Appendix to Behavioral Learning Equilibria Appendix on Learning with constant gains (not for publication)

Cars Hommes[a], Mei Zhu[b,a] *

[a] CeNDEF, School of Economics, University of Amsterdam

Roetersstraat 11, 1018 WB Amsterdam, Netherlands

[b] Institute for Advanced Research, Shanghai University of Finance and Economics,

and the Key Laboratory of Mathematical Economics(SUFE), Ministry of Education, Shanghai 200433, China

August 1, 2013

1 Learning with constant gains

In the paper *Behavioral Learning Equilibria*, Hommes and Zhu (2012), stability of SCEE under SAC-learning with decreasing gains was examined. Decreasing gains allow for the possibility of full convergence to SCEE, and are thus convenient for studying local stability under learning. It is known that constant gains are appropriate if agents want to allow for the possibility of structural change of an unknown form. When agents use constant-gain econometric learning, which to some extent discounts past dynamics, the learning system is time-invariant with an ergodic distribution that can be studied, as shown in Branch and Evans (2011). In this Appendix we study the learning dynamics with constant gains in the two applications by some real-time learning simulations.

Following Evans and Honkapohja (2001), we restrict attention to the case without the so-called complementary term. Based on equations (2.8), SAC-learning with constant

* *E-mail addresses:* C.H.Hommes@uva.nl, M.Zhu@uva.nl

gain is given by[1]

$$\begin{cases} \alpha_t = \alpha_{t-1} + \kappa(x_t - \alpha_{t-1}), \\ \beta_t = \beta_{t-1} + \kappa R_t^{-1}[(x_t - \alpha_{t-1})(x_{t-1} - \alpha_{t-1}) - \beta_{t-1}(x_t - \alpha_{t-1})^2], \end{cases} \tag{1.1}$$

where

$$R_t = R_{t-1} + \kappa[(x_t - \alpha_{t-1})^2 - R_{t-1}]$$

is the estimated state variance. The parameter κ is a constant gain governing the rate at which agents adjust their econometric estimates in a real-time learning algorithm, describing the sensitivity of agents' beliefs to recent data. Smaller values of κ imply more smoothing in estimating mean, autocorrelation and variance.

We now turn to simulations of the real-time versions of the two applications, the asset pricing model and the New Keynesian Philips curve.

1.1 Asset pricing model under constant gain learning

Branch and Evans (2011) demonstrate that under constant gains learning, if agents' beliefs are displaced away from the REE, the transitional path may include agents temporarily believing tha stock prices follow a random walk without drift and such random-walk beliefs are almost self-fulfilling. We show here the random-walk beliefs are not only the result of transitory learning dynamics, but also they do not vanish after the economy converges to the equilibrium for some plausible parameters.

In the special case $\rho = 0$, very similar to the asset pricing model in Branch and Evans (2011), $\beta^* = 0$ as discussed in Section 3 in our revised manuscript. We are interested in the case with nonzero β^*. For convenience of comparison with the model in Branch and Evans (2011), we first take a small $\rho (= 0.1)$. In this case $\beta^* = 0.112$[2]. Figure 1 illustrates the dynamics under constant-gain learning with $p_0 = 1.5$, $\kappa = 0.01$, $R = 1.05$. When there is a sequence of shocks that drives the system far away from the SCEE, with relatively small constant gain agents predict the mean α of prices in the following periods also high as shown in the top panel in Figure 1 (the dotted line), which reinforces the following prices. One the other hand, the path back to the SCEE governed by the equations (3.13) in our revised manuscript induces the prices to decrease [3]. Hence the

[1] If the SAC-learning is modified with $R_t = S_{t-1}$, the following results still hold.

[2] If β is not limited in $[-1, 1]$, in fact, there is another $\beta^* (= 1.065)$.

[3] For the related theoretical analysis under constant-gain learning, the readers are referred to Evans and Honkapohja (2001). We leave this to future work.

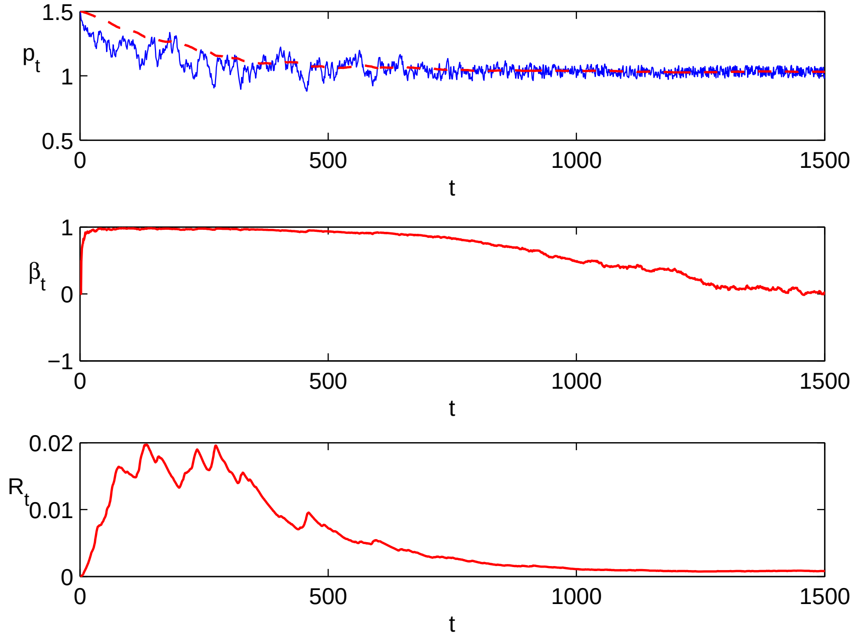

Figure 1: Asset pricing model under constant gain learning with $\kappa = 0.01$ and $\rho = 0.1$.

combination of the beliefs and the mean dynamics leads prices to decrease gradually with high persistence, as shown in the top panel in Figure 1 (the real line). This results in increasing serial correlations up to random-walk beliefs. As shown in the second panel in Figure 1, the autocorrelation β_t rapidly increases up to approximately 1 and persists for some period. That is, random-walk beliefs are almost self-fulfilling. Note that, for $\beta_t \approx 1$, $\alpha_t(1-\beta_t) \approx 0$, hence the random-walk process is without drift. Moreover, in the ODE (3.13) in our revised manuscript governing the mean dynamics under constant-gain learning, for $\beta = 1$, $F(\beta) - \beta = -0.0391$, which is very close to zero. Thus if a sequence of random shocks leads agents to have random-walk beliefs, these beliefs are likely to be sustained for a substantial period of of time. The beliefs do eventually return to fluctuate around the SCEE ($\beta^* = 0.112$). The bottom panel in Figure 1 plots the estimated state variance R_t, which implies that large R_t corresponds to random-walk beliefs ($\beta_t \approx 1$).

For small ρ, as shown in Figure 2a in our revised manuscript, β^* is also relatively small, and random-walk beliefs are the result of transitory learning dynamics. We now show that for some plausible parameters random-walk beliefs do not vanish after the economy converges to the equilibrium. Corresponding to Figure 1 in our revised manuscript, set $\rho = 0.9$ and then the learning dynamics is displayed in Figure 2. The top panel plots the

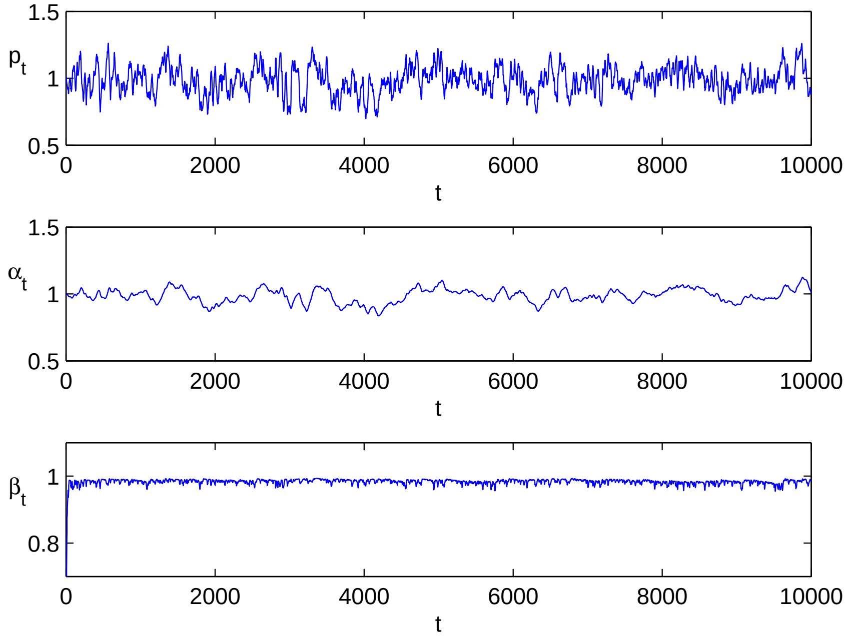

Figure 2: Constant gain learning with $\kappa = 0.01$ and $\rho = 0.9$

stock prices p_t, while the bottom two panels plot the estimated mean α_t and the estimated first-order autocorrelation coefficient β_t, illustrating that the beliefs parameters α_t and β_t are time-varying and fluctuate around the corresponding SCEE values. As shown in Figure 1b in our revised manuscript, if $\beta < \beta^*$, then the actual autocorrelation $F(\beta)$ implied by this belief is much larger than β. Especially with the perception of large autocorrelation, this belief will be reinforced by the actual law of motion (4.8) in our revised manuscript. Hence the autocorrelation β_t increases rapidly up to approximate one and fluctuates around $\beta^* = 0.997$ after some periods, see the bottom panel in Figure 2. Thus agents hold self-fulfilling random walk beliefs and the self-fulfilling random-walk beliefs are not the result of transitory learning dynamics, which instead describe a long-run learning equilibrium. In addition, if κ is increased to $\kappa = 0.1$, the plots (not shown here) for p_t exhibit a smooth process, and α_t and β_t display more volatility.

After a sequence of shocks drives the system far away from the SCEE equilibrium, it is anticipated that the dynamics under constant-gain learning would consist of recurrent episodes of high, then low, persistence and volatility that is similar to what is seen in Branch and Evans (2010, 2011) and in the data. We would like to illustrate this more clearly in the second example of New Keynesian Philips curve with constant gain learning.

1.2 New Keynesian Philips curve under constant gain learning

Under real-time learning agents do not have fixed beliefs. In the case of multiple equilibria, time-varying parameter estimates make it possible that a sequence of shocks could move the economy from one equilibrium to another. As shown in Branch and Evans (2010), a real-time learning formulation is capable of reproducing regime-switching returns and volatilities. We show that in the New Keynesian Philips curve with multiple SCEE under constant gain SAC-learning, inflation is likely to switch between different regimes with high and low persistence.

In Figure 4 in Section 4.2 in our revised manuscript, there are three SCEE. We have shown that for different initial states the learning dynamics will converge to different stable SCEE. Here we show that under constant-gain learning unanticipated shocks push the inflation from one basin of attraction to another with different persistence. As shown in Branch and Evans (2010), the intuition for why the economy may switch from one equilibrium to another revolves around the interaction between the exogenous shocks and the gain parameter κ. A particularly large shock, mediated through beliefs via κ, may induce agents to hold a belief on first-order autocorrelation that falls within another basin of attraction. The economy will then jump between a high and a low persistence equilibrium. Because κ is positive, there are repeated realizations of shocks sufficiently large to cause the economy to switch between equilibria. Beliefs reinforce the direct effect of shocks on inflation. These beliefs are almost self-confirming and tend to last for a period until new shocks push the system into another basin of attraction. The frequency with which these regime switches occur, are governed by a complicated interaction between the gain parameter κ and the stochastic shocks. Figure 3 illustrates these switches between different basins of attraction for $\kappa = 0.008$, where other parameters are taken as in Figure 4 in our revised manuscript. The inflation π_t clearly exhibits regime switching, with the autocorrelation coefficient β_t switching between phases of high persistence with near unit root behaviour and phases of low persistence with close to 0 autocorrelation.

References

[1] Branch, W.A., Evans, G.W., 2010. Asset return dynamics and learning. Review of Financial Studies 23(4), 1651-1680.

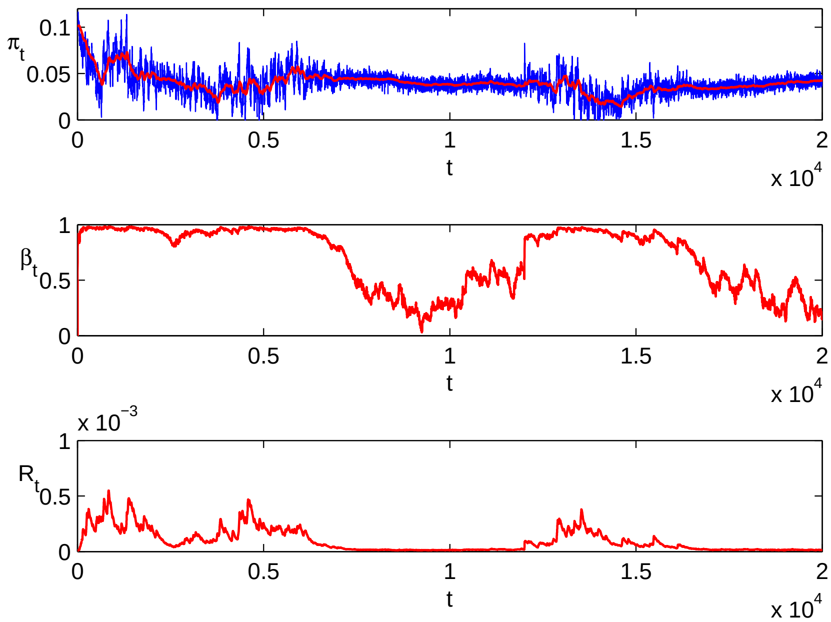

Figure 3: New Keynesian Philips curve under constant gain learning with $\kappa = 0.008$.

[2] Branch, W.A., Evans, G.W., 2011. Learning about risk and return: a simple model of bubbles and crashes. American Economic Journal: Macroeconomics 3(3), 159-191.

[3] Evans, G.W., Honkapohja, S., 2001. Learning and Expectations in Macroeconomics. Princeton University Press, Princeton.