Adaptive Learning and Business Cycles

(Preliminary and Incomplete)

Noah Williams

Department of Economics, Princeton University and NBER
E-mail: noahw@princeton.edu

Revised January 9, 2003

This paper analyzes the quantitative importance of adaptive learning in business cycle fluctuations. We first introduce adaptive learning in a real business cycle model and a New Keynesian model, using specifications drawn from the literature which assume that agents learn about the equilibrium laws of motion. We consider a variety of learning rules, and find that in both environments learning has very minor effects on the volatility and the persistence of the key economic variables. However we discuss some potential theoretical drawbacks to this formulation of learning, and consider an alternative formulation in which agents learn about the structural features of the economy. In some simplified settings, we show that structural learning has much greater effects. We also illustrate how learning with misspecified beliefs can lead to fluctuations of a different kind, as agents "escape" from an equilibrium. Overall, our results show that the importance of learning depends greatly on the specification of beliefs.

1. INTRODUCTION

Many theoretical economic models have difficulty in matching the volatility and persistence observed in economic data. The inherent amplification and propagation mechanisms in many models are relatively weak, so that most of the persistence in the theoretical models is inherited from the persistence of exogenous stochastic shocks. In order to better match the data, many models have incorporated adjustment costs, decision delays, or additional frictions in order to prolong the effects of a shock. In this paper we study the importance of adaptive learning as a source of amplification and propagation in some standard economic models.

When agents are uncertain about their economic environment, exogenous shocks lead to revisions of beliefs over time, which may draw out the effects of a shock. Moreover the belief updating process may lead to more overall volatility, as agents' subjective views of the economy change over time and alter their decisions. We seek to quantify the size of these potential effects in explaining business cycle fluctuations. In broad terms, the idea of this paper follows Barsky and DeLong (1993), Timmermann (1993), and Timmermann (1996) who looked at effects of learning on stock market. Timmermann (1996) in particular showed how adaptive learning can lead to the excess volatility and predictability of returns. This suggests the possibility that learning in a business cycle framework could have similar effects.

We focus throughout on adaptive learning models, which have been widely studied in macroeconomics. A distinct but related literature in macroeconomics focuses on "learning" in the sense of filtering or signal extraction in partial information environments. Models in this line of research have dealt more explicitly with the issues we focus on here. In contrast, typical papers in the adaptive learning literature focus on the question of whether adaptive agents, who base their actions on simple learning rules, could eventually learn a rational expectations equilibrium. This provides a foundations for rational expectations models, and limits focus to equilibria which are "learnable".[1] The focus of this paper however is not on the stability of equilibria, but on the quantitative importance of learning for outcomes.

We first follow much of the learning literature by supposing that instead of rational expectations, agents forecast future outcomes using a statistical model. They update this model over time as they observe data. In particular, we initially assume that agents form beliefs about the equilibrium reduced form law of motion for the economy. We introduce adaptive learning in this form into a real business cycle model and a New Keynesian model. The stability of rational expectations under learning has been established in the literature, with Evans and Honkapohja (2001), Packalen (2000), and Bullard and Duffy (2002) studying the RBC model and Evans and Honkapohja (2002b) and Bullard and Mitra (2002) the New Keynesian model. However, to our knowledge, the quantitative importance of learning has not been addressed.[2] A number of modifications have needed to be introduced in both models (see King and Rebelo (1999) and Christiano, Eichenbaum, and Evans (2001) for example) in order to improve their empirical fit, largely by increasing the persistence of responses to external shocks. Learning thus may add to or substitute for some of these modifications in improving the performance of the models.

While theoretical results from the theory of stochastic approximation (see Kushner and Yin (1997)) may help to characterize the effects of learning dynamics analytically, we follow the "calibration" literature by conducting some simulation studies. We find that the effects of learning are very small in both environments. Learning only slightly increases the volatility and persistence of the key economic variables over a horizon comparable to the post-war period. However even over the very short run, assuming that agents have very relatively little prior information, the effects of learning are quite modest in both models. Agents who learn in this manner rather quickly learn to have rational expectations, and so the learning process has very little consequence for outcomes.

However we discuss below some potential difficulties in the interpretation of this specification of learning. By learning about equilibrium laws of motion, the distinction between beliefs and decisions is blurred and the learning rules are difficult to interpret. We then develop a natural alternative, where agents learn about the structure of the economy instead of its reduced form. We show in some simplified settings that in this case learning leads to much more substantial effects. By learning about the structural parameters, we make more explicit how beliefs influence decisions. Decision rules are typically nonlinear functions of the underlying structural parameters, which amplifies

[1]See Evans and Honkapohja (2001) for a monograph on these issues, with a wide variety of applications.

[2]Bullard and Duffy (2002) study the quantitative importance of learning for the trend-cycle decomposition in the RBC model.

the effect of variations in structural beliefs. Moreover, by more explicitly modeling agents' decision problems, we are naturally led consider what happens if agents' beliefs are misspecified. We develop a simple model with an aggregate production externality that illustrates that misspecified learning rules may lead to an alternate source of long-run fluctuations. Agents' beliefs occasionally, but recurrently, "escape" from their equilibrium values, which leads to significant changes in outcomes. We apply the results of Williams (2002a) to characterize these fluctuations, showing that they take a particular, predictable form in which agents internalize the external effect.

While we argue that the structural learning specification is more natural and has more significant effects in some simple settings, we do not quantify its importance in a more reasonably calibrated model. As we discuss below, there are difficulties in formulating structural learning in the baseline RBC model which remain to be addressed.[3] More generally, our results illustrate the importance of explicitly modeling agents' beliefs and decisions, and show that the specification of learning rules can have a substantial effect on outcomes. Clearly more work is needed to sort through these issues and to arrive at a learning specification which is consistent with individual behavior and with the aggregate data.

2. THE EFFECTS OF LEARNING: REDUCED FORMS

2.1. Overview

In this section we investigate the importance of introducing learning in two well-known "canonical" models: the real business cycle model (RBC), as discussed in Cooley (1995), and the baseline "New Keynesian" (NK) monetary model, as discussed in Woodford (2002). Both models are known to have relatively weak propagation mechanisms, and modifications such as adjustment costs have needed to be introduced in both in order to improve their empirical fit. To fit the data reasonably well, both specifications we study require persistent economic shocks. The goal of this section is to see if learning may add to or substitute for some of these modifications in improving the performance of the models.

In dynamic macroeconomic settings, the most widely used specification of learning has focused on Euler equations. In particular, most studies have applied the methodology described in Evans and Honkapohja (2001). In this framework, the structural model of the economy is used to formulate a system of linear expectational difference equations. As is a standard practice in much of macroeconomics, this is accomplished by log-linearizing the Euler equations and laws of motion which summarize the dynamic equilibrium in the model.[4] This leads to systems of the form:

$$\begin{aligned} y_t &= A_0 + A_1 E_t^* y_{t+1} + A_2 y_{t-1} + A_3 z_t \\ z_t &= B_0 + B_1 z_{t-1} + w_t, \end{aligned} \tag{1}$$

[3]Evans and Honkapohja (2002b) and Evans and Honkapohja (2002a) discuss a form of structural learning by policymakers in the New Keynesian model.

[4]See Williams (2002b) for results justifying linearizations for some of the dynamic properties of models.

where y_t is a vector of endogenous variables, z_t is a vector of exogenous variables, and w_t is a martingale difference sequence of stochastic shocks. As is well-known, many dynamic macro models can be put into this form, and there are efficient solution methods for computing equilibria when E_t^* corresponds to rational expectations.

The learning literature focuses on cases in which the E_t^* is instead associated with a *perceived law of motion* (PLM), which in most cases takes the form of the reduced form equilibrium law of motion for the economy. In this example, a natural PLM is:

$$y_t = a_0 + a_1 y_{t-1} + a_2 z_t. \tag{2}$$

If agents use (2) to form expectations, then substituting these into (1) leads to the *actual law of motion*:

$$y_t = T_0 + T_1 y_{t-1} + T_2 z_t, \tag{3}$$

where the T_i constants depend on the perceptions a_i and the true structural parameters A_i and B_i. This "T-map" is the key to the notion of "E-stability" and also governs the (local) stability of adaptive learning rules (see Marcet and Sargent (1989) and Evans and Honkapohja (2001)).

Under adaptive learning, the parameters a_i in (2) are updated according to a statistical algorithm each period as agents observe data. We focus on different specifications of the recursive least squares algorithm, which have the following structure. As of date t, agents' beliefs are summarized by the vector of coefficients $a_t = [a_{0t}', a_{1t}', a_{2t}']'$, and we denote the regressors $x_t = [1, y_t', z_t']'$. The algorithm accounts for the volatility of the regressors by updating a matrix R_t which is an estimate of the second moment matrix of x_t. The updating rule is:

$$a_{t+1} = a_t + \varepsilon_t R_t^{-1} x_t \left(y_t - a_t' x_t\right) \tag{4}$$
$$R_{t+1} = R_t + \varepsilon_t \left(R_t - x_t x_t'\right). \tag{5}$$

Here ε_t is a sequence known as the *gain*. We consider the least squares (LS) case in which $\varepsilon_t = \frac{1}{t+1}$ and we also consider constant gain (CG) settings in which $\varepsilon_t = \varepsilon$. Constant gain learning rules make learning a persistent process, as agents discount past data and continually pay equal attention to new observations.

A key question of interest is whether adaptive agents who use such learning rules would eventually converge to a rational expectations equilibrium. While this is an important issue, here we focus not on the stability of equilibria, but on the quantitative importance of learning for outcomes. Stability is a prerequisite for our analysis, and this has been addressed in the RBC model by Evans and Honkapohja (2001), Packalen (2000), and Bullard and Duffy (2002) and in the NK model by Bullard and Mitra (2002) and Evans and Honkapohja (2002b). We use the specifications in these papers to simulate the models and analyze the effects of introducing learning. We study different specifications of the learning rule that agents are assumed to employ, and we summarize the ability of learning to add additional persistence and volatility to the models.

We find that this formulation of learning, focusing on the reduced form laws of motion, has very modest effects. Later in the paper we then study the effects of *structural* learning in simpler settings, which has much more sizable effects. We also discuss there some difficulties in the interpretation of these reduced-form learning rules as a model of agent learning. This opens the possibility that more plausible specifications lead to different results, as our results there suggest.

2.2. A Real Business Cycle Model

A natural place to begin the study of the quantitative importance of learning is with a real business cycle (RBC) model. RBC models are arguably the benchmark for all dynamic macro models, and the New Keynesian model discussed below developed as an extension of them. While the empirical success or failure of the RBC approach has always been a contentious issue, it has resurfaced with recent papers by Gali (1999) and Francis and Ramey (2001), who show that estimated technology shocks do not produce business cycle patterns in the data. Further, as mentioned above the RBC model has been extended to incorporate additional propagation mechanisms to improve its match to the data. We address whether learning can add additional amplification and propagation mechanisms, which could potentially make the model consistent with smaller technology shocks and eliminate some of the ad-hoc adjustment costs. However, as we discuss below, we find that in our model that this does not appear to be the case. Significant effects from learning seem to require a different formulation. One possibility is studied later in the paper.

We study a benchmark calibrated RBC model from Cooley and Prescott (1995). In the planning version of the model, a social planner chooses sequences of consumption, capital, and employment to solve:

$$\max \; E\sum_{t=0}^{\infty}\left(\beta\frac{1+\eta}{1+\gamma}\right)^{t}\left[\log C_t+\theta\log(1-h_t)\right] \tag{6}$$

$$\begin{aligned}\text{s.t. } (1+\gamma)(1+\eta)K_{t+1} &= (1-\delta)K_t+e^{z_t}K_t^{\alpha}h_t^{1-\alpha}-C_t\\ z_t &= \rho z_{t-1}+w_t\end{aligned}$$

Here γ is the growth rate of technology, η is the growth rate of the labor force, and all variables except hours h_t are measured in efficiency units, normalizing by the level of technology and the labor force. The rest of the notation is standard: C_t is consumption, K_t capital, h_t hours (taken as a fraction between zero and one), and z_t the AR(1) technology shock. Throughout we use the parameter values calibrated in Cooley and Prescott (1995).

As is well-known, this model is nearly log-linear, so we employ the standard log-linearization. Letting lower case letters denote logarithms, this allows us to represent the equilibrium of the economy via the expectational difference equation system:

$$\begin{aligned}c_t &= A_{10}+A_{11}E_t^*c_{t+1}+A_{12}E_t^*k_{t+1}+A_{13}E_t^*z_{t+1}\\ k_t &= A_{20}+A_{21}c_{t-1}+A_{22}k_{t-1}+A_{23}z_{t-1}\\ z_t &= \rho z_{t-1}+w_t.\end{aligned} \tag{7}$$

Details of the linearization are given in Bullard and Duffy (2002) who pay special attention to the constant terms, and provide explicit formulas for the constants A_{ij} in terms of the underlying parameters.

As Evans and Honkapohja (2001) note, there two main alternatives for formulating a reduced-form PLM for this economy. The first natural candidate, which is that used by Bullard and Duffy (2002), takes the form:

$$\begin{aligned} c_t &= a_{10} + a_{11}c_{t-1} + a_{12}k_{t-1} + a_{13}z_{t-1} \\ k_t &= a_{20} + a_{21}c_{t-1} + a_{22}k_{t-1} + a_{23}z_{t-1}. \end{aligned} \tag{8}$$

There is a problem with learning based on this specification, however, as consumption is an endogenous variable which is a linear combination of capital and the technology shock. Thus there is perfect collinearity which causes the learning algorithms to explode.[5] A conceptually simple way around this problem is to add a exogenous shock to the consumption equation in (7), which breaks the collinearity. If the shock is sufficiently small its importance is minor. However, the fact that the system remains nearly collinear leads to some fragility in its behavior. An alternative PLM proposed by Evans and Honkapohja (2001) avoids this problem. It uses only the state (or predetermined) variables as regressors and forecasts consumption within a period:

$$\begin{aligned} c_t &= b_{10} + b_{11}k_t + b_{12}z_t \\ k_t &= b_{20} + b_{21}k_{t-1} + b_{22}z_{t-1}. \end{aligned} \tag{9}$$

In this section, we focus on the effects of learning based on the first specification (8). Similar results in this model obtained using the alternative (9). However when we consider a simplified RBC model later in the paper, we use the PLM (9) which proved slightly more stable in that environment.

To analyze the impact of learning in the RBC model, we ran a number of simulations of different specifications. For each specification, we first ran the economy in the rational expectations equilibrium for a small number (twenty) of "training" periods and estimate the parameters of the PLM (non-recursively) via OLS. This served as the initial condition for agent's beliefs. We then ran the simulation for 150 periods under adaptive learning, to roughly correspond to a typical post-war quarterly data set. We considered both the recursive least squares (LS) learning rule and several different constant gain (CG) settings, repeating each specification 1000 times. For some of the simulation runs, the learning algorithm led to explosive outcomes. Rather than set up a "projection facility" as in Marcet and Sargent (1989) to guide beliefs back to a stable region, we instead simply re-started the simulation run. Thus the results to be follow are conditioned on non-explosion, which seems a reasonable restriction.

Table 1 summarizes our results. For each specification, the table lists the mean across simulations of the different time series statistics. As is clear from the table, the effects of learning in this model are very small. Most specifications result in slightly

[5]The recursive least squares learning algorithm requires an inversion of the second moment matrix of the regressors. The collinearity makes this matrix singular, and thus the inversion becomes unstable.

TABLE 1.
Summary statistics from 1000 simulations of 150 observations of the RBC model under rational expectations and different learning rules.

	Rational Expectations	Least Squares	Constant Gain $\varepsilon = 0.1$	Constant Gain $\varepsilon = 0.05$	Constant Gain $\varepsilon = 0.03$
$std(y_t)$	3.0459	3.0172	3.0969	3.1062	3.1481
$std(c_t)$	1.9105	1.9099	1.9503	1.9262	1.9207
$std(k_t)$	2.4447	2.4268	2.5217	2.5524	2.6201
$corr(y_t, y_{t-1})$	0.9125	0.9096	0.9152	0.9158	0.9166
$corr(c_t, c_{t-1})$	0.9779	0.9766	0.9787	0.9798	0.9794

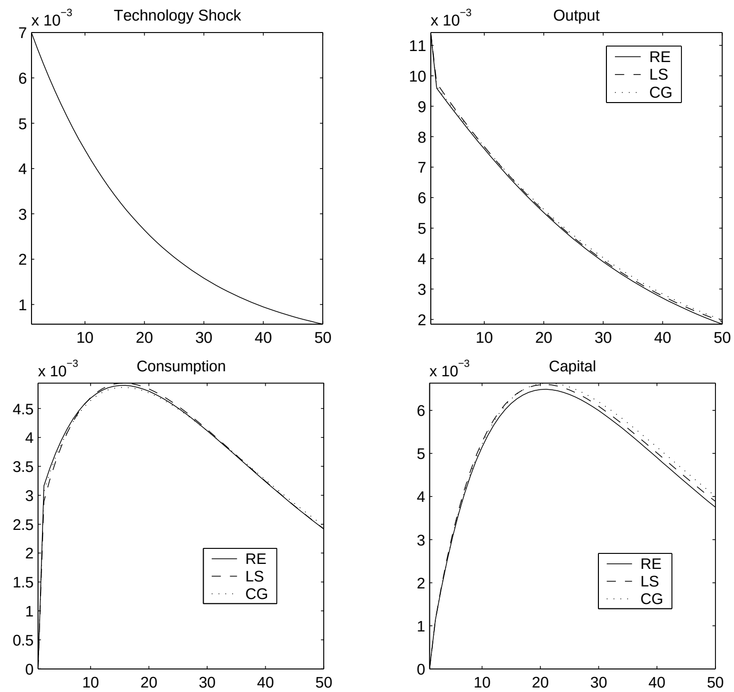

FIGURE 1. Mean impulse and pseudo-impulse responses of different variables in the RBC model to a technology shock. RE= rational expectations, LS=least squares, CG= constant gain with $\varepsilon = 0.03$.

higher volatility for most of the variables, and a slight increase in the persistence of output and consumption, but the effects are very small. Further, for output and the capital stock there is an increase in volatility as the size of the gain ε shrinks in the constant gain specifications. This suggests that most of the minor changes that we do find are due to the predictable path of convergence toward the rational expectations equilibrium, not to fluctuations around it. Smaller gain settings imply slower learning, but smaller fluctuations around a limit point. The results indicate that the fluctuations around the limit are very small, even for relatively large gain settings. This point was also confirmed by directly initializing the simulations at the limit point instead of using the training data. In this case, the differences from rational expectations were incredibly minute (roughly the order of machine precision for my computer).

The results in Table 1 suggest that in this model, agents learn to have rational expectations very quickly. The overall picture is one of very slight changes in volatility, mostly due to the convergence of beliefs. However it is possible that at least at the beginning of a sample period, learning may increase the persistence of the response of the economy to exogenous shocks. To investigate this, we ran an additional set of simulations in which we initialized agents' beliefs with some training data as above. Then, rather than simulate the economy, we simply input a one standard deviation technology shock and traced out the pseudo-impulse responses of the different variables. These are only pseudo-impulse responses because the learning process adds some nonlinearities to the model. However these nonlinearities proved to be slight, as nearly identical results obtained if the shock size was increased to two standard deviations.

Figure 1 summarizes the results. Shown in the figure are the impulse responses under rational expectations, and the mean across simulations of the pseudo-impulse responses for the LS learning rule and the CG setting with $\varepsilon = 0.03$. Again we see that, even at the beginning of the sample and with a relatively short set of training data (five years), there is very little impact of learning. The only noticeable change comes in the response of the capital stock, where the learning rules add a bit more persistence on average, but even this effect is quite small.

2.3. A New Keynesian Model

In this section we analyze the impact of learning in a baseline New Keynesian (NK) model following Woodford (2002) and Clarida, Gali, and Gertler (1999), which is the basis for much current monetary policy research. A number of papers have studied the properties of different policy rules models of this type, and much attention has focused on the determinacy and more recently, as in Bullard and Mitra (2002) and Evans and Honkapohja (2002b), the "learnability" of equilibria in this environment. However it is also known that the simplest specifications of this model, which are completely forward looking, do not match important features of the data. To lead to better empirical fit, in one of the original formulations of this class of models Rotemberg and Woodford (1997) added decision delays, while more recently Christiano, Eichenbaum, and Evans (2001) have added additional sources of stickiness and adjustment costs. Thus again, it is of interest to know whether learning could potentially lead to increased volatility and persistence in this environment. While the results of the previous section suggest that the effect of reduced-form learning in the RBC model is small, it is unclear how broadly these results apply. In this section, we show that in the NK model as well the effects of reduced-from learning are very modest.

The baseline NK model assumes monopolistic competition and sticky prices in the form of staggered price setting, and introduces money through currency in the utility function. It is laid out in Woodford (2002) among other places, and we do not reproduce the derivation here, but instead proceed directly to the expectational difference equations governing the rational expectations equilibrium.[6] These take the

[6] This may not be without consequence, as Preston (2002) argues.

form of an "IS equation" (log-linearized Euler equation) and a forward-looking Phillips curve:

$$\begin{aligned} y_t &= E_t^* y_{t+1} + \psi(r_t - r_t^n - E_t^* \pi_{t+1}) \\ \pi_t &= \kappa y_t + \beta E_t^* \pi_{t+1} \\ r_t^n &= \rho r_{t-1}^n + w_t. \end{aligned} \tag{10}$$

Here y_t is the output gap, π_t is the inflation rate, r_t is the nominal interest rate, and r_t^n is the "natural rate of interest," which is the driving shock process. All variables are in percentage deviations from their steady state values. We supplement the model with a policy rule, which we take to be a simple Taylor rule. In fact, for simplicity we focus on "the" Taylor (1993) rule:

$$r_t = 1.5\pi_t + 0.5y_t.$$

Substituting this specification into (10) eliminates r_t. We take the remaining parameters of the model from the calibration of Woodford (1999), which in turn is based on the estimation results of Rotemberg and Woodford (1997).

Bullard and Mitra (2002) studied the stability of rational equilibria under adaptive learning in this model for a variety of simple policy rules, including the Taylor rule we adopt. We follow them in analyzing the natural reduced-form PLM:

$$\begin{aligned} y_t &= a_{10} + a_{11} r_t^n \\ \pi_t &= a_{20} + a_{21} r_t^n. \end{aligned} \tag{11}$$

As in the previous section, we run a number of simulations in order to analyze the impact of learning in this model. We proceed exactly as above by generating a small sample to initialize agents' beliefs, and then run 1000 simulations of 150 observations each. Table 2 summarizes our results.

TABLE 2.

Summary statistics from 1000 simulations of 150 observations of the New Keynesian model under rational expectations for different learning rules.

	Rational Expectations	Least Squares	Constant Gain $\varepsilon = 0.1$	$\varepsilon = 0.05$	$\varepsilon = 0.03$
$std(y_t)$	6.1007	6.1188	6.1138	6.0727	6.0655
$std(\pi_t)$	0.2241	0.2257	0.2204	0.2137	0.2071
$std(r_t)$	3.3864	3.3966	3.3868	3.3560	3.3424
$corr(\pi_t, \pi_{t-1})$	0.3381	0.3399	0.3399	0.3402	0.3402
$corr(r_t, r_{t-1})$	0.3381	0.3389	0.3378	0.3381	0.3376

As in the RBC model, we find that the effects of learning in this model are very small. Here the largest effects are in the LS specification, which produces slightly higher mean volatility and slightly more persistence. The CG specifications actually lead to slightly less volatility than under rational expectations, with the lower gain

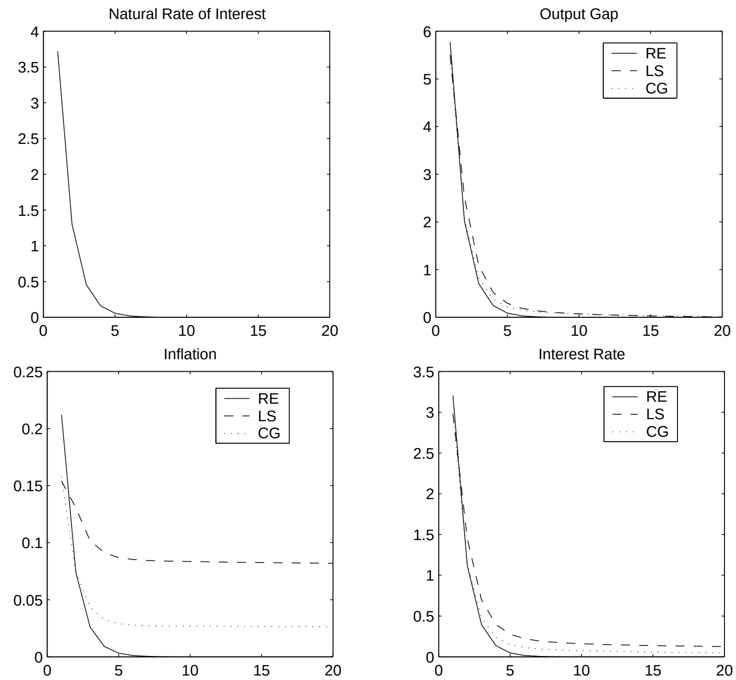

FIGURE 2. Mean impulse and pseudo-impulse responses of different variables in the New Keynesian model to a shock to the natural interest rate. RE= rational expectations, LS=least squares, CG= constant gain with $\varepsilon = 0.1$.

settings leading to lower volatility. However none of these effects are at all substantial. Thus, as in the RBC case, we turn to the analysis of the pseudo-impulse responses of the model under learning, which are shown in Figure 2. As above, this figure summarizes the average responses of the variables to a shock (now to the natural interest rate) at the beginning of a sample. Here we see that learning has more noticeable effects, particularly for inflation. In this case, under learning the shock has a nearly permanent mean effect on the inflation rate, and a slightly more persistent output gap response. These combine to make noticeable, if mild, increase in the interest rate response. As in the table, all of these effects are largest in the LS case.

Taken together, these results suggest that over short horizons with relatively little prior knowledge, learning may lead to additional volatility and persistence. But over longer horizons, its effects are very mild and may lead to slightly less volatility. These results are further illustrated in Figure 3, which shows some time series from a representative sample of the model. The top two panels plot the time series of the output gap and inflation under learning, here subtracting off the rational expectations levels. We clearly see that the largest effects of learning occur within the first twenty periods, after which only the slow-learning specification (the CG setting with $\varepsilon = 0.03$) displays any noticeable fluctuations. However while the slow-learning case displays larger fluctuations *around rational expectations*, it leads to smaller overall volatility. The beliefs underlying these simulations are shown in the bottom two panels, which plot the time series of the estimated slope coefficients from the IS and Phillips curve

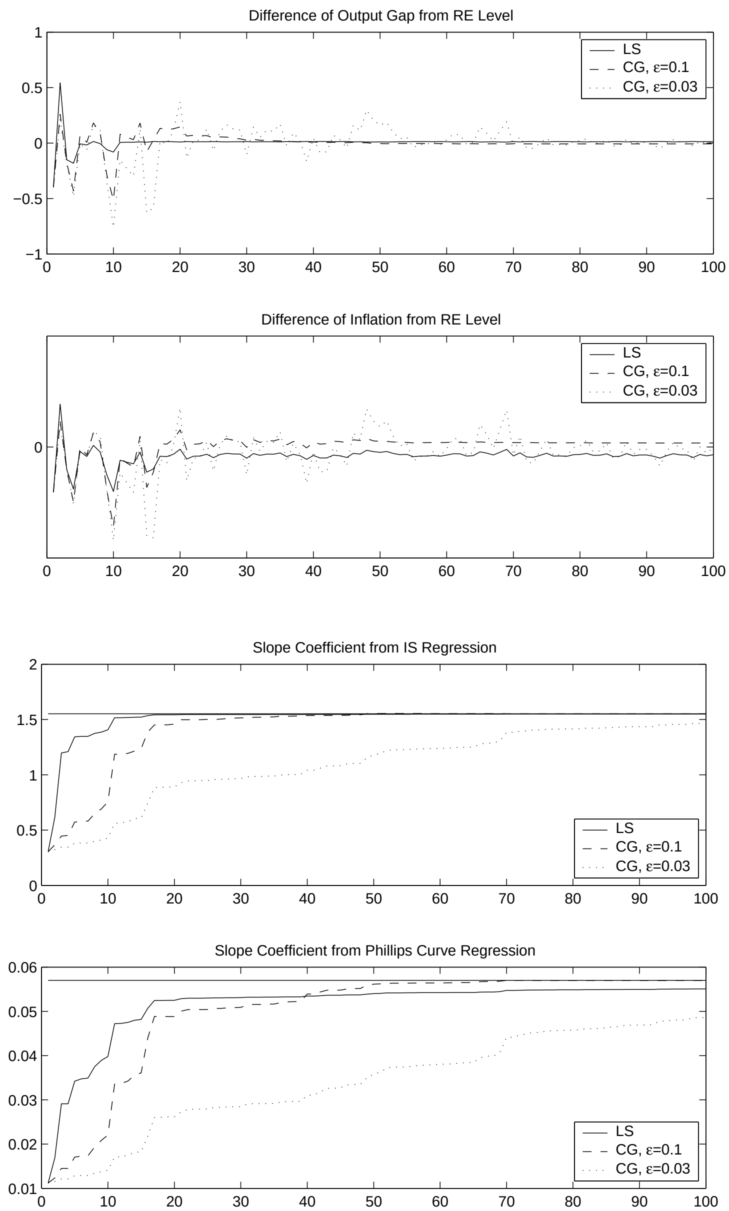

FIGURE 3. Simulated outcomes and beliefs in the New Keynesian model under different learning rules. LS=least squares, CG= constant gain.

regressions, $a_1 1$ and $b_1 1$ from (11). The straight lines give the rational expectations values. We see that for the LS and fast-learning CG case, convergence happens quite quickly, while the slow-learning CG case takes the whole sample to converge. The fact that both slopes remain below the rational expectations levels for an extended period of time, means that agents' perceived law of motion displays significantly less responsiveness to fluctuations in the natural rate of interest. This in turn explains why there is less overall volatility when learning is slow.

3. THE EFFECTS OF LEARNING: STRUCTURAL UNCERTAINTY

3.1. Overview

It may be tempting to conclude from our results above that learning has little quantitative impact in the models that we studied. Instead, the results should be read as suggesting that learning of the form specified above, in which agents learn reduced form equilibrium laws of motion, has little impact. While this is bad news from the vantage of trying to increase the amplification and propagation mechanisms in the models, it is good news for the plausibility of rational expectations—nearly the same results obtain when agents use boundedly rational learning rules. But there are clearly many other ways in which agents may learn about the economy, some of which may be arguably more plausible specifications. Here we consider a natural alternative, in which agents are uncertain about the structure of the economy and learn about this over time. We argue in the next section that this may be a more reasonable and theoretically coherent formulation. Moreover, we show that it makes the learning process have much more substantial effects on outcomes.

While it would ultimately be of interest to introduce structural learning directly in the models above, this proved more difficult than it might seem. Thus we focus for now on simpler models in the same vein. While this makes the analysis more transparent, and establishes some points which are relatively general, we are not able to make as strong a quantitative case. We focus on a simple special case of the RBC model which has an analytical solution (full depreciation, i.i.d. technology, and inelastic labor) and for which there is a natural structural learning specification. We then extend this model to consider the effects of aggregate externalities, as in Romer (1986). This example was introduced in Williams (2002a), and it provides a simple framework in which to study misspecified learning rules. We show that misspecified learning gives rise to additional low frequency fluctuations of a qualitatively different nature. With correctly specified beliefs, agents converge to an equilibrium and learning (especially constant gain versions) leads to fluctuations around the equilibrium. But with misspecified beliefs, learning leads to occasional, but recurrent, large movements away from an equilibrium. These are driven by the *escape dynamics* as studied by Sargent (1999), Williams (2002a), and Cho, Williams, and Sargent (2002). We include the results of Williams (2002a) which characterize the escape dynamics, and show that they are driven by agents recurrently (but unknowingly) internalizing the externality.

3.2. Structural Learning in a Simplified RBC model

The learning rules formulated above inherit a typical problem of reduced form estimates: they are difficult to interpret. Agents' decisions have in effect been substituted out, so that the only role left for agents is forecasting. The specification does not explicitly state any separation between agents' beliefs and their decisions, as agents forecast variables which are under their control. By contrast, if agents learn about the structural parameters of the economy, there is a clear separation between beliefs and decisions. Focusing on the structural features also shifts the focus from forecasting, to analyzing how agents' beliefs affect their decisions, and how this may matter for outcomes. By more explicitly considering agents' decision problems and beliefs, it is also more straightforward to consider what happens when agents' beliefs are incorrect. In the next section we consider such an example of misspecified learning.

But being more explicit about agents' uncertainty also leads to difficulties. Simply put, there is not a natural way to introduce structural learning in a well-posed manner in the full RBC model. The leading sources of uncertainty may be the features of the technology accumulation process and the production process. But the level of technology and the technology shock are observed, and coupled with the observations of output and inputs, the technology parameters could be inferred exactly. A natural extension would be to re-formulate the model with partial observations, so that agents need to infer the state of technology from data they observe. However the problem of coupling signal extraction with learning about the underlying parameters was beyond the scope of this paper. Instead, as mentioned above, we focus on a well-known simple special case of the RBC model. In this case, since the technology process is i.i.d. no signal extraction is necessary and there is a direct way to formulate structural learning.

Formally, we consider a special case of the RBC model (6), in which technology is i.i.d. ($\rho = 0$), there is no growth ($\gamma = \eta = 0$), labor is inelastically supplied ($\theta = 0$), and there is full depreciation of capital ($\delta = 1$). We keep the remaining parameters constant, except we increase the innovation standard deviation σ to maintain the same unconditional technology shock variance. As is well-known, under rational expectations the Euler equation is the following:

$$\frac{1}{C_t} = E_t \left[\beta \frac{\alpha e^{z_{t+1}} K_{t+1}}{C_{t+1}} \right], \tag{12}$$

where $K_{t+1} = e^{z_t} K_t^{\alpha} - C_t$. As is also well-known, this can be solved for the optimal consumption policy, which is to consume a constant fraction of output:

$$C_t = (1 - \beta\alpha) Y_t = (1 - \beta\alpha) e^{z_t} K_t^{\alpha}. \tag{13}$$

As before, we now relax the assumption of rational expectations and consider adaptive learning. It is straightforward to formulate the perceived law of motion (9) in this model, in which agents estimate a log-linear consumption rule within a period. Taking logarithms in (13), we see that the rational expectations values of the parameters are:

$$b_{10} = \log(1 - \beta\alpha), \; b_{11} = \alpha, \; b_{12} = 1.$$

However the interpretation of (9) is unclear. Consumption is a choice variable, so presumably agents should know the consumption rule that they are using and do not need to learn its coefficients. Further (9) does not directly determine consumption, but is only used to link forecasts of future consumption to forecasts of the other variables in determining the actual level of consumption. This seems to introduce an inconsistency between how agents perceive that they will behave and how they actually behave, one that is resolved only once they converge to rational expectations. The specification (8) encounters similar conceptual difficulties.[7]

It is possible that a more fully specified model, one which does not take the representative agent/social planner specification as literally, could justify the specifications (8)-(9). Such a "little k—big K" model would treat the perceived laws of motion as forecasts of aggregate behavior which are used in an individual optimization problem. But this does not immediately alleviate the problem: agents care about the marginal utility of their individual consumption, which is a choice variable for them. The more natural specification in a decentralized model would be to have agents forecast their future wages and rates of return.

The difficulties with the reduced form specification relate to the failure to separate beliefs and decisions. As discussed above, a natural alternative is to suppose that agents are uncertain about structural features of the economy, but optimize given their beliefs. In particular, we suppose here that agents to not know the parameters of the production function, and do not separately observe the technology shock. They estimate the return on their capital inputs via the regression equation:

$$y_t = a_0 + a_1 k_t + \xi_t. \tag{14}$$

Here ξ_t is a regression error which is due to the technology shock and the deviations of the parameters from their true values, which clearly are:

$$a_0 = 0, \ a_1 = \alpha.$$

Notice that the i.i.d. technology assumption is crucial in making (14) well-posed, as it insures that k_t and ξ_t are independent.

Following the literature on adaptive learning, we suppose that at each date agents treat their estimates as if they were true and would be constant forever. We thus abstract from parameter uncertainty and any potential experimentation motives. We assume that agents know the law of motion for capital, which seems reasonable given their direct observations of capital, output, and consumption. Agents then optimize given their beliefs, leading to the optimal decision rule as in (13) which simply replaces α by its estimate:

$$C_t = (1 - \beta a_1) Y_t. \tag{15}$$

[7]These points are related to but distinct from the issues raised by Preston (2002). He likewise considers the consistency of learning models with agent optimization, but his concerns dealt with the derivation of log-linearized models under potentially non-rational expectations. However he still considered reduced form learning specifications.

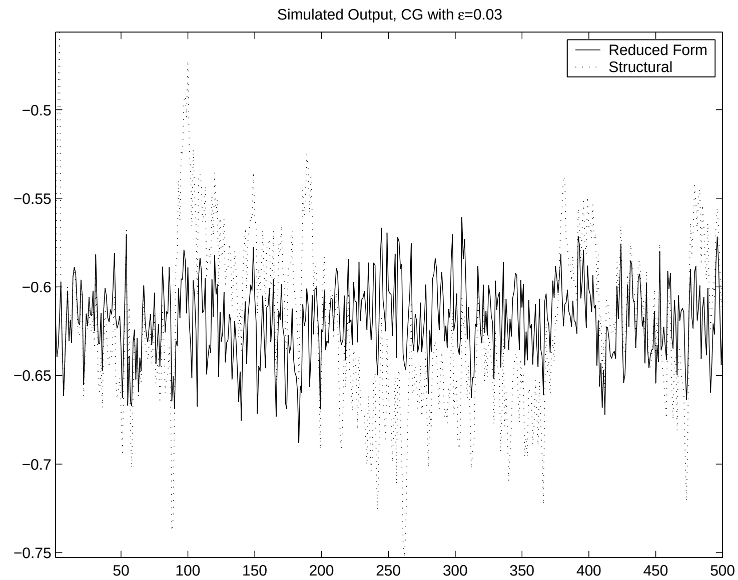

FIGURE 4. Simulated output time series in the simplified RBC model, for the reduced form and structural learning specifications. Both use CG with $\epsilon = 0.03$.

Based on their observations, agents then update their beliefs using a learning rule of the form (4).

TABLE 3.
Summary statistics from 1000 simulations of 150 observations of the simplified RBC model under rational expectations for different learning rules.

	Rational Expectations	Least Squares	Constant Gain $\varepsilon = 0.1$	$\varepsilon = 0.05$	$\varepsilon = 0.03$
		Reduced-Form Learning (9)			
$std(y_t)$	2.4229	2.4268	2.4310	2.4235	2.4302
$std(c_t)$	2.4244	2.4372	2.4421	2.4346	2.4412
$std(k_t)$	2.4186	2.4427	2.4455	2.4379	2.4453
$corr(y_t, y_{t-1})$	0.3873	0.3874	0.3819	0.3853	0.3829
$corr(c_t, c_{t-1})$	0.3878	0.3790	0.3732	0.3772	0.3747
		Structural Learning			
$std(y_t)$	2.4229	5.0775	8.8232	6.8955	5.7279
$std(c_t)$	2.4244	6.6774	11.7096	8.8261	7.4200
$std(k_t)$	2.4186	11.0567	21.2654	16.1966	13.0332
$corr(y_t, y_{t-1})$	0.3873	0.5603	0.6891	0.6816	0.6374
$corr(c_t, c_{t-1})$	0.3878	0.2571	0.1503	0.2306	0.2190

While we have expressed some theoretical reservations about the reduced form specification (9), we now compare its quantitative impact to the structural specification (14). As in the previous examples, we run 1000 simulations of 150 periods each, initializing beliefs in each case via a training sample (which we now increase to fifty periods). The results are summarized in Table 3. Once again, we find that the reduced form learning specifications have very little effect, leading to very slight increases in volatility and decreases in persistence relative to rational expectations. However structural learning has an enormous impact. The volatility of each variable is increased dramatically. The largest effect is on capital (which is the same as investment in this model), whose standard deviation increases by a factor of 5 to 7 across the different learning specifications. Further, the persistence of output is greatly increased, while that of consumption is decreased.

The increased volatility that we find is due to substantial fluctuations around the rational expectations equilibrium, as Figure 4 illustrates. The figure plots a time series of output from a single simulation run under reduced form and structural learning. The figure shows that the mean level of output is roughly equal in each case (as it was in the full suite of simulations), but that the structural rule displays persistent fluctuations around this mean. These fluctuations are driven by the revisions of beliefs and, as the table shows, their magnitude decreases with the weight that agents place on new information. When agents heavily discount past data (having larger gain settings), they are more willing to believe that the underlying parameters have changed. Since the parameter estimates affect decisions in a nonlinear way, this revision process can lead to substantial changes in outcomes.

3.3. An Extended RBC Model with External Effects

Thus far we have considered subjective models which, although not rational, were correctly specified. However once we retreat from rational expectations, there is no clear reason to suppose that agents always rely on models of the correct form. In this section we analyze the effects of misspecification in a simple extension of the previous RBC model which was analyzed in Williams (2002a). We extend the model by incorporating an aggregate production externality, as in Romer (1986), which leads to external increasing returns. Similar models have been widely used in the literature on indeterminacy and sunspots, as discussed in Benhabib and Farmer (1999). However under our assumptions, there is a unique, determinate, equilibrium and therefore no room for exogenous sunspots to play a role. We show that, as in the previous section, learning produces substantial fluctuations. However when agents' beliefs are misspecified, there are occasional fluctuations of a qualitatively different nature. Rather than fluctuating around an equilibrium, agents "escape" from the equilibrium for an extended period. These escape dynamics are further considered in the next section, which reports some results of Williams (2002a).

We extend the model by supposing that there are two different types of firms, which operate in segmented markets and are owned by separate representative agents who supply their capital. We consider different firm types because with a single representative firm there would be no distinction between individual and aggregate behavior. Further, the segmentation of markets allows us to dispense with prices, and focus on the determination of quantities. Neither of these assumptions are compelling, but they serve to keep the model simple and close to the preceding RBC model. We can interpret the different firm types as reflecting different sectors or closed economies that do not trade, but are only linked by productive externalities. Each firm uses capital and inelastically-supplied labor inputs in production, but because of external effects, each firm's output also depends on the capital input of the other firm. Thus we specify that firms produce according to:

$$\begin{aligned} Y_{1t} &= e^{z_{1t}} K_{1t}^{\alpha} K_{2t}^{\nu} \\ Y_{2t} &= e^{z_{2t}} K_{2t}^{\alpha} K_{1t}^{\nu}, \end{aligned} \tag{16}$$

for $\alpha, \nu > 0$. Firms have constant internal returns to scale (with the inelastic labor input which is suppressed), but there are social increasing returns of $1+\nu$. We continue to assume that the technology shock process is i.i.d. over time, but we now suppose that the two types have independent shocks z_{it} with common variance. We retain the parameter values of the previous section, and set the external effect at the relatively small value of $\nu = 0.08$.

We first consider the extension of the structural specification above, in which agents learn their internal returns via the regression:

$$y_{it} = a_{0i} + a_{1i} k_{it} + a_{2i} k_{jt} + \xi_{it}, \tag{17}$$

for $i, j = 1, 2$ with $i \neq j$. However we also study the case when agents' beliefs are misspecified. Here we assume that agents do not consider the possibility of productive spillovers. This could be justified either by assuming that agents do not observe the output of the other type of firm (and do not try to account for this fact), or simply that they focus on estimating their own internal returns. We assume that agents learn adaptively about their firm's production function according to the regression:

$$y_{it} = \gamma_{0i} + \gamma_{1i} k_{it} + \xi_{it}^{*}. \tag{18}$$

Now ξ_{it}^{*} also absorbs the effects other firm's capital input which has been omitted from (18). Agents optimize under either belief specification, and the optimal decision rule is again of the form (15), where the fraction of output saved depends on their estimate of the return on their own capital (a_{i1} or γ_{1i}).

Standard results as in Evans and Honkapohja (2001) show that with the correct belief specification, (under some technical conditions) agents will converge to the rational expectations (competitive) equilibrium. However with the misspecified learning rule their beliefs continually differ from rational expectations. Williams (2002a) shows that agents converge to a *self-confirming equilibrium*, with outcomes close to, but distinct from, rational expectations. In a self-confirming equilibrium agents optimize, and their beliefs are not contradicted by their observations (see Fudenberg and Levine (1998) and Sargent (1999)). Under (18), agents' beliefs have an omitted variable which is correlated with the regressor k_{it}. This induces some bias in the limiting beliefs, and therefore in agents' decisions, which causes a wedge between the self-confirming and rational expectations equilibrium outcomes. But we focus on parameterizations with small increasing returns, so this bias is small and has relatively little effect.

We once again analyze the effects of learning through some simulations, now focusing on structural learning and comparing the correctly specified beliefs (17) with the misspecified beliefs (18). For reasons that will be clear, we now focus on longer

TABLE 4.
Summary statistics from 400 simulations of 2000 observations of the externality model under rational expectations for different learning rules.

	Rational Expectations	Least Squares	Constant Gain $\varepsilon = 0.05$	$\varepsilon = 0.03$	$\varepsilon = 0.01$
		Correctly Specified			
$std(y_t)$	2.4447	5.2939	10.5821	9.2550	8.1892
$std(c_t)$	2.4464	6.3331	10.7224	9.1124	7.9879
$std(k_t)$	2.4464	10.1516	23.7765	20.4553	17.7038
$corr(y_t, y_{t-1})$	0.3878	0.6941	0.8030	0.7581	0.6696
$corr(c_t, c_{t-1})$	0.3884	0.3738	0.4644	0.4405	0.3442
		Misspecified			
$std(y_t)$	2.4447	4.2588	10.5221	7.4573	9.1732
$std(c_t)$	2.4464	3.7078	8.0159	4.8799	6.1286
$std(k_t)$	2.4464	7.6811	23.0556	15.7852	19.6129
$corr(y_t, y_{t-1})$	0.3878	0.6395	0.7777	0.7271	0.6974
$corr(c_t, c_{t-1})$	0.3884	0.2703	0.3122	0.2789	0.3269

simulated time series of 2000 observations, and we doing 400 runs of each specification. Table 4 summarizes the results. As we see, the results are fairly similar, as both correct and misspecified structural beliefs lead to an increase in the volatility of all variables and an increase in the persistence of output. For all but the smallest constant gain setting, the misspecified case leads to lower volatility than the correct specification, and the misspecification results in consumption being less persistent than under rational expectations. In all but the smallest gain setting, the volatility decreases with the amount of weight put on new information, which again points to the fact that the volatility is driven by fluctuations around the (rational expectations or self-confirming) equilibrium. However it is interesting to note that under misspecified beliefs, the volatility increases when the gain is cut from 0.03 to 0.01. This represents the influence of the *escape dynamics* (see Sargent (1999), Williams (2002a), and Cho, Williams, and Sargent (2002)), which lead to occasional episodes in which agents' beliefs move away from the equilibrium.

The emergence of the escape dynamics is shown in Figure 5, which plots simulated output series for one firm type under the correctly specified and misspecified beliefs. The two series track each other rather closely at the beginning of the sample, as they bounce around the equilibrium levels. But then after a period of time, we see that output increases significantly in the misspecified case, after which it gradually drifts back down to the equilibrium level. These "escapes" always occur in a similar manner, leading to a rapid expansion of output. The figure also suggests that, at least in this version and parameterization of the model, the escape dynamics happen at very low frequencies. For larger gain settings, escapes may occur, but their effect is swamped by the overall volatility of the model. There may be ways of separating these effects and speeding up the escape dynamics, perhaps by considering a more general learning rule which allows more flexibility in the specification of beliefs, as Sargent and Williams (2002) study. However at present our results suggest that escape dynamics are not

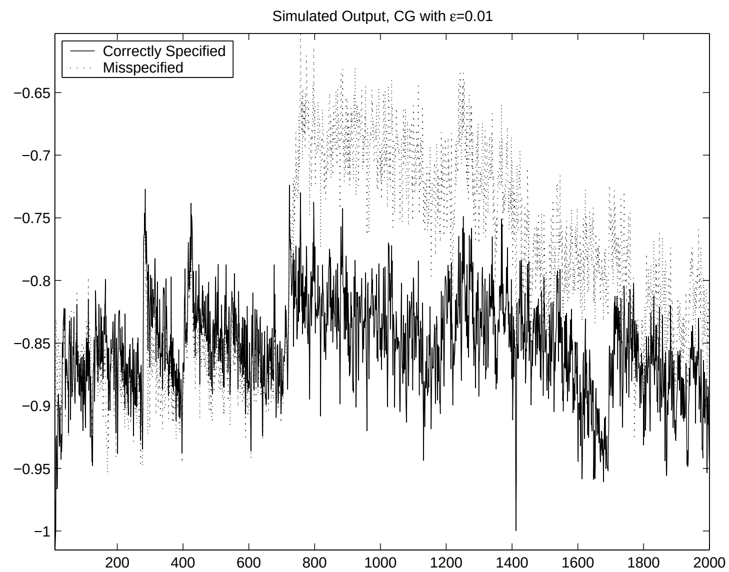

FIGURE 5. Simulated output time series in the externality model, with correctly specified and misspecified learning specifications. Both use CG with $\epsilon = 0.01$.

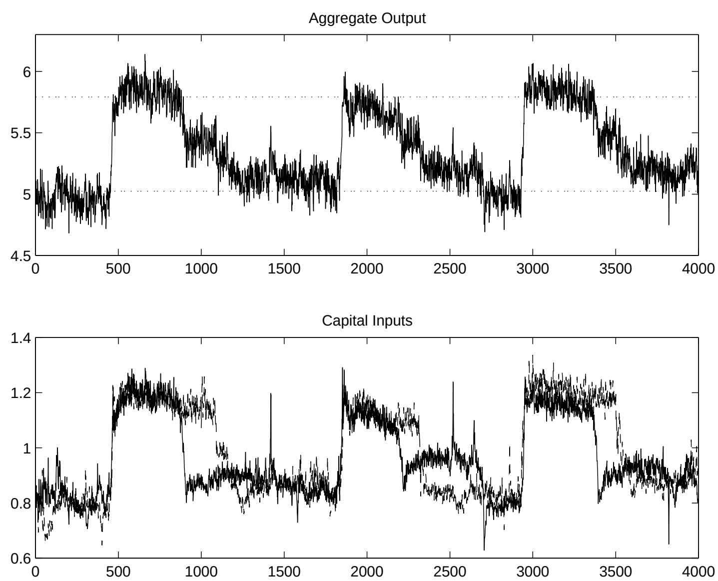

FIGURE 6. Simulated time paths of aggregate output and individual capital inputs from the externality model. The self-confirming equilibrium output is 5.07, and the symmetric social optimum output is 5.84.

likely to be a source of business cycles, but rather may lead to longer term changes. In the next section we explore these escape dynamics further.

3.4. Escape Dynamics

In order to focus more fully on the escape dynamics, we now work with a parameterization which makes them more apparent. We add a mean level to the productivity factors, which are now of the form $\exp(1 + z_{it})$, and we set $\alpha = 0.33$, $\nu = 0.08$,

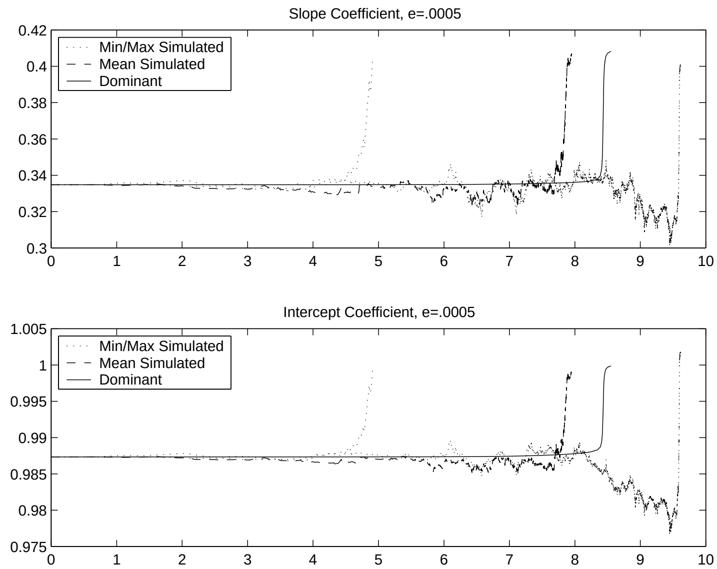

FIGURE 7. Dominant and simulated time paths of a firm's beliefs from the externality model, logarithmic time scale.

$\sigma = 0.02$, and $\beta = 0.995$. The production parameters reflect small increasing returns and the discount factor is set relatively high to reflect a short interval between periods. Under this parameterization, the self-confirming equilibrium beliefs are $(\gamma_{0i}, \gamma_{1i}) = (0.9873, 0.3348)$. This differs only slightly from the beliefs $(\gamma_{1i} = \alpha)$ that support the competitive equilibrium, so that the corresponding SCE outcomes are close to the competitive equilibrium levels.

In Figure 6 we plot the time paths of aggregate output (adding up the two firms) and the two firms' capital inputs from a simulation. The figure plots the variables in levels, not logarithms. We initialize agents' beliefs at the SCE, and use a gain setting of $\varepsilon = 0.01$. The figure is essentially a less noisy version of Figure 5, and shows that the increases in output recur repeatedly in the model. The time paths are characterized by rapid expansions, in which firms act nearly in unison to increase output, and slower reductions in output as the firms stagger their return to the self-confirming equilibrium level. Moreover, as figure shows, output always increases to nearly the same level. This corresponds to the level of output in the social optimum in which agents internalize the external effect. Although the time paths look similar to shifts between distinct equilibria, it is important to recall that there is a single equilibrium in the model.

The escapes from the self-confirming equilibrium are driven by random occurrences with small probability, due to an unlikely sequence of shock realizations. But as Figure 6 shows, when an escape occurs with very high probability it happens in a particular, predictable way. Williams (2002a) provides a complete characterization of the escape dynamics, determining their frequency and identifying the most likely path that beliefs follow when they escape from a self-confirming equilibrium. There it is shown that the most likely or *dominant* path can be found by solving a simple dynamic control problem, whose solution also determines the frequency of escapes. We now apply these results in this model.

Figure 7 shows some results for the beliefs of one firm type (the other type is similar) from 1000 simulations from the model. The top panel plots the slope coefficient from the regression (18) and the bottom panel plots the constant coefficient, where we show the dominant escape paths and the minimum, maximum, and mean simulated escape paths. We see that the escapes are very regular, and that our predicted dominant path provides a reasonable prediction of the mean path of beliefs. The time series are characterized by a variable period near the self-confirming equilibrium level, followed by a rapid increase in both the estimated slope and intercept coefficients of both firm types. These changes in beliefs lead agents to increase savings, leading to increases in output.

These plots suggest that the escapes are caused by agents internalizing the externality. Along an escape the slope coefficient increases from its SCE level of $\overline{\gamma}_1 = 0.3348$ to $\alpha + \nu = 0.413$, which corresponds to the social returns to capital. The driving shocks in the model are the independent idiosyncratic technology shocks. This independence implies that on average agents' capital stocks are nearly independent, and thus agents converge to the SCE. However, occasional correlated shock realizations cause agents to synchronize their actions. The correlation in the capital inputs of the different firm types in these episodes leads agents to discover the external return. Agents effectively internalize the externality, which they interpret as an increase in the productivity of their capital inputs. During these episodes, the complementarity of firms' capital inputs overrides the effect of the idiosyncratic shocks, and agents learn that it is beneficial to jointly increase savings and expand production. However once they reach the social optimum, they do not increase savings further and the idiosyncratic shocks break the correlation between agents' actions. Thus agents again only perceive the effect of their internal returns on their outcomes, and this leads them back to the SCE. Thus the escape dynamics and mean dynamics drive these endogenous long-run cycles of the expansion and contraction of output.

4. CONCLUSION

The goal of the paper was to determine the importance of adaptive learning as a source of amplification and propagation in standard economic models. While the results are far from definitive, we have shown that under the most common learning specification, in which agents learn about reduced-form equilibrium laws of motion, learning has very little effect in either a calibrated RBC model or a New Keynesian monetary model. However we discussed some difficulties with the interpretation of this specification, and formulated an alternative in which agents learn about structural features of the economy. In a simplified setting, we showed that structural learning has a much more sizeable impact, and that it may lead to increased volatility and some increases in persistence. Further, we have illustrated that if agents have misspecified beliefs, then learning may add a further source of fluctuations via escape dynamics. The quantitative importance of these escape dynamics and the increased volatility due to structural learning still remain to be established, but we have illustrated in a simple setting that they may matter greatly. A more complete quantitative assessment requires a more complete model, which is a task of future research.

More generally, our results point toward the importance of explicitly modeling agents' beliefs, including the possibility that their subjective models may be misspecified. If learning is an important issue, which many would agree, then we need to take a closer look at individuals' beliefs, decisions, and learning procedures. In so doing, there is the potential to bridge the extensive literature on learning in macroeconomics with extensive literature on learning in games, which tackles precisely these issues.

References

Barsky, R. B. and J. B. DeLong (1993). Why Does the Stock Market Fluctuate? *Quarterly Journal of Economics 108*, 291–311.

Benhabib, J. and R. E. A. Farmer (1999). Indeterminacy and Sunspots in Macroeconomics. In J. Taylor and M. Woodford (Eds.), *Handbook of Macroeconomics*, pp. 387–448. Amsterdam: North-Holland.

Bullard, J. and J. Duffy (2002). Learning to Detrend Macroeconomic Data. Working paper, Federal Reserve Bank of St. Louis.

Bullard, J. and K. Mitra (2002). Learning about Monetary Policy Rules. *Journal of Monetary Economics 49*, 1105–1129.

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

Christiano, L. J., M. Eichenbaum, and C. Evans (2001). Nominal Rigidities and the Dynamic Effects of a Shock to Monetary Policy. Working paper, Northwestern University.

Clarida, R., J. Gali, and M. Gertler (1999). The Science of Monetary Policy: A New Keynesian Perspective. *Journal of Economic Literature 37*, 1661–1701.

Cooley, T. and E. C. Prescott (1995). Economic Growth and Business Cycles. See Cooley (1995), pp. 1–38.

Cooley, T. F. (Ed.) (1995). *Frontiers of Business Cycle Research*. Princeton University Press.

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

Evans, G. W. and S. Honkapohja (2002a). Adaptive Learning and Monetary Policy Design. Working paper, University of Oregon.

Evans, G. W. and S. Honkapohja (2002b). Expectations and the Stability Problem for Optimal Monetary Policies. Working paper, University of Oregon.

Francis, N. and V. Ramey (2001). Is the Technology-Driven Real Business Cycle Hypothesis Dead? Shocks and Aggregate Fluctuations Revisited. Working paper, UCSD.

Fudenberg, D. and D. Levine (1998). *The Theory of Learning in Games*. MIT Press.

Gali, J. (1999). Technology, Employment, and the Business Cycle: Do Technology Shocks Explain Aggregate Flucutations? *American Economic Review 89*, 249–271.

King, R. G. and S. T. Rebelo (1999). Resuscitating Real Business Cycles. In J. Taylor and M. Woodford (Eds.), *Handbook of Macroeconomics*, pp. 927–1007. Amsterdam: North-Holland.

Kushner, H. J. and G. G. Yin (1997). *Stochastic Approximation Algorithms and Applications*. Springer-Verlag.

Marcet, A. and T. J. Sargent (1989). Convergence of Least Squares Learning Mechanisms in Self Referential Linear Stochastic Models. *Journal of Economic Theory 48*, 337–368.

Packalen, M. (2000). On the Learnability of Rational Expectations Equilibria in Three Business Cycle Models. Research Report No. 87, University of Helsinki Department of Economics.

Preston, B. (2002). Learning About Monetary Policy Rules when Long-Horizon Expectations Matter. Working paper, Princeton University.

Romer, P. M. (1986). Increasing Returns and Long-run Growth. *Journal of Political Economy 94*, 1002–37.

Rotemberg, J. and M. Woodford (1997). Optimization-Based Econometric Framework for the Evaluation of Monetary Policy. In B. S. Bernanke and J. J. Rotemberg (Eds.), *NBER Macroeconomics Annual 1997*, pp. 297–346. Cambridge: MIT Press.

Sargent, T. J. (1999). *The Conquest of American Inflation*. Princeton University Press.

Sargent, T. J. and N. Williams (2002). Impacts of Priors on Convergence and Escape from Nash Inflation. Working paper, Princeton University.

Taylor, J. B. (1993). Discretion versus Policy Rules in Practice. *Carnegie-Rochester Conference Series on Public Policy 39*, 195–214.

Timmermann, A. G. (1993). How Learning in Financial Markets Generates Excess Volatility and Predictability in Stock Prices. *Quarterly Journal of Economics 108*, 1135–4563.

Timmermann, A. G. (1996). Excess Volatility and Predictability of Stock Prices in Autoregressive Dividend Models with Learning. *Review of Economic Studies 63*, 523–557.

Williams, N. (2002a). Escape Dynamics in Learning Models. Working paper, Princeton University.

Williams, N. (2002b). Small Noise Asymptotics for a Stochastic Growth Model. Working paper, Princeton University.

Woodford, M. (1999). Optimal Monetary Policy Inertia. National Bureau of Economic Research Working Paper: 7261.

Woodford, M. D. (2002). *Interest and Prices: Foundations of a Theory of Monetary Policy*. Princeton: Princeton University Press. Forthcoming.