First draft. Please do not quote without authors' permission.

Stochastic Gradient versus Recursive Least Squares Learning

Sergey Slobodyan, Anna Bogomolova, and Dmitri Kolyuzhnov*

CERGE-EI†
Politických vězňů 7, 111 21 Praha 1,
Czech Republic

This version: March 7, 2006

Abstract

In this paper we perform an in–depth investigation of relative merits of two adaptive learning algorithms with constant gain, Recursive Least Squares (RLS) and Stochastic Gradient (SG), using the Phelps model of monetary policy as a testing ground. The behavior of the two learning algorithms is very different. RLS is characterized by a very small region of attraction of the Self–Confirming Equilibrium (SCE) under the mean, or averaged, dynamics, and "escapes", or large distance movements of perceived model parameters from their SCE values.

On the other hand, the SCE is stable under the SG mean dynamics in a large region. However, actual behavior of the SG learning algorithm is divergent for a wide range of constant gain parameters, including those that could be justified as economically meaningful. We explain the discrepancy by looking into the structure of eigenvalues and eigenvectors of the mean dynamics map under the SG learning.

As a result of our paper, we express a warning regarding the behavior of constant gain learning algorithm in real time. If many eigenvalues of the mean dynamics map are close to the unit circle, Stochastic Recursive Algorithm which describes the actual dynamics under learning might exhibit *divergent* behavior despite *convergent* mean dynamics.

JEL Classification: C62, C65, D83, E10, E17

Keywords: constant gain adaptive learning, E–stability, recursive least squares, stochastic gradient learning

* *{Sergey.Slobodyan,Dmitri.Kolyuzhnov, Anna.Bogomolova}@cerge-ei.cz.*

† CERGE-EI is a joint workplace of the Center for Economic Research and Graduate Education, Charles University, and the Economics Institute of the Academy of Sciences of the Czech Republic.

1 Introduction

In this paper we perform an in–depth investigation of relative merits of two adaptive learning algorithms with constant gain, Recursive Least Squares (RLS) and Stochastic Gradient (SG). Properties of RLS as a learning algorithm are reasonably well understood, as it has been used extensively in the adaptive learning literature. For an extensive review, see Evans and Honkapohja (2001). SG learning received a more limited attention in the past, but the situation is changing: Evans, Honkapohja and Williams (2005) promote the constant gain SG, together with a generalized SG, as a robust learning rule which is well suited to the situation of time–varying parameters.

A different motivation to studying the properties of the SG learning comes from a recent interest in the heterogeneous learning, cf. Honkapohja and Mitra (2005) or Giannitsarou (2003). Several types of agents use different adaptive rules to learn the parameter values in the model. Often, some of the groups are using RLS while the others employ SG. A desirable property of such a model is its stability under all implemented types of learning.

Finally, our interest is focused on the properties of the learning algorithm which are not strictly local. It is known that E–stability of the rational expectations equilibrium (REE), which implies local stability under RLS with decreasing gain learning, does not automatically imply local stability under SG with decreasing gain, see Giannitsarou (2005). In contrast, we work in a situation when both RLS and SG are, indeed, locally stable, but the behavior of the constant gain versions of the two methods differ substantially away from the equilibrium.

As a testing ground for comparison we use the Phelps problem of a government controlling inflation while adaptively learning the approximate Phillips curve, studied previously by Sargent (1999) and Cho, Williams and Sargent (2002) (CWS in the sequel). A phenomenon known as "escape dynamics" can be observed in the model under the constant gain RLS learning. In Kolyuzhnov, Bogomolova and Slobodyan (2006) we applied a continuous–time version of the large deviations theory to study the escape dynamics, and argued that a simple approximation by a one–dimensional Brownian motion can be better suited for description of the escape dynamics in a large interval of values of the constant gain. Here we derive an even better one–dimensional approximation and discuss Lyapunov function–based approach to establishing limits of applicability of this approximation. We also extend our analysis to the SG constant gain learning.

The rest of the paper is organized as follows. We briefly describe the dynamic and static versions of the model of CWS in Section 2. Section 3 is devoted to describing the local stability (mean dynamics) results for the RLS and SG learning. In Section 4, we present and contrast the non–local effects arising under constant gain versions of these algorithms and discuss possible explanations for the difference in behavior of the mean dynamics and the actual real–time learning algorithm. Section 6 concludes.

2 The model

The economy consists of the government and the private sector. The government uses the monetary policy instrument x_n to control inflation rate π_n and attempts to minimize losses from inflation and unemployment U_n. It believes (in general, incorrectly) that an exploitable trade-off between π_n and U_n (the Phillips curve) exists. The true Phillips curve is subject to random shifts and contains this trade-off only for unexpected inflation shocks. The private sector possesses rational expectations $\widehat{x}_n = x_n$ about the inflation rate, and thus unexpected inflation shocks come only from monetary policy errors.

$$U_n = u - \chi\left(\pi_n - \widehat{x}_n\right) + \sigma_1 W_{1n},\ u > 0,\ \theta > 0, \tag{1a}$$

$$\pi_n = x_n + \sigma_2 W_{2n}, \tag{1b}$$

$$\widehat{x}_n = x_n, \tag{1c}$$

$$U_n = \gamma_1 \pi_n + \gamma_{-1}^T X_{n-1} + \eta_n. \tag{1d}$$

Vector $\gamma = \left(\gamma_1, \gamma_{-1}^T\right)^T$ represents government's beliefs about the Phillips curve. W_{1n} and W_{2n} are two uncorrelated Gaussian shocks with zero mean and unit variance. η_n is the Phillips curve shock as perceived by the government, believed to be a white noise uncorrelated with regressors π_n and X_{n-1}. In the "dynamic" version of the model, vector X_{n-1} contains two lags of inflation and unemployment rates and a constant,

$$X_{n-1} = \left(\ U_{n-1},\ \ U_{n-2},\ \ \pi_{n-1},\ \ \pi_{n-2},\ \ 1\ \right)^T, \tag{2}$$

while only the constant is present in X_{n-1} in the "static" version.

Given beliefs γ, the government solves

$$\min_{\{x_n\}_{n=0}^{\infty}} E\sum_{n=0}^{\infty}\beta^n\left(U_n^2+\pi_n^2\right), \tag{3}$$

subject to (1b) and (1d). This linear–quadratic problem produces a linear monetary policy rule

$$x_n = h(\gamma)^T X_{n-1}. \tag{4}$$

CWS identify three particular beliefs in model, which replicate different equilibria of the correctly specified version (government believes in [1a]) of the model. Under Belief 1, $\gamma = \left(-\chi, 0, 0, 0, 0, u(1+\chi^2)\right)^T$, policy function is $x_n = \chi u$. This is the Nash, or discretionary equilibrium of Sargent (1999). Beliefs 2 of the form $\gamma = (0,0,0,0,0,u^*)^T$ lead to $x_n = 0$ and zero average inflation for any u^*: Ramsey, or the optimal time–inconsistent equilibrium of Kydland and Prescott (1977). Finally, Beliefs 3 where $\gamma_1+\gamma_4+\gamma_5 = 0$ asymptotically lead to $x_n = 0$: this is an "induction hypothesis" belief, see Sargent (1999). In the misspecified model where the government believes in (1d), the equilibrium is defined as a vector of beliefs at which the government's assumptions about orthogonality of η_n to the space of regressors are consistent with observations:

$$E\left[\eta_n \cdot (\pi_n, X_{n-1})^T\right] = 0. \tag{5}$$

CWS call this point a *self–confirming equilibrium*, or SCE. Williams (2001) shows that the only SCE in the model is Belief 1.

3 RLS and SG

In a period n, the government uses its current vector of beliefs γ_n to solve (3), assuming the beliefs will never change. The generated monetary policy action x_n is correctly anticipated by the private sector and produces U_n according to (1a). Then the government adjusts its beliefs about the Phillips curve coefficients γ_n and, if RLS learning is used, their second moments matrix R_n in an adaptive learning step. Define $\xi_n = \begin{bmatrix} W_{1n} & W_{2n} & X_{n-1}^T \end{bmatrix}^T$, $g(\gamma_n, \xi_n) = \eta_n \cdot \left(\pi_n, X_{n-1}^T\right)^T$, and $M_n(\gamma_n, \xi_n) =$

$\left(\pi_n, X_{n-1}^T\right)^T \cdot \left(\pi_n, X_{n-1}^T\right)$. Next period's beliefs γ_{n+1} and R_{n+1} are given by

$$\gamma_{n+1} = \gamma_n + \epsilon R_n^{-1} g(\gamma_n, \xi_n), \tag{6a}$$

$$R_{n+1} = R_n + \epsilon \left(M_n(\gamma_n, \xi_n) - R_n\right), \tag{6b}$$

under RLS learning and by

$$\gamma_{n+1} = \gamma_n + \epsilon g(\gamma_n, \xi_n) \tag{7}$$

under the SG learning.

The evolution of the state vector ξ_n can be written as

$$\xi_{n+1} = A(\gamma_n)\xi_n + BW_{n+1}, \tag{8}$$

where $W_{n+1} = \left[\begin{array}{cc} W_{1n+1} & W_{2n+1} \end{array}\right]^T$, for some matrices $A(\gamma_n)$ and B. The parameter vector $\theta_n^{\epsilon,SG}$ is given as γ_n for SG and as

$$\theta_n^{\epsilon,RLS} = \left[\begin{array}{cc} \gamma_n^T, & vech^T(R_n) \end{array}\right]^T \tag{9}$$

for the RLS case. Finally, define $H^{SG}(\theta_n^\epsilon, \xi_n) = g(\gamma_n, \xi_n)$ and

$$H^{RLS}(\theta_n^\epsilon, \xi_n) = \left[\begin{array}{cc} \left(R_n^{-1} g(\gamma_n, \xi_n)\right)^T, & vech^T\left(M_n(\gamma_n, \xi_n) - R_n\right) \end{array}\right]^T \tag{10}$$

to obtain the Stochastic Recursive Algorithm (SRA) in the standard form:

$$\theta_{n+1}^{\epsilon,j} = \theta_n^{\epsilon,j} + \epsilon H^j(\theta_n^\epsilon, \xi_n),\ j = \{RLS, SG\}, \tag{11a}$$

$$\xi_{n+1} = A(\gamma_n)\xi_n + BW_{n+1}. \tag{11b}$$

Finally, the approximating ordinary differential equations corresponding to the above SRA are given by

$$\dot{\theta^j} = E[H^j(\gamma, \xi_n)]. \tag{12}$$

The SCE, which consists of the vector $\overline{\gamma}$ that forms Belief 1, and corresponding 2^{nd} moments matrix $\overline{R}$ if RLS is used, is the only equilibrium of the above ODE. The SCE is stable for both RLS and SG in the dynamic and static versions of the model. Under some assumptions, this means that the continuous–time process $\theta_t^{\varepsilon,j}$ defined as $\theta_t^{\varepsilon,j} = \theta_n^{\epsilon,j}$ for $t \in [n\varepsilon, (n+1)\varepsilon)$ converges weakly (in distribution) to $\theta^j(t,a)$, solution of the ODE (12) with the initial condition $a = \theta(0)$ which is also a starting point of the process $\theta_t^{\varepsilon,j}$. This solution is also called the "mean dynamics trajectory" of the SRA (11), with the right–hand side of (11) being the "mean dynamics". A variant of the mean dynamics approximation is the following difference equation obtained from (11a):

$$\theta_{n+1}^{\epsilon,j} = \theta_n^{\epsilon,j} + \epsilon \cdot E[H^j(\gamma, \xi_n)]. \tag{13}$$

The difference between the above approximation and (12) is that ϵ is not assumed to be approaching zero asymptotically. For details and derivations, see Evans and Honkapohja (2001).

When the gain is constant, the convergence of θ_n^{ϵ} to the mean dynamics trajectory $\theta(t)$ is only in distribution. Evans and Honkapohja (2001, Prop. 7.8) show that as $\varepsilon \to 0$ the process $U_t^{\varepsilon,j} = \frac{\theta_t^{\varepsilon,j} - \theta^j(t,a)}{\sqrt{\varepsilon}}$ converges (weakly) to a following diffusion:

$$dU_t^{\epsilon,j} = D_\theta p(\theta^j(t,a)) U_t^{\epsilon,j} dt + \Sigma^{1/2}\left(\theta^j(t,a)\right) dW_t, \tag{14}$$

where W_t is a multi–dimensional Brownian process with dimensionality equal to that of θ. $p(\theta)$ is the mean dynamics vector, and Σ the matrix whose elements are covariances of different components of the mean dynamics vector, both with respect to the unique invariant probability distribution $\Gamma_\theta(dy)$ of the state vector X:[1]

$$p(\theta^j) = \int H^j(\theta^j, y)\Gamma_\theta(dy), \tag{15}$$

$$\Sigma_{ik} = \sum_{k=-\infty}^{\infty} Cov\left[H_i^j(\theta^j, X_k(\theta^j)), H_k^j(\theta^j, X_0(\theta^j))\right]. \tag{16}$$

This result is used to get continuous–time approximation of SRA, given any initial condition:

$$d\theta_t^{\varepsilon,j} = D_\theta p(\theta^j(t,a))\left[\theta_t^{\varepsilon,j} - \theta^j(t,a)\right] dt + \sqrt{\varepsilon}\Sigma^{1/2}\left(\theta^j(t,a)\right) dW_t. \tag{17}$$

For the RLS case Williams (2001, Theorem 3.2) shows that the above results can be used to derive a local continuous–time approximation of the SRA around the SCE $\overline{\theta}$:

$$d\varphi_t^{RLS} = D_\theta p(\overline{\theta}^{RLS})\varphi_t dt + \sqrt{\epsilon}\Sigma^{1/2}(\overline{\theta}^{RLS}) dW_t, \tag{18}$$

where $\varphi_t = \theta_t^{RLS} - \overline{\theta}^{RLS}$ are deviations from the SCE. Similar result is easily obtained in the SG case.

[1] State vector ξ_n has a unique invariant probability distribution: it contains stationary Gaussian random variables W_{1n} and W_{2n}, a constant, and a stationary 4–dimensional (in the dynamics model) or 1–dimensional (in the static model) AR(1) variable. This distribution can be calculated explicitly.

4 Behavior of Simulations

The discussion below refers to the model as parametrized in CWS: $\sigma_1 = \sigma_2 = 0.3$, $u = 5$, $\chi = 1$, $\beta = 0.98$.

4.1 Recursive Least Squares

It is well known that under the constant gain RLS learning beliefs in the Phelps problem can exhibit "escapes": after a number of periods spent in the neighborhood of the SCE, the beliefs vector γ suddenly deviates from the SCE towards "induction hypothesis" plane $\gamma_1 + \gamma_4 + \gamma_5 = 0$ ($\gamma_1 = 0$ axis for the static model), see CWS, in particular Figs. 6 and 7. During such an escape, the inflation rate falls from its Nash equilibrium value equal to u and approaches 0, see Fig. 1 in CWS.

In Kolyuzhnov et al. (2006), we have studied these escapes extensively and described the following sequence of events. If the constant gain parameter ϵ is not too small, the behavior of equation (6a) is almost one–dimensional. It is well known that in this model the region of attraction of the SCE is very small, see Fig. 1 reprinted from Kolyuzhnov et al. (2006). Outside of the immediate neighborhood of the SCE, the mean dynamics trajectories point away from it and towards the "induction hypothesis" plane. These trajectories linger in the neighborhood of the plane for a relatively long time and then start a slow return to the SCE. The largest eigenvalue of $\overline{R}^{-1}$ is λ_1 =3083.8 and the next largest λ_2 =29.1, less that 1% of λ_1. The projection of $g(\gamma_n, \xi_n)$ onto v_1, the dominant eigenvector of $\overline{R}^{-1}$, is magnified 100 times as strongly as the projection onto the second eigenvector. As a result, simulation runs with escapes tend to contain a set of points aligned along the dominant eigenvector of $\overline{R}^{-1}$ all the way towards the "induction hypothesis" plane, which is clearly demonstrated in the Figure 2 reprinted from Kolyuzhnov et al. (2006).

In the Figure, 6–dimensional vector of beliefs γ is presented in the space of $(\widetilde{\gamma}_1, \widetilde{\gamma}_2)$, defined as $\gamma_1 + \gamma_4 + \gamma_5$ and $u \cdot (\gamma_2 + \gamma_3) + \gamma_6$. For a discussion on this choice of variables, see that paper. The significant disbalance of eigenvalues of $\overline{R}^{-1}$ is inherited by the matrix Σ, and the eigenvector v_1 is essentially collinear to first 6 components of $\widetilde{v}_1$, the dominant eigenvector of Σ.

We use this essential one–dimensionality to derive the following approximation of (18). Write $\varphi_t \approx x_t \cdot \widetilde{v}_1$, and multiply (18) by $\widetilde{v}_1^T$ from the left. The resulting

1–dimensional approximation is then given by

$$\begin{aligned} dx_t &\approx \widetilde{v}_1^T D_\theta p(\bar{\theta}^{RLS})\widetilde{v}_1 \cdot x_t dt + \sqrt{\epsilon\widetilde{\lambda}_1} \cdot \widetilde{v}_1^T \cdot dW_t, \\ dx_t &\approx A \cdot x_t dt + \sqrt{\epsilon\widetilde{\lambda}_1} dW_t. \end{aligned} \tag{19}$$

Note that dW_t in the second line is now a one–dimensional standard Brownian motion. (19) is an Ornstein–Uhlenbeck process with well–known properties. In particular, one could easily derive the expected time until the process leaves any interval of the real line, see Borodin and Salminen (1996). A equals -0.41 and $\widetilde{\lambda}_1$ =277.58.[2]

To estimate the region of applicability of the approximation (19), take x_t^2 as the Lyapunov function and calculate LV for one–dimensional diffusion (19):[3]

$$LV = 2 \cdot \left(Ax_t^2 + \epsilon\widetilde{\lambda}_1 \right).$$

Clearly, LV is positive for small x, and thus $V(x_t) = x_t^2$ is expected to *increase.* In other words, in a small neighborhood of the SCE the Stochastic Recursive Algorithm (11) is expected to be locally divergent on the average. We would call values of ϵ "small" if for x_t corresponding to the boundary of the SCE's stability region under the mean dynamics, the value of LV is negative: once the SRA approaches this boundary, it is expected to turn back towards the SCE. If such behavior is observed, one expects the invariant distribution derived along the lines of Evans and Honkapohja (2001, Ch. 14.4) to be valid, and other methods of describing escape dynamics are needed, such as the Large Deviations Theory, see CWS and Kolyuzhnov et al. (2006). For "not small" ϵ, the approximation (19) could be used to derive expected escape time. In the dynamic model, values of ϵ below $2 \cdot 10^{-5}$ are "small".

Dynamics of the static model under the constant gain RLS learning is qualitatively similar: a move out of the immediate region of attraction of the SCE, followed

[2] Ornstein–Uhlenbeck approximation could also be useful in case one is interested in selecting the value of ϵ such that for a given time period the probability of observing an escape is below some given threshold (dynamics under learning is empirically stable).

[3] The operator L defined for a function V has the following meaning: Under certain conditions, the expected value of $V(t, X(t)) - V(s, X(s))$ is given as an integral from s to t over LV, see Khasminskii (1980, Ch. 3). In some sense, in stochastic differential equations LV plays the role of time derivative of the Lyapunov function $\frac{dV}{dt}$ for the deterministic system.

by a long trek to the Ramsey equilibrium outcome with zero average inflation. The dynamics is essentially one–dimensional. However, the radius of the region of attraction is slightly larger in the dominant direction than in the dynamic model, and the diffusion is less powerful: A =-0.52 and $\widetilde{\lambda}_1$ equals 26.09. The combined effect of the stronger drift, weaker diffusion, and larger stability region is obvious: a significantly larger expected number of periods until the simulations escape neighborhood of the SCE. Table 1 compares empirically observed average time needed to escape with the theoretically predicted values for different choices of the constant gain parameter ϵ. For relatively large values of $\epsilon \geq 4 \cdot 10^{-4}$, the agreement is rather good, especially for the static model. The theory starts to overpredict for smaller ϵ; again, this effect is more pronounced in the static model, because "smallness" of ϵ starts earlier: Gains below $3 \cdot 10^{-4}$ are "small" in the static model.

TABLE 1. A comparison of the theoretically derived values of expected escape time and empirically observed average escape times

	Dynamic model		**Static model**	
ϵ	Simulations	Theory	Simulations	Theory
$2 \cdot 10^{-5}$	$1.10 \cdot 10^5$	$1.86 \cdot 10^5$		
$3 \cdot 10^{-5}$	$5.10 \cdot 10^4$	$7.21 \cdot 10^4$	$4.40 \cdot 10^7$	$9.40 \cdot 10^8$
$5 \cdot 10^{-5}$	$1.88 \cdot 10^4$	$2.34 \cdot 10^4$	$1.93 \cdot 10^6$	$9.90 \cdot 10^6$
$1 \cdot 10^{-4}$	$4.84 \cdot 10^3$	$5.43 \cdot 10^3$	$1.50 \cdot 10^5$	$2.75 \cdot 10^5$
$2 \cdot 10^{-4}$	$1.26 \cdot 10^3$	$1.31 \cdot 10^3$	$2.38 \cdot 10^4$	$2.97 \cdot 10^4$
$4 \cdot 10^{-4}$	336.96	321.5	$5.06 \cdot 10^3$	$5.26 \cdot 10^3$
$1 \cdot 10^{-3}$	64.59	50.9	733.57	701.5
$2 \cdot 10^{-3}$	21.49	12.68	189.98	165.7
$3 \cdot 10^{-3}$	12.50	5.63	87.00	72.27
$4 \cdot 10^{-3}$	8.77	3.16	52.08	40.28
$5 \cdot 10^{-3}$	6.79	2.02	34.39	25.64
$6 \cdot 10^{-3}$	5.99	1.40	24.76	17.74
$7 \cdot 10^{-3}$	4.98	1.03	19.14	13.00
$8 \cdot 10^{-3}$	4.49	0.79	15.02	9.93
$9 \cdot 10^{-3}$	4.12	0.62	13.32	7.84
$1 \cdot 10^{-2}$	3.70	0.51	11.16	6.34

4.1.1 What is the right ϵ and the time scale?

How should one approach the problem of choosing ϵ? Putting aside any considerations related to the stability of learning in a particular model, two rules of thumb for selecting ϵ seem sensible. The first is based on the fact that constant gain adaptive learning is well suited to situations with time–varying parameters or structural breaks. In this case, $1/\epsilon$ should be related to the typical time which is needed to observe a break, or for the time variation to become "significant". Alternatively, one could imagine that the initial value of parameters is obtained through some method of statistical estimation such as OLS. In this case, it is natural to assign to every point in the initial estimation a weight equal to $1/N$. If there is no reason to believe that subsequent points are in some sense superior to those used to derive an initial estimate, the constant gain ϵ should be comparable to $1/N$. Given the nature of the Phelps problem where inflation might be available on the monthly basis but the output gap could be obtained only at the quarterly basis, values of ϵ not much larger or smaller than 0.01 seem empirically justified.

Notice that a period in the Phelps model could not be shorter than a quarter (or a month). At economically relevant time scale (at most a hundred years), there are no escapes for $\epsilon < 1 \cdot 10^{-4}$ in the dynamic model and $\epsilon < 4 \cdot 10^{-4}$ in the static one. An important caveat to this statement is that both the theoretical and simulation results are obtained by imposing the SCE as the starting point. In other words, one starts from a situation of a learning which is completed in the sense that the government and the private sector are playing Nash equilibrium, and is interested in the expected time until the economy "unlearns" Nash equilibrium given a particular constant gain learning rule. If, instead of the SCE, initial beliefs are given by a point which is closer to the stability region's boundary, one would expect smaller escape times.

4.2 Stochastic Gradient Learning

4.2.1 Dynamic Model

The behavior of the dynamic Phelps model under the SG learning is dramatically different. In the approximation (13), the matrix

$$F(\epsilon) = I + \epsilon D_\theta p(\bar{\theta}^{SG})$$

is stable but only just: for $\epsilon = 0.01$, its eigenvalues range from 0.2447 to 0.9988 to 0.99999862. Five out of six eigenvalues are almost unitary. Under the mean dynamics (13), any deviation from the SCE results in a fast movement along x_1, the eigenvector which corresponds to 0.2447 eigenvalue, and then an extremely slow convergence back to the SCE along the remaining five directions, see Figure 3. On the other hand, simulations behave very differently. Figure 4 plots a norm of deviations from the SCE: there is a clearly distinguishable movement away from the SCE which seems almost deterministic.[4] Figure 5 plots values of γ_6 and $\widetilde{\gamma}_2$, which both exhibit a clear divergence. For this value of ϵ, inflation rate will drop below 4 (its mean equals 5 at the SCE) in a couple hundred iterations, which is definitely the time scale one should be concerned with. How could one explain the discrepancy between the mean dynamics (13) and the simulations?

Fig. 6 plots a projection of $\frac{\gamma_n - \overline{\gamma}}{\|\gamma_n - \overline{\gamma}\|}$ onto the subspace spanned by five eigenvectors of $\digamma(\epsilon)$ which correspond to the almost unitary eigenvalues for a typical simulation run with $\epsilon = 0.01$. Within the first hundred simulation periods, this projection becomes very close to unity: average value for the first ten (hundred) periods is 0.69 (0.80). Thus, simulation run very fast approaches some neighborhood of the subspace and does not leave it for any extended period of time. This behavior is natural: any initial deviation along x_1 will shrink to 0.2447^3 ~1.5% of its initial size in just 3 steps. On the other hand, deviations along five other eigenvectors will take at least $\frac{\ln(0.5)}{\ln(0.9988)}$ ~577 periods to reach 50% of their initial magnitude.

Another feature of the matrix $\digamma(\epsilon)$ which helps to explain the behavior of simulations is the presence of directions along which deviations are expected to *increase* before declining. Such directions exist because the matrix $\digamma(\epsilon)+\digamma(\epsilon)^T$ is not stable. In this case, one could find a unit vector w such that $w^T \cdot \digamma(\epsilon) \cdot w$ is greater than one. A deviation in the direction w is thus expected to *increase* its projection onto w and thus to *increase* its norm, at least initially. We define the vector w as the unit vector which maximizes $w^T \cdot \digamma(\epsilon) \cdot w$ at 1.103 (this value equals 1.01 at $\epsilon = 0.001$ and 1.001 at $\epsilon = 1 \cdot 10^{-4}$). A projection of $\frac{\gamma_n - \overline{\gamma}}{\|\gamma_n - \overline{\gamma}\|}$ onto w is plotted in the Figure 7 (only the absolute value of the projection matters, not its sign). It becomes large very fast, in about one hundred simulation periods or less. A system (13) is expected to demonstrate a locally divergent behavior whenever this projection is large. To support further the crucial importance of the projection onto w, Figure 8 presents

[4] If we observe the simulations for larger number of periods, the belief vector γ eventually reaches values at which the state vector process loses stationarity, and the simulation breaks down.

the norm of deviation from the SCE for the mean dynamics trajectory which started from a point γ that lies in the direction w. There is a steep initial increase in the norm, followed by a long decline which is still far from complete after 2000 periods. To overcome the initial increase and return the system to the norm of deviation equal to its initial value, 150 periods are needed.

Notice that the norm of the projection of w onto the subspace spanned by the five eigenvectors is rather large and equals 0.95. When the dynamics of (11) is restricted almost exclusively to this subspace, mean dynamics plays almost no role in the short run. Random disturbances are then very likely to produce value of $\gamma_n - \overline{\gamma}$ which has a significant projection onto w during the 150 periods which are needed to eliminate the effect of the previous shock in this direction. Once such shock happens, the projection is not likely to disappear given a very weak stabilizing force of the mean dynamics on the subspace.

As a final piece of evidence connecting the vector w with the divergent behavior of simulations, consider Figure 9. In the periods when the projection of $\frac{\gamma_n - \overline{\gamma}}{\|\gamma_n - \overline{\gamma}\|}$ onto w (crosses) is particularly large, the distance between the beliefs γ_n and the SCE $\overline{\gamma}$ (solid line) grows the fastest; a relative decline in the projection is correlated with a temporary stop or even reversal of the divergent behavior.

Summarizing the discussion, we could say that a clear instability observed in the behavior of the SRA for SG learning in the dynamic Phelps problem is caused by a particular geometric structure. The subspace spanned by the almost unitary eigenvalues' eigenvectors of the mean dynamics map is almost parallel to the direction along which the mean dynamics is expanding in the short run rather than contracting. Given that any random deviation away from the subspace is likely to be very short–lived, and that contracting mean dynamics within the subspace is very weak, random vectors with a relatively large projection onto the expansive direction are likely to appear. Once such a projection appears, it is unlikely to be averaged away by the mean dynamics.

We checked the behavior of the algorithm for other values of ϵ. Qualitatively, the picture does not change: there is still an apparent divergence of the vector of government's beliefs γ_n away from the SCE. One could still observe a very fast convergence towards the subspace spanned by the five almost unitary eigenvalues' eigenvectors and a significant projection onto the expanding direction w. The direction w remains almost parallel to the subspace. Only for very small values of $\epsilon \leq 8 \cdot 10^{-6}$ we start observing a different behavior; the system (11) does not systematically diverge and fluctuates in some neighborhood of the SCE.

4.2.2 Static Model

Taking into account that under the RLS learning the static model was much more stable (it took much longer for the escape to the "induction hypothesis" plane to happen), we expect this feature to be preserved under the SG learning as well. This is what is indeed observed. Clearly unstable behavior is observed only for relatively large values of ϵ above $3 \cdot 10^{-2}$. This instability could take two forms: either a convergence to a quasi–stable stochastic steady state where $\|\gamma - \overline{\gamma}\|$ is about 3 for ϵ between approximately $6.5 \cdot 10^{-2}$ and $7.9 \cdot 10^{-2}$ (above $\epsilon \sim 7.9 \cdot 10^{-2}$, the mean dynamics map $\mathcal{F}(\epsilon)$ has a real eigenvalue which is less than -1, and the SCE is thus unstable under the mean dynamics), or a divergence of simulations from the SCE for $3.5 \cdot 10^{-2} \lesssim \epsilon \lesssim 6.5 \cdot 10^{-2}$. When ϵ equals $3.5 \cdot 10^{-2}$ or less, empirically relevant time scales are characterized by what seems to be a stable dynamics. The speed of divergence significantly depends on the value of ϵ: while at $\epsilon = 5 \cdot 10^{-2}$ less than 100 iterations are typically needed to observe a deviation from the SCE such that $\|\gamma - \overline{\gamma}\| \geq 0.1$, such large excursions are not likely to be observed before 500^{th} iteration for $\epsilon = 4 \cdot 10^{-2}$. As in the dynamic model, the eventual outcome of divergent simulations is the value of γ which leads to at least one eigenvalue of the matrix $A(\gamma)$ in (11b) being outside of the unit circle and thus to non–stationary state process.

Applying the reasoning demonstrated above to the dynamics of the static model under SG learning in real time, we could say the following. The map $\mathcal{F}(\epsilon)$ has two eigenvalues. One is always close to one (0.9999 for $\epsilon = 3 \cdot 10^{-2}$). Another is a linearly decreasing function of ϵ. It equals -1 when $\epsilon \sim 7.9 \cdot 10^{-2}$ and approaches 1 as $\epsilon \to 0$. It is still true that the divergent behavior is related to the movement along the almost unitary eigenvalue's eigenvector: projection of w onto this eigenvector equals 0.9988, and the fastest divergence of beliefs from their SCE values occurs when $\gamma - \overline{\gamma}$ is in the closest alignment with w ($w^T \cdot \frac{\gamma_n - \overline{\gamma}}{\|\gamma_n - \overline{\gamma}\|}$ is close to one). There are two crucial differences with the dynamics model, however: first, the direction w is very weakly expansive, as $w^T \cdot \mathcal{F}(\epsilon) \cdot w$ equals only 1.0018 when $\epsilon \sim 3 \cdot 10^{-2}$ and becomes even smaller as ϵ decreases. At the same time, the dominant eigenvalue of $\mathcal{F}(\epsilon)$ equals 0.23 for $\epsilon \sim 3 \cdot 10^{-2}$ and is decreasing in ϵ. Thus, for smaller values of ϵ the dynamics of (13) loses its essentially one–dimensional nature, and the expansive movement in the direction w is not too strong (compare 1.0018 to the 1.103 reported for the dynamic model). Instead of 150 periods needed to to start reversing a deviation in the direction of w which we reported for the dynamic model at $\epsilon = 0.01$, only 3-4 iterations are needed to achieve the same result in the static model at similar values of ϵ. It is not a big surprise, then, that the static model under the SG learning stops diverging at much larger values of the constant gain.

5 Conclusion

We compared the performance of two methods of adaptive learning with constant gain, Recursive Least Squares and Stochastic Gradient learning, in a Phelps model of a monetary policy which has been extensively studied previously. For the values of ϵ which might be justified for the problem, it is a well–known fact that the RLS adaptive learning could lead to "escapes": large deviations of the government's beliefs about the Phillips curve from the Self–Confirming Equilibrium where inflation level is set at high levels towards the beliefs which lead the policymaker to set inflation close to zero. We approximated the discrete–time Stochastic Recursive Algorithm which describes RLS constant gain learning by a one–dimensional continuous–time Ornstein–Uhlenbeck process and derived expected escape times out of a small neighborhood of the SCE. The theoretical prediction works rather well when compared with the simulation results.

Turning our attention to the SG learning, we showed that the dynamics behaves in a divergent way for a large interval of values of ϵ. The divergence is especially pronounced when SG learning is used in the dynamic version of the Phelps problem. This behavior is caused by existence of eigenvalues of the SRA which are very close to the unit circle, and thus deviations in the direction of corresponding eigenvectors contract very slowly. Moreover, the mean dynamics of the SRA has directions which are expected to expand in the short run rather than contract, and these directions are almost parallel to the subspace spanned by the slowly contracting eigenvectors. Such combination leads to a divergent behavior of the SRA, which is reversed only for the very small ϵ values when the expansion rate reduces to very small values. Behavior of the static model exhibits similar features, with a crucial difference of the expansion rate: for the empirically relevant values of ϵ, it is less than 1.02 instead of 1.1 as in the dynamic model. This difference means that the SRA stops exhibiting divergent behavior for much larger values of the constant gain parameter in the static than in the dynamic model.

Comparing the two variants of the model under two types of constant gain adaptive learning, we could say that only SG learning in the static model demonstrates an absence of large excursions of beliefs from the SCE at empirically relevant time scale and for constant gain values likely to be used in practice ("stability"). Following Evans et al. (2005), one could thus endorse using this adaptive learning method for the static model. The overall result, however, cannot be judged as very good, as three out of four modifications produce an "unstable" result.

A very unbalanced nature (large differences between the dominant eigenvalue and the rest) of the second moments matrix $\overline{R}$ plays a significant role in the results, making the stochastic dynamics strongly one–dimensional in the RLS case and leading to almost unitary eigenvalues in the SG case. Whether this feature is caused by the fact that the government uses a misspecified model in the Phelps problem warrants further investigation.

The behavior of the SRA under SG learning in real time leads us to express a warning. Checking that all the eigenvalues of the mean dynamics map are stable is not enough to guarantee "stable" behavior of the constant gain learning algorithm in real time; moreover, checking that the mean dynamics trajectories are stable in a large region is not enough either. If many eigenvalues of the mean dynamics map for a constant gain learning algorithm are close to the unit circle, Stochastic Recursive Algorithm might exhibit *divergent* behavior despite *convergent* mean dynamics.

References

Borodin, A. N. and Salminen, P.: 1996, *Handbook of Brownian Motion — Facts and Formulae*, Birkhauser Verlag, Basel Boston Berlin.

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

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

Evans, G. W., Honkapohja, S. and Williams, N.: 2005, Generalized stochastic gradient learning. NBER Technical Working Paper 317.

Giannitsarou, C.: 2003, Heterogeneous learning, *Review of Economic Dynamics* **6**, 885–906.

Giannitsarou, C.: 2005, E–stability does not imply learnability, *Macroeconomic Dynamics* **9**, 276–287.

Honkapohja, S. and Mitra, K.: 2005, Learning stability in economies with heterogeneous agents. mimeo.

Khasminskii, R. Z.: 1980, *Stochastic Stability of Differential Equations*, Sijthoff & Noordhoof, Alphen aan den Rijn, The Netherlands.

Kolyuzhnov, D., Bogomolova, A. and Slobodyan, S.: 2006, Escape dynamics: A continuous–time approximation. CERGE–EI Working Paper 285.

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

Williams, N.: 2001, *Escape Dynamics in Learning Models*, PhD thesis, University of Chicago.

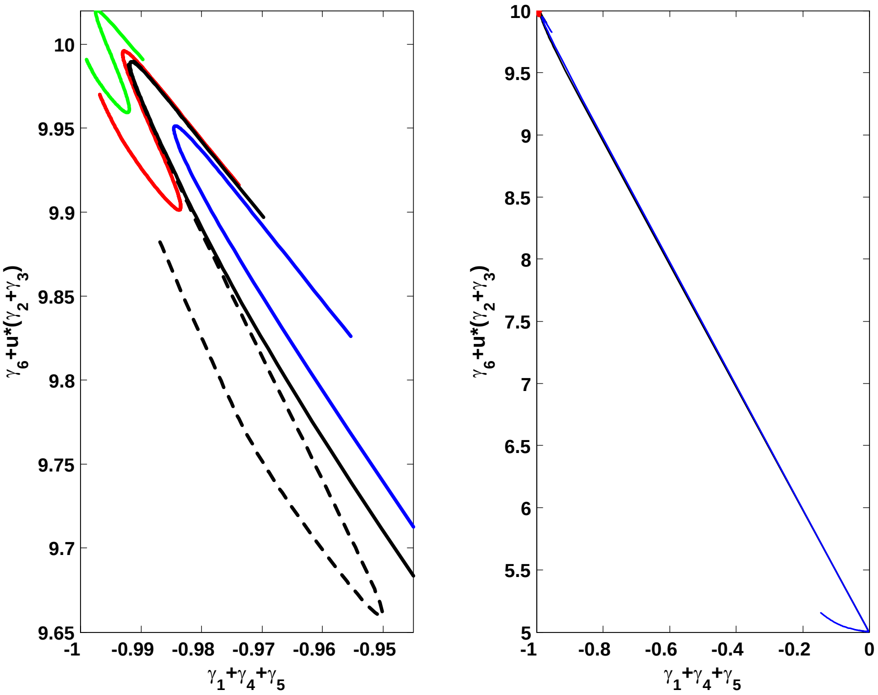

Figure 1: The mean dynamics trajectories under RLS

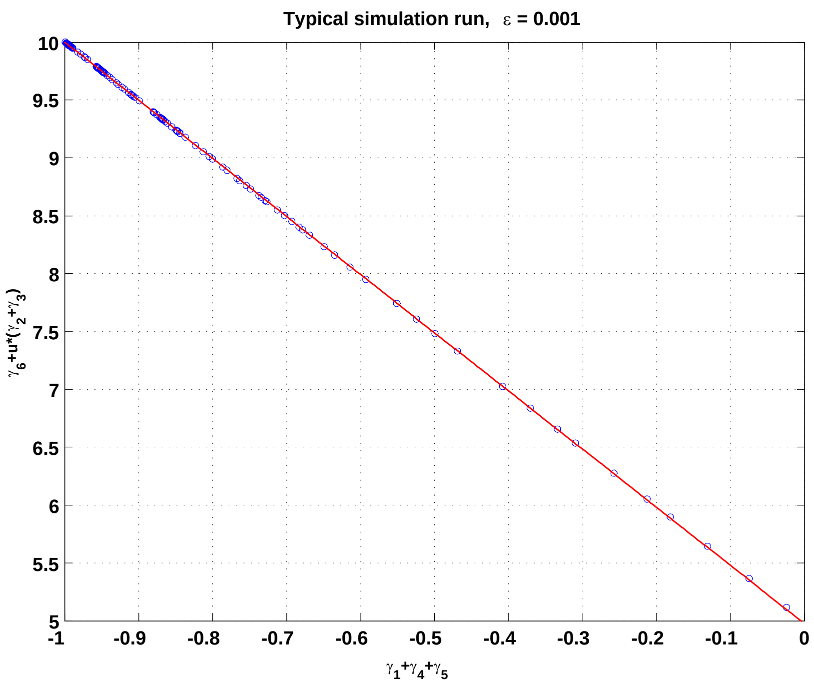

Figure 2: Typical simulation run and the “largest” eigenvector of R^{-1}.

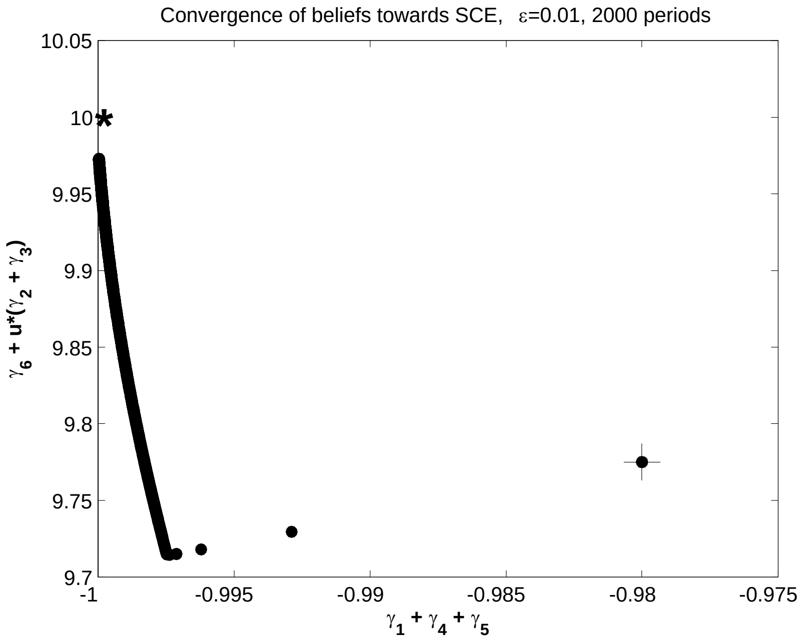

Figure 3: '+' sign: start of the simulation. '*' - the SCE location

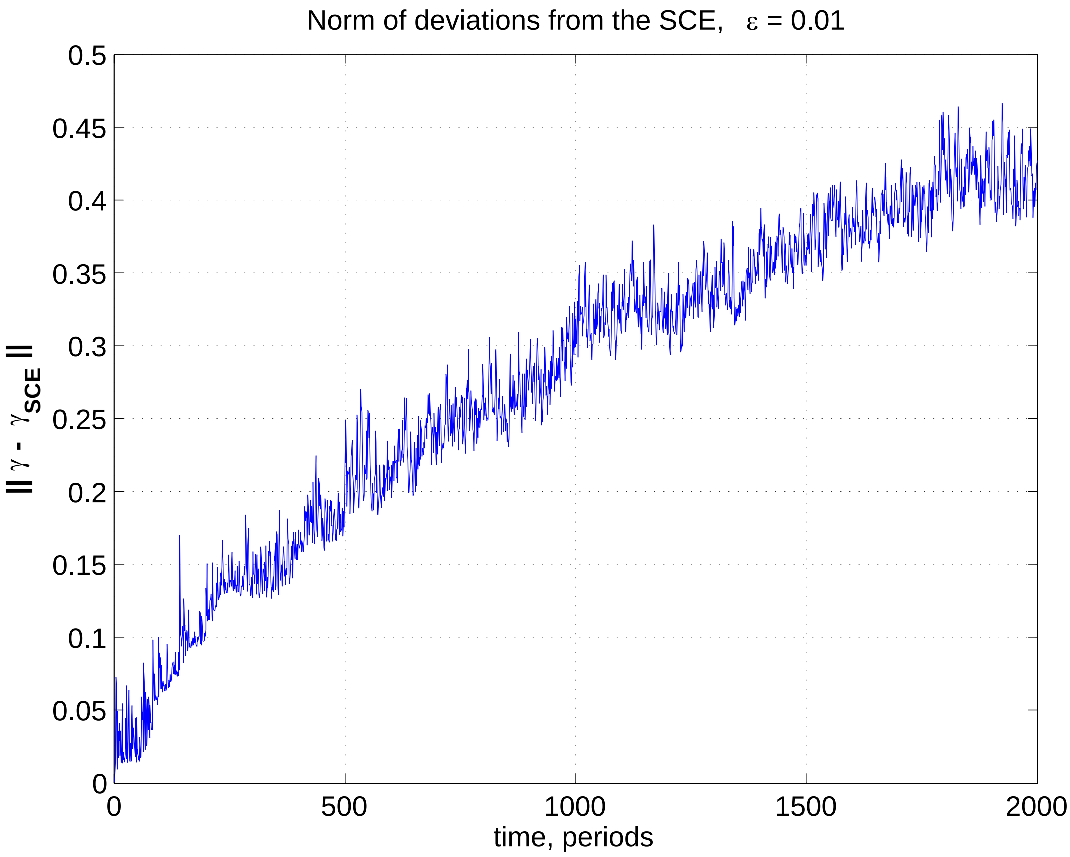

Norm of deviations from the SCE, ε = 0.01
0.5
0.45
0.4
0.35
0.3
0.25
0.2
0.15
0.1
0.05
0
|| γ - γ_SCE ||
0
500
1000
1500
2000
time, periods

Figure 4:

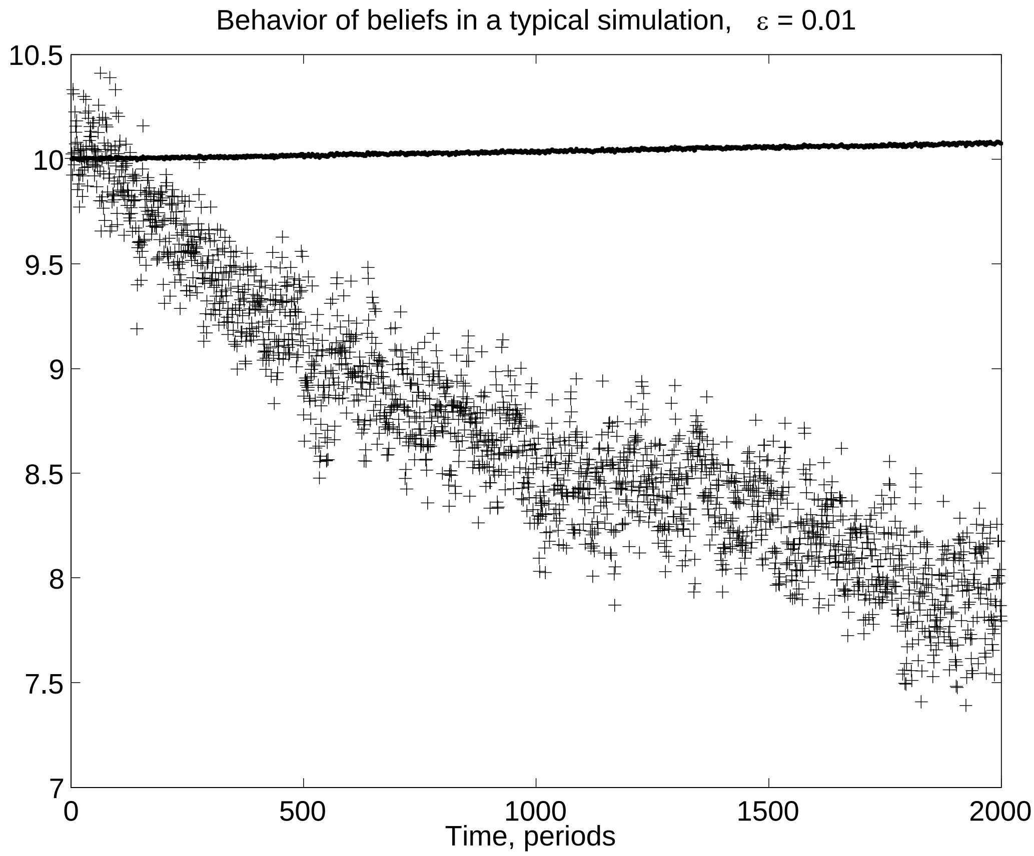

Figure 5: Behavior of beliefs during a typical simulation run. Dots: γ_6, crosses: $\widetilde{\gamma}_2$.

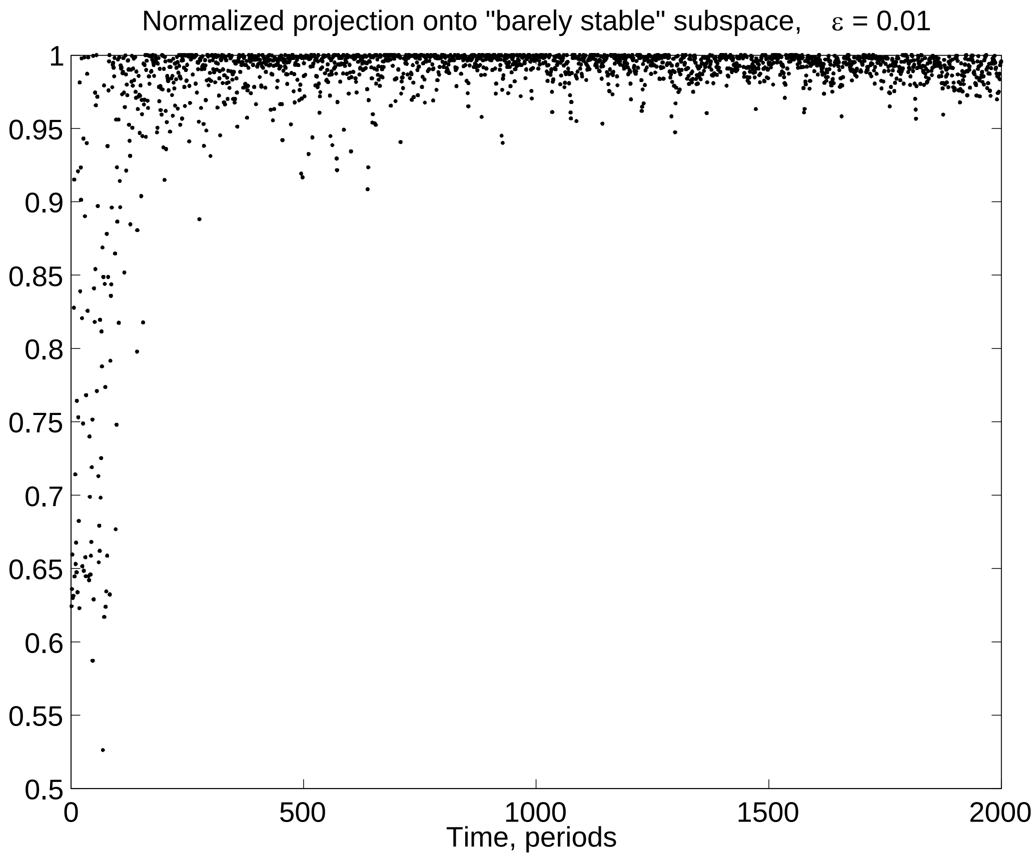
Normalized projection onto "barely stable" subspace, ε = 0.01
1
0.95
0.9
0.85
0.8
0.75
0.7
0.65
0.6
0.55
0.5
0
500
1000
1500
2000
Time, periods

Figure 6:

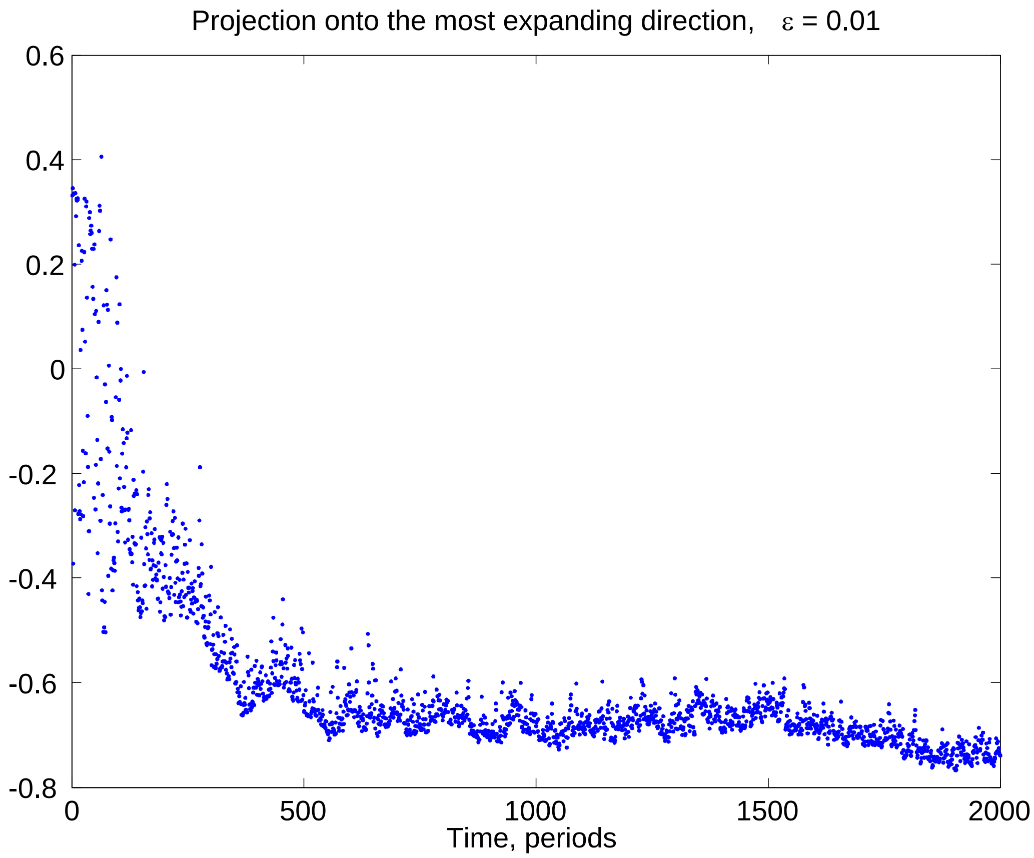
Projection onto the most expanding direction, ε = 0.01
0.6
0.4
0.2
0
-0.2
-0.4
-0.6
-0.8
0
500
1000
1500
2000
Time, periods

Figure 7:

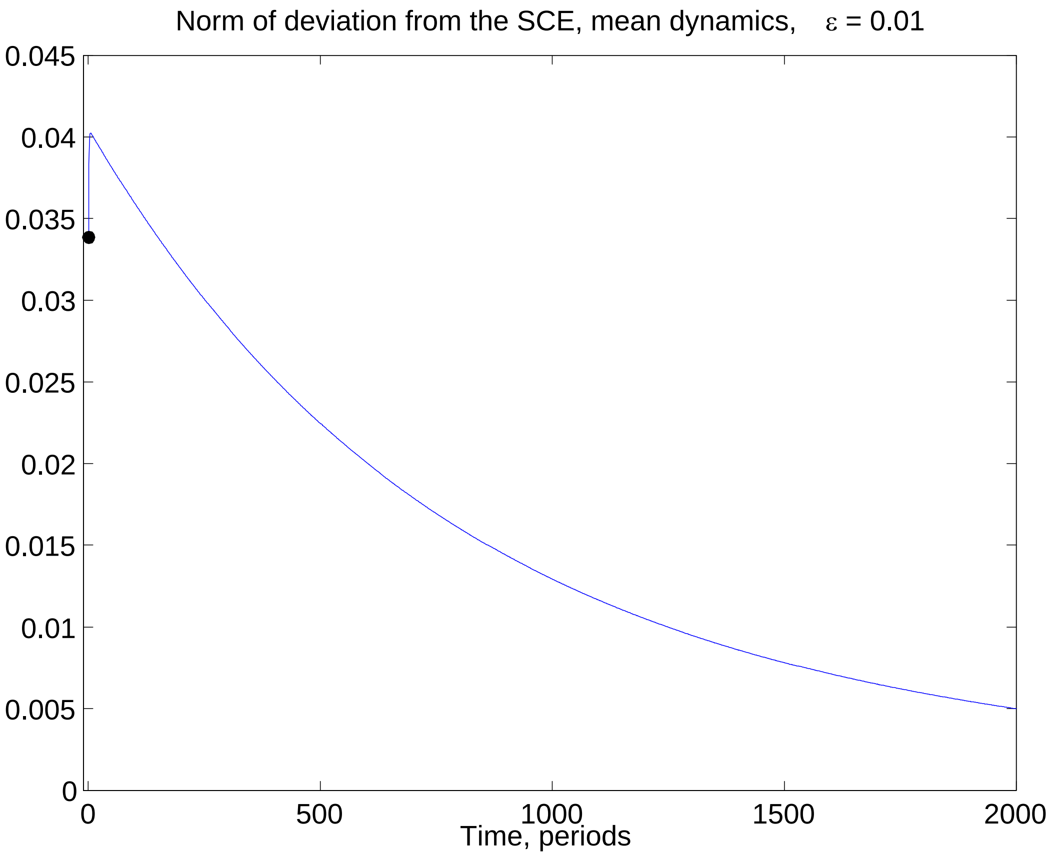
Norm of deviation from the SCE, mean dynamics, ε = 0.01
0.045
0.04
0.035
0.03
0.025
0.02
0.015
0.01
0.005
0
0
500
1000
1500
2000
Time, periods

Figure 8:

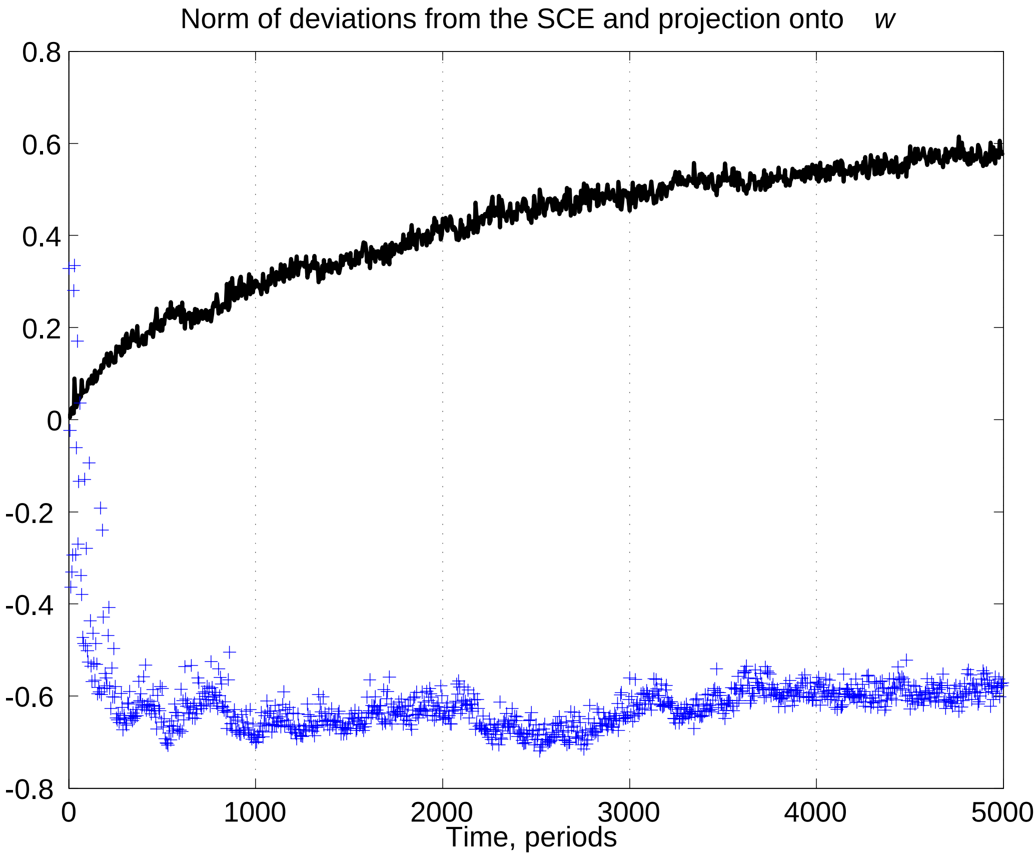
Norm of deviations from the SCE and projection onto w
0.8
0.6
0.4
0.2
0
-0.2
-0.4
-0.6
-0.8
0
1000
2000
3000
4000
5000
Time, periods

Figure 9: