat{\pi}], \; E\pi_t = \bar{\pi}. \tag{15}$$

In (15) the latter requirement states that the (unconditional) mean of π_t is equal to the forecast $\bar{\pi}$ of the agents.

Note that, if the min and max in (15) are not binding, we have $E\pi_t = Ef^{-1}(\beta^{-1}\bar{\pi} - 1)/\theta_t)$. Moreover, if the support of the shock θ_t is small, the stochastic steady states will be approximately equal to the non-stochastic steady states that satisfy the equation $\pi = f^{-1}(\beta^{-1}\pi - 1)$. We assume that, corresponding to each non-stochastic steady state π', there exists a unique $\bar{\pi}$ in a neighborhood of π' satisfying (15). Note that, due to the nonlinearities, the mean of a stochastic steady state $\bar{\pi}$ is not in general equal to a nonstochastic steady state. For convenience, we will nevertheless refer to the corresponding nearby non-stochastic steady state when we discuss the theoretical and simulation results below.

The derivation of the conditions for convergence of the learning rule (14) under decreasing gain can be studied using standard techniques for SRAs. In particular, convergence of adaptive learning to an REE is governed by E-stability conditions for the REE. Moreover, provided the range of variation (support) of θ_t is sufficiently small, the E-stability conditions for the corresponding nonstochastic steady state will determine the stability under learning of the stochastic steady state process (15).[9] We will derive the relevant E-stability condition below.

3.1.2 Constant Gain Learning

Another natural formulation of learning is to assume a constant gain, i.e. $\phi_t = \phi$, where ϕ is a small positive constant in (14).[10] In this case (14) becomes a time-autonomous stochastic difference equation, and the parameter

[9] See Chapters 11 and 12 of (Evans and Honkapohja 2001) and (Evans and Honkapohja 1995) for analysis of learning of steady states in stochastic models.

[10] For learning of steady states constant gain is formally the same as classic adaptive expectations. This feature does not, however, hold for richer settings in which the estimated parameters are coefficients, e.g., in a regression.

estimate π^f_{t+1} of the steady state mean no longer converges to a constant value. Instead, for a small gain the estimate can converge to a random variable such that the mean of this random variable is approximately equal to the nonstochastic steady state. E-stability of the steady state is necessary for this convergence to take place. See Chapter 14 of (Evans and Honkapohja 2001) for an introduction to learning under constant gain.

E-stable steady states in the nonstochastic model thus provide guidance to the possible convergence properties of constant gain learning. The economy often spends considerable periods of time in a neighborhood of an E-stable steady state. However, the random fluctuations under constant gain learning can have rich patterns of dynamics. The economy may, for example, occasionally experience relatively sudden deviations that move the economy far away from the neighborhood of a steady state. These are called escape dynamics. Along such escape routes the economy moves far away from a neighborhood of an E-stable steady state.[11] In such an event the economy may possibly settle in a neighborhood of another E-stable steady state (if a second E-stable steady state exists) for a period of time after which it may move back, along a new escape path, to a neighborhood of the former E-stable equilibrium. Simulations of such escape dynamics are shown in (Evans and Honkapohja 1993) and Chapter 14 of (Evans and Honkapohja 2001) for an overlapping generations model with multiple E-stable steady states, by (Kasa 2002) for a model of currency crises and by (Williams 2002b) in game theoretic settings. We remark that this is not the only case in which escape dynamics can occur. In some models escape routes can also exist when the equilibrium is unique, see (Sargent 1999), (Cho, Williams, and Sargent 2002), (Bullard and Cho 2002), (Williams 2002a) and (Cho and Kasa 2002). Later we will examine the present model for escape dynamics.

3.2 VAR Learning

Instead of employing steady state learning, the agents may view the economy as following a more complex stochastic process. In section 2 we saw that locally near a steady state the model can be linearized as shown by equations (12)-(13). Introducing the notation $y_t = (\pi_t, b_t)'$, $v_t = (\theta_t, \psi_t)'$, we write the

[11] The terminology is due to (Sargent 1999). Theoretical analysis of escape routes is developed in (Cho, Williams, and Sargent 2002) and (Williams 2002a).

linearized model near steady state i in vector form as

$$y_t = \mathcal{K}_i + \mathcal{M}_i y^e_{t+1} + \mathcal{N}_i y_{t-1} + \mathcal{P}_i v_t + \mathcal{R}_i v_{t-1} + \mathcal{S}_i \varepsilon_t, \tag{16}$$

where

$$\mathcal{M}_i = \begin{pmatrix} (\alpha_i\beta)^{-1} & 0 \\ -\varphi_{i,1}(\alpha_i\beta)^{-1} & 0 \end{pmatrix}, \mathcal{N}_i = \begin{pmatrix} 0 & 0 \\ -\varphi_{i,2} & \beta^{-1} - \gamma \end{pmatrix},$$

$$\mathcal{P}_i = \begin{pmatrix} -\alpha_i^{-1} & 0 \\ \varphi_{i,1}\alpha_i^{-1} - \varphi_{i,3} & -1 \end{pmatrix}, \mathcal{R}_i = \begin{pmatrix} 0 & 0 \\ -\varphi_{i,4} & 0 \end{pmatrix}.$$

Near a steady state i the linearized model (16) has a unique stationary REE of the form

$$y_t = A_i + B_i y_{t-1} + C_i v_t + D_i v_{t-1} + F_i \varepsilon_t \tag{17}$$

in the two cases LAM/PF (with $i = H$) and LPM/AF (with $i = L$), see Proposition 3 above. We remark that in the LAM/PF case the first rows of B_H and D_H are zero, so that in this case the REE solution reduces to one in which inflation is a noisy steady state. In contrast, for the LPM/AF case both rows of B_L are nonzero. See (Evans and Honkapohja 2002) for details.

To study the stability of these REE under learning, we assume that agents think that the stochastic process for inflation has the form (17), where the parameters A_i, B_i, C_i and D_i are to be estimated from past data. We are here assuming that the exogenous shocks $v_t = (\theta_t, \psi_t)'$ are observable, so that agents are estimating and updating the coefficients of a VAR with exogenous variables v_t and v_{t-1}. For this estimation agents might use either recursive least squares or stochastic gradient learning.[12] For computational simplicity, we used stochastic gradient algorithms in the simulations of VAR learning. In both cases, there are also decreasing and constant gain versions of learning and the general remarks made in connection with the same two versions of steady state learning also apply here. In computing the expectations agents make forecasts using the perceived law of motion (17) and, in simulations, the actual expectations π^e_{t+1} take account of both the bounds on inflation and possible expectation shocks.

The system under VAR learning can again be written as a stochastic recursive algorithm and analyzed using standard techniques. As discussed above, it can in general be shown that parameter estimates under least

[12] See (Evans and Honkapohja 2001) for formal details on these two algorithms.

squares learning with decreasing gain converge to an REE if and only if the REE is E-stable, and E-stability tends also to govern stability under stochastic gradient learning.[13] If constant gain is used instead, the parameter estimates do not converge to a fixed point but instead remain random even asymptotically. For a small gain the mean of the parameter estimates is approximately equal to the REE values, but escape dynamics can occur with VAR learning under constant gain. We will examine this possibility below.

4 Theoretical Stability Results

4.1 Steady State Learning

The relevant E-stability condition can be obtained as follows. Suppose that $\pi^e_{t+1} = \pi$, for some constant π, and set the shock θ_t equal to its mean $E\theta_t = 1$. Then the temporary equilibrium value of π_t (in the corresponding nonstochastic model) is given by

$$T(\pi) = \min[\max(f^{-1}(\beta^{-1}\pi - 1)), \tilde{\pi}), \hat{\pi}].$$

It can be shown that the local asymptotic stability of the ordinary differential equation

$$\frac{d\pi}{du} = T(\pi) - \pi$$

provides the relevant E-stability criterion for the stochastic model, under steady state learning, when the shock is small, i.e. θ_t has small bounded support around its mean. Here u denotes notional time. For π_H, π_L and $\tilde{\pi}$ we have:

Proposition 4 *Assume that $\tilde{\pi} < \pi_L < \pi_H < \hat{\pi}$ and that the support of θ_t is sufficiently small. Under steady state learning we have:*
(i) π_H and $\tilde{\pi}$ are E-stable, and
(ii) π_L is not E-stable.

[13] Stability under stochastic gradient learning can be sensitive to details of the algorithm, and there exist models in which stability under stochastic gradient learning is not in all cases governed by E-stability. However, our simulations of the current model appear fully consistent with the predictions of the E-stability principle.

Proof. Since π_H and π_L are interior and satisfy $f'(\pi_H) > \beta^{-1}$, $f'(\pi_L) < \beta^{-1}$ we get

$$T'(\pi_H) = \beta^{-1}\frac{1}{f'(\pi_H)} < 1 \text{ and } T'(\pi_L) = \beta^{-1}\frac{1}{f'(\pi_L)} > 1.$$

Thus π_H is and π_L is not E-stable. On the lower boundary $\exists\breve{\pi}$ defined by $1 + f(\tilde{\pi}) = \beta^{-1}\breve{\pi}$ such that $T(\pi) = \tilde{\pi}$ for $\pi \leq \breve{\pi}$, which implies that $\tilde{\pi}$ is E-stable. ■

Figure 2 illustrates the mapping $T(\pi)$ and the different steady states π_H, π_L and $\tilde{\pi}$. We emphasize once more the key result that E-stability properties of the nonstochastic steady states determine E-stability and hence stability under learning of the stochastic steady state when the range of variation of the shock θ_t is small. This feature will be crucial in interpreting some of the simulations below.

FIGURE 2 HERE

A notable feature of steady state learning in this model is that the evolution of government debt, i.e. real bonds does not influence the dynamics of inflation. If steady state learning of inflation is convergent, then in the long run the evolution of bonds is approximately determined by equations (7) and (8) with $m_t = m(\bar{R})$, $\pi_t \approx \bar{\pi}$, $\bar{R} = \beta^{-1}\bar{\pi}$ and $b_t = (\beta^{-1} - \gamma)b_{t-1} + \psi_t + \varepsilon_t +$ constant. Thus the process for real bonds is stationary if fiscal policy is passive. In the case of steady state learning and active fiscal policy, π_H or $\tilde{\pi}$ remain stable, but the bond path becomes explosive until it reaches the bounds set on government debt.

4.2 VAR Learning

In this case agents are assumed to have a linear PLM of the form (17) and we can consider the E-stability of the REE that are local to the steady states and for which the dynamics are given approximately by the linearized model (16). It is possible to adapt the results in (Evans and Honkapohja 2002), which lead to the following proposition:

Proposition 5 *Under VAR learning the linearized REE that are local to a steady state have the following E-stability properties:*
(i) The local REE associated with π_H is E-stable when fiscal policy is passive.
(ii) The local REE associated with π_L is E-stable (resp. E-unstable) when fiscal policy is active (resp. passive).

Proof. (i) Monetary policy is locally active near π_H, so that by Proposition 4 in (Evans and Honkapohja 2002) the REE is E-stable with PF.

(ii) Monetary policy is locally passive near π_L, and the results follow from Proposition 4 in (Evans and Honkapohja 2002). ■

The results of Proposition 5 indicate that the nature of fiscal policy is crucial for the possibility of a liquidity trap at π_L and provide important conclusions about the reasonableness of the REE considered in the literature.

First, with PF the usual solution is the E-stable REE near π_H. If fiscal policy becomes active at π_H, then the situation is less clear as the system under RE is locally explosive and analysis based on linearization does not yield a full answer. (Evans and Honkapohja 2002) study incipient tendencies for the explosive case and suggest that, depending on monetary and fiscal policy parameters there can be solutions that are stable under learning.[14]

Second, if fiscal policy is passive the solutions near π_L are not stable under learning.[15] (Benhabib, Schmitt-Grohe, and Uribe 2001) assumed PF and suggested that the indeterminacy of π_L indicates the possibility of perfect foresight paths converging to π_L. Part (ii) of Proposition 5 shows that these paths are not interesting if one adopts the learning viewpoint. However, there exists an E-stable REE near π_L if fiscal policy is active. This raises the question of whether convergence to π_L can arise from initial points near π_H when fiscal policy shifts from PF to AF regime.

Third, we remark that simulations below will indicate that the low boundary stochastic steady state $\tilde{\pi}$ is stable under VAR learning under PF.

5 Numerical Analysis

In this section we will present several numerical simulations illustrating the preceding theoretical results about convergence of learning to the different types of equilibria. We will also examine further issues for which analytical answers are not available. For the most part, we will present only simulations with constant gain algorithms. For a learnable equilibrium we should then anticipate that, for the most part, the economy fluctuates near the equilib-

[14] A full study of learning in this explosive case has not yet been completed.

[15] (Eusepi 2002) has recently studied the implications of forward-looking global Taylor rules. Under such a rule learnable cycles and sunspots can exist even if fiscal policy is passive. Eusepi assumes that either (i) money and consumption are complements in the utility function or (ii) real balances affect the production function.

rium but, as a result of specific sequences of random shocks, with occasional escape paths that move the economy further away from the equilibrium and possibly to a neighborhood of another equilibrium whenever a second learnable equilibrium exists.

We specify the following baseline numerical values for the parameters (we will report new values only if they deviate from these):
(i) utility function and output: $\beta = 0.95$, $\sigma_1 = \sigma_2 = 0.95$, $\mathcal{A} = 0.1$, $y = 10$;
(b) fiscal policy: $g = 1.5$, $\gamma_0 = 0.5$, $\gamma = \beta^{-1} - 1 + 0.15$;
(c) interest rate rule: $\pi^* = 1.1$, $A = 1.2$;
(d) inflation bounds: $\tilde{\pi} = 1$, $\hat{\pi} = 2\pi^*$.
The shocks η_t and ψ_t are assumed to be normal with standard deviations $\sigma_\eta = 0.02$ and $\sigma_\psi = 0.01$. θ_t is assumed to be log-normal with $\sigma_\theta = 0.1$ for the corresponding normal variate. The mean of θ_t is set at one.

These parameter values are relatively unsurprising, though we have not tried to do any calibrations to data. We remark that these values imply the existence of a low inflation steady state at $\pi_L = 1.0477$. Monetary policy is locally active at $\pi_H = \pi^* = 1.1$ and it is locally passive at π_L. (The value of π^* is chosen purely for convenient presentation of the numerical results). Fiscal policy is passive under the baseline parameter values since $\gamma = \beta^{-1} - 1 + 0.15$ satisfies the definition of PF, compare Definition 1 in Section 2. We will vary fiscal policy from PF regime to AF regime or vice versa in some simulations.

5.1 Convergence to Equilibria

In the simulation shown in Figure 3 we use the basic parameter settings given above and we assume that agents do steady state learning with the constant gain parameter set at $\phi = 1/10$. The initial values for the economy are assumed to be near the desired steady state π_H. The dynamics were run for 20000 periods. The three panels show the rate of inflation, the quantity of bonds and expectations of inflation, respectively.

FIGURE 3 HERE

The simulations confirms the stability results for π_H and $\tilde{\pi}$ above. With constant gain learning, the economy remains near the high steady state for much of the time but occasionally moves along an escape path to the vicinity of the low boundary steady state $\tilde{\pi}$, which is also E-stable in the PF regime

with steady state learning. Having stayed near $\tilde{\pi}$ for a period of time, the economy eventually moves to the vicinity of π_H along a different escape path. The economy continues to occasionally shift between these regions.

We make some further remarks. First, simulations of learning with decreasing gain (not shown) indicate local stability of π_H and $\tilde{\pi}$ under both steady state and VAR learning in the PF regime. Second, under AF, simulations (not shown) indicate local convergence of steady state learning inflation to π_H or $\tilde{\pi}$, with explosive Euler paths for bonds. However, under VAR learning there is divergence from both π_H and $\tilde{\pi}$.

The low interior steady state π_L was examined for learning in the AF regime for fiscal policy. Propositions 4 and 5 indicate that the stationary REE near π_L is unstable for steady state learning but stable for VAR learning. Figure 4 illustrates that latter result under the parameter setting $\gamma = (\beta^{-1} - 1)/2$, $\sigma_\theta = 0.0001$, $\sigma_\eta = 0.0001$, $\sigma_\psi = 0.001$ and with other parameters at their base values given above. For initial conditions we set $b_0 = b_L + 0.01$, $\pi_0 = \pi_L + 0.001$, $R_0 = R_L + 0.001$ and $m_0 = m_l + 0.01$. (Here the subscript L refers to the steady state value of the corresponding variable at the steady state π_L.) Learning was assumed to use a small constant gain $\phi = 5000^{-1}$ and the simulation was run for 20000 periods. The results of the simulation accord with the theoretical analysis. The top panels of Figure 4 show that the economy fluctuates near the steady state values $\pi_L = 1.0477$ and $b \approx -23$.[16] The regression errors are relatively small, as indicated by the lower panels.

FIGURE 4 HERE

Further experiments in this case showed that stability of the stationary REE near π_L under learning is quite sensitive to assuming small shocks (note the standard deviations used) and requires initial conditions very close to the low interior steady state values. The experiments showed that larger constant gain increases in the shock variability or use of initial conditions further away from the equilibrium can each lead to divergence of learning from π_L. This suggests that, at least for the parameter values studied, convergence to the stationary REE near π_L is "very local" in some sense and that this solution is not easily reached from initial conditions that do not lie in a small neighborhood of this steady state.

[16]The steady state value for bonds is negative meaning that the government is a net lender. It can also be checked that the comparative dynamic properties of this equilibrium are non-intuitive.

The fragility of the stability of the REE near π_L can be contrasted to the stability properties of the REE near π_H and $\tilde{\pi}$ when PF prevails. Simulations (not shown) indicate that the stability of these solutions is relatively robust, i.e. it occurs also from initial conditions that are relatively far away from the corresponding equilibrium.

5.2 Policies for Avoiding Liquidity Traps

Since the stability in some key cases depends on fiscal policy it is of interest to examine the consequences of changes in the fiscal regime. In the first experiment it is assumed that the economy is initially near the low boundary equilibrium $\tilde{\pi}$ with passive fiscal policy. In this case $\tilde{\pi}$ is a locally stable equilibrium with no inflation, and we think of the fluctuations near $\tilde{\pi}$ as a liquidity trap of, say, the Japanese economy.

If fiscal policy is made active, i.e. is no longer geared towards keeping the public debt under control, the equilibrium is no longer stationary. In particular, the switch in fiscal policy leads to build up of debt (until the upper bound on debt is reached) with no essential upward movement in inflation. In other words, the liquidity trap cannot be cured by active fiscal policy that is not geared towards control of public debt.

Elimination of the liquidity trap can instead be achieved by a reformulation of monetary policy. The obvious remedy is to raise the minimum permissible level of inflation $\tilde{\pi}$ above the low interior steady state π_L. In terms of Figure 2 this leads to a shift up of the horizontal portion of the $T(\pi)$ map sufficient to ensure a unique steady state at $\pi = \pi_H$. Figure 5 illustrates the $T(\pi)$ map after such a change.[17] It is easily seen that in this case π_H continues to be E-stable under both the steady state and VAR learning.

Figure 6 presents simulation results showing the dynamics after this type of change in monetary policy.

FIGURES 5 AND 6 HERE

The three panels respectively show the rate of inflation, the quantity of bonds and inflation expectations. This simulation assumes that the shocks and

[17] This policy changes the number of steady states. It is conceptually similar to the tightening of a fiscal constraint in a monetary inflation model; see (Evans, Honkapohja, and Marimon 2001).

the constant gain parameter are all small (so that escape paths are highly unlikely): $\sigma_\theta = 0.0004$, $\sigma_\eta = 0.0001$, $\sigma_\psi = 0.0001$ and $\phi = 100^{-1}$. The run is for 10000 periods. The economy is initially assumed to be near π_L, which is unstable since we now maintain passive fiscal policy. The economy converges to the liquidity trap with π near $\tilde{\pi} = 1.00$. In period 2000 the low boundary is shifted up to $\tilde{\pi} = 1.05$ ($> \pi_L$). As seen from the figure, the economy gradually converges to the desired equilibrium at π_H.

As discussed in Section 2, this method of eliminating the liquidity trap is achieved by an implementation of monetary policy that switches to a money supply rule if private sector inflation expectations are too low. Operationality of this switch requires sufficiently accurate information on private inflation expectations.[18]

There are other ways to diminish the likelihood of the liquidity trap by policy design. Evidently, the nature of the fluctuations under constant gain learning depends on the "sizes" of the basins of attraction of the desired steady state π_H and the floor steady state $\tilde{\pi}$, as well as on the strength of dynamic adjustments. A change in the interest rate rule that shifts down the value of the unstable steady state π_L will alter the basins of attraction and therefore the likelihood of escape paths. Figure 7 presents the same setup as in Figure 3 but with $A = 1.35$ in place of 1.2. With this value for A, the intermediate steady state inflation is $\pi_L = 1.0234$. The figure shows that the likelihood of the escape paths from π_H to $\tilde{\pi}$ is greatly reduced, since the basin of attraction of π_H is made larger by the shift in the interest rate rule.

FIGURE 7 HERE

We conclude this section by reemphasizing the important role played by fiscal policy. A passive fiscal regime, in which taxation responds appropriately to the level of public debt, has several key properties. First, it helps to ensure convergence to the desired target inflation rate. This is clear from part (i) of Proposition 5. Furthermore active fiscal policy leads to locally explosive paths near π_H and it can also be shown that these paths may be unstable under learning. Second, part (ii) of Proposition 5 shows the theoretical risk of learning leading the economy to equilibria near π_L when fiscal

[18]We remark that the upper bound on inflation, the implementation of which also requires an analogous switch to money supply rule, is less critical than the lower bound as the upper bound does not represent a new steady state to the economy.

policy is active. Finally, active fiscal policy leads to explosive debt paths at $\tilde{\pi}$ as well as at π_H.

Our results provide an interesting contrast to those of (Benhabib, Schmitt-Grohe, and Uribe 2002). They propose two possible policies, either of which they view as sufficient to avoid the possibility of a liquidity trap: (i) a fiscal rule that reduces tax revenue at low inflation rates, or (ii) a monetary policy that switches to money growth rules at low interest rates in combination with a suitable non Ricardian fiscal rule. Their results are based squarely, and rely heavily, on the perfect foresight assumption (in a nonstochastic model). In our approach we replace perfect foresight (or fully rational expectations) by adaptive econometric learning rules. Although our learning dynamics allow for the possibility of convergence to a rational stationary solution under active (non Ricardian) fiscal policy, they appear more likely to lead instead to explosive debt paths.[19] Our results thus suggest the desirability of relying on a "second pillar" targeting of inflation, through money supply rules, at sufficiently low inflation rates, in combination with a passive (Ricardian) fiscal policy.

6 Conclusions

Recent research by (Benhabib, Schmitt-Grohe, and Uribe 2001) and (Benhabib, Schmitt-Grohe, and Uribe 2002) has brought to economists' attention the possibility of the economy sliding into a liquidity trap when monetary policy is conducted using Taylor-type interest rate rules. Their analysis was conducted under the perfect foresight assumption and we have re-examined this issue under the assumption that agents form expectations using econometric learning rules. One major finding is that although the interior liquidity trap equilibrium π_L can be stable under learning when fiscal policy is active, the basin of attraction for this equilibrium appears small. A greater concern is the economy slipping even further to inflation rates below π_L, with adaptive learning dynamics pushing the economy towards some lower boundary $\tilde{\pi}$ established by monetary authorities. $\tilde{\pi}$ then becomes a low level inflation trap from which it is difficult to escape.

[19] However, the policy combinations that we consider do differ in detail from those (Benhabib, Schmitt-Grohe, and Uribe 2002), and a more careful analysis of their specific policies, under learning, seems warranted.

Without the floor at $\tilde{\pi}$ the economy under learning would slide into cumulative deflation. This floor can be interpreted as a "second pillar" of monetary policy that gives primacy to money supply rules when inflation is sufficiently low. The second pillar inflation rate can be achieved for any given inflation expectations, money demand function and inherited price level by sufficiently increasing the money supply. Active fiscal policy alone is unable to push the economy out of the low boundary inflation trap $\tilde{\pi}$, and leads instead to an explosive build-up of debt with little change in inflation.

The required policy is instead a switch to a more aggressive monetary policy in which the second pillar inflation rate $\tilde{\pi}$ is increased above π_L. This leads to a cumulative increase in inflation to the desired inflation target π_H, achieved through implementation of the interest rate rule once inflation and inflation expectations are above $\tilde{\pi}$. In fact, the policy of relying more aggressively on the money supply component of the rule by setting a higher $\tilde{\pi} > \pi_L$ will also insulate the economy against liquidity traps and ensure global convergence to the desired inflation target π_H.

The results of this paper indicate the need for further research in several directions. Within the current model framework there are a number of open issues. In particular, are there specifications in which the interior low inflation steady state π_L has more robust stability properties under learning so that it might plausibly emerge as an outcome of these dynamics? More generally, our analysis has been cast in terms of a flexible price economy with constant output. Extending the model to one with sticky prices and variable output would be considerably more complicated, but clearly desirable, since the main concern of liquidity traps is their association with low output and stagnation. Such an extension is planned for the near future.

References

Benhabib, J., S. Schmitt-Grohe, and M. Uribe (2001): "The Perils of Taylor Rules," *Journal of Economic Theory*, 96, 40–69.

——— (2002): "Avoiding Liquidity Traps," *Journal of Political Economy*, 110, 535–563.

Bullard, J., and I.-K. Cho (2002): "Escapist Policy Rules," Working paper.

Cho, I.-K., and K. Kasa (2002): "Learning Dynamics and Endogenous Currency Crises," mimeo.

Cho, I.-K., N. Williams, and T. J. Sargent (2002): "Escaping Nash Inflation," *Review of Economic Studies*, 69, 1–40.

European Central Bank, E. U. (2001): *The Monetary Policy of the ECB*. European Central Bank, www.ecb.int.

Eusepi, S. (2002): "Forecast-Based vs. Backward-Looking Taylor Rules: A Global Analysis," mimeo.

Evans, G. W., and S. Honkapohja (1993): "Adaptive Forecasts, Hysteresis and Endogenous Fluctuations," *Federal Reserve Bank of San Francisco Economic Review*, 1993(1), 3–13.

——— (1995): "Local Convergence of Recursive Learning to Steady States and Cycles in Stochastic Nonlinear Models," *Econometrica*, 63, 195–206.

——— (2001): *Learning and Expectations in Macroeconomics*. Princeton University Press, Princeton, New Jersey.

——— (2002): "Policy Interaction, Learning and the Fiscal Theory of Prices," mimeo, www.valt.helsinki.fi/RAKA/seppo.hmtl, and http://darkwing.uoregon.edu/∼gevans/.

Evans, G. W., S. Honkapohja, and R. Marimon (2001): "Convergence in Monetary Inflation Models with Heterogeneous Learning Rules," *Macroeconomic Dynamics*, 5, 1–31.

Kasa, K. (2002): "Learning, Large Deviations, and Recurrent Currency Crises," *International Economic Review, forthcoming*.

Leeper, E. M. (1991): "Equilibria under 'Active' and 'Passive' Monetary and Fiscal Policies," *Journal of Monetary Economics*, 27, 129–147.

McCallum, B. (2001a): "Indeterminacy, Bubbles, and the Fiscal Theory of Price Level Determination," *Journal of Monetary Economics*, 47, 19–30.

McCallum, B. T. (2001b): "Inflation Targeting and the Liquidity Trap," mimeo, Carnegie Mellon University.

PRESTON, B. (2003): "Learning about Monetary Policy Rules when Long-Horizon Forecasts Matter," mimeo, Princeton University.

SARGENT, T. J. (1993): *Bounded Rationality in Macroeconomics*. Oxford University Press, Oxford.

——— (1999): *The Conquest of American Inflation*. Princeton University Press, Princeton NJ.

WILLIAMS, N. (2002a): "Escape Dynamics in Learning Models," mimeo.

——— (2002b): "Stability and Long Run Equilibrium in Stochastic Fictitious Play," mimeo.

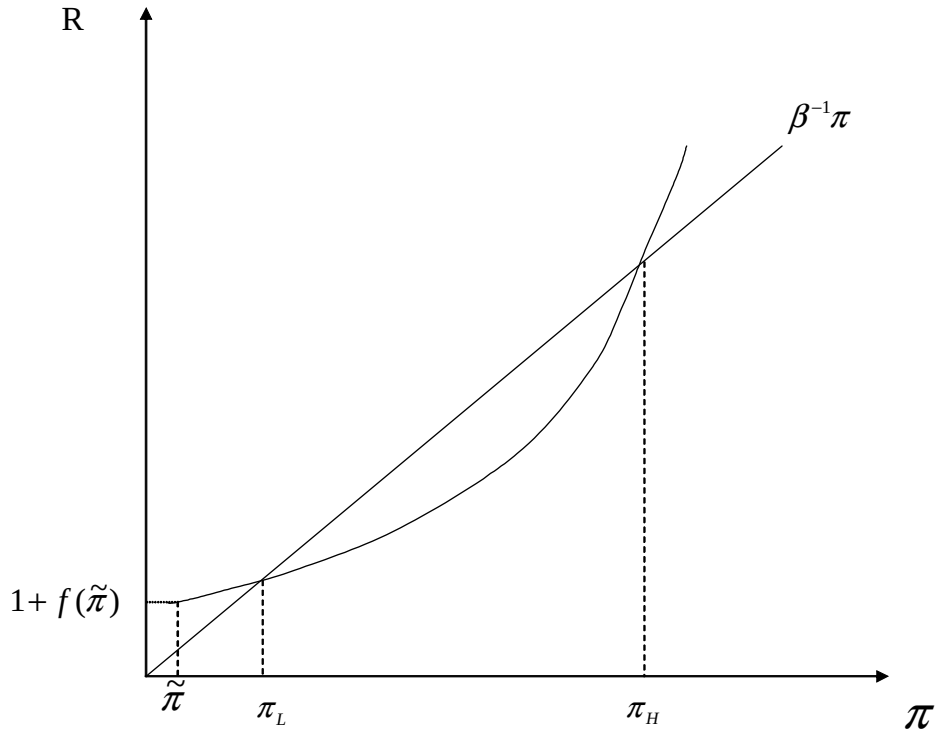

Figure 1: Steady states

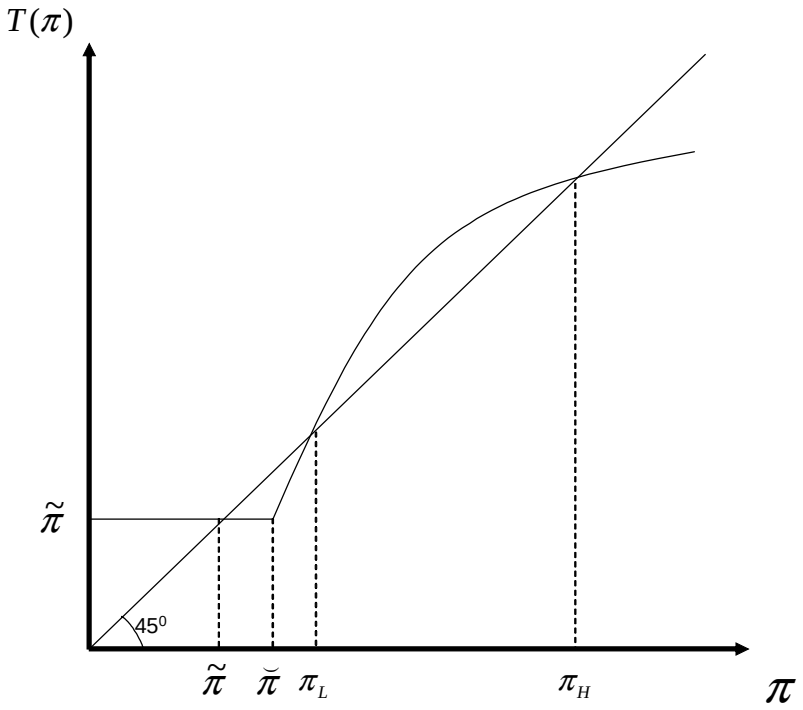

Figure 2: T-map governing E-stability

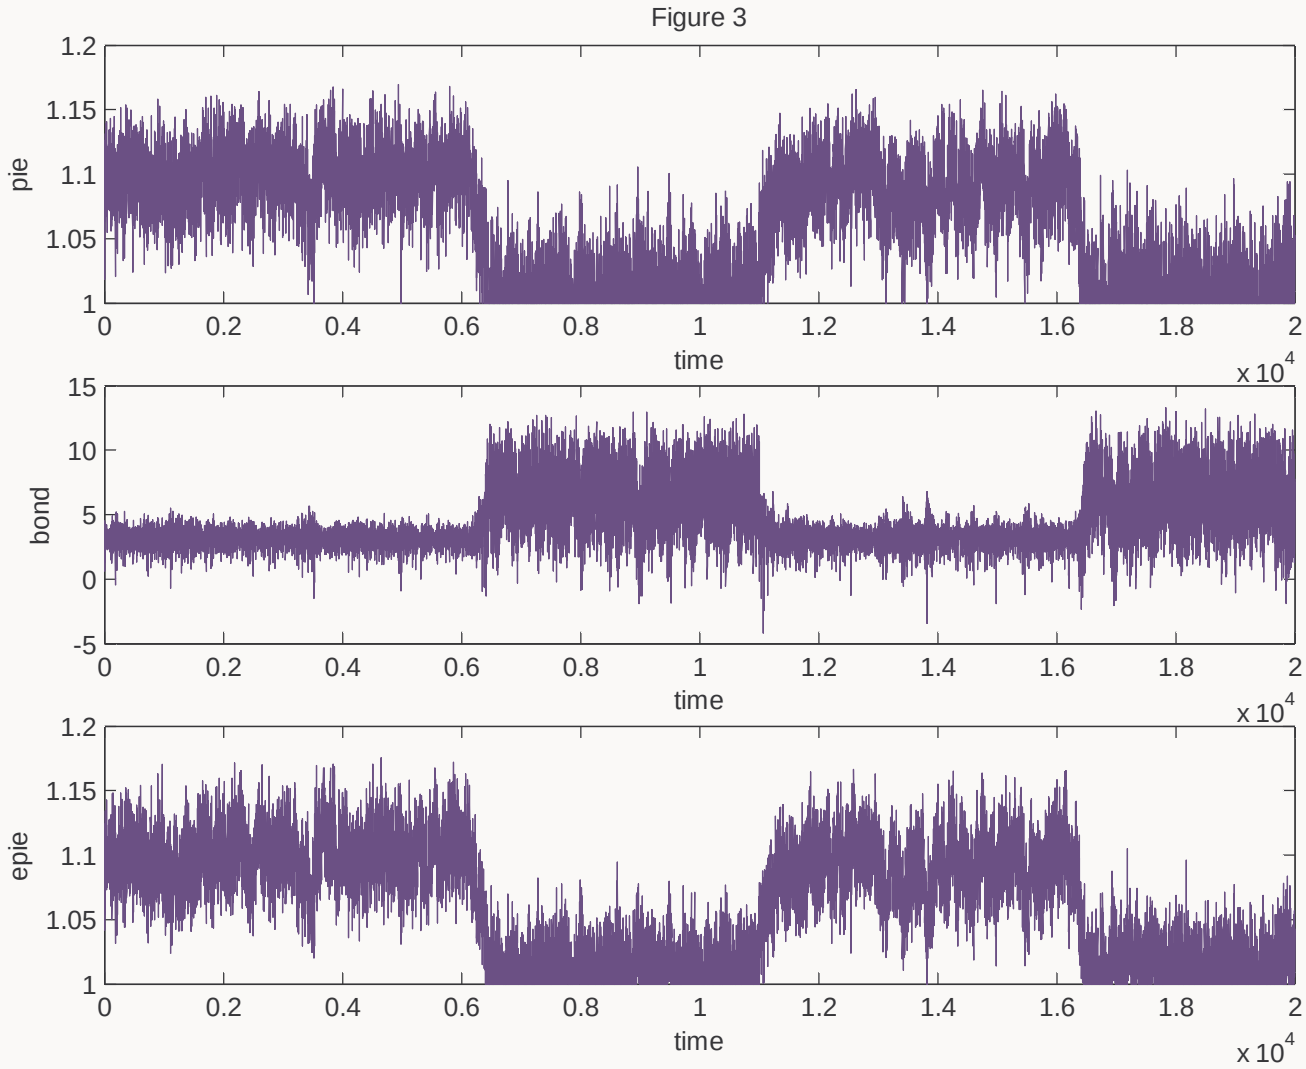
Figure 3
pie
1.2
1.15
1.1
1.05
1
0
0.2
0.4
0.6
0.8
1
1.2
1.4
1.6
1.8
2
time
x 10^4
bond
15
10
5
0
-5
0
0.2
0.4
0.6
0.8
1
1.2
1.4
1.6
1.8
2
time
x 10^4
epie
1.2
1.15
1.1
1.05
1
0
0.2
0.4
0.6
0.8
1
1.2
1.4
1.6
1.8
2
time
x 10^4

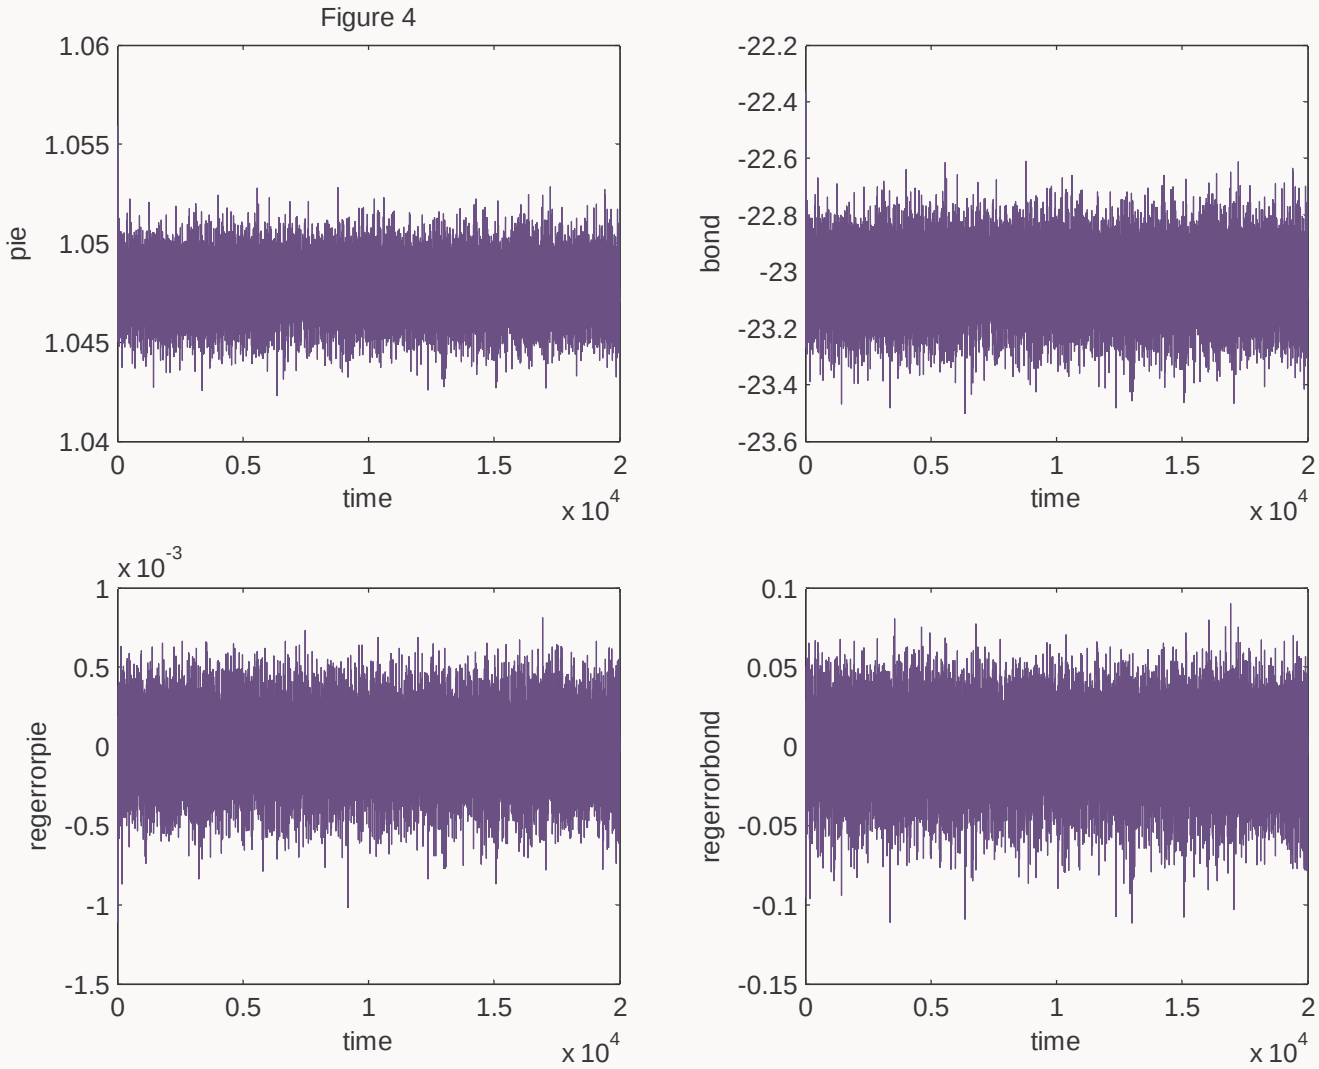
Figure 4
pie
1.06
1.055
1.05
1.045
1.04
0
0.5
1
1.5
2
time
x 10^4
bond
-22.2
-22.4
-22.6
-22.8
-23
-23.2
-23.4
-23.6
0
0.5
1
1.5
2
time
x 10^4
regerrorpie
x 10^-3
1
0.5
0
-0.5
-1
-1.5
0
0.5
1
1.5
2
time
x 10^4
regerrorbond
0.1
0.05
0
-0.05
-0.1
-0.15
0
0.5
1
1.5
2
time
x 10^4

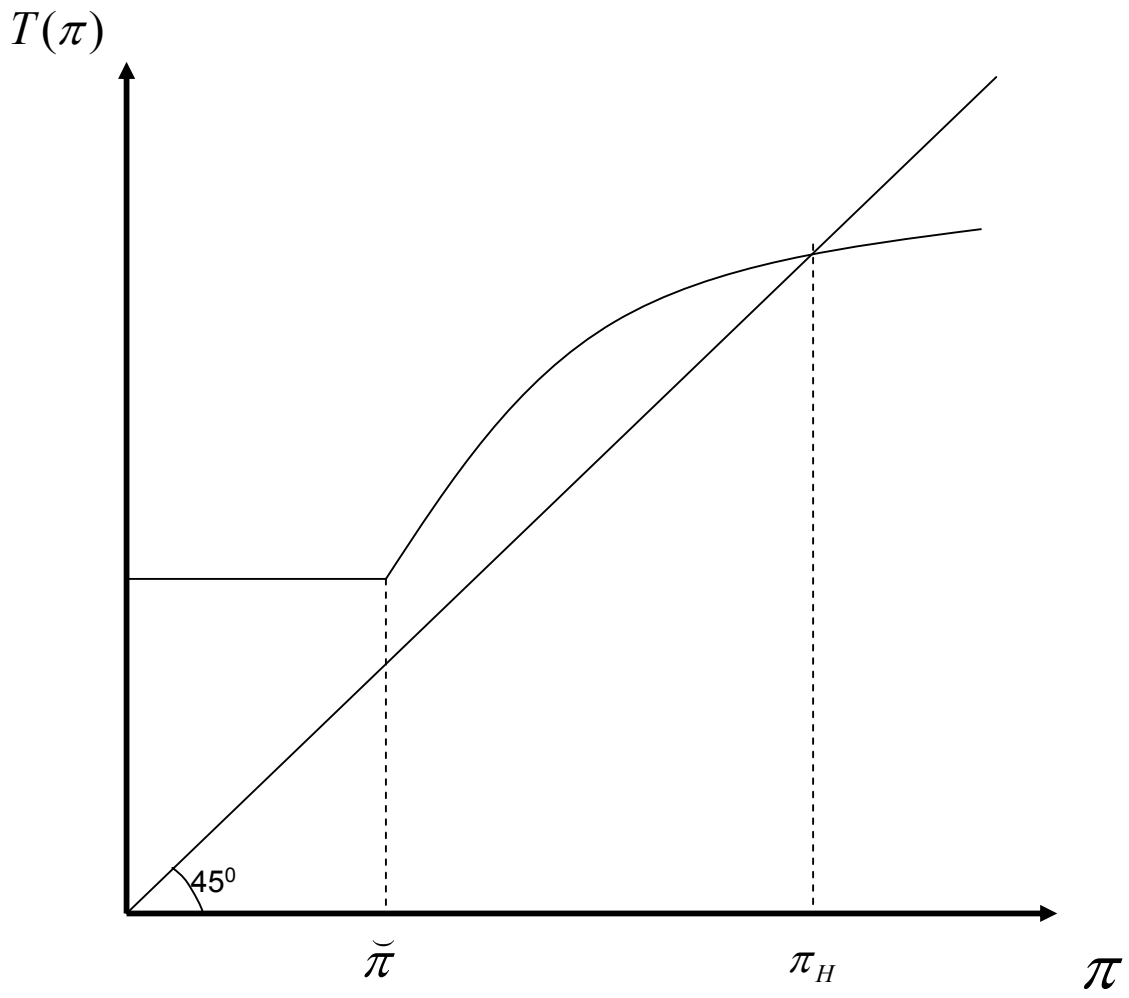

Figure 5: *T*-map with aggressive money rule

Figure 6

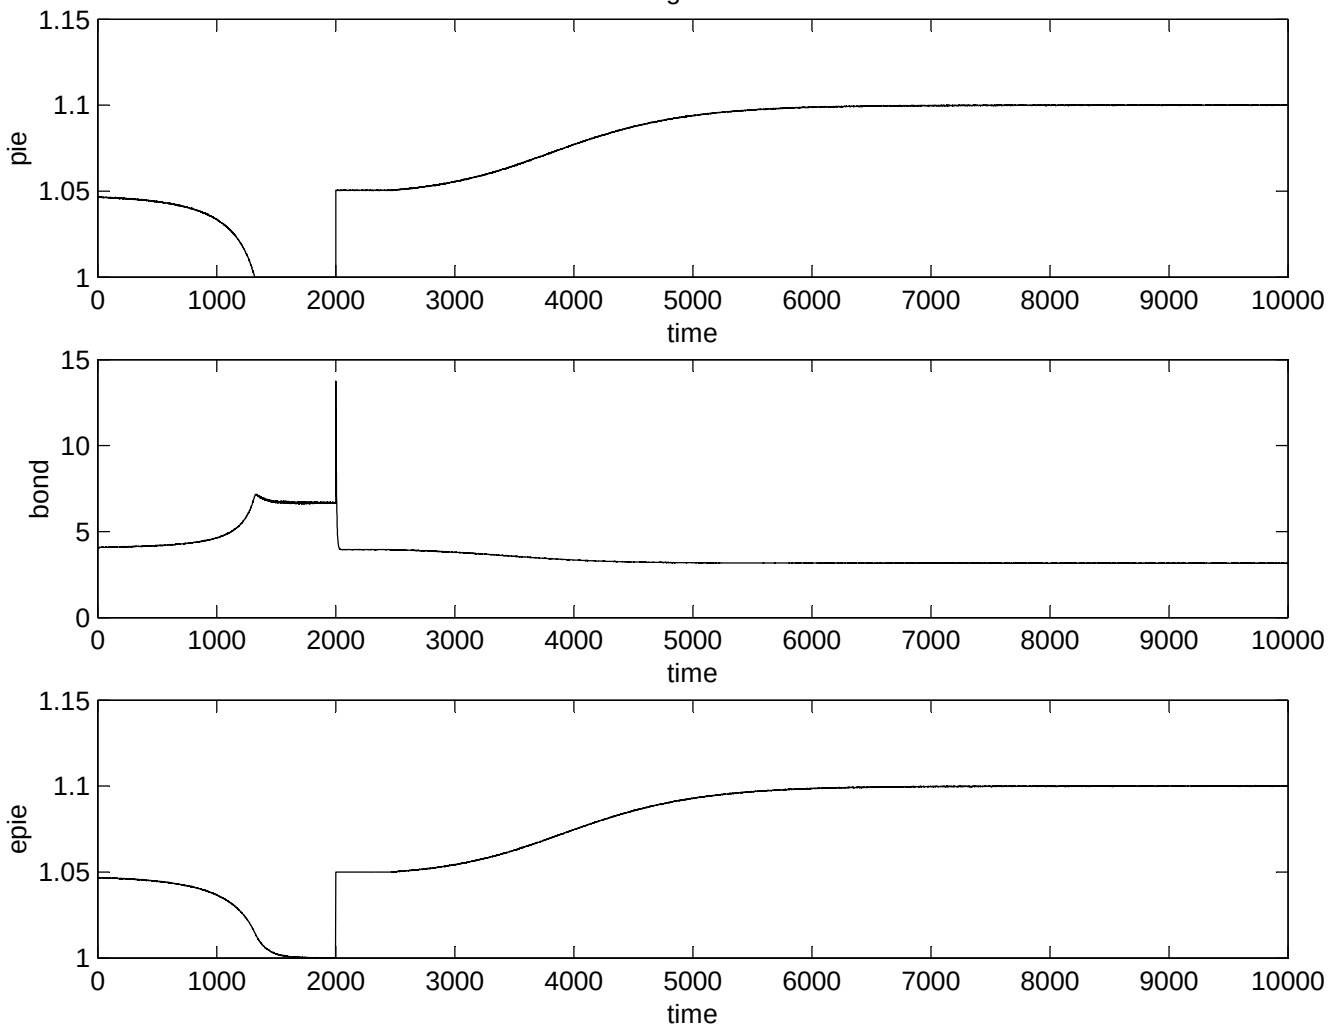
pie
1.15
1.1
1.05
1
0
1000
2000
3000
4000
5000
6000
7000
8000
9000
10000
time
bond
15
10
5
0
time
epie
1.15
1.1
1.05
1
time

Figure 7

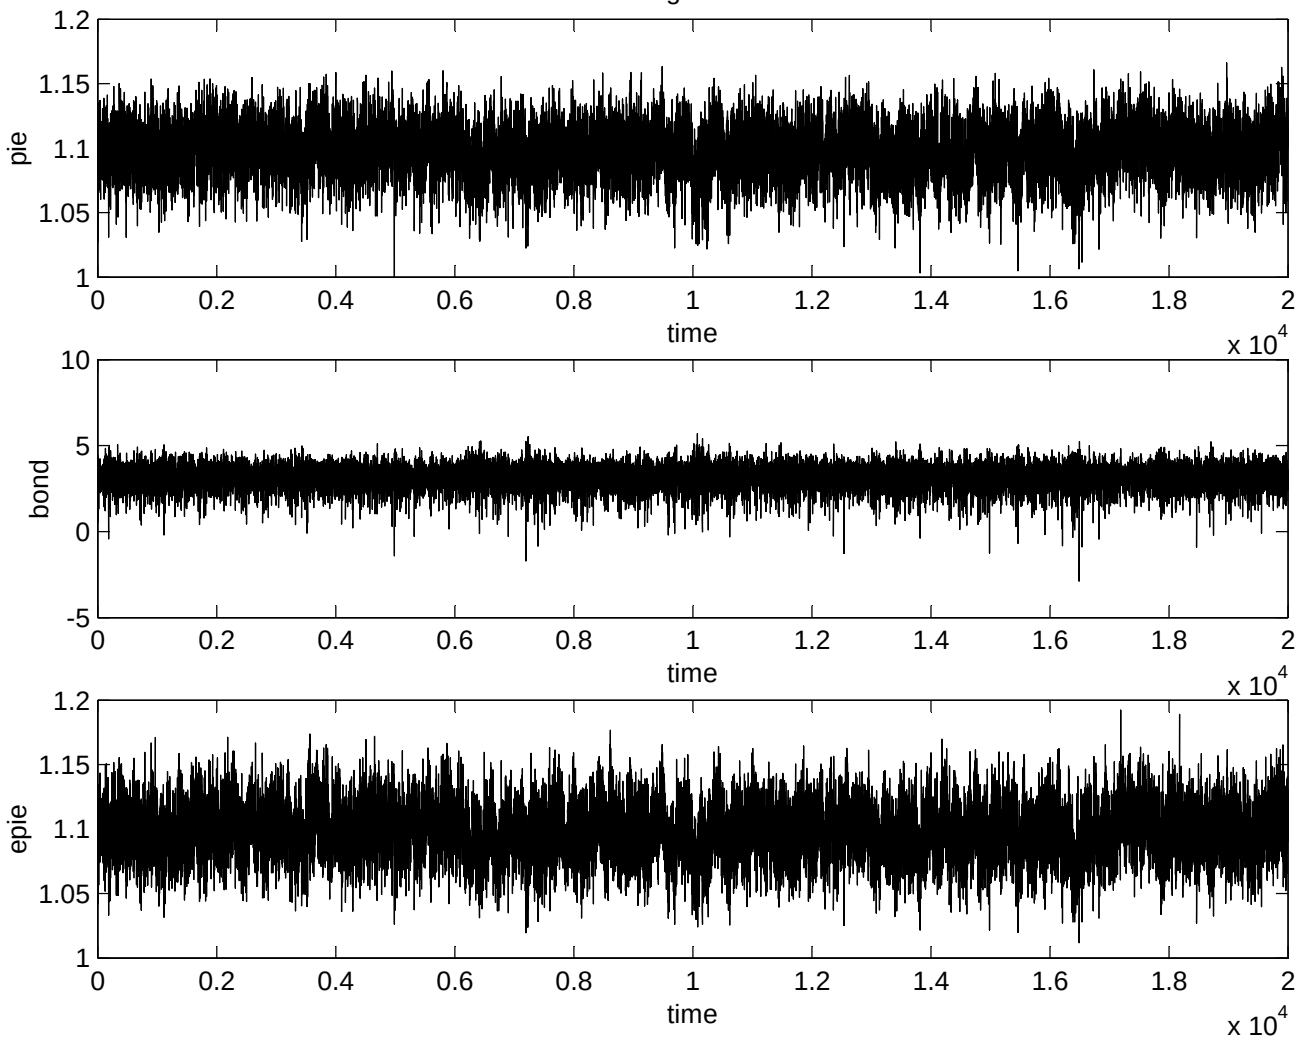
pie
1.2
1.15
1.1
1.05
1
0
0.2
0.4
0.6
0.8
1
1.2
1.4
1.6
1.8
2
time
x 10^4
bond
10
5
0
-5
0
0.2
0.4
0.6
0.8
1
1.2
1.4
1.6
1.8
2
time
x 10^4
epie
1.2
1.15
1.1
1.05
1
0
0.2
0.4
0.6
0.8
1
1.2
1.4
1.6
1.8
2
time
x 10^4