Bayesian Model Averaging, Learning and Model Selection*

George W. Evans, University of Oregon and
University of St. Andrews
Seppo Honkapohja, Bank of Finland
Thomas Sargent, New York University
Noah Williams, University of Wisconsin, Madison

January 26, 2012

Abstract

Agents have two forecasting models, one consistent with the unique rational expectations equilibrium, another that assumes a time-varying parameter structure. When agents use Bayesian updating to choose between models in a self-referential system, we find that learning dynamics lead to selection of one of the two models. However, there are parameter regions for which the non-rational forecasting model is selected in the long-run. A key structural parameter governing outcomes measures the degree of expectations feedback in Muth's model of price determination.

1 Introduction

In the past two decades there has been a significant amount of macroeconomic research studying the implications of adaptive learning in the formation of expectations. This approach replaces rational expectations with the

*Any views expressed are those of the authors and do not necessarily reflect the views of the Bank of Finland. Corresponding author: Seppo Honkapohja, email: seppo.honkapohja@bof.fi.

assumption that economic agents employ a statistical forecasting model to form expectations and update the parameters of their forecasting model as new information becomes available over time. One goal of this literature is to find the conditions under which the economy with this kind of learning converges to a rational expectations equilibrium (REE).

The basic learning setting presumes that the agents' perceptions take the form of a forecasting model with fixed unknown parameters, estimates of which they update over time.[1] Such a setting does not explicitly allow for parameter uncertainty or the use of averaging across multiple forecasting models.[2] In this paper, we postulate that economic agents use Bayesian estimation and Bayesian model averaging to form their expectations about relevant variables.

We investigate this issue both to study the robustness of some existing convergence results in the learning literature, and also to provide some further justification for recursive updating scheme models. It is widely understood that if agents learn by updating what they believe to be a fixed parameter, in environments with feedback their beliefs are only correct asymptotically. That is, as agents change the coefficients in their perceived laws of motion, their actions influence economic outcomes in ways that make the law of motion actually generating the data change over time. After beliefs have ultimately converged to a REE, agents' beliefs are correctly specified, but along the transition path to the REE the data generating process has drifting coefficients. It seems natural to suppose that agents would allow the coefficients of their forecasting models to drift over time. Heuristically, that reasoning motivated the body of papers on "constant gain" learning in which agents attempt to track drifting coefficients and learning is perpetual.[3] In this paper, we suppose that agents subjectively entertain *two* possibilities – one that says that the data generating process is constant, and another that says that it drifts over time. Our agents update the probabilities that they place on these two possibilities. We study the long-run behavior of this process of jointly updating models and probabilities over models.

More precisely, we study a setting in which the pair of models used by

[1] See Evans and Honkapohja (2001) for the earlier literature; for recent critical overviews see Sargent (2008) and Evans and Honkapohja (2009).

[2] A few papers have incorporated model averaging in a macroeconomic learning setting. For examples, see Cogley and Sargent (2005) and Slobodyan and Wouters (2008).

[3] See Cho, Williams, and Sargent (2002), Sargent and Williams (2005), and Evans, Honkapohja, and Williams (2010), for examples.

agents includes a “grain of truth”in the sense that the functional form of one of the models is consistent with the REE of the economy while the other model is misspecified relative to the REE.[4] In particular, as above we assume that agents also employ a time-varying parameter (TVP) model as a second available forecasting model. The analysis is carried out using a standard general set-up, discussed, e.g., in chapter 2 of Evans and Honkapohja (2001). It is known that for this model there is convergence of usual least-squares (LS) learning to a unique REE unless the expectations feedback is positive and more than one-to-one.

We thus consider a set-up with Bayesian model averaging over a constant parameter model that nests the REE and a TVP model. The parameters of the models and the probability weight over models are updated in Bayesian fashion as new data becomes available. Does learning converge to REE?

- Convergence occurs for a range of structural parameters in which the influence of expectations on the current outcome is not too strong and positive.

- The set of structural parameters for which convergence occurs is smaller than the less than the one-to-one feedback that is crucial for LS learning *without* model averaging.

- A striking result is that there can be convergence to the (non-REE) TVP forecasting model even though the prior puts an atom on the RE forecasting model. The happens when the expectations feedback parameter is positive and sufficiently strong but less than one-to-one.

- Learning via Bayesian model averaging usually leads to model selection. The proportion of cases of no selection in the long run is small.

One version of our general setup applies to the Muth market (or “cobweb”) model in which expectations feedback is negative. For the Muth model, learning by Bayesian model averaging converges to the REE. Our setup also covers a version of the Lucas “island” model in which the feedback of expectations on current outcomes is positive. For that setting, the strength of the response of output to expected inflation determines the convergence

[4] We note that there is also a game-theory literature on convergence of Bayesian learning and the issue of the “grain of truth”, see Young (2004) for an overview. Here we have a set-up in which the prior of agents includes a grain of truth on the REE.

outcome. If the feedback is sufficiently strong, learning by Bayesian model averaging may converge to a situation in which agents perpetually use the TVP forecasting model.

2 Muth Model with Bayesian Learning

We consider the Muth market model

$$p_t = \mu + \alpha E^*_{t-1} p_t + \delta z_{t-1} + \eta_t, \quad (1)$$

where p_t is the market price, $E^*_{t-1} p_t$ denotes expectations of p_t conditioned on information at date $t-1$, z_{t-1} is an exogenous observable variable following a stationary $AR(1)$ process $z_t = \rho z_{t-1} + w_t$ with $w_t \sim iid(0, \sigma^2_w)$ and η_t is an unobserved white noise shock with $E\eta^2_t = \bar{\sigma}^2_\eta$. We normalize $\mu = 0$. We denote the subjective expectations E^*_t to highlight that they are not necessarily the rational (mathematical) expectation.

We remark that the Muth model (1) can be obtained from aggregating firm supply curves that depend on $E^*_{t-1} p_t$ and a market demand curve depending on p_t, with each also depending on aggregate shocks. The firm supply curves in turn are derived from maximization of expected profits and quadratic costs. If the demand and supply curves are, respectively, downward- and upward-sloping, the Muth market model has the parameter restriction $\alpha < 0$, so that there is a negative feedback from expectations to outcomes. More generally, as noted by Bray and Savin (1986), $\alpha < 1$ provided the demand curve crosses the supply curve from above.

The setting (1) also arises for a version of the Lucas aggregate supply model, in which supply depends on price surprises and aggregate demand is given by a quantity-theory type of equation. For the Lucas-type macro model the parameter restriction is $0 < \alpha < 1$, so that there is positive feedback from expectations to outcomes. See Evans and Honkapohja (2001), Chapter 2, Sections 2.2 and 2.3 for more details on the Muth and Lucas models.

The REE for model (1) is

$$p_t = \bar{\beta} z_{t-1} + \eta_t, \text{ where } \bar{\beta} = (1-\alpha)^{-1}\delta.$$

We begin with the case in which agents have a constant parameter forecasting model, which they estimate using Bayesian techniques. The beliefs of the agents are

$$p_t = \beta z_{t-1} + \eta_t,$$

where $\eta_t \perp z_{t-1}$ and $\eta_t \sim N(0, \sigma_\eta^2)$. The forecasting model of the agents at the end of period $t-1$, also called the perceived law of motion (PLM), is

$$p_t = b_{t-1} z_{t-1} + \eta_t,$$

where b_{t-1} is the time $t-1$ estimate of β. Note that in general we allow $\sigma_\eta^2 \neq \bar{\sigma}_\eta^2$. There is a prior distribution $\beta \sim N(b_0, V_0)$, which implies a posterior distribution of $f(\beta \mid y^{t-1})$, where $y^t = (y_t, y_{t-1}, y_{t-2}, ...)$ and $y_t' = (p_t, z_t)$, of the form $N(b_t, V_t)$. Here the updating of parameters b_t, V_t is given by

$$\begin{aligned} b_t &= b_{t-1} + \frac{V_{t-1} z_{t-1}}{\sigma_\eta^2 + V_{t-1} z_{t-1}^2} (p_t - b_{t-1} z_{t-1}) \\ V_t &= V_{t-1} - \frac{z_{t-1}^2 V_{t-1}^2}{\sigma_\eta^2 + V_{t-1} z_{t-1}^2}. \end{aligned}$$

using the Kalman filter.

The dynamics of the system can be formulated as a stochastic recursive algorithm (SRA) as indicated in the appendix, where it is shown that we have the result:

Proposition 1 *There is convergence to the REE with probability* 1 *if* $\alpha < 1$. *Moreover, we get*

$$V_t = \frac{\sigma_\eta^2}{(t+1) S_t - z_t^2} \rightarrow 0$$

with probability 1 *for all* σ_η^2, *irrespective of whether* σ_η^2 *is correct or not.*[5]

Bayesian learning was already considered in a somewhat different formulation by Bray and Savin (1986), who assumed that agents have heterogeneous expectations and there is a continuum of initial priors $b_0(i), i \in [0,1]$ with the same initial precision. Our setting could handle a finite number of classes of agents with different priors.

[5] In particular, we get convergence to the REE whether or not the actual η_t is normal.

3 Bayesian Learning with Subjective Model Averaging

3.1 Priors on parameter variation

In the preceding section it was assumed that agents' beliefs treat the parameter β as an unknown constant that does not vary over time. An alternative set-up would be to allow time variation in β. Papers by Bullard (1992), McGough (2003), Sargent and Williams (2005), and Evans, Honkapohja, and Williams (2010) look at this issue in models with learning. Cogley and Sargent (2005) look at empirical time-varying parameter models without learning. In our self-referential set-up with learning, we adopt a formulation where agents entertain multiple forecasting models and form the final forecast as a weighted average of the forecasts from the different models.

Although other extensions may be useful, we consider a simple example of multiple forecasting models below. We assume that agents have a prior that puts a weight $\pi_0 > 0$ on β constant over time and $1-\pi_0 > 0$ on the TVP model $\beta_t = \beta_{t-1} + v_t$, where $v_t \overset{iid}{\sim} N(0, \sigma_v^2)$. In general, σ_v^2 could be unknown, but we assume that it is known. The next steps are (i) the computation of the model-weighted forecast and (ii) the updating of the parameters in the forecasting models and of the models weights as new information becomes available.

We now develop these ideas in a simple setting using model (1) and the assumption that agents employ two different forecasting models.

3.2 Model averaging

Thus there are just two forecasting models in use: a constant coefficient model, and a TVP model with $\beta_t = \beta_{t-1} + v_t$, where $v_t \overset{iid}{\sim} N(0, \sigma_v^2)$ and $\sigma_v^2 > 0$ is known.[6] More specifically, and changing notation somewhat, the PLMs of the agents are

$$\begin{aligned} p_t &= \beta_t(i) z_{t-1} + \sigma_p(i)\varepsilon_{pt}, \text{ for } i = 0, 1 \\ \beta_t(i) &= \beta_{t-1}(i) + \sigma_\beta(i)\varepsilon_{\beta t}, \text{ for } i = 0, 1, \end{aligned}$$

[6] We adopt this formulation for simplicity. Clearly a useful extension will be to have finitely many subjective models.

where z_t is an exogenous observable. Here the first equation is the PLM for p_t of subjective model i. Various assumptions about $\sigma_p(i)$ are possible. They can be assumed known or unknown and equal or allowed to be different for $i = 0, 1$. The second equation specifies the perceived parameter drift in each subjective model. We will assume $0 \leq \sigma_{\beta(0)} < \sigma_{\beta(1)}$ are known. A third equation, specified below, gives the subjective probability weight for subjective model one, with a *prior* placing (say) an equal weight on the two models. The ε_{jt}, $j = p, \beta$ are i.i.d. standard normal and mutually independent. Agents have normal priors on the $\beta_0(i)$ and a given prior probability π_0 that model 0 is correct. We will usually take $\pi_0 = 0.5$.

Expectations of agents are given by subjective model averaging, i.e.

$$E^*_{t-1}p_t = \left(\pi_{t-1}(0)\hat{\beta}_{t|t-1}(0) + \pi_{t-1}(1)\hat{\beta}_{t|t-1}(1)\right) z_{t-1},$$

where $\hat{\beta}_{t|t-1}(0)$ and $\hat{\beta}_{t|t-1}(1)$ are the means of the posterior distribution for $\beta_t(0)$ and $\beta_t(1)$ and where $\pi_{t-1}(i)$ is the posterior probability that model i is correct, i.e.

$$\pi_{t-1}(i) = \Pr(i|p^{t-1}, z^{t-1}]),$$

for $i = 0, 1$. The actual evolution of price p_t is given by

$$p_t = \alpha E^*_{t-1}p_t + \delta z_{t-1} + \sigma_a \varepsilon_{at},$$

where the exogenous observables z_{t-1} are a stationary AR(1) process, i.e.,

$$z_t = \rho z_{t-1} + \sigma_z \varepsilon_{zt}.$$

For simplicity, we have set all the intercepts to be zero. Otherwise, subjective intercepts for each model would also need to be estimated and agents would plausibly also allow for parameter drift for intercepts. As noted above, in the Muth market model we usually have $\alpha < 0$, but in "Lucas-type" models $0 < \alpha < 1$.

We are especially interested in how the weight $\pi_t(1)$ of the TVP model evolves. Suppose that $\sigma_\beta(0) = 0$ and $\sigma_\beta(1) > 0$. Will we have $\pi_t(1) \to 0$ and $\hat{\beta}_t \to \bar{\beta} = \delta(1-\alpha)^{-1}$ as $t \to \infty$, with probability one, so that there is convergence to REE? We suspect that this will depend on the magnitudes of both α and $\sigma_\beta(1)$. We venture an initial guess that the most "stable" cases are possibly in the range $-1 < \alpha < 0.5$. The basis for this guess that in the standard LS learning setting parameter values in the range $0.5 < a < 1$ may yield slow convergence to REE and in the case $\alpha < -1$ a possible problem of overshooting can emerge when agents overparameterize the PLM under LS learning. For $\alpha = 0$ the p_t process is exogenous and here we certainly expect $\pi_t(1) \to 0$ and $\hat{\beta}_t \to \delta$ with probability one. We would therefore expect $\pi_t \to 0$ and convergence to REE also for α near 0. However, there does seem to be the possibility with $\alpha \neq 0$ that $\pi_t(1)$ remains near 1 for long periods or even that $\pi_t(1) \to 1$. In what follows we examine these issues by means of numerical simulations.

We now give the recursive updating equations for $i = 0, 1$. The Kalman filter, see e.g. Hamilton (1994) p. 399 and 380, gives the updating equations for the mean $\hat{\beta}_{t+1|t}(i)$ of the (Gaussian) posterior distribution of $\beta_{t+1}(i)$ as follows:

$$\begin{aligned}
\hat{\beta}_{t+1|t}(i) &= \hat{\beta}_{t|t-1}(i) + \frac{V_{t|t-1}(i)z_{t-1}}{\sigma_p^2(i) + V_{t|t-1}(i)z_{t-1}^2}(p_t - \hat{\beta}_{t|t-1}(i)z_{t-1}) \\
V_{t+1|t}(i) &= V_{t|t-1}(i) - \frac{z_{t-1}^2 V_{t|t-1}(i)^2}{\sigma_p^2(i) + V_{t|t-1}(i)z_{t-1}^2} + \sigma_\beta^2(i).
\end{aligned}$$

Here $\text{var}(\hat{\beta}_{t|t-1}(i) - \beta_t(i)) = V_{t|t-1}(i)$. We will also need the mean and variance of the conditional distribution for β_t conditional on information through t, which are given by $\hat{\beta}_{t|t}(i) = \hat{\beta}_{t+1|t}(i)$ and $V_{t|t}(i) = V_{t+1|t}(i) - \sigma_\beta^2(i)$.

3.3 Updating formulae for model probabilities

Finally, we need the updating formula for $\pi_t(i) = \Pr(i|p^t, z^t)$. We will make use of Cogley and Sargent (2005), Appendix A to get the recursion. Writing $\pi_t(i) = \Pr(i|p^t, z^t)$, we have

$$\begin{aligned}
\pi_t(i) &= \frac{f(p^t, z^t|i)\pi_0(i)}{f(p^t, z^t)} \propto m_{it}\pi_0(i), \text{ where} \\
m_{it} &= f(p^t, z^t|i) \\
&= \int L(p^t, z^t|\beta_0(i); i) f(\beta_0(i)) d\beta_0(i).
\end{aligned}$$

Here $f(\beta_0(i))$ denotes the prior distribution for $\beta_0(i)$, $f(p^t, z^t|i)$ denotes the probability density for (p^t, z^t) conditional on the model, $L(p^t; z^t; \beta_0(i)|i)$ is the likelihood function for model i, and $\pi_0(i)$ is the prior probability for model $i = 0, 1$ (with $\pi_0(1) = \pi_0$ and $\pi_0(0) = 1 - \pi_0(1)$). Moreover, $f(p^t, z^t)$

is the marginal distribution of the sample across models. Thus, m_{it} is the marginalized likelihood for model i. Since $\sigma_{\beta(0)}$ and $\sigma_{\beta(1)}$ are assumed known, we have not made the dependence of the distributions on them explicit.

First note that

$$\begin{aligned} f(\beta_t(i)|p^t, z^t; i) \times m_{it} &= L(p_t, z_t|p^{t-1}, z^{t-1}, \beta_t(i); i) \times f(\beta_t(i)|p^{t-1}, z^{t-1}; i) \\ &\quad \times f(p^{t-1}, z^{t-1}|i), \text{ or} \end{aligned}$$

$$\frac{m_{it}}{m_{i,t-1}} = \frac{L(p_t, z_t|p^{t-1}, z^{t-1}, \beta_t(i); i) \times f(\beta_t(i)|p^{t-1}, z^{t-1}; i)}{f(\beta_t(i)|p^t, z^t; i)} \equiv \tilde{A}_t(i).$$

Here $f(\beta_t(i)|p^{t-1}, z^{t-1}; i)$ denotes the normal density with mean $\hat{\beta}_{t|t-1}(i)$ and variance $V_{t|t-1}(i)$, i.e.

$$f(\beta|p^{t-1}, z^{t-1}; i) = (2\pi V_{t|t-1}(i))^{-1/2} \exp\left\{-\frac{1}{2}\frac{(\beta - \hat{\beta}_{t|t-1}(i))^2}{V_{t|t-1}(i)}\right\}.$$

Similarly,

$$f(\beta|p^t, z^t; i) = (2\pi V_{t|t}(i))^{-1/2} \exp\left\{-\frac{1}{2}\frac{(\beta - \hat{\beta}_{t|t}(i))^2}{V_{t|t}(i)}\right\},$$

and

$$L(p_t, z_t|p^{t-1}, z^{t-1}, \beta; i) = f(p_t|z_{t-1}, \beta; i) \times f(z_t|z_{t-1}),$$

where

$$\begin{aligned} f(p_t|z_{t-1}, \beta; i) &= (2\pi\sigma_p^2(i))^{-1/2} \exp\left\{-\frac{1}{2}\frac{(p_t - z_{t-1}\beta)^2}{\sigma_p^2(i)}\right\}, \text{ and} \\ f(z_t|z_{t-1}) &= (2\pi\sigma_z^2)^{-1/2} \exp\left\{-\frac{1}{2}\frac{(z_t - z_{t-1}\rho)^2}{\sigma_z^2}\right\}. \end{aligned}$$

It can be verified that $m_{it}/m_{i,t-1}$ does not depend on $\beta_t(i)$ even though each of the three terms in the expression does.[7] In fact,

$$\tilde{A}_t(i) = f(z_t|z_{t-1})A_t(i), \text{ where}$$

[7] A Mathematica routine for this is available on request.

$$A_t(i) = \exp\left\{-\frac{(p_t - \hat{\beta}_{t|t-1}(i)z_{t-1})^2}{2(\sigma_p^2(i) + V_{t|t-1}(i)z_{t-1}^2)}\right\}\left(2\pi(\sigma_p^2(i) + V_{t|t-1}(i)z_{t-1}^2)\right)^{-1/2}.$$

Since

$$\pi_t(i) = \Pr(i|p^t, z^t) = \frac{f(p^t, z^t|i)\pi_0(i)}{f(p^t, z^t)},$$

the posterior odds ratio for the two models is given by

$$r_t = \frac{\pi_t(1)}{\pi_t(0)} = \frac{f(p^t, z^t|i=1)\pi_0(1)}{f(p^t, z^t|i=0)\pi_0(0)} = \frac{m_{1t}}{m_{0t}}$$

assuming $\pi_0(1) = \pi_0(0)$ for the prior of the two models. (More generally, the prior odds ratio would come in.). We thus have

$$\frac{m_{1,t+1}}{m_{0,t+1}} = \frac{m_{1,t}}{m_{0,t}}\frac{A_{t+1}(1)}{A_{t+1}(0)}.$$

We then use the fact that $m_{1,t}/m_{0,t} = \pi_t(1)/(1 - \pi_t(1))$ in the last equation. Solving for $\pi_{t+1}(1)$ then gives:

$$\pi_{t+1}(1) = \frac{\pi_t(1)A_{t+1}(1)}{A_{t+1}(0) - \pi_t(1)A_{t+1}(0) + \pi_t(1)A_{t+1}(1)}.$$

This equation describes the updating of the model weights over time and completes our specification of the formulae for the posteriors of the parameters of both forecasting models and for the posterior probabilities of the two models.

4 Simulation Results

We now present simulation results for our set-up. Key parameters are the model parameters α and ρ and the belief parameter $\sigma_\beta(1)$. We set $\sigma_\beta(0) = 0$. Other model parameters are set at values $\delta = 0.5, \sigma_z = 1$ and $\sigma_a = 1$ in the simulations. We assume that agents set $\sigma_p(1) = \sigma_p(0) = 1$. Their priors at $t = 1$ are assumed to be

$$\begin{aligned} V_{1|0}(i) &= 0.2 \text{ for } i = 0, 1 \\ \hat{\beta}_{1|0}(i) &= \bar{b} + \varepsilon_b\sqrt{V_{1|0}(i)} \text{ for } i = 0, 1, \end{aligned}$$

α	% model 1	% model 0	% unselect
-0.5	0	100.0	0
0.1	0	100.0	0
0.4	0.2	95.8	4.0
0.5	5.7	82.7	11.6
0.6	32.7	60.5	6.8
0.7	56.0	43.3	0.7
0.85	70.0	30.0	0
0.95	61.6	38.2	0.2
0.99	49.3	50.7	0

Table 1: The role of expectations feedback in model selection. Percentage of 10,000 simulation runs selecting either model 1, model 2, or neither after 40,000 periods.

where ε_b is a standard normal random variable. In addition we set the model priors as

$$\pi_0(i) = 0.5 \text{ for } i = 0, 1.$$

Except where otherwise stated, we simulate for $T = 40,000$ periods and do $N = 10,000$ simulations. In our tables we report the proportion of the simulations in which each model is selected. To assess this we say that model 1 is selected if $\pi_T(1) > 0.99999999$ and model 0 is selected if $\pi_T(1) < 0.00000001$. In our benchmark table the role of α is studied, and the other key parameters are set at $\sigma_\beta(1) = 0.005$ and $\rho = 0$. The results are in Table 1:

The results in Table 1 are intriguing, and are further illustrated in Figure 1. For $\alpha \ll 0.5$ learning with Bayesian model averaging converges to the REE with high (empirical) probability. As the value of α gets closer to 0.5 the probability starts to fall below one and for values near $\alpha = 0.5$ both cases of selection of the TVP model and of non-selection have small positive probabilities. As the value of α is raised above 0.5 the frequency of selection of the TVP model increases but in a non-monotonic way as $\alpha \to 1$.

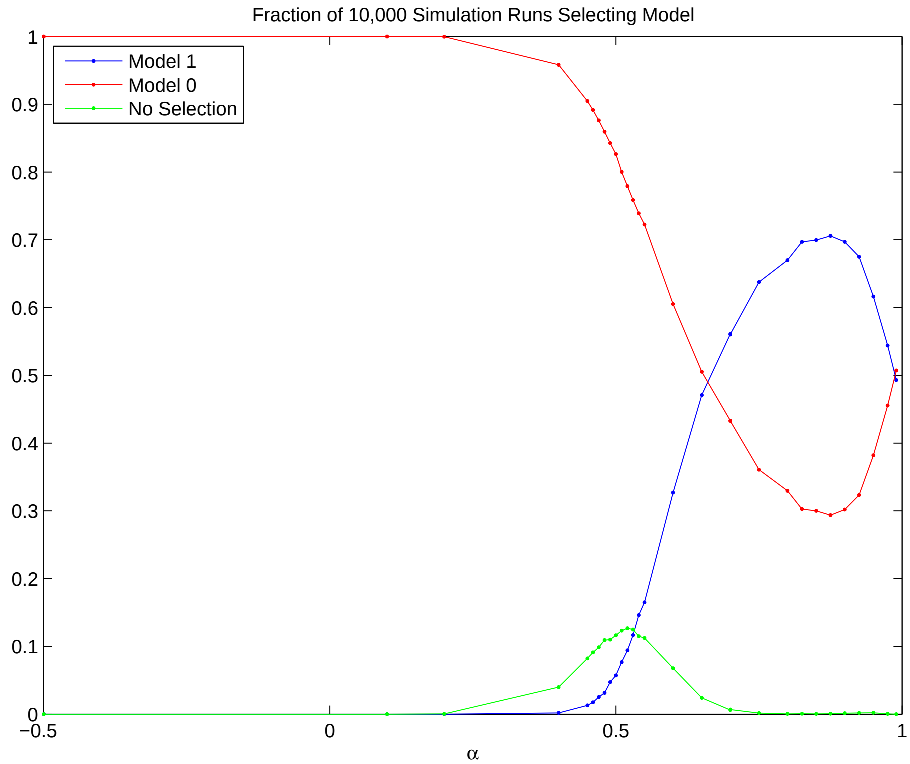

Figure 1: Proportions of selections of models 0 and 1.

It can be seen from Figure 1 that a fairly strong positive expectational feedback creates the possibility that agents come to believe that the economy is generated by the time-varying parameter model 1. When there is negative expectations feedback, as in the Muth cobweb model with normal supply and demand slopes, then agents learn asymptotically the REE. Convergence to the REE also occurs when there is positive expectational feedback that is not too strong, i.e. for $\alpha \ll 0.5$. However for $\alpha > 0.5$ there is a clear possibility of convergence to a non-RE equilibrium in which agents believe that the endogenous variable p_t is generated by a time-varying model of the form $p_t = b_t z_{t-1} + \varepsilon_{pt}$ where b_t follows a random walk.

Interestingly, the dependence of the selection result on α is non-monotonic. As α increases from 0.1, the proportion of simulations selecting model 1 increases until around $\alpha = 0.85$. At that point, further increases in α lead to reductions in selection of model 1, and as α gets very close to 1 the proportions are nearly 50-50. Thus, sufficiently strong expectations feedback makes

ρ	% model 1	% model 0	% unselect
0.99	60.8	39.2	0
0.90	54.5	45.4	0.1
0.75	50.5	48.8	0.6
0.25	46.7	51.0	2.3
0	46.8	50.6	2.6
-0.25	47.2	50.2	2.6
-0.75	50.3	49.2	0.5
-0.90	54.0	45.9	0.1
-0.99	60.6	39.5	0

Table 2: Robustness of results with respect to autocorrelation of observable shocks.

agents more likely to believe in the drifting parameter model, but with very strong feedback they are just as likely to believe in the constant parameter model.

We remark that for cases in which model 1 is selected at $T = 40,000$ we have found that in longer simulations of $T = 100,000$ the convergence is to $\pi_T(1) = 1$ up to computer accuracy. This implies that such a numerical simulation would deliver $\pi_t(1) = 1$ for all $t > T$.

Next we consider the sensitivity of the outcomes to the parameters ρ and $\sigma_\beta(1)$. We fix $\alpha = 0.65$ and $\sigma_\beta(1) = 0.005$, and first consider different values of ρ. The results are reported in Table 2. It can be seen that impact of ρ, the correlation in z_t, is fairly small. Larger values of ρ, either positive or negative, increase the likelihood of model 1 being selected. However this impact is not very substantial.

Finally, for $\alpha = 0.65$ and $\rho = 0$ we consider variations in $\sigma_\beta(1)$. The results are shown in Table 3. As the perceived parameter variation in the forecasting model 1 falls, the proportion of simulations converging to model 0 decreases and, apparently falls to zero for $\sigma_\beta(1)$ sufficiently small. However, for small values of $\sigma_\beta(1)$ Table 3 suggests the possibility of non-selection between the two models. To study this further we considered longer simulations of $T = 1,000,000$ for the cases when $\sigma_\beta(1) < 0.005$. These results are shown in the second section of the table. There we see that the instances of non-selection eventually resolve, and the proportion of simulations converging to model 1 continues to increase as the variability in the its random walk innovation decreases. Intuitively, for $\sigma_\beta(1)$ small it is more difficult for agents

$\sigma_\beta(1)$	% model 1	% model 0	% unselect
1.0000	3.2	96.8	0
0.5000	7.0	93.0	0
0.2500	14.1	85.9	0
0.1000	21.3	78.7	0
0.0500	29.6	70.4	0
0.0100	41.4	58.5	0.1
0.0050	47.3	50.7	2.0
0.0025	42.2	30.1	27.7
0.0010	31.1	0	68.9
0.0005	22.8	0	77.2
0.0050	48.4	51.7	0
0.0025	52.8	47.2	0
0.0010	59.1	40.9	0
0.0005	62.9	36.9	0.2

Table 3: Role of standard deviation of random walk in model selection. The first section is $T = 40,000$, while the second is for $T = 1,000,000$.

to distinguish between the two models, which is why the selection results require longer samples. Nevertheless, for large enough samples, model 1 is increasingly selected. For this case of $\alpha = 0.65$ there is sufficient dependence on expectations to make the time-varying parameter model a possible outcome, but the likelihood of this outcome increases as the drift in coefficients becomes smaller. That is, a slowly varying random walk seems to be a better fit for the actual law of motion than a model with substantial parameter drift. [8]

We normally do not see switching between $\pi_t(1)$ near 0 and near 1 within a simulation. There may possibly be such cases for α near 0.5, since we have observed a few cases when $\alpha = 0.4$ and $\sigma_\beta(1) = 0.001$ in which $\pi_t(1)$ is near 1 at $t = 40,000$, which might resolve to $\pi_t(1) = 0$ eventually.

[8] For $\alpha = 0.4$ it is also the case that for $\sigma_\beta(1)$ sufficiently small the possibility of non-selection increases with relatively shorter T. But as T is increased the proportion of non-selecting simulations falls and agents increasingly select model 0.

5 Additional Interpretation and Analysis

The preceding results are not straightforward to interpret, because a number of forces are at play. In this section we provide two auxiliary results that will help in obtaining a partial interpretation of the surprising result that non-RE forecasting model 1 has fair chance to get selected when the expectational feedback is sufficiently strong.

The first result is about comparison of rational and non-rational mean forecast errors. At any moment of time, for each model $i = 0, 1$ agents have estimates $\hat{\beta}_{t|t-1}(i)$ of β_t, which are used to form forecasts of p_t. For convenience, we use the temporary notation

$$b = \pi_{t-1}(0)\hat{\beta}_{t|t-1}(0) + \pi_{t-1}(1)\hat{\beta}_{t|t-1}(1).$$

Then the following result will be useful in interpreting our results.

Lemma 2 *Suppose at time t agents believe in the PLM*

$$p_t = bz_{t-1} + \varepsilon_{pt}$$

with probability 1, where $b \neq \bar{b} = \delta/(1-\alpha)$. (Here $\bar{b}$ is the REE coefficient value). Then under the resulting actual law of motion (ALM), the forecast $E^*_{t-1}p_t = bz_{t-1}$ has lower conditional MSE (mean square forecast error, conditional on z_{t-1}) than the "REE forecast" $E^*_{t-1}p_t = \bar{b}z_{t-1}$, provided $1/2 < \alpha < 1$.

Suppose that model 0 converges to $\bar{b}$, i.e. $\hat{\beta}_{t|t-1}(0) \to \bar{b}$ as $t \to \infty$. We know this is the case if only model 0 is used by agents, but it plausibly holds more generally when both models are in play. In that case, the forecasts from model 1 will be more accurate, in the conditional MSE sense, if and only if $\alpha > 0.5$. On average the more accurate forecasting model will lead to an upward revision in its probability. Thus, for $\alpha > 0.5$ one would expect $\pi_t(1)$ to tend to increase over time. This suggests that one may have $\pi_t(1) \to 1$ when $\alpha > 0.5$.[9]

Special case: agents set $\pi_t(1) = 1$ for all t. Here the "small constant gain" approximation (i.e. small $\sigma_\beta(1)$) is employed, making the framework

[9]The Lemma also holds for $\alpha \geq 1$, but $\alpha < 1$ is a maintained assumption (the standard condition for stability under LS learning).

similar to that used in Evans, Honkapohja, and Williams (2010). Formally, the setting is now

$$\begin{aligned} p_t &= \beta_t z_{t-1} + \sigma_p \varepsilon_{pt}, \text{ and} \\ \beta_t &= \beta_{t-1} + \sigma_\beta \varepsilon_{\beta t} \end{aligned}$$

for the PLM, where for brevity the index to model 1 has been omitted in β_t, σ_p and σ_β. The rest of the system is

$$\begin{aligned} p_t &= \alpha E^*_{t-1} p_t + \delta z_{t-1} + \sigma_a \varepsilon_{at}, \\ E^*_{t-1} p_t &= \hat{\beta}_{t|t-1} z_{t-1}, \\ \hat{\beta}_{t+1|t} &= \hat{\beta}_{t|t-1} + \frac{V_{t|t-1} z_{t-1}}{\sigma_p^2 + V_{t|t-1} z_{t-1}^2} (p_t - \hat{\beta}_{t|t-1} z_{t-1}) \\ V_{t+1|t} &= V_{t|t-1} - \frac{z_{t-1}^2 V_{t|t-1}^2}{\sigma_p^2 + V_{t|t-1} z_{t-1}^2} + \sigma_\beta^2. \end{aligned}$$

The ALM is $p_t = (\hat{\beta}_{t|t-1} + \delta) z_{t-1} + \sigma_a \varepsilon_{at}$. We analyze this system in the Appendix. The analysis implies that asymptotically $\hat{\beta}_{t+1|t}$ is approximately an AR(1) process, with mean equal to the RE value, and with (i) a variance proportional to $\sigma_\beta^2(1)$ and (ii) a first-order autocorrelation parameter that tends to 1 as $\sigma_\beta^2(1)$ tends to 0.

The result for the special case and the above Lemma suggest the reason for convergence of Bayesian updating to the TVP model 1, provided $\sigma_\beta^2(1) > 0$ is small enough. A sketch of the argument is as follows (we are assuming $\sigma_\beta^2(1) > 0$, $\sigma_\beta^2(0) = 0$). Suppose that $\pi_t(1) = 1$, all t, i.e. agents believe in the TVP model 1 w.p.1 for all t. Under model 0, we expect $\beta_{t+1|t}(0) \to \bar{b}$ as $t \to \infty$, based on the approximation results. Under model 1 we will have $\beta_{t+1|t}(1)$ converging to a distribution centered on $\bar{b}$, with deviations that are strongly and positively serially correlated. Hence by the Lemma the average forecast error under model 1 will be less than under model 0 if $0.5 < \alpha < 1$. Since actual squared errors strongly impact the evolution of $\pi_t(1)$ this strongly suggests that $\pi_t(1) = 1$ can be a stable outcome. (However other factors influence the $\pi_t(1)$ updating.) This argument also suggests that for $\alpha < 0$ and for $0 \leq \alpha < 0.5$ model 1 will not be a fixed point asymptotically.

6 Conclusions

It is natural to assume that in situations with imperfect knowledge economic agents try to use multiple models and weighted averages of forecasts when they form expectations about the future. We consider the consequences of such practices in a simple self-referential model in which expectations affect outcomes and agents learn using appropriate Bayesian techniques. At the outset we impose the assumption that one of the forecasting models employed by contains a "grain of truth", i.e., for particular parameter values that model corresponds to the correct forecasting model in an REE.

The central result in our paper shows that convergence of learning with Bayesian model averaging to an REE occurs only when the feedback of agents' expectations on actual outcomes is relatively weak, less than 0.5. See Table 1 and Figure 1. This observation should be contrasted with Proposition 1, where it is showed that when agents only use the correctly specified forecasting model Bayesian learning converges to the REE provided that expectations feedback has coefficient less than 1.

More generally, it is seen from Table 1 and Figure 1 that learning by Bayesian model averaging leads to selection of a unique forecasting model with very high probability. However, the selection can be a misspecified forecasting model when the expectations feedback parameter has sufficiently high value. We allow agents to consider a drifting coefficients model and use it to form expectations. If the feedback from expectations to outcomes is sufficiently high, then the resulting drift in the actual data generating process may justify selecting the drifting coefficients model. Thus, even though a constant parameter model is correctly specified asymptotically, the process of learning may make agents doubt that model.

7 Appendix

Proof of Proposition 1 (outline): Define

$$S_{t-1}^{-1} = \frac{tV_{t-1}}{\sigma_\eta^2 + V_{t-1}z_{t-1}^2}.$$

Then we have

$$S_{t-1} = \frac{1}{t}(\sigma_\eta^2 + V_{t-1}z_{t-1}^2)/V_{t-1}, \text{ or} \tag{2}$$

$$\sigma_\eta^2/V_{t-1} = tS_{t-1} - z_{t-1}^2 \tag{3}$$

and

$$\begin{aligned} V_t &= V_{t-1}\left(1-\frac{z_{t-1}^2 V_{t-1}}{\sigma_\eta^2+V_{t-1}z_{t-1}^2}\right)=V_{t-1}\frac{\sigma_\eta^2}{\sigma_\eta^2+V_{t-1}z_{t-1}^2} \\ &= V_{t-1}\frac{\sigma_\eta^2/V_{t-1}}{\sigma_\eta^2/V_{t-1}+z_{t-1}^2}=V_{t-1}\frac{tS_{t-1}-z_{t-1}^2}{tS_{t-1}}. \end{aligned}$$

Using the last expression and (3), we also have

$$V_t=\frac{\sigma_\eta^2}{(t+1)S_t-z_t^2}=\frac{\sigma_\eta^2}{tS_{t-1}-z_{t-1}^2}\frac{tS_{t-1}-z_{t-1}^2}{tS_{t-1}}$$

and so

$$\begin{aligned} \frac{1}{(t+1)S_t-z_t^2} &= \frac{1}{tS_{t-1}} \text{ or} \\ (t+1)S_t-z_t^2 &= tS_{t-1}, \end{aligned}$$

from which we get

$$\begin{aligned} S_t &= \frac{t}{t+1}S_{t-1}+\frac{1}{t+1}z_t^2 \\ &= S_{t-1}+\frac{1}{t+1}(z_t^2-S_{t-1}). \end{aligned}$$

Since $E_{t-1}^* p_t = b_{t-1}z_{t-1}$ note that $p_t=(\alpha b_{t-1}+\delta)z_{t-1}+\eta_t$.

Collecting the results together, the system under Bayesian learning is

$$\begin{aligned} b_t &= b_{t-1}+t^{-1}S_{t-1}^{-1}z_{t-1}(p_t-b_{t-1}z_{t-1}) \\ S_t &= S_{t-1}+t^{-1}\left(\frac{t}{t+1}\right)(z_t^2-S_{t-1}) \\ p_t &= (\alpha b_{t-1}+\delta)z_{t-1}+\eta_t, \\ z_t &= \rho z_{t-1}+w_t. \end{aligned}$$

Since

$$b_t=b_{t-1}+t^{-1}S_{t-1}^{-1}z_{t-1}(z_{t-1}(\alpha-1)b_{t-1}+\delta+\eta_t),$$

we can apply standard results on stochastic recursive algorithms to show convergence. Note that the starting point is $S_0=\frac{V_0}{\sigma_\eta^2+V_0z_0^2}$. As usual,

$$\lim_{t\to\infty}(b_t,S_t)=((1-\alpha)^{-1}\delta,\bar{S}).$$

Q.E.D.

Proof of Lemma 2: The actual model is $p_t = \alpha E^*_{t-1}p_t + \delta z_{t-1} + \varepsilon_{at}$. For this PLM the ALM is

$$p_t = (\alpha b + \delta)z_{t-1} + \varepsilon_{at}.$$

Thus, the forecast $E^*_{t-1}p_t = bz_{t-1}$ has lower conditional MSE than $E^*_{t-1}p_t = \bar{b}z_{t-1}$ when

$$\begin{aligned} |(\alpha b + \delta) - b| &< \left|(\alpha b + \delta) - \bar{b}\right|, \text{ i.e.} \\ \left|(\alpha - 1)(b - \bar{b})\right| &< \left|\alpha(b - \bar{b})\right|, \text{ or} \\ |\alpha - 1| &< |\alpha|, \end{aligned}$$

which holds for $\alpha > 0.5$ and fails to hold for $\alpha < 0$ and for $0 \leq \alpha < 0.5$. Q.E.D.

Analysis of the case $\pi_t(1) = 1$: Letting $\hat{P}_t = P_t/\sigma^2_\beta$, we get

$$\begin{aligned} \hat{\beta}_{t+1|t} &= \hat{\beta}_{t|t-1} + \sigma^2_\beta \frac{\hat{P}_t z_{t-1}}{1 + \sigma^2_\beta \hat{P}_t z^2_{t-1}}[(p_t - \hat{\beta}_{t|t-1}z_{t-1}) + \sigma_a \varepsilon_{at}] \\ \hat{P}_{t+1} &= \hat{P}_t + \sigma^2_\beta \left[\frac{1}{\sigma^2_p} - \frac{z^2_{t-1}\hat{P}^2_t}{1 + \sigma^2_\beta \hat{P}_t z^2_{t-1}}\right], \end{aligned}$$

which is a constant-gain stochastic recursive algorithm (SRA) when σ^2_β is treated as the gain. The associated differential equation is

$$\frac{d\beta}{d\tau} = \hat{P}\sigma^2_z[(\alpha - 1)\beta + \delta] \quad (4)$$

$$\frac{d\hat{P}}{d\tau} = \frac{1}{\sigma^2_p} - \hat{P}^2\sigma^2_z \quad (5)$$

with fixed point

$$\begin{aligned} \bar{\beta} &= \delta(1 - \alpha)^{-1} \\ \bar{P} &= \frac{1}{\sigma_z \sigma_p}. \end{aligned}$$

We consider the stochastic differential equation approximation to the algorithm. Introduce the notation

$$\theta = \begin{pmatrix} \beta \\ \hat{P} \end{pmatrix},$$

The mean dynamics $\tilde{\theta}(\tau, a)$ are given by (4)-(5) and we write the system in vector form as

$$\frac{d\theta}{d\tau} = h(\theta). \tag{6}$$

We also define

$$U^{\sigma_\beta^2}(\tau) = \sigma_\beta^{-1}[\theta^{\sigma_\beta^2}(\tau) - \tilde{\theta}(\tau, a)],$$

where a is the initial condition for (6). As $\sigma_\beta^2 \to 0$, the normalized deviation $U^{\sigma_\beta^2}(\tau)$ for $0 \leq \tau \leq \bar{\tau}$ converges weakly to the solution $U(\tau)$ of

$$dU(\tau) = D_\theta h(\tilde{\theta}(\tau, a))U(\tau)d\tau + R^{1/2}(\tilde{\theta}(\tau, a))dW(\tau) \tag{7}$$

with initial condition $U(0) = 0$.[10] Here

$$A = Dh(\bar{\theta}) = \begin{pmatrix} \hat{P}\sigma_z^2(\alpha - 1) & 0 \\ 0 & -2\frac{\sigma_z}{\sigma_p} \end{pmatrix},$$

where $\bar{\theta}' = (\bar{\beta}, \bar{P})'$. Writing the SRA in vector form as

$$\theta_t = \theta_{t-1} + \gamma H(\theta_{t-1}, X_t),$$

where $X_t' = (z_{t-1}, \varepsilon_{at})$ we consider (7) from the starting point $\bar{\theta}$, i.e.

$$dU(\tau) = D_\theta h(\bar{\theta})U(\tau)d\tau + \bar{R}^{1/2}dW(\tau), \text{ where } \bar{R} = R(\bar{\theta}).$$

The elements of $\bar{R}$ are given by

$$\bar{R}^{ij} = \sum_{k=-\infty}^{\infty} cov[H_i(\bar{\theta}, X_k), H_j(\bar{\theta}, X_0)].$$

In particular, for

$$\bar{R}^{11} = \sum_{k=-\infty}^{\infty} cov[H_1(\bar{\theta}, X_k), H_1(\bar{\theta}, X_0)]$$

we get

$$\begin{aligned} H_1(\bar{\theta}, X_t) &= \frac{\bar{P}z_{t-1}}{1 + 1 + \sigma_\beta^2 \bar{P} z_{t-1}^2}[((\alpha - 1)\bar{\beta} + \delta)z_{t-1} + \sigma_a \varepsilon_{at}] \\ &= \frac{\bar{P}\sigma_a z_{t-1}}{1 + \sigma_\beta^2 \bar{P} z_{t-1}^2} z_{t-1}\varepsilon_{at}. \end{aligned}$$

[10] We use the results in Section 7.4 of Evans and Honkapohja (2001).

It follows that

$$cov[H_1(\bar{\theta}, X_k), H_1(\bar{\theta}, X_0)] = 0 \text{ for } k \neq 0$$

since $\{\varepsilon_{ak}\}$ is independent of $\{z_t\}$ and $E\varepsilon_{ak}\varepsilon_{a0} = 0$. Thus,

$$\bar{R}^{11} = var(H_1(\bar{\theta}, X_t)) = \bar{P}\sigma_a^4 var\left(\frac{z_{t-1}}{1 + \sigma_\beta^2 \bar{P} z_{t-1}^2}\right).$$

In particular,

$$\bar{R}^{11} \to \bar{P}\sigma_a^4 var(z_t) \text{ as } \sigma_\beta^2 \to 0,$$

where $\bar{P} = \sigma_z^{-1}\sigma_p^{-1}$. If $\sigma_a = \sigma_p$ then $\bar{R}^{11} \to \sigma_p^3\sigma_z$.

Next, we compute

$$cov[U(\tau), U(\tau - \hat{\tau})] \equiv \rho(\tau, \tau - \hat{\tau}) = \exp(\hat{\tau}A).C,$$

where

$$C = \int_0^\infty \exp(vA)\bar{R}\exp(vA)dv.$$

Note also that $\bar{R}^{12} = 0$ as $H_2(\bar{\theta}, X_0)$ does not depend on ε_{at}. It follows that $\exp(vA)\bar{R}\exp(vA)$ is diagonal and

$$\begin{aligned} C^{11} &= \bar{R}^{11}\int_0^\infty \exp(v\sigma_z\sigma_p^{-1}(\alpha - 1)2)dv \\ &= \bar{R}^{11}\frac{\sigma_p}{2(1-\alpha)\sigma_z} = \frac{\sigma_p^4}{2(1-\alpha)} \end{aligned}$$

since $\bar{R}^{11} \to \sigma_p^3\sigma_z$ as noted above. This implies that the autocorrelation function of $U(\tau)$ is $\tilde{r}(\hat{\tau}) = \exp(-\hat{\tau}\frac{\sigma_z}{\sigma_p}(1-\alpha))$. As $\bar{\theta} + \sigma_\beta^2 U^{\sigma_\beta^2}(\tau) = \theta^{\sigma_\beta^2}(\tau)$ we have the approximation $\theta^{\sigma_\beta^2}(\tau) \approx \theta_t$, where $\tau = \sigma_\beta^2 t$. Using the notation $\hat{\beta}_t = \hat{\beta}_{t+1|t}$ we have

$$\begin{aligned} cor(\hat{\beta}_t, \hat{\beta}_{t-k}) &= cor[U_\beta^{\sigma_\beta^2}(\tau)(\sigma_\beta^{-2}\tau), U_\beta^{\sigma_\beta^2}(\tau)(\sigma_\beta^{-2}\tau - k)] \\ &= \tilde{r}_\beta(\sigma_\beta^2 k) = \exp[-\sigma_\beta^2\frac{\sigma_z}{\sigma_p}(1-\alpha)k]. \end{aligned}$$

Thus, for any $k > 1$, $cor(\hat{\beta}_t, \hat{\beta}_{t-k}) \to 1$ as $\sigma_\beta^2 - 0$.

References

Bray, M., and N. Savin (1986): "Rational Expectations Equilibria, Learning, and Model Specification," *Econometrica*, 54, 1129–1160.

Bullard, J. (1992): "Time-Varying Parameters and Nonconvergence to Rational Expectations under Least Squares Learning," *Economics Letters*, 40, 159–166.

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

Cogley, T., and T. J. Sargent (2005): "The Conquest of US Inflation: Learning and Robustness to Model Uncertainty," *Review of Economic Dynamics*, 8, 528–563.

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

——— (2009): "Learning and Macroeconomics," *Annual Review of Economics*, 1, 421–451.

Evans, G. W., S. Honkapohja, and N. Williams (2010): "Generalized Stochastic Gradient Learning," *International Economic Review*, 51, 237–262.

Hamilton, J. D. (1994): *Time Series Analysis*. Princeton University Press, Princeton, NJ.

McGough, B. (2003): "Statistical Learning and Time Varying Parameters," *Macroeconomic Dynamics*, 7, 119–139.

Sargent, T. J. (2008): "Evolution and Intelligent Design," *American Economic Review*, 98, 5–37.

Sargent, T. J., and N. Williams (2005): "Impacts of Priors on Convergence and Escapes from Nash Inflation," *Review of Economic Dynamics*, 8, 360–391.

Slobodyan, S., and R. Wouters (2008): "Estimating a Medium-Scale DSGE Model with Expectations Based on Small Forecasting Models," mimeo.

Young, H. P. (2004): *Strategic Learning and Its Limits*. Oxford University Press, Oxford UK.