

STOCHASTIC GRADIENT VERSUS RECURSIVE LEAST SQUARES LEARNING

Sergey Slobodyan
Anna Bogomolova
Dmitri Kolyuzhnov

CERGE-EI

Charles University
Center for Economic Research and Graduate Education
Academy of Sciences of the Czech Republic
Economics Institute

Working Paper Series **309**
(ISSN 1211-3298)

**Stochastic Gradient versus Recursive
Least Squares Learning**

Sergey Slobodyan
Anna Bogomolova
Dmitri Kolyuzhnov

CERGE-EI
Prague, October 2006

ISBN 80-7343-105-X (Univerzita Karlova. Centrum pro ekonomický výzkum a doktorské studium)
ISBN 80-7344-094-6 (Akademie věd České republiky. Národohospodářský ústav)

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: May 23, 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. Under the mean (averaged) RLS dynamics, the Self-Confirming Equilibrium (SCE) is stable for initial conditions in a very small region around the SCE. Large distance movements of perceived model parameters from their SCE values, or “escapes”, are observed.

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 SG learning.

Results of our paper hint that caution is needed when constant gain learning algorithms are used. If the mean dynamics map is stable but not contracting in every direction, and most eigenvalues of the map are close to the unit circle, the constant gain learning algorithm might diverge.

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

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

*{*Sergey.Slobodyan, Anna.Bogomolova, Dmitri.Kolyuzhnov*}@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.

Abstrakt

V této práci podrobně zkoumáme relativní výhody dvou algoritmů adaptivního učení s konstantním výnosem, Rekurzivního algoritmu nejmenších čtverců (Recursive Least Squares, RLS) a Algoritmu stochastického gradientu (Stochastic Gradient, SG), na Phelpsově modelu monetární politiky. Chování těchto dvou algoritmů je velmi odlišné. Uvažujeme-li průměrnou dynamiku RLS, pak je stav Sebe-potvrzující rovnováhy (Self-Confirming Equilibrium, SCE) stabilní pro počáteční podmínky v malé oblasti okolo SCE. Jsou ale pozorovány velké pohyby sledovaných parametrů od SCE hodnot, tzv. “úniky”. Na druhé straně, SCE je stabilní pro průměrnou dynamiku SG v rozsáhlé oblasti. Nicméně chování SG jako algoritmu adaptivního učení je divergentní pro velký rozsah parametrů konstantních výnosů, včetně těch které mohou být vnímány jak ekonomicky realistické. Vysvětlujeme tyto nesrovnalosti pomocí analýzy vlastních hodnot (eigenvalues) a vlastních vektorů (eigenvectors) rozdělení průměrné dynamiky v SG algoritmu adaptivního učení. Výsledky naší práce naznačují, že při používání konstantních výnosů učícího algoritmu je potřeba opatrnosti. Jestliže je rozdělení průměrné dynamiky stabilní, nikoli však ve všech směrech, a většina vlastních hodnot je blízko k jednotkové kružnici, konstantní výnos algoritmu adaptivního učení může divergovat.

1 Introduction

In this paper, we perform an in–depth investigation of the 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 more limited attention in the past, but the situation is changing: Evans, Honkapohja, and Williams (2005) promote the constant gain SG (and generalized SG) as a robust learning rule, which is well suited to the situation of time–varying parameters.

A different motivation for studying the properties of the SG learning comes from recent interest in heterogeneous learning, cf. Honkapohja and Mitra (2005) or Giannitsarou (2003). In this literature, several types of agents use different adaptive rules to learn the parameter values of 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 not restricted to the dynamics of the learning algorithm in a small neighborhood of the rational expectations equilibrium (REE) which motivates our focus on constant gain learning. It is known that E–stability of the REE, which implies local stability under RLS learning with decreasing gain, does not automatically imply local stability under SG with decreasing gain, see Giannitsarou (2005). Here the equilibrium is E–stable under both RLS and SG learning, but the behavior of the constant gain versions of the two methods is substantially different 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 hereafter). 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 describing the escape dynamics in a large interval of values of the constant gain. Here, we derive an even better one-dimensional approximation and discuss the Lyapunov function-based approach in establishing the 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 and define the RLS and SG learning in Section 2. In Section 3, we present and contrast the non-local effects arising under the constant gain versions of these algorithms and discuss the possible explanations for the difference in behavior of the mean dynamics and the actual real-time learning algorithm. Section 4 concludes.

2 The model and learning algorithms

The economy consists of the government and the private sector. The government attempts to minimize losses from inflation π_n and unemployment U_n :

$$\min_{\{x_n\}_{n=0}^{\infty}} E \sum_{n=0}^{\infty} \beta^n (U_n^2 + \pi_n^2), \quad (1)$$

It uses the monetary policy instrument x_n to control π_n , Eq. (2b). It believes (in general, incorrectly) in the Phillips curve (2c). The true Phillips curve is given by (2a): Unemployment is affected only by unexpected inflation. The private sector possesses rational expectations $\hat{x}_n = x_n$ about the inflation rate, and thus unexpected inflation shocks come only from monetary policy errors. The whole model is presented below.

$$U_n = u - \chi(\pi_n - \hat{x}_n) + \sigma_1 W_{1n}, \quad u > 0, \theta > 0, \quad (2a)$$

$$\pi_n = x_n + \sigma_2 W_{2n}, \quad (2b)$$

$$U_n = \gamma_1 \pi_n + \gamma_{-1}^T X_{n-1} + \eta_n. \quad (2c)$$

In the “static” version of the model, X_{n-1} contains only a constant, while two lags of π and U are added to X_{n-1} in the “dynamic” version. W_{1n} and W_{2n} are zero mean, unit-variance independent Gaussian shocks. Vector $\gamma = (\gamma_1, \gamma_{-1}^T)^T$ represents a government’s beliefs about the Phillips curve; it is 6-dimensional in the “dynamic”

and 2-dimensional in the “static” model. η_n is perceived by the government as a white noise uncorrelated with regressors π_n and X_{n-1} .

The equilibrium is defined as a vector of beliefs $\bar{\gamma}$ at which the government’s assumptions about orthogonality of η_n to the space of regressors are consistent with observations:

$$E \left[\eta_n \cdot (\pi_n, X_{n-1})^T \right] = 0. \quad (3)$$

CWS call this point a *self-confirming equilibrium*, or SCE. Williams (2001) shows that at the SCE, $\gamma = (-\chi, 0, 0, 0, 0, u(1 + \chi^2))^T$, and the average inflation is $x_n = \chi u$. For a detailed description of the model, see CWS.

In a period n , the government solves (1), subject to (2b) and (2c), assuming that current beliefs γ_n will never change. The monetary policy action x_n is correctly anticipated by the private sector. U_n is generated according to (2a), and the government’s beliefs are adjusted in a constant gain adaptive learning step. Let $\xi_n = [W_{1n} \ W_{2n} \ X_{n-1}^T]^T$; $g(\gamma_n, \xi_n) = \eta_n \cdot (\pi_n, X_{n-1}^T)^T$; and $M_n(\gamma_n, \xi_n) = (\pi_n, X_{n-1}^T)^T \cdot (\pi_n, X_{n-1}^T)$. The next period’s beliefs γ_{n+1} and R_{n+1} are given by

$$\gamma_{n+1} = \gamma_n + \epsilon R_n^{-1} g(\gamma_n, \xi_n), \quad (4a)$$

$$R_{n+1} = R_n + \epsilon (M_n(\gamma_n, \xi_n) - R_n), \quad (4b)$$

under RLS learning and by

$$\gamma_{n+1} = \gamma_n + \epsilon g(\gamma_n, \xi_n) \quad (5)$$

under the SG learning.¹

Set the parameter vector $\theta_n^{\epsilon, SG}$ equal to γ_n for the SG and $\theta_n^{\epsilon, RLS} = [\gamma_n^T, \text{vech}^T(R_n)]^T$ for the RLS case.² Define $H^{RLS}(\theta_n^\epsilon, \xi_n) = [(R_n^{-1} \cdot g(\gamma_n, \xi_n))^T, \text{vech}^T(M_n(\gamma_n, \xi_n) - R_n)]^T$ and $H^{SG}(\theta_n^\epsilon, \xi_n) = g(\gamma_n, \xi_n)$ to write the Stochastic Recursive Algorithm (SRA) in the standard form:

$$\theta_{n+1}^{\epsilon, j} = \theta_n^{\epsilon, j} + \epsilon H^j(\theta_n^{\epsilon, j}, \xi_n), \quad j = \{RLS, SG\}, \quad (6a)$$

$$\xi_{n+1} = A(\gamma_n) \xi_n + B \cdot [W_{1n+1} \ W_{2n+1}]^T. \quad (6b)$$

¹ R_n is the current estimate of the 2nd moments matrix of the regressors.

²Following the notation of Lütkepohl (1996), *vech* denotes a column vector in which abridged columns (the main diagonal and below) of a square matrix are stacked.

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

$$\dot{\theta}^j = E[H^j(\theta^{\epsilon,j}, \xi_n)]. \quad (7)$$

The SCE (vector $\bar{\gamma}$ and corresponding 2^{nd} moments matrix \bar{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: the SCE is E-stable under both algorithms. Solution of the ODE (7) is called the “mean dynamics trajectory” of the SRA (6), with the right-hand side of (7) being the “mean dynamics”. For details and derivations, see Evans and Honkapohja (2001). Another local continuous-time approximation of the SRA around the SCE $\bar{\theta}$ can be derived in the constant gain case, as shown by Evans and Honkapohja (2001, Prop. 7.8) and Williams (2001, Theorem 3.2),

$$d\varphi_t^{RLS} = P\varphi_t dt + \sqrt{\epsilon}\Sigma^{1/2}(\bar{\theta}^{RLS})dW_t, \quad (8)$$

where $\varphi_t = \theta_t^{RLS} - \bar{\theta}^{RLS}$ are deviations from the SCE, see Kolyuzhnov, Bogomolova, and Slobodyan (2006). We use the approximation (8) to study behavior of the model when when RLS learning is employed.

Still another variant of the mean dynamics approximation is the following difference equation obtained from (6a):

$$\theta_{n+1}^{\epsilon,j} = \theta_n^{\epsilon,j} + \epsilon \cdot E[H^j(\theta_n^{\epsilon,j}, \xi_n)]. \quad (9)$$

The difference between the above approximation and (7) is that ϵ is not assumed to be approaching zero asymptotically. This approximation turns out to be useful when we consider the learning dynamics in the SG case.

3 Behavior of Simulations

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

3.1 Recursive Least Squares

3.1.1 Dynamic Model

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 neighbor-

hood of the SCE, the beliefs vector γ suddenly deviates from the SCE towards the “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 χu and approaches 0, see Fig. 1 in CWS.

In Kolyuzhnov, Bogomolova, and Slobodyan (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 (4a) is almost one-dimensional because the two largest eigenvalues of \bar{R}^{-1} , λ_1 and λ_2 , equal 3083.8 and 29.1. As a result, the projection of $g(\gamma_n, \xi_n)$ onto v_1 , the dominant eigenvector of \bar{R}^{-1} , is amplified about 100 times as strongly as the projection onto the second largest eigenvector. It is also well known that in this model, the region of attraction of the SCE is very small, see Fig. 1 reprinted from Kolyuzhnov, Bogomolova, and Slobodyan (2006) or Figs. 8 and 9 in CWS. Outside of the immediate neighborhood of the SCE, the mean dynamics point away from it and towards the “induction hypothesis” plane in the direction which is very close to v_1 . These trajectories linger in the neighborhood of the plane for a relatively long time and then start a slow return to the SCE. As a result, simulation runs with escapes tend to contain a set of points aligned along the dominant eigenvector of \bar{R}^{-1} all the way towards the “induction hypothesis” plane, which is clearly demonstrated in the Figure 2 reprinted from Kolyuzhnov, Bogomolova, and Slobodyan (2006).³

We use this essential one-dimensionality to derive the following approximation of (8). Write $\varphi_t \approx x_t \cdot \tilde{v}_1$, and multiply (8) by \tilde{v}_1^T from the left. The resulting 1-dimensional approximation is then given by

$$dx_t \approx \tilde{v}_1^T D_{\theta p}(\bar{\theta}^{RLS}) \tilde{v}_1 \cdot x_t dt + \sqrt{\epsilon \tilde{\lambda}_1} \cdot \tilde{v}_1^T \cdot dW_t = A \cdot x_t dt + \sqrt{\epsilon \tilde{\lambda}_1} dW_t, \quad (10)$$

where $\tilde{\lambda}_1$ is the dominant eigenvalue of Σ . Note that $\tilde{v}_1^T \cdot dW_t$ is a one-dimensional standard Brownian motion. (10) is then an Ornstein-Uhlenbeck process with well-known properties. In particular, one could easily derive the expected time until the

³In the Figure, 6-dimensional vector of beliefs γ is presented in the space of $(\tilde{\gamma}_1, \tilde{\gamma}_2)$, defined as $\gamma_1 + \gamma_4 + \gamma_5$ and $u \cdot (\gamma_2 + \gamma_3) + \gamma_6$. A government’s beliefs about the influence of past and current inflation on U_n are given by $\tilde{\gamma}_1$, while $\tilde{\gamma}_2$ represents the beliefs about the effect of past unemployment (and a constant). The significant disbalance of eigenvalues of \bar{R}^{-1} is inherited by the matrix Σ in (8), and the eigenvector v_1 is essentially collinear to the first 6 components of \tilde{v}_1 , the dominant eigenvector of Σ .

process leaves any interval of the real line, see Borodin and Salminen (1996).⁴

To estimate the region of applicability of the approximation (10), take x_t^2 as the Lyapunov function and calculate LV for one-dimensional diffusion (10):⁵

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

Clearly, LV is positive for small x_t , 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 (6) is expected to be locally divergent on 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, Bogomolova, and Slobodyan (2006). For values of ϵ which are not “small”, the approximation (10) could be used to derive expected escape time. In the dynamic model, values of ϵ below $2 \cdot 10^{-5}$ are “small”.

3.1.2 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

⁴Ornstein–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).

⁵The 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.

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 a monthly basis but the output gap could be evaluated only quarterly, values of ϵ not much larger or smaller than 0.01 seem empirically justified. In a recent paper, Orphanides (2006) considers values of ϵ between 0.01 and 0.03 as fitting the data in a model with constant gain RLS learning. He also uses $\epsilon = 0.005$ for the SG constant gain learning of a natural real rate and a natural unemployment rate.

Notice that the period in the Phelps model could not be shorter than a quarter (or a month). As Table 1 shows, for $\epsilon < 1 \cdot 10^{-4}$ in the dynamic model and $\epsilon < 4 \cdot 10^{-4}$ in the static one, the expected time until escape becomes larger than an economically relevant time scale (say, a hundred years); probability of observing an escape within this time becomes negligible as ϵ decreases even further. An important caveat to this statement is that both the theoretical and simulation results are obtained by imposing the SCE as the starting point of learning. In other words, one starts from a situation of a completed learning, where 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.

3.1.3 Static Model

Dynamics of the static model under the constant gain RLS learning is qualitatively similar to that of the dynamic one: a move out of the immediate region of attraction of the SCE, followed 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.⁶ As a result, in the static model ϵ starts to be “small” at about $3 \cdot 10^{-4}$.

The combined effect of the stronger drift, weaker diffusion, and larger stability region is obvious: a significantly larger than in the dynamic model expected num-

⁶In the dynamic (static) model, $A = -0.41$ (-0.52) and $\tilde{\lambda}_1 = 278$ (26).

ber of periods until the simulations escape the 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 values of ϵ which are not “small”, the agreement is rather good, especially for the static model. In agreement with our estimate of the Ornstein–Uhlenbeck approximation’s applicability, it starts to overpredict for “small” ϵ . This effect is especially pronounced for the static model.

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

ϵ	Dynamic model		Static model	
	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

3.2 Stochastic Gradient Learning

It is necessary to note that in the SG case, the dependence of the learning dynamics on ϵ is dramatically different from the RLS case. In a nutshell, simulations are divergent for a rather wide interval of ϵ . On the other hand, the term R_n^{-1} does not multiply the right-hand-side in Eq. (5), which prevents usage of a one-dimensional approximation which proved to be so successful in the RLS case.

3.2.1 Dynamic Model

In the approximation (9), 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 $\lambda_1 = 0.2447$ to $\lambda_2 = 0.9988$ to $\lambda_6 = 0.99999862$. Five out of six eigenvalues are almost unitary. Under the mean dynamics (9), any deviation from the SCE results in a fast movement along x_1 , the eigenvector which corresponds to λ_1 , and then an extremely slow convergence back to the SCE along the remaining five directions, see Figure 3. On the other hand, simulations of (6) behave very differently. Figure 4 plots a norm of deviations from the SCE and $\gamma_6 - \bar{\gamma}_6$: There is a clearly distinguishable movement away from the SCE which seems almost deterministic.⁷ For this value of ϵ , the inflation rate will drop below 4 (at the SCE, its mean equals 5) in a couple of hundred periods, which is definitely the time scale with which one should be concerned. How could one explain the discrepancy between the mean dynamics (9) and the simulations?

Fig. 5 plots a projection of $\frac{\gamma_n - \bar{\gamma}}{\|\gamma_n - \bar{\gamma}\|}$ onto the sub-space spanned by five eigenvectors of $F(\epsilon)$ 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, a simulation run quickly approaches some neighborhood of the sub-space and does not leave it for any extended period of time. This behavior is natural: Any initial deviation along x_1 will shrink to $0.25^3 \sim 1.5\%$ of its initial size in just 3 steps. On the other hand, deviations along five other eigenvectors will take at least $\frac{\ln(0.5)}{\ln(0.9988)} \sim 577$ periods to reach 50% of their initial magnitude.

Another feature of the matrix $F(\epsilon)$ 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 symmetric part of $F(\epsilon)$, $F^{sym}(\epsilon) = \frac{F(\epsilon) + F(\epsilon)^T}{2}$, is not stable. After one iteration of the map $F(\epsilon)$, initial deviation in the direction w , the unstable eigenvector of $F^{sym}(\epsilon)$, is expected to *increase* its projection onto w and thus to *increase* its norm, at least initially.⁸ The

⁷If we observe the simulations for a larger number of periods, the belief vector γ eventually reaches values at which the state vector process loses stationarity, and the simulation breaks down.

⁸Suppose an initial deviation is given by w . After one period, this deviation is transformed

largest eigenvalue of $F^{sym}(\epsilon)$ equals 1.103 at $\epsilon = 0.01$, 1.01 at $\epsilon = 0.001$ and 1.001 at $\epsilon = 1 \cdot 10^{-4}$. A projection of $\frac{\gamma_n - \bar{\gamma}}{\|\gamma_n - \bar{\gamma}\|}$ onto w is plotted in Figure 6 (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 (9) 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 7 presents 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.

The norm of the projection of w onto the sub-space spanned by the five eigenvectors is rather large and equals 0.95. When the dynamics of (6) 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 the value of $\gamma_n - \bar{\gamma}$ 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 sub-space.

As a final piece of evidence connecting the vector w with the divergent behavior of simulations, consider Figure 8. In the periods when the projection of $\frac{\gamma_n - \bar{\gamma}}{\|\gamma_n - \bar{\gamma}\|}$ onto w (crosses) is particularly large, the distance between the beliefs γ_n and the SCE $\bar{\gamma}$ (solid line) grows the fastest; a relative decline in the projection is correlated with a temporary stop or even a 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 structure of the mean dynamics map $F(\epsilon)$. The sub-space spanned by the almost unitary eigenvalues' eigenvectors of $F(\epsilon)$ 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

into Fw . Projection $w^T Fw$ then gives a measure of expansion or contraction in the direction of w after one iteration of map F . But for any vector w , $w^T Fw = w^T F^{sym} w$. Therefore, in order to find expanding (after one iteration) directions of F , one could look at eigenvalues of F^{sym} .

be very short-lived, and that a contracting mean dynamics within the sub-space 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 a government’s beliefs γ_n away from the SCE. One could still observe a very fast convergence towards the sub-space spanned by the five almost unitary eigenvalues’ eigenvectors and a significant projection onto the expanding direction w . Only for very small values of $\epsilon \leq 8 \cdot 10^{-6}$ we start observing a different behavior, when the system (6) does not systematically diverge and fluctuates in some neighborhood of the SCE.

3.2.2 Static Model

Taking into account that under 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 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 $\|\gamma - \bar{\gamma}\|$ 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 $F(\epsilon)$ has a real eigenvalue which is less than -1 making the SCE unstable), 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 $\|\gamma - \bar{\gamma}\| \geq 0.1$; such large excursions are not likely to be observed before the 500th 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 (6b) being outside of the unit circle and thus to a 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 $F(\epsilon)$ has two eigenvalues. One is always close to one (0.9999 for $\epsilon = 3 \cdot 10^{-2}$). The other is a linearly decreasing function of ϵ . It equals -1 when $\epsilon \sim 7.9 \cdot 10^{-2}$ and approaches 1 as $\epsilon \rightarrow 0$. 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 $\gamma - \bar{\gamma}$ is in the closest alignment with w ($w^T \cdot \frac{\gamma_n - \bar{\gamma}}{\|\gamma_n - \bar{\gamma}\|}$ is close to one). There are two crucial differences with the dynamics model, however. First, the direction w is very weakly expansive: The unstable eigenvalue of $F^{sym}(\epsilon)$ 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 $F(\epsilon)$ equals 0.23 for $\epsilon \sim 3 \cdot 10^{-2}$ and is decreasing in ϵ . Thus, for smaller values of ϵ , the dynamics of (9) loses its essentially one-dimensional nature in the expanding direction, 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 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.

3.2.3 Reasons for difference with the RLS case

Why do we observe the diverging behavior documented above only in the SG case? RLS case differs from the SG one in three respects. First, the mean dynamics is very weak relative to the stochastic dynamics especially in the direction of dominant eigenvector of \bar{R}^{-1} , as documented in Kolyuzhnov, Bogomolova, and Slobodyan (2006). Second, the mean dynamics map $F(\epsilon)$ does not contain strongly contracting eigenvalues. And third, those eigenvalues of $F(\epsilon)$ that are closest to the unit circle are much further from it than in the SG case. A combination of these three factors assures that even though simulation runs under RLS learning do exhibit relatively large projections in the expanding direction of $F^{sym}(\epsilon)$, these projections are not correlated with episodes of particularly fast deviations from the SCE.

4 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 force the government's beliefs about the Phillips curve to "escape", or deviate significantly, from the neighborhood of the Self-Confirming Equilibrium where the 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 model dynamics is divergent 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 the existence of eigenvalues of the SRA mean dynamics map which are very close to the unit circle, and thus, deviations in the direction of corresponding eigenvectors contract very slowly. Moreover, the SRA mean dynamics map has directions which are expected to expand in the short run rather than contract, and these directions are almost parallel to the sub-space spanned by the slowly contracting eigenvectors. Such a 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 an empirically relevant time scale and for constant gain values likely to be used in practice ("stability").

Additionally, the expected escape time rises very steeply as ϵ decreases. Following Evans, Honkapohja, and Williams (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 \bar{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 mis-specified 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 the mean dynamics map is asymptotically 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, and the mean dynamics map is not contracting in every direction, the Stochastic Recursive Algorithm might exhibit *divergent* behavior despite *convergent* mean dynamics.

References

- BORODIN, A. N., AND P. SALMINEN (1996): *Handbook of Brownian Motion — Facts and Formulae*. Birkhauser Verlag, Basel Boston Berlin.
- CHO, I.-K., N. WILLIAMS, AND T. J. SARGENT (2002): “Escaping Nash Inflation,” *Review of Economic Studies*, 69(1), 1–40.
- EVANS, G. W., AND S. HONKAPOHJA (2001): *Learning and Expectations in Macroeconomics*. Princeton University Press, Princeton, NJ.
- EVANS, G. W., S. HONKAPOHJA, AND N. WILLIAMS (2005): “Generalized Stochastic Gradient Learning,” NBER Technical Working Paper 317.
- GIANNITSAROU, C. (2003): “Heterogeneous Learning,” *Review of Economic Dynamics*, 6, 885–906.
- (2005): “E-Stability Does Not Imply Learnability,” *Macroeconomic Dynamics*, 9, 276–287.
- HONKAPOHJA, S., AND K. MITRA (2005): “Learning Stability in Economies with Heterogeneous Agents,” mimeo.
- KHASHMINSKII, R. Z. (1980): *Stochastic Stability of Differential Equations*. Sijthoff & Noordhoff, Alphen aan den Rijn, The Netherlands.
- KOLYUZHNOV, D., A. BOGOMOLOVA, AND S. SLOBODYAN (2006): “Escape Dynamics: A Continuous-Time Approximation,” CERGE–EI Working Paper 285.
- LÜTKEPOHL, H. (1996): *Handbook of Matrices*. John Wiley & Sons, Chichester.
- ORPHANIDES, A. (2006): “Inflation Targeting under Imperfect Knowledge,” FRB San Francisco Working Paper 2006-14.
- SARGENT, T. J. (1999): *The Conquest of American Inflation*. Princeton University Press.
- WILLIAMS, N. (2001): “Escape Dynamics in Learning Models,” Ph.D. thesis, University of Chicago.




Figure 1. The mean dynamics trajectories under RLS




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

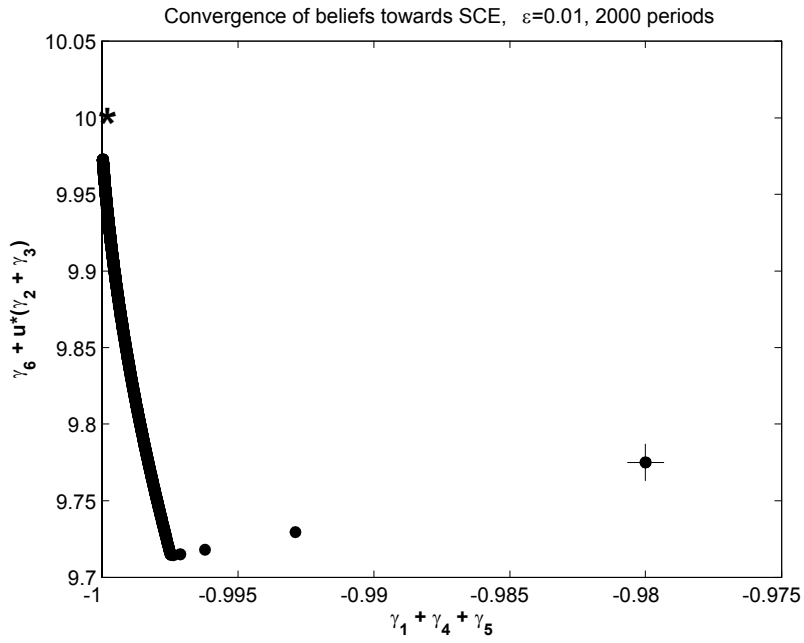


Figure 3. Iterations of the mean dynamics map. '+' sign: start of the simulation.
 '*' - the SCE location

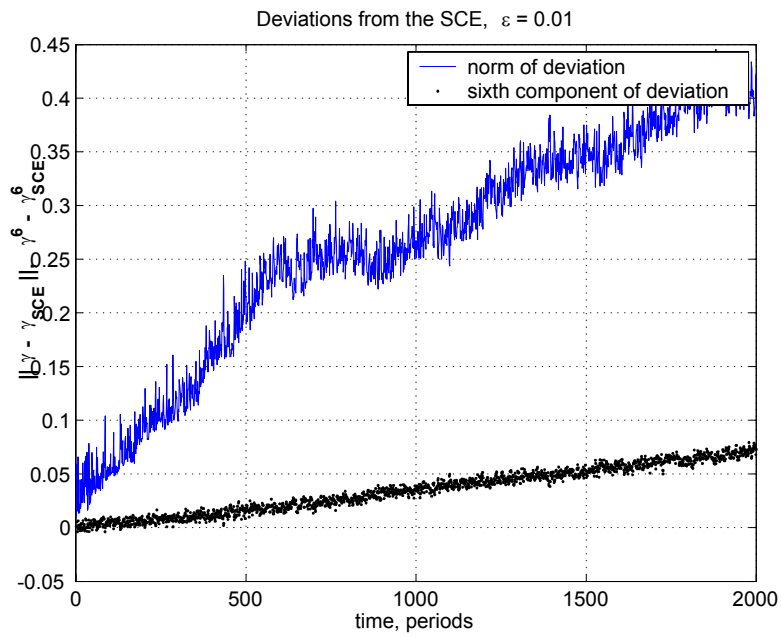


Figure 4. Divergence of a simulation run from the SCE.




Figure 5. Projection of beliefs onto the sub-space spanned by the mean dynamics map's almost unitary eigenvalues' eigenvectors in a typical simulation run.




Figure 6. Projection of beliefs onto the expansive direction of the mean dynamics map in a typical simulation run.

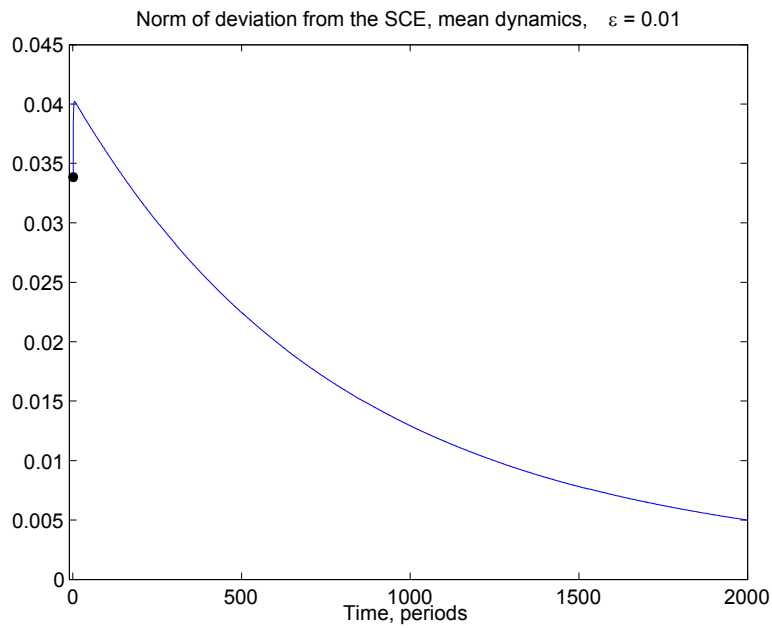


Figure 7. Evolution of beliefs under iterations of the mean dynamics map. Initial deviation in the expanding direction.




Figure 8. The largest is the projection of beliefs onto the expanding direction of the mean dynamics map ('+' sign), the faster the beliefs deviate from the SCE ('.' sign).

Individual researchers, as well as the on-line and printed versions of the CERGE-EI Working Papers (including their dissemination) were supported from the following institutional grants:

- Center of Advanced Political Economy Research [Centrum pro pokročilá politicko-ekonomická studia], No. LC542, (2005-2009),
- Economic Aspects of EU and EMU Entry [Ekonomické aspekty vstupu do Evropské unie a Evropské měnové unie], No. AVOZ70850503, (2005-2010);
- Economic Impact of European Integration on the Czech Republic [Ekonomické dopady evropské integrace na ČR], No. MSM0021620846, (2005-2011);

Specific research support and/or other grants the researchers/publications benefited from are acknowledged at the beginning of the Paper.

(c) Sergey Slobodyan, Anna Bogomolova, Dmitri Kolyuzhnov, 2006

All rights reserved. No part of this publication may be reproduced, stored in a retrieval system or transmitted in any form or by any means, electronic, mechanical or photocopying, recording, or otherwise without the prior permission of the publisher.

Published by

Charles University in Prague, Center for Economic Research and Graduate Education (CERGE) and

Economics Institute (EI), Academy of Sciences of the Czech Republic

CERGE-EI, Politických vězňů 7, 111 21 Prague 1, tel.: +420 224 005 153, Czech Republic.

Printed by CERGE-EI, Prague

Subscription: CERGE-EI homepage: <http://www.cerge-ei.cz>

Editors: Directors of CERGE and EI

Managing editors: Deputy Directors for Research of CERGE and EI

ISSN 1211-3298

ISBN 80-7343-105-X (Univerzita Karlova. Centrum pro ekonomický výzkum a doktorské studium)

ISBN 80-7344-094-6 (Akademie věd České republiky. Národohospodářský ústav)



CERGE-EI
P.O.BOX 882
Politických vězňů 7
111 21 Praha 1
Czech Republic
<http://www.cerge-ei.cz>