Adaptive Learning in Infinite Horizon Decision Problems

Bruce Preston
Columbia University*

September 22, 2005

Preliminary and Incomplete

Abstract

Building on Marcet and Sargent (1989) and Preston (2005) this paper shows that for infinite horizon decision problems in which agents optimize but have arbitrary subjective expectations about the evolution of state variables beyond their control optimal decision rules necessarily depend on infinite-horizon expectations. This contrasts with most work on adaptive learning in macroeconomics since Marcet and Sargent (1989) which posit Euler equations as decision rules in which only one-period-ahead forecasts matter. Using the Townsend (1983) investment model and the canonical consumption model various pathologies of the Euler equation approach are adduced. The two modeling approaches are shown to give very different conclusions about policy in a simple model of output gap and inflation determination.

JEL Classifications: E52, D83, D84

Keywords: Adaptive learning, Multi-period decision problems

*Department of Economics, Columbia University, 420 West 118th Street, Rm 1022, New York NY 10027. E-mail: bruce.preston@columbia.edu

1 Introduction

Fundamental to a rational expectations analysis of multi-period decision problems is the notion that to form an optimal decision today agents must form expectations about future economic conditions at all dates of the decision horizon. By taking appropriate account of the future uncertainty about the evolution of state variables relevant to the agents' decision problem resources can be optimally allocated. Building on Marcet and Sargent (1989a) this paper argues that under arbitrary beliefs that satisfy standard probability laws – and specifically adaptive learning – optimal decision rules in infinite horizon decision problems will in general require forecasts of state variables into the indefinite future. This contrasts with many recent analyses of non-rational expectations in the macroeconomics literature which posit decision rules in which only one-period-ahead forecasts matter.

Two dynamic intertemporal decision models – the Townsend (1983) model of investment and the canonical consumption model – are analyzed under the assumption of non-rational expectations. Careful attention is given to the proper specification of beliefs: though agents have non-rational expectations like an analysis under rational expectations they form beliefs about the evolution of state variables that are beyond their control though relevant to their decision problems. Under this assumption, the optimal decision rules are shown to depend on infinite horizon expectations of these same state variables. The logic inherent in the derivation of such optimal decisions is identical to that of a rational expectations analysis, and, indeed, these solutions are necessarily valid under the specific assumption of rational expectations.

The optimal decision rules are then related to learning analyses that depend only on one-period-ahead forecasts which have been prevalent in the recent macroeconomics literature. This literature typically commences with a system of Euler equations – recursive representations of optimal actions mapping expectations about tomorrow into actions today – that characterize aggregate dynamics under the rational expectations assumption. As an example, canonical consumption theory yields a consumption Euler equation which predicts current consumption to be a linear function of expectations of consumption tomorrow and

the real interest rate. It is then asked what are the implications of replacing the rationally formed expectations with the alternative adaptive expectations assumption. Of particular interest are the conditions under which the learning dynamics converge to those predicted by a rational expectations equilibrium analysis.

This paper demonstrates that such Euler equations implied by a rational expectations equilibrium analysis will in general not represent optimal decision rules under adaptive learning for the infinite horizon decision problems considered. This departure from the conclusions of a rational expectations analysis arises due to arbitrary specification of the conditional distributions of future endogenous decision variables with respect to which expectations are taken. Only when expectations are taken with respect to the correct distribution – the distribution induced by the optimal decision rules, beliefs and the properties of state variables – will the Euler equation approach represent an optimal decision rule. Moreover, because Euler equation approaches require agents to forecast endogenous decision variables using arbitrary statistical rules, such approaches will be revealed to be difficult to square with primitive assumptions inherent in microfounded multiperiod decision problems.

The analysis presented here is not intended to suggest that it is inappropriate to postulate Euler equation-type decision procedures as plausible models of bounded rationality. However, to the extent that researchers are interested in studying microfounded dynamic stochastic general equilibrium models under learning, then Euler equation approaches will not lead to optimal decisions. Indeed, given that agents face multiperiod decision problems it seems unnecessary to restrict attention to models in which agents base their decisions only on one-period-ahead expectations. Evidence is also adduced underscoring a number of pathologies that arise under the Euler equation approach. For instance, in a variant of the canonical consumption model, the Euler equation approach can lead households to consume permanently too much or too little and therefore violate their intertemporal budget constraint. This difficulty arises because Euler equation-based methods do not account for wealth in any way whatsoever.

The paper concludes with a discussion of Preston (2005) which considers a simple New

Keynesian model of output gap and inflation determination of the kind that has been used in recent analyses of monetary policy – see, for example, Bernanke and Woodford (1997), Clarida, Gali, and Gertler (1999) and Woodford (1999). The logic underpinning the analyses of the Townsend investment model and the canonical consumption model is shown to apply directly and tractably to this dynamic stochastic general equilibrium model (and more generally to any class of such models including real business cycle theory). As such, the paper owes much to Marcet and Sargent (1989a). Importantly, the discussion reveals that the Euler equation approach attributes much greater knowledge about the underlying rational expectations equilibrium – about which agents are attempting to learn – than do the optimal decision rules. A direct implication is that many policies recently argued to be desirable to stabilizing economic fluctuations in the Euler equation-based learning literature in fact lead to instability in the model based on optimal decision rules.

2 Infinite Horizon Learning

This section first recapitulates the treatment of the Townsend (1983) model of investment by Marcet and Sargent (1989a) and Sargent (1993) under the assumption of non-rational expectations. To the author's knowledge this is the first and only example of infinite horizon learning in the macroeconomics literature. An example of the permanent income theory is then discussed and shown to imply that infinite horizon forecasts are required by optimal decision rules.

2.1 The Townsend Investment Model

Consider a representative firm facing the following decision problem:

$$\max_{\{k_T\}} \hat{E}_t \sum_{T=t}^{\infty} \beta^{T-t} \left[p_T \gamma k_T - \frac{\delta}{2} \left(k_T - k_{T-1} \right)^2 \right]$$

where $0 < \beta < 1$ and $\gamma,\ \delta > 0$. The firm therefore chooses a sequence of capital stocks $\{k\}_{T=t}^{\infty}$ to maximize output (γk_T) of the economy's only good whose price (p_t) is determined

in a competitive market and satisfies the demand system

$$\begin{aligned} p_t &= A_0 - A_1 \gamma K_t + u_t \\ u_t &= \rho u_{t-1} + \varepsilon_t \end{aligned} \tag{1}$$

where K_t is the aggregate capital stock of the economy, taken as given by the representative firm. u_t is an exogenous disturbance process with ε_t a mean zero, bounded, i.i.d. disturbance and $0 < \rho < 1$.

The representative firm's beliefs are denoted $\hat{E}_t$ and are determined as follows. Because the firm faces an intertemporal decision problem – indeed an infinite-horizon decision problem – it must forecast all state variables relevant to its decision problem in all future dates and states of uncertainty. As the firm is assumed only to have knowledge of its own objectives and the constraints it faces it does not have a complete economic model of the determination of aggregate variables. It therefore does not have knowledge of the relation (1) or the equilibrium condition $k_t = K_t$. Thus the set of state variables – the set of variables beyond the control of the representative firm and whose statistical properties firms are attempting to deduce – are given by $z_t \equiv \{p_t,\ K_t,\ u_t\}$. In contrast to a rational expectations analysis where the equilibrium probability laws are known and coincide with the predictions of the economic model it is assumed that inferences about the evolution of future state variables are made by extrapolation from observed patterns in historical data. Agents therefore recursively estimate the vector autoregression

$$z_t = a_t + b_t z_{t-1} + \eta_t \tag{2}$$

where η_t is vector white noise and a_t and b_t are coefficient matrices of obvious dimension. Given estimates, forecasts of future state variables at any horizon T are determined as:

$$\hat{E}_t z_T = (I_3 - \beta)^{-1} \left(I_3 - b_t^{T-t}\right) a_t + b_t^{T-t} z_t. \tag{3}$$

This completes the specification of the firm's decision problem. It should be underscored that the arguments presented in the following generalize to arbitrary beliefs satisfying standard probability laws. The assumption of adaptive learning is made to facilitate analytical tractability and to render comparison with existing work transparent.

The optimal decision rule that solves the firm's infinite horizon decision problem is

$$k_t = k_{t-1} + \frac{\gamma}{\delta}\hat{E}_t \sum_{T=t}^{\infty} \beta^{T-t} p_T \tag{4}$$

for all $t \geq 0$. Thus to solve for the optimal level of the capital stock at time t the firm requires forecasts of goods prices into the indefinite future. This will be shown to be a general characteristic of optimal decision rules in infinite horizon decision problems. Using (3) to evaluate expectations in (4) provides the optimal allocation of capital given observed aggregate data. The adaptive learning literature typically concerns itself with the conditions under which learning dynamics converge to the model's associated rational expectations equilibrium. Application of the convergence criterion detailed in appendix A.1 to the present problem provides the following proposition.

Proposition 1 *The Townsend investment model implies optimal allocation of capital is determined in a rational expectations equilibrium according to*

$$k_t = \mu K_{t-1} + \frac{\gamma}{\delta(\beta\mu - \rho)} u_t$$

where $0 < \mu < 1$. *Under learning dynamics, optimal decisions converge to this allocation rule with probability 1.*

The proof is found in the appendix. Given the paucity of examples of long-horizon learning in the macroeconomics literature the question naturally arises as to whether there are other equally valid representations of the optimal decision rule that may have been adopted by the literature on learning dynamics. To this end note that the beliefs (2) imply conditional distributions for z_t that satisfy standard probability laws. It follows that the law of iterated expectations is necessarily satisfied. Thus the optimal decision rule can be reformulated as

$$\begin{aligned} k_t &= k_{t-1} + \frac{\gamma}{\delta} p_t + \frac{\gamma}{\delta}\hat{E}_t \hat{E}_{t+1} \sum_{T=t+1}^{\infty} \beta^{T-t} p_T \\ &= k_{t-1} + \frac{\gamma}{\delta} p_t + \frac{\gamma}{\delta}\beta \hat{E}_t \left(k_{t+1} - k_t\right) \\ &= (\delta + \gamma\beta)^{-1} \left[k_{t-1} + \gamma p_t + \gamma\beta \hat{E}_t k_{t+1}\right]. \end{aligned} \tag{5}$$

The final equality gives an Euler equation of the kind central to rational expectations equilibrium analyses of dynamic macroeconomic models. It determines the current capital stock as a function of the lagged capital stock, the current price and the expectations of tomorrow's capital stock. This remains a valid formulation of the optimal decision rule so long as the expectation $\hat{E}_t k_{t+1}$ is properly interpreted. In particular, expectations, denoted by $\hat{E}_t$, must be taken with respect to the correct probability distribution – that is the conditional distribution of k_{t+1} that is induced by the optimal decision rule, (4), and the specified beliefs, (2). Hence the optimal program in conjunction with the assumed beliefs imply the firm's optimal forecast to be

$$\hat{E}_t k_{t+1} = k_t + \frac{\gamma}{\delta} \hat{E}_t \sum_{T=t+1}^{\infty} \beta^{T-t} p_T \tag{6}$$

formed by leading (4) forward one period and taking expectations. Thus interpreted, the Euler equation represents an optimal decision rule.

To push this discussion further, suppose relation (5) were taken as a primitive of the model environment. Does this decision rule lead to the same allocations as the optimal rule (4)? To answer this endow the representative firm with the beliefs

$$k_t = \tilde{a}_t + \tilde{b}_t K_{t-1} + \tilde{c}_t u_t + \tilde{\eta}_t \tag{7}$$

so that the future firm's capital stock is conjectured to depend on the lagged aggregate capital stock and current demand shock. Beliefs therefore nest the model's minimum-state-variable rational expectations solution as derived in Proposition 1. The required forecast can be determined as

$$\hat{E}_t k_{t+1} = \tilde{a}_t + \tilde{b}_t K_t + \tilde{c}_t \rho u_t \tag{8}$$

assuming for simplicity that ρ is known. It is immediate that the forecasts (6) and (8) will in general not coincide and therefore that the decision rule (5) combined with the beliefs (7) does not represent an optimal decision rule given the microfoundations described above.

This discussion underscores that expectations of endogenous choice variables must be taken with respect to the probability distribution induced by (i) the beliefs regarding the evolution of state variables that are beyond the control of the firm and (ii) the optimal

decision rule. Yet much of the recent work in learning and macroeconomics has commenced with an Euler equation such as (5) coupled with the beliefs (7). Examples include Bullard and Mitra (2000, 2002), Evans and Honkapohja (2002, 2003), Euseppi, Honkapohja and Mitra (2003), Ferrero, Marcet and Sargent (1989b). While these papers do not generally commence analysis with a microfounded decision problem some certainly use such models to motivate the particular economy being studied. The central point of the present analysis is to underscore that to the extent we are interested in studying microfounded dynamic models the Euler equation approach – as it will be referred to subsequently – will not provide a theoretically consistent and coherent approach to modeling learning in these problems.

The present analysis builds on Preston (2005) and represents an appeal to model learning dynamics with infinite horizon forecasts as first proposed by Marcet and Sargent (1989a). It will be argued that long-horizon learning provides a more general and appealing approach to modeling learning dynamics. Such an approach requires agents to learn about more dimensions of the rational expectations equilibrium, giving agents must less information ex ante about the set of cross-variable restrictions that hold under rational expectations. Furthermore, Euler equation approaches may omit economic variables relevant to optimal decisions, leading to permanently sub-optimal decisions even in the case that learning dynamics converge. Optimal decisions rules, by virtue of being an implication of all relevant constraints and optimality conditions, necessarily imply that decisions are optimal and are conditioned on all relevant information. Finally, to the extent that we are interested in studying multi-period decision problems in macroeconomics it seems unappealing to restrict attention to decision procedures that only depend on one-period-ahead forecasts.

2.2 The Canonical Consumption Model

The discussion now turns to another fundamental intertemporal allocation problem: a version of the permanent income theory. The optimal consumption decision rule is again shown to depend on infinite horizon expectations and the Euler equation approach is shown to be inconsistent with optimality. Moreover, because the Euler equation approach fails to

take into account wealth dynamics in any way whatsoever, households may violate their intertemporal budget constraint even when learning dynamics converge. As a result they will permanently consume too much or too little.

A continuum of households indexed by $i \in [0,1]$ solve the decision problem

$$\max_{\{C_T^i\}} \hat{E}_t^i \sum_{T=t}^{\infty} \beta^{T-t} U\left(C_T^i\right)$$

subject to the flow budget constraint

$$W_{t+1}^i = R_t\left(W_t^i + Y_t - C_t^i\right)$$

and the usual No-Ponzi condition. These constraints together imply the intertemporal budget constraint

$$\sum_{T=t}^{\infty} R_{t,T} C_t^i = W_t^i + \sum_{j=0}^{\infty} R_{t,T} Y_T$$

where

$$R_{t,T} = \prod_{j=1}^{T-t} R_{t+j}^{-1}$$

and $R_{t,t} = 1$. C_t^i is the consumption of household i, R_t the period return on the economy's only available asset. W_t^i is the stock of wealth at the beginning of period t of household i. Aggregate wealth is assumed to satisfy the constraint $\int_0^1 W_t^i di = 0$ so that some households are lenders and some borrowers. The model is interpreted as a simple general equilibrium model that determines the real interest rate R_t. As such the market clearing condition $\int_0^1 C_t^i di = Y_t$ is satisfied. All households receive the same wages, Y_t, assumed to follow the exogenous process

$$Y_t = M Y_{t-1}^{\rho} e^{\nu_t}$$

where $\log M > 0$ and $0 < \rho < 1$ and $\{\nu_t\}$ is an i.i.d. mean zero bounded stochastic disturbance. To close the model, the beliefs of each household need to be specified. As in the investment problem, it is assumed that agents forecast state variables that are beyond their control using a simple vector autoregression. Here the state variables are $\{Y_t,\ R_t\}$. However, because we will employ a log-linear approximation to the optimality conditions of

this simple non-linear model, beliefs will not be specified in terms of $\{Y_t,\ R_t\}$ but rather log-linear approximations to these variables. This is simply to minimize notation. However, it should be noted that beliefs are necessarily specified as a primitive of the problem.

A log-linear approximation to the model's optimality conditions provides the Euler equation

$$c_t^i = \hat{E}_t^i c_{t+1}^i - \sigma r_t \tag{9}$$

and intertemporal budget constraint

$$\hat{E}_t^i \sum_{T=t}^{\infty} \beta^{T-t} c_T^i = w_t^i + \hat{E}_t^i \sum_{T=t}^{\infty} \beta^{T-t} y_T \tag{10}$$

where

$$c_t^i \equiv \ln(C_t^i/\bar{Y})\ ;\quad y_t^i \equiv \ln(Y_t^i/\bar{Y});\quad w_t^i = W_t^i/\bar{Y};\ \text{and}\ r_t \equiv \ln(R_t/\bar{R}) = \ln(R_t\beta).$$

$\bar{Y}$ is the steady state value of the real wage and $\sigma \equiv -U_{cc}/U_c\bar{Y} > 0$ is the intertemporal elasticity of substitution.[1] The income process satisfies the approximation

$$y_t = \rho y_{t-1} + \nu_t.$$

Given this linear structure beliefs can now be specified directly. Defining $z_t \equiv \{r_t,\ y_t\}$ agents recursively estimate the vector autoregression (2) with the appropriate re-definition of $\{a_t,\ b_t,\ \eta_t\}$. Given estimates, forecasts of future states at any horizon T are again determined by (3).[2]

Combining (9) and (10) yields the optimal decision rule:

$$c_t^i = (1-\beta)\, w_t^i + \hat{E}_t^i \sum_{T=t}^{\infty} \beta^{T-t} \left[(1-\beta)\, y_T - \sigma\beta r_T\right]. \tag{11}$$

[1] The log-linearization imposes the restriction $\beta^{-1} = \bar{R}$. This is done so that the steady state coincides with the usual assumptions of the rational expectations permanent income hypothesis.

[2] Note that each household has no knowledge of the preferences or beliefs of other households in this economy. There are therefore no higher order beliefs. It is implicitly being assumed that each household adopts a statistical model which they believe is the best forecasting technology available to them and do not concern themselves with the forecasting models of others – even though these happen to be the same in the example being considered.

Once again the optimal decision rule requires each household to forecast general macroeconomic conditions, as captured by the prices $\{r_t,\ y_t\}$ into the indefinite future. Moreover, application of the law of iterated expectations provides

$$\begin{aligned} c_t^i &= (1-\beta)\, w_t^i + (1-\beta)\, y_t - \sigma\beta r_t + \beta \hat{E}_t[c_{t+1} - (1-\beta)\, w_{t+1}^i] \\ &= \hat{E}_t^i c_{t+1}^i - \sigma r_t \end{aligned} \tag{12}$$

where the second equality follows directly from a log-linear approximation to the flow budget constraint. This Euler equation is a valid decision rule in the same way that equation (5) is for the investment problem. So long as the expectation $\hat{E}_t c_{t+1}$ is taken with respect to the correct probability distribution the Euler equation will provide the optimal allocation of consumption expenditures. Thus expectations are optimally determined according to

$$\begin{aligned} E_t^i c_{t+1} &= (1-\beta)\, E_t^i w_{t+1}^i + \hat{E}_t \sum_{T=t+1}^{\infty} \beta^{T-t} \left[(1-\beta)\, y_T - \sigma\beta r_T\right] \\ &= (1-\beta)\, \beta^{-1}(w_t^i + y_t - c_t^i) + \hat{E}_t \sum_{T=t+1}^{\infty} \beta^{T-t} \left[(1-\beta)\, y_T - \sigma\beta r_T\right] \end{aligned} \tag{13}$$

obtained by leading (11) one period and taking expectations. Expectations must be taken with respect to the conditional distribution of next-period consumption that is induced by the optimal decision rule (11) and beliefs (2).

Finally, aggregate consumption dynamics are determined by integrating (11) over $i \in [0,1]$ to give

$$c_t = \hat{E}_t \sum_{T=t}^{\infty} \beta^{T-t} \left[(1-\beta)\, y_T - \sigma\beta r_T\right] \tag{14}$$

defining $c_t = \int_0^1 c_t^i di$ and $\hat{E}_t = \int_0^1 \hat{E}_t^i di$ where the latter gives the average beliefs of all households. It is crucial to underscore that $\hat{E}_t$ has no behavioral content. One might be tempted to conclude that average beliefs necessarily satisfy standard probability laws. This is true in the present example, though is unique to the case of homogeneous beliefs – given heterogeneity of household beliefs average beliefs would not satisfy the law of iterated expectations. But suppose we proceed, noting that homogeneity of beliefs implies that the law of iterated

expectations holds, with the following manipulations:

$$
\begin{aligned}
c_t &= \hat{E}_t \sum_{T=t}^{\infty} \beta^{T-t} \left[(1-\beta) y_T - \sigma \beta r_T \right] && (15) \\
&= (1-\beta) y_t - \sigma \beta r_t + \beta \hat{E}_t \hat{E}_{t+1} \sum_{T=t+1}^{\infty} \beta^{T-t-1} \left[(1-\beta) y_T - \sigma \beta r_T \right] \\
&= (1-\beta) y_t - \sigma \beta r_t + \beta \hat{E}_t c_{t+1} \\
&= \hat{E}_t c_{t+1} - \sigma r_t && (16)
\end{aligned}
$$

where the final equality follows from application of the goods market clearing condition $c_t = \int_0^1 c_t^i di = y_t$. Note that this final relation could be obtained directly by integrating over (12). But on the ground of the same logic adduced above, this aggregate Euler equation only represents an optimal decision rule if expectations are taken with respect to the correct distribution induced by the aggregate decision rule (14) – and hence its constituent elements, the optimal household decision rules – and beliefs about state variables beyond the control of agents.

Honkapohja, Mitra, and Evans (2002) in a recent discussion of Preston (2005) have argued that because (16) can be derived from the optimal decision rule (14) – as shown in the above calculations – the Euler equation necessarily represents a decision rule that is consistent with the model's microfoundations. However, this is optimal only if expectations are taken with respect to the correct conditional distribution of future aggregate consumption. A direct implication is assuming beliefs of the form

$$\hat{E}_t^i c_{t+1}^i = a_t + b_t y_t \tag{17}$$

or the aggregate equivalent

$$\hat{E}_t c_{t+1} = a_t + b_t y_t$$

as specified by Honkapohja, Mitra, and Evans (2002) is inconsistent with the underlying microfoundations. Such an analysis runs aground because beliefs are never specified as a primitive in any of the examples discussed in their paper. Their analysis emphasizes the law of iterated expectations as being central to the consistency of the Euler equation approach.

But the law of iterated expectations necessarily holds by assumption given that beliefs satisfy standard probability laws and is not central to emerging differences. What is important is the consistent and correct interpretation of what the notation $\hat{E}_t^i$ or $\hat{E}_t$ summarizes. By not specifying beliefs as a primitive of the decision problem difficulties arise through assumptions of the kind represented by the beliefs (17). Indeed, Honkapohja, Mitra, and Evans (2002) in discussing a version of the canonical consumption model presented here, when deriving the Euler equation from the optimal decision rule as done in equations (15) through (16), transform the decision problem from one in which households forecast state variables beyond their control to one in which they forecast their own future consumption decisions using an arbitrary statistical rule – and as a result form expectations with respect to a probability distribution that is not implied by the decision problem. It follows that the Euler equation fails to deliver optimal allocations of consumption over time. This failure to consistently specify beliefs and the resulting inference problems regarding optimal decision rules in infinite horizon decision pervades all their discussion and in particular the simple model of output gap and inflation determination of Preston (2005) recapitulated in the sequel.

A direct implication is that it is difficult to square the Euler equation approach with the primitive decision problem. Why would a household solve a dynamic programming problem for the optimal decision rule and then form forecasts of their own future consumption decisions using an arbitrary statistical rule? One might argue that this difficulty can be avoided once it is noted that market clearing implies that $c_t = y_t$ in all periods in this model and therefore propose an analysis of the Euler equation

$$y_t = \hat{E}_t y_{+1} - \sigma r_t.$$

But such logic only compounds the nature of the difficulty. Market clearing conditions are part of the very equilibrium about which agents are attempting to learn. There does not seem to be any particularly reason why agents ought to know anything about this aspect of the equilibrium.

2.3 Alternative Approaches to Modeling Learning

Given the primacy of the Euler equation approach in the recent literature it is worth considering the convergence properties of such a model of learning relative to the optimal decision rule to further understand differences in the two modeling approaches. To this end, consider the stability properties of the long horizon expectations model given by (11) and (2) and the Euler equation model given by (12) with the beliefs (17). A learning analysis gives the following proposition.

Proposition 2 *Under the maintained assumptions both the long-horizon and Euler equation approaches lead to convergence of aggregate dynamics to the same equilibrium. However, the Euler equation approach has the property that household decisions do not converge to the optimal decision rule. Thus households fail to make optimal decisions both in transition to rational expectations equilibrium and in the rational expectations equilibrium itself. Households therefore violate their intertemporal budget constraint at all times.*

The proof is contained in the appendix. This is an important point: not only does the Euler equation approach imply households make sub-optimal decisions while learning, but also when learning dynamics converge – despite the fact that aggregate dynamics converge to the same rational expectations equilibrium in both approaches. The difference arises because the Euler equation approach fails to take into account household wealth in any way whatsoever. As a result, households will permanently consume too much or too little depending on whether they are a lender or borrower in this economy, and therefore plan to violate their intertemporal budget constraint. Indeed, the appendix demonstrates that in the case of the Euler equation approach individual decision rules converge to

$$c_t^i = y_t$$

In contrast, the limiting decision rules under long-horizon learning take the form

$$c_t^i = (1-\beta)\, w_t^i + y_t.$$

It is immediate that Euler equation-based decision procedures lead to consumption permanently too high or too low as w_t^i is negative and positive respectively.

In a rational expectations analysis the equilibrium probability laws are an implication of, among other things, the intertemporal budget constraint being satisfied. A consequence is that the equilibrium conditional distribution of future consumption with respect to which expectations are taken appropriately accounts for household wealth. It is for this reason that the Euler equation represents an optimal decision rule. But this is not true in a model of non-rational expectations. Indeed, in the case of adaptive learning, beliefs about state variables are exogenously specified and therefore cannot in anyway directly account for wealth.[3]

Finally, it should be emphasized that this paper – and the analyses of Preston (2003, 2004a, 2004b) – does not claim it to be unreasonable to posit directly Euler equation-type decision procedures as plausible models of bounded rationality. What is contested is to the extent that we are interested in multi-period decision problems with non-rational expectations, the Euler equation approach will not give rise to optimal decisions and is therefore inconsistent with the primitive model assumptions. Finally, note that the Euler equation approach does not necessarily give a more plausible model of bounded rationality than the infinite horizon learning framework – the latter extrapolates current beliefs into the indefinite future even though such beliefs will be updated in the next period. The fact that households are hyper-rational in one dimension of their decision rule – i.e. having an infinite horizon – serves to magnify the effect of misspecified beliefs on current consumption. In contrast, the Euler equation approach uses only one-period-ahead expectations and therefore limits the impact of misspecified beliefs on current consumption decisions. Moreover, in the next section it will be demonstrated that the long-horizon expectations approach typically requires agents to learn much more about the underlying rational expectation equilibrium of interest. This should already be clear from the above example. The optimal decision rule requires forecasts of the entire path of future real interests rates, which is not required in the Euler equation approach. As a result, the stability conditions under learning dynamics can be considerably different relative to the Euler equation approach.

[3] The expectations function (17) could be augmented with a term in w_t^i. But for the same reasons adduced earlier, the resulting allocations for consumption would remain suboptimal in transition to rational expectations equilibrium.

3 A Simple DSGE Model

Preston (2005) analyzes the microfoundations used in several recent studies of monetary policy rules – see, for example, Bernanke and Woodford (1997), Clarida, Gali, and Gertler (1999) and Woodford (1999) – under a specific non-rational expectations assumption. This leads to important differences in the model's implied aggregate dynamics when agents are learning relative to the predictions of a rational expectations equilibrium analysis and the Euler equation approach to modeling learning dynamics. This section recapitulates the central assumptions and results of that paper and the reader is encouraged to consult it for details.

The model based on the optimal decision rules is contrasted with the associated Euler equation approach. The two models are shown to give different conclusions about the desirability of a number of monetary policy rules for stabilizing economic fluctuations. This is true even though both approaches are consistent with the same underlying rational expectations equilibrium.

3.1 Household and Firm Decision Problems

The economy is populated by a continuum of households which seek to maximize future expected discounted utility

$$\hat{E}_t^i \sum_{T=t}^{\infty} \beta^{T-t} \left[U(C_T^i; \xi_T) - \int_0^1 v(h_T^i(j); \xi_T) dj \right] \tag{18}$$

where utility depends on a consumption index, C_t^i, of the economy's available goods (to be specified), a vector of aggregate preference shocks, ξ_t, and the amount of labor supplied for the production of each good j, $h^i(j)$. The second term in the brackets captures the total disutility of labor supply. The consumption index, C_t^i, is the Dixit-Stiglitz constant-elasticity-of-substitution aggregator of the economy's available goods and has an associated

price index written, respectively, as

$$C_t^i \equiv \left[\int_0^1 c_t^i(j)^{\frac{\theta-1}{\theta}} dj \right]^{\frac{\theta}{\theta-1}} \quad \text{and} \quad P_t \equiv \left[\int_0^1 p_t(j)^{1-\theta} dj \right]^{\frac{1}{1-\theta}}$$

where $\theta > 1$ is the elasticity of substitution between any two goods and $c_t^i(j)$ and $p_t(j)$ denote household i's consumption and the price of good j.[4]

$\hat{E}_t^i$ denotes the subjective beliefs of household i about the probability distribution of the model's state variables: that is, variables that are beyond agents' control though relevant to their decision problems – i.e. prices and exogenous variables. Beliefs are assumed to be homogenous across households, though each household has no knowledge of the beliefs of other households. Agents therefore do not have a complete economic model of the determination of aggregate state variables and it is about the evolution of such variables that agents are attempting to learn. The learning algorithm is discussed in the appendix. The discount factor is assumed to satisfy $0 < \beta < 1$.

Asset markets are assumed to be *incomplete*: there is a single one-period riskless non-monetary asset available to transfer wealth intertemporally. Under this assumption, the household's flow budget constraint can be written as

$$M_t^i + B_t^i \leq \left(1 + i_{t-1}^m\right) M_{t-1}^i + \left(1 + i_{t-1}\right) B_{t-1}^i + P_t Y_t^i - T_t - P_t C_t^i \tag{19}$$

where M_t^i denotes the household's end-of-period holdings of money, B_t^i the household's end-of-period nominal holdings of risk-less bonds, i_t^m and i_t are the nominal interest rates paid on money balances and bonds held at the end of period t, Y_t^i the period income (real) of households and T_t denotes lump sum taxes and transfers. The household receives income in the form of wages paid, $w(j)$, for labor supplied in the production of each good, j. Furthermore, all household's i are assumed to own an equal part of each firm and therefore receive a common share of profits $\Pi_t(j)$ from the sale of each firm's good j. Period nominal

[4] The absence of real money balances from the period utility function (18) reflects the assumption that there are no transaction frictions that can be mitigated by holding money balances. However, agents may nonetheless choose to hold money if it provides comparable returns to other available financial assets.

income is therefore determined as $P_t Y_t^i = \int_0^1 [w_t(j) h_t^i(j) + \Pi_t(j)] dj$ for each household i. Fiscal policy is assumed to be Ricardian.

To summarize, the household's problem in each period t is to choose $\{c_t^i(j), h_t^i(j), M_t^i, B_t^i\}$ for all $j \in [0,1]$ so as to maximize (18) subject to the constraint (19) taking as parametric the variables $\left\{p_T(j),\ w_T(j),\ \Pi_T,\ i_{T-1},\ i_{T-1}^m,\ \xi_T\right\}$ for $T \geq t$.

Firms face a Calvo-style price-setting problem in period t and maximize the expected present discounted value of profits

$$\hat{E}_t^i \sum_{T=t}^{\infty} \alpha^{T-t} Q_{t,T} \left[\Pi_T^i(p_t(i))\right] \tag{20}$$

where

$$\Pi_T^i(p) = y_t(i)\, p_t(i) - w_t(i) h_t(i) \tag{21}$$

and $y_t(i) = Y_t(p_t(i)/P_t)^{-\theta}$ is the demand curve faced by producer i and output is produced using technology $y_t(i) = A_t f(h_t(i))$ where A_t is an exogenous technology shock. The factor α^{T-t} in the firm's objective function is the probability that the firm will not be able to adjust its price for the next $(T-t)$ periods.

Finally firms are assumed to value future streams of income at the marginal value of aggregate income in terms of the marginal value of an additional unit of aggregate income today. That is, a unit of income in each state and date T is valued by the stochastic discount factor $Q_{t,T} = \beta^{T-t} P_t U_c(Y_T, \xi_T)/(P_T U_c(Y_t; \xi_t)$. This simplifying assumption is appealing in the context of the symmetric equilibrium that is examined in this model. To summarize, the firm's problem is to choose $\{p_t(i)\}$ to maximize (20) taking as given $\{Y_T,\ P_T,\ w_T(j),\ A_T,\ Q_{t,T}\}$ for $T \geq t$ and $j \in [0,1]$.

3.2 Optimal Decision Rules

Preston (2005) shows that a log-linear approximation to the first order conditions of the above infinite horizon decision problems lead to following optimal decision rules. Each household

allocate consumption according to

$$\hat{C}_t^i - Y_t^n = (1-\beta)\,\varpi_t^i + \hat{E}_t^i \sum_{T=t}^{\infty} \beta^{T-t}\left[(1-\beta)x_T - \beta\sigma(\hat{\imath}_T - \hat{\pi}_{T+1}) + \beta r_T\right].$$

where

$$\hat{C}_t^i \equiv \ln(C_t^i/\bar{Y});\ \hat{Y}_t \equiv \ln(Y_t/\bar{Y});\quad \hat{\imath}_t \equiv \ln[(1+i_t)/(1+\bar{\imath})]; \tag{22}$$

$$\hat{\pi}_t = \ln(P_t/P_{t-1});\quad \varpi_t^i = W_t^i/(P_t\bar{Y})$$

and for any variable g, $\bar{g}$ denotes the variables steady state value. The output gap $x_t = \hat{Y}_t - \hat{Y}_t^n$ is the difference between output and the economy's natural rate of output that obtains in this model with fully flexible prices. It is a function of model disturbances as is the natural rate of interest r_t. $\sigma = -U_{cc}/U_c\bar{Y} > 0$ is the intertemporal elasticity of substitution.

The connection of this relation to the predictions of permanent income theory and the canonical consumption model of the previous section is immediate. The first two terms capture precisely the basic insight of the permanent income hypothesis that agents should consume a constant fraction of the expected future discounted wealth, given a constant real interest rate equal to $\beta^{-1} - 1$. The third term arises from the assumption of a time-varying real interest rate, and represents deviations from this constant real rate due to either variation in the nominal interest rate or inflation. The final term results from allowing stochastic disturbances to the economy.

Each firm determines its optimal price according to

$$\widehat{p}_t^*(i) = \hat{E}_t^i \sum_{T=t}^{\infty} (\alpha\beta)^{T-t}\left[\frac{1-\alpha\beta}{1+\omega\theta}\cdot(\omega+\sigma^{-1})x_T + \hat{\pi}_T\right]$$

where $\widehat{p}_t^*(i) \equiv \ln(p_t^*(i)/P_t)$ and $\omega > 0$ is the elasticity of real marginal costs with respect to own output. Here the presence of long-horizon expectations arise due to the pricing frictions induced by Calvo pricing. When a firm has the opportunity to change its price in period t there is probability α^{T-t} that the firm will not get to change its price in the subsequent $T-t$ periods. The firm must therefore concern itself with macroeconomic conditions relevant to marginal costs into the indefinite future when deciding the current price of its output.

Finally, to complete the specification of household and firm behavior, beliefs are specified as

$$z_t = a_t + b_t z_{t-1} + \eta_t$$

where $z_t \equiv [x_t \ \pi_t \ i_t \ r_t]'$ and a_t and b_t coefficient matrices of appropriate dimension and η_t vector white noise.[5]

For precisely the same reasons articulated earlier, neither of these optimal decision rules can be written in a way that requires only one-period-ahead forecasts. Optimality requires long-horizon expectations of general macroeconomic conditions. All points made in regards to the permanent income model discussed in Section 2 apply equally well to the household decision problem presented here. Similarly the discussion of the Townsend investment model extends directly and naturally to the Calvo-price setting problem.

Integrating over the optimal decision rules of the above household and firm decision problems, applying the market clearing condition $\int_0^1 \hat{C}_t^i = \hat{Y}_t$ and the definition of x_t, gives a simple New Keynesian sticky-price model of output gap and inflation determination, comprising aggregate demand and supply relations of the form:

$$x_t = \hat{E}_t \sum_{T=t}^{\infty} \beta^{T-t} \left[(1-\beta)x_{T+1} - \sigma(i_T - \pi_{T+1}) + r_T\right] \tag{23}$$

and

$$\pi_t = \kappa x_t + \hat{E}_t \sum_{T=t}^{\infty} (\alpha\beta)^{T-t} \left[\kappa\alpha\beta \cdot x_{T+1} + (1-\alpha)\beta \cdot \pi_{T+1} + u_T\right] \tag{24}$$

where x_t is the output gap, π_t the inflation rate, i_t the nominal interest rate and r_t and u_t are exogenous disturbance terms, with all variables being properly interpreted as log-deviations from steady state values.[6] $\sigma > 0$ is the intertemporal elasticity of substitution and $\kappa > 0$. $\hat{E}_t$ denotes the average expectations of individuals in this economy.

Neither the aggregate demand relation (23) nor the Phillips curve (24) can be simplified under arbitrary subjective expectations about the evolution of state variables if agents are

[5] As for the canonical consumption model, beliefs are a primitive of the decision problem. They are specified here in terms of log-linearized variables purely to save on notational and computational complexity.

[6] The cost-push shock u_t has been added to ensure a non-trivial stabilization for the central bank detailed below. The '^' notation has been dropped for simplicity.

optimizing, as demonstrated in Preston (2005). However, under the assumption of rational expectations these relations simplify to give:

$$x_t = E_t x_{t+1} - \sigma^{-1}(i_t - E_t \pi_{t+1}) + r_t \tag{25}$$

$$\pi_t = \beta E_t \pi_{t+1} + \kappa x_t + u_t. \tag{26}$$

A rational expectations equilibrium analysis implies that once the true probability laws are known, only one-period-ahead expectations matter for aggregate dynamics. It is clear that learning has important implications for aggregate economic dynamics: with subjective expectations agents optimally require long-horizon expectations of macroeconomic conditions into the indefinite future. The presence of these expectational variables is important for the study of monetary policy, as expectations represent an important source of instability.

The recent work of Bullard and Mitra (2000, 2002) and Evans and Honkapohja (2002, 2003) take these relations and replace the rational expectations assumption with the adaptive learning assumption to give the system

$$x_t = \hat{E}_t x_{t+1} - \sigma(i_t - \hat{E}_t \pi_{t+1}) + r_t \tag{27}$$

$$\pi_t = \kappa x_t + \beta \hat{E}_t \pi_{t+1} + u_t. \tag{28}$$

where $\hat{E}_t$ indicates expectations are taken with respect to distribution implied by whatever beliefs may be specified. Honkapohja, Mitra, and Evans (2002) argue that (27) and (28) are implied (23) and (24) so long as i) the law of iterated expectation holds and ii) agents are endowed with knowledge of the market clearing conditions today and in all future periods. The analysis of the previous sections should make clear that these conditions are not sufficient for (27) and (28) to be consistent with the microfoundations presented here and in Preston (2003).

Note also that the assumption that agents know the market clearing conditions in all future periods is questionable. It might seem to have appeal because of the trivial nature of this condition in the present model – i.e. that $C_t = Y_t$. But its apparent simplicity derives from the fact that we are analyzing an endowment economy. In more general modeling

contexts, such as when there are non-trivial wealth dynamics and agents have incentives to trade assets in equilibrium, this market clearing condition will take a more complicated form and there is no obvious reason why agents ought to be informed about the nature of aggregate resource constraints or market clearing conditions. These are objects of the economic environment that agents are attempting to learn about.

The next section turns to a discussion of a number of policies for which their desirability for stabilizing economic fluctuations differs markedly depending on whether the Euler equation or long-horizon expectations approach is adopted.

3.3 Policy Implications

Having laid out a theory for modeling learning in infinite horizon decision problems and contrasted the approach with the more common Euler equation approach it is natural to ask whether these modeling differences lead to important differences in conclusions about stability of model economies. We now turn to the specific issue of stability of learning dynamics and monetary policy.

Consider a central bank seeking minimize the loss function

$$W = E_{t_0} \sum_{t=t_0}^{\infty} \beta^{t-t_0} \left[\pi_t^2 + \lambda x_t^2 \right] \tag{29}$$

where $0 < \beta < 1$ is the household's discount factor subject to the constraints (25), (26) and the additional constraint that $\pi_{t_0} = \bar{\pi}_{t_0}$ where

$$\bar{\pi}_{t_0} = (1-\mu) \frac{\lambda}{\kappa} x_{t-1} + \frac{\mu}{1-\beta\mu\gamma} u_t.$$

Thus the central bank wishes to stabilize variation in inflation and the output gap, and $\lambda > 0$ determines the relative importance of these stabilization objectives. Woodford (2003, chap. 7) shows that this decision problem is completely characterized by the relation

$$\pi_t = -\frac{\lambda}{\kappa} (x_t - x_{t-1}) \tag{30}$$

in the sense that if the central bank can arrange for this target criterion to be satisfied at all times it will successfully implement the optimal policy associated with the above decision problem.

Evans and Honkapohja (2002) show that the policy rule

$$i_t = \frac{1}{\sigma}\left[\hat{E}_t x_{t+1} - \frac{\lambda}{\lambda+\kappa^2}x_{t-1} + \left(\frac{\beta\kappa}{\lambda+\kappa^2}+\sigma\right)\hat{E}_t\pi_{t+1} + \frac{\kappa}{\lambda+\kappa^2}u_t + r_t^n\right] \tag{31}$$

is consistent with implementing the target criterion (30) in an economy given by (27) and (28). Beliefs are then assumed to be given by

$$\hat{E}_t z_{t+1} = a_t + b_t r_t + c_t u_t$$

with obvious notation and $z_t = \{\pi_t,\ x_t\}$. The instrument rule (31) guarantees convergence in learning dynamics to the optimal rational expectations equilibrium for all maintained parameter assumptions. However, in the context of the model given by the structural equations (23) and (24) and beliefs (2) where the state vector z_t is redefined to be$\{\pi_t,\ x_t,\ i_t\}$ such rules give rise to divergent learning dynamics for many parameter values. These results are summarized in the following proposition.

Proposition 3 *In an economy given by (27) and (28) the instrument rule (31) guarantees convergence to the associated rational expectations equilibrium for all parameter values under the maintained assumptions. In contrast, in an economy given (23) and (24) such an instrument gives rises to divergent learning dynamics for many parameter values, even though the two models are equivalent under the rational expectations assumption.*

Proof. See Evans and Honkapohja (2002) for the first result and Preston (2004) for the second. ■

This is a striking result given that both models are equivalent under the rational expectations assumption. Hence the instrument rule (31) gives rise to a different kind of out-of-equilibrium behavior in the long-horizon forecast model than in the Euler equation-based model. Indeed, the differences arise because the long-horizon forecast model requires agents to learn about a greater number of restrictions required by rational expectations equilibrium. In particular they need to forecast the path of interest rates which is not a property of the model given by (27) and (28). Agents therefore know less about the set of restrictions that must be satisfied in a rational expectations equilibrium. These instrument rules have the property that because they depend on agent's learning behavior they make

the central banks own decisions depend on such instability and hence promote the possibility of divergent learning dynamics.

Preston (forthcoming) demonstrates further examples of policy rules that can lead to markedly different conclusions about their desirability for stabilizing economic fluctuations depending the whether the Euler equation model of the long-horizon forecast model is adopted. It demonstrates that recently popular forecast-based instrument rules are similarly afflicted by instability under learning dynamics in a model with long-horizon learning. For instance, consider an instrument rule of the form

$$i_t = \psi_\pi \hat{E}_t \pi_{t+1} + \psi_x \hat{E}_t x_{t+1}$$

where $\psi_\pi, \psi_x > 0$. This Taylor-type rule is of the general form proposed by Batini and Haldane (1999) and Levin, Wieland, and Williams (2003) as being a desirable approach to implement policy on the ground of robustness and also given empirical support by Clarida, Gali and Gertler (1998, 2000).

Bullard and Mitra (2002) consider the performance of such rules in an economy given by (27) and (28). They find that the so-called Taylor principle

$$\kappa (\psi_\pi - 1) + (1 - \beta) \psi_x > 0$$

is necessary and sufficient for stability under learning dynamics. In contrast, Preston (forthcoming) shows in the model of long-horizon expectations given by (23) and (24) that the Taylor principle is no longer sufficient for stability. In fact, for many plausible parameter values, such rules require implausible policy responses to forecasted inflation for stability.

4 Conclusion

This paper represents an appeal to model learning dynamics in multi-period decision problems as first proposed by Marcet and Sargent (1989a) and more recently by Preston (2005). In multiperiod decision problems in which agents have arbitrary subjective beliefs about the evolution of state variables relevant to their decision problems optimal decision rules depend

on long-horizon forecasts. The logic of this result is identical to a rational expectations analysis of such models: optimal allocations today require forecasts of state variables in all states of uncertainty at all future dates.

The approach contrasts with many recent analyses of learning in macroeconomics in which decision rules are assumed to depend only on one-period-ahead forecasts. Such decision rules are often loosely based on Euler equations that characterize microfounded multiperiod decision problems under rational expectations, though assuming forecasts are formed under the learning assumption. Importantly, this approach is shown to imply sub-optimal decisions. An example is provided where agents violate their intertemporal budget constraint (transversality condition) even when learning dynamics converge. The optimal decision rule is not prone to such problems by virtue of appropriately accounting for all relevant constraints faced by the decision maker.

In a simple model useful for policy evaluation it is shown that many policies that have been argued to be desirable for stabilizing economic fluctuations when evaluated in one-period-ahead forecast models in fact lead to significant instability in the model based on optimal decision rules. This arises because the former attributes much greater knowledge about the underlying rational expectations equilibrium – about which agents are attempting to learn – than do the optimal decision rules. This might cast doubt on the motivation of the one-period-ahead forecast models being a more plausible model of boundedly rational agents.

A Appendix

This appendix first outlines the general approach to analyzing learning dynamics in the context of the model of this framework. It then turns to sketching the proofs of the central results which are all applications of this general methodology. Since the algebra underpinning these results is at times tedious, it is largely omitted. Most calculations were performed in Mathematica.

A.1 Expectational Stability

Consider the Towsend investment model. If agents know the form of the minimum-state-variable solution they estimate a linear model given by (2) where η_t is vector white noise and a_t and b_t are coefficient matrices of obvious dimension. Given estimates, forecasts of future state variables at any horizon T are determined according to (3) for $T \geq t$. The estimation procedure makes use of the entire history of available data in period t, $\{1, z_t\}_0^{t-1}$. As additional data become available, agents update their estimates of the coefficients (a_t, b). This is neatly represented as the recursive least squares formulation

$$\phi_t = \phi_{t-1} + t^{-1} R_t^{-1} w_{t-1}(z_t - \phi'_{t-1} w_{t-1}) \tag{32}$$

$$R_t = R_{t-1} + t^{-1}(w_{t-1} w'_{t-1} - R_{t-1}) \tag{33}$$

where the first equation describes how the forecast coefficients, $\phi_t = (a'_t, b'_t)'$, are updated with each new data point and the second the evolution of the matrix of second moments of the appropriately stacked regressors $w_t \equiv \{1, z_{t-1}\}_0^{t-1}$.

To obtain the actual law of motion, substitute (3) into the optimal decision rule and applying the market clearing condition $k_t = K_t$ together with () and () give

$$z_t = \bar{a}_t + \bar{b}_t z_{t-1} + \bar{c}_t \varepsilon_t. \tag{34}$$

where $\{\bar{a}_t, \bar{b}_t, \bar{c}_t\}$ are functions of the current private forecast parameters $\{a_t, b_t, c_t\}$ and model primitives. Comparison with (2) makes clear that agents are estimating a misspecified model of the economy – agents assume a stationary model when in fact the true model has time

varying coefficients. Leading this expression one period and taking expectations (rational) provides

$$E_t z_{t+1} = \bar{a}_t + \bar{b}_t z_t$$

which describes the optimal rational forecast conditional on private-sector behavior. Taken together with (34) at $T = t + 1$ it defines a mapping that determines the optimal forecast coefficients given the current private-sector forecast parameters (a'_t, b'_t), written as

$$T(a_t, b_t) = (\bar{a}_t, \bar{b}_t). \tag{35}$$

A rational expectations equilibrium (REE) is a fixed point of this mapping. For such REE, we are then interested in asking under what conditions does an economy with learning dynamics converge to this equilibrium. Using stochastic approximation methods, Evans and Honkapohja show that the conditions for convergence of the learning algorithm (32) and (33) are neatly characterized by the local stability properties of the associated ordinary differential equation

$$\frac{d}{d\tau}(a, b) = T(a, b) - (a, b), \tag{36}$$

where τ denotes "notional" time. The REE is said to be expectationally stable, or E-Stable, if this differential equation is locally stable in the neighborhood of the REE. From standard results for ordinary differential equations, a fixed point is locally asymptotically stable if all eigenvalues of the Jacobian matrix $D\left[T(a, b) - (a, b)\right]$ have negative real parts (where D denotes the differentiation operator and the Jacobian is understood to be evaluated at the rational expectations equilibrium of interest.) See Evans and Honkapohja (2001) for further details on expectational stability.

A.2 Proof of Proposition 1

[TO BE ADDED]

A.3 Proof of Proposition 2

[TO BE ADDED]

References

Batini, N., and A. G. Haldane (1999): "Forward-Looking Rules for Monetary Policy," in *Monetary Policy Rules*, ed. by J. Taylor. University of Chicago Press, Chicago.

Bernanke, B. S., and M. Woodford (1997): "Inflation Forecasts and Monetary Policy," *Journal of Money, Credit, and Banking*, 29(4), 653–686.

Bullard, J., and K. Mitra (2000): "Determinacy, Learnability and Monetary Policy Inertia," Working Paper no. 2000-030A, Federal Reserve Bank of St. Louis. [Available online at http://www.stls.frb.org/docs/research/wp/2000-001.pdf].

——— (2002): "Learning About Monetary Policy Rules," *Journal of Monetary Economics*, 49(6), 1105–1129.

Clarida, R., J. Gali, and M. Gertler (1998): "Monetary Policy Rules in Practice: Some International Evidence," *European Economic Review*, 42, 1033–1067.

——— (1999): "The Science of Monetary Policy: A New Keynesian Perspective," *Journal of Economic Literature*, 37, 1661–1707.

——— (2000): "Monetary Policy Rules and Macroeconomic Stability: Evidence and Some Theory," *Quarterly Journal of Economics*, 115, 147–180.

Evans, G. W., and S. Honkapohja (2001): *Learning and Expectations in Economics*. Princeton, Princeton University Press.

——— (2002): "Monetary Policy, Expectations and Commitment," *Working Paper*.

——— (2003): "Expectations and the Stability Problem for Optimal Monetary Policies," *Review of Economic Studies*, 70(4).

Honkapohja, S., and K. Mitra (2003): "Performance of Monetary Policy with Internal Central Bank Forecasting," *Journal of Economic Dynamics and Control*, Forthcoming.

Honkapohja, S., K. Mitra, and G. W. Evans (2002): "Notes on Agents' Behavioral Rules Under Adaptive Learning and Recent Studies of Monetary Policy," Unpublished working paper. Available at: http//www.valt.helsinki.fi/raka/seppo.htm.

Levin, A., V. Wieland, and J. C. Williams (2003): "Robustness of Forecast-Based Monetary Policy Rules under Model Uncertainty," *American Economic Review*, forthcoming.

Marcet, A., and T. J. Sargent (1989a): "Convergence of Least-Squares Learning in Environments with Hidden State Variables and Private Information," *Journal of Political Economy*, pp. 1306–1322.

——— (1989b): "Convergence of Least Squares Learning Mechanisms in Self-Referential Linear Stochastic Models," *Journal of Economic Theory*, (48), 337–368.

Preston, B. (2004): "Adaptive Learning and the Use of Forecasts in Monetary Policy," unpublished, Columbia University.

——— (2005): "Learning About Monetary Policy Rules when Long-Horizon Expectations Matter," *International Journal of Central Banking*, 1(2), 81–126.

——— (forthcoming): "Adaptive Learning, Forecast-Based Instrument Rules and Monetary Policy," *Journal of Monetary Economics*.

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

Townsend, R. M. (1983): "Forecsating the Forecasts of Others," *The Journal of Political Economy*, pp. 546–588.

Woodford, M. (1999): "Optimal Monetary Policy Inertia," NBER Working Paper 7261.

——— (2003): *Interest and Prices: Foundations of a Theory of Monetary Policy*. Princeton University Press.