

STOCHASTIC GRADIENT LEARNING AND INSTABILITY: AN EXAMPLE

SERGEY SLOBODYAN

CERGE-EI

ANNA BOGOMOLOVA

CERGE-EI

and

Novosibirsk State University

DMITRI KOLYUZHNOV

CERGE-EI

and

*Institute of Economics and Industrial Engineering of the Siberian Branch of
the Russian Academy of Sciences*

In this paper, we investigate real-time behavior of constant-gain stochastic gradient (SG) learning, using the Phelps model of monetary policy as a testing ground. We find that whereas the self-confirming equilibrium is stable under the mean dynamics in a very large region, real-time learning diverges for all but the very smallest gain values. We employ a stochastic Lyapunov function approach to demonstrate that the SG mean dynamics is easily destabilized by the noise associated with real-time learning, because its Jacobian contains stable but very small eigenvalues. We also express caution on usage of perpetual learning algorithms with such small eigenvalues, as the real-time dynamics might diverge from the equilibrium that is stable under the mean dynamics.

Keywords: Constant Gain, Adaptive Learning, E-stability, Stochastic Gradient Learning

1. INTRODUCTION

Recently, asymptotic stability criteria such as E-stability or stochastic gradient (SG) and generalized SG stability were put forward as criteria a good monetary policy rule should satisfy, making possible anchoring of agents' inflation expectations on a desired equilibrium through a learning process; see Bullard and Mitra (2002) and Evans et al. (2010). A number of papers [Orphanides and Williams (2007), Milani (2008), Bullard et al. (2010), Slobodyan and Wouters

A previous version of this paper was circulated as Slobodyan, Bogomolova, and Kolyuzhnov, "Stochastic Gradient versus Recursive Least Squares Learning," CERGE-EI WP 309, Oct 2006. The first author acknowledges financial support from GACR Grant P402/11/J018. Address correspondence to: Sergey Slobodyan, CERGE-EI, a joint workplace of Charles University in Prague and the Economics Institute of the Academy of Sciences of the Czech Republic, Politických veznu 7, 111 21 Prague, Czech Republic; e-mail: Sergey.Slobodyan@cerge-ei.cz.

(2012), Hommes and Lux (2013)] have simulated or estimated macro models with expectations formed using adaptive learning algorithms. This development puts the problem of looking for models and adaptive learning algorithms that are well behaved not only *asymptotically* but also in *real time* at the forefront of the research agenda in the monetary policy literature.

Studies of the Phelps problem of a government controlling inflation while adaptively learning the approximate Phillips curve using constant-gain recursive least squares (CG RLS) have discovered a phenomenon known as “escape dynamics”; see Sargent (1999), Cho et al. (2002) (CWS hereafter), and Berardi (2013). “Escape” is a sudden deviation of government beliefs from the neighborhood of the equilibrium, in which a nonzero inflation–unemployment tradeoff is perceived to exist, toward vertical Phillips curve beliefs. During the escape, average inflation drops from a high (Nash equilibrium) value to zero. Escapes are observed despite the fact that the equilibrium is E-stable and thus asymptotically stable under the mean, or average, dynamics of beliefs updated by RLS.

It is clearly interesting to study real-time properties of other learning algorithms in the Phelps model. In this note we study constant-gain (CG) SG learning, which is simpler than RLS, and for certain priors could be robust to agents’ uncertainty about their beliefs, as well as appropriate to the case of time-varying parameters; see Evans et al. (2010). We find that the mean dynamics under SG learning remains stable in a large region, and that there are no escapes, but real-time learning diverges for all but very low values of the gain. Low convergence speed of the mean dynamics in a high-dimensional space is a major cause of such behavior. Thus, a policy maker conducting monetary policy that is optimal given current beliefs would lead the economy along a divergent path.

To establish our results, we use a stochastic Lyapunov function approach that is novel in the adaptive learning literature. Although our analytical results are limited, this tool turns out to be useful in understanding the sources of instability of CG SG learning in the Phelps model.

We briefly summarize the model of CWS and define SG learning in Section 2. In Section 3, we present the nonlocal effects arising under CG SG learning and discuss the possible explanations for the difference in behavior between the mean dynamics and the actual real-time learning algorithm. Section 4 concludes.

2. THE MODEL AND LEARNING ALGORITHMS

The CWS model of a government adaptively learning a misspecified Phillips curve can be summarized as follows. The economy consists of the government and the private sector. The government attempts to minimize losses from inflation π_n and unemployment U_n :

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

The true Phillips curve is given by

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

The government uses the monetary policy instrument x_n to control π_n , as shown in

$$\pi_n = x_n + \sigma_2 W_{2n}. \tag{3}$$

It holds misspecified beliefs about the Phillips curve,

$$U_n = \gamma_1 \pi_n + \gamma_{-1}^T X_{n-1} + \eta_n, \tag{4}$$

whereas in (2) 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.

In a “dynamic” version of the model, which we consider in this paper, X_{n-1} contains a constant and two lags of π and U . W_{1n} and W_{2n} are zero-mean, unit-variance independent Gaussian shocks.¹ A 6-dimensional vector $\gamma = (\gamma_1, \gamma_{-1}^T)^T$ represents the government’s beliefs about the Phillips curve. η_n , the residual in the misspecified Phillips curve defined by (4), 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 assumption about orthogonality of η_n to the space of regressors is indeed consistent with observations:

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

CWS call this point a *self-confirming equilibrium*, or SCE. Williams (2009) shows that at the SCE, $\gamma = \bar{\gamma} = [-\chi, 0, 0, 0, 0, u(1 + \chi^2)]^T$, and thus the government perceives a nonzero inflation-unemployment tradeoff. As a result, it sets the policy instrument x_n equal to $\chi u > 0$ (which is also the average inflation at the SCE) because costs of higher unemployment are perceived to be too high. For a detailed description of the model, see CWS.

In period n , the government solves (1), subject to (3) and (4), 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 (2), and the government’s beliefs are adjusted in a constant-gain adaptive learning step. Let $\xi_n = [W_{1n} \ W_{2n} \ X_{n-1}^T]^T$ and $g(\gamma_n, \xi_n) = \eta_n \cdot (\pi_n, X_{n-1}^T)^T$. Under the SG learning the next period’s beliefs γ_{n+1} are given by

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

which together with the law of motion for the state variable ξ_n ,

$$\xi_{n+1} = A(\gamma_n)\xi_n + B \cdot [W_{1n+1} \ W_{2n+1}]^T, \tag{7}$$

with suitably defined matrices A and B constitutes a stochastic recursive algorithm (SRA). Finally, the approximating ordinary differential equation corresponding to

this SRA is given by

$$\dot{\gamma} = \lim_{n \rightarrow \infty} E[g(\gamma, \bar{\xi}_n)] = h(\gamma). \tag{8}$$

The right-hand side of this ordinary differential equation (ODE) is obtained as follows. For every parameter vector γ , construct the state vector process $\bar{\xi}_n$ corresponding to this γ as in (7), with $A(\gamma_n)$ becoming a fixed matrix $A(\gamma)$. $\bar{\xi}_n$ is thus a VAR with constant coefficients. The mathematical expectation in (8) is then taken with respect to the invariant distribution induced by this VAR. Solution of (8) is called the “mean dynamics trajectory” of the real-time learning dynamics (6) and (7), with the right-hand side of (8) being the “mean dynamics.” For details and derivations, see Evans and Honkapohja (2001), referred to as EH in the remainder of the paper. The SCE $\bar{\gamma}$ is the only equilibrium of this ODE.²

Another local continuous-time approximation of the SRA around the SCE $\bar{\gamma}$ can be derived in the constant-gain case, as shown by EH, Proposition 7.8:

$$d\gamma_t = J \cdot (\gamma_t - \bar{\gamma}) dt + \sqrt{\epsilon} \Sigma^{1/2}(\bar{\gamma}) dW_t. \tag{9}$$

The matrix $J = D_\gamma h|_{\gamma=\bar{\gamma}}$ is the Jacobian of the mean dynamics linearized around the SCE. The matrix Σ contains covariances of elements of $g(\gamma, \bar{\xi}_n)$ summed over all leads and lags of $\bar{\xi}_n$. See EH (Sections 7.4 and 14.4) for details.³ We also consider a generalized diffusion approximation to the SRA,

$$d\gamma_t = h(\gamma_t)dt + \sqrt{\epsilon} \Sigma^{1/2}(\gamma_t) dW_t. \tag{10}$$

Notice that (9) is a linear approximation of (10) around the SCE.

3. BEHAVIOR OF SIMULATIONS

The discussion that follows refers to the model as parameterized in CWS: $\sigma_1 = \sigma_2 = 0.3$, $u = 5$, $\chi = 1$, $\beta = 0.98$. We also explore the sensitivity of the results to parameter values.

For these parameter values the SCE is stable under the flow defined by (8), or SG-learnable using a term from Evans et al. (2010): J has one “fast” eigenvalue, $\lambda_1 = -75.5$, three “slow” ones, ranging from -0.12 to -0.07 , and two “extremely slow” eigenvalues as small as -1.4×10^{-4} .⁴ As a consequence, any deviation from the SCE results in a fast movement along the dominant or “fast” eigenvector v_1 , and then an extremely slow convergence back to the SCE along the remaining five “slow” directions.

The convergent dashed line in the left panel of Figure 1 presents the norm of deviations of beliefs from the SCE, $\|\gamma_t - \bar{\gamma}\|$, along one such trajectory. Note that stability of the mean dynamics (8) is observed for initial conditions in a very wide region. However, simulations of the real-time learning algorithm (6) and (7) show that the beliefs move away from the SCE in a fashion that looks like a deterministic trend with a strong stochastic component, as evidenced by the left panel of Figure

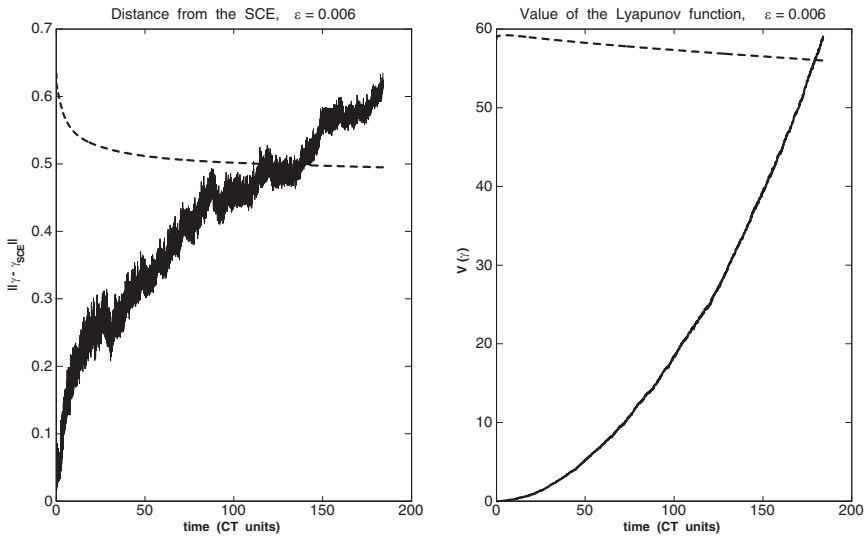


FIGURE 1. Evolution of the distance from the SCE and value of the Lyapunov function for mean dynamics trajectory (dashed line) and simulated SRA (black dots) under SG learning.

1, black dots.^{5,6} The divergence in Figure 1 is documented for $\epsilon = 0.006$, but similar behavior was observed for ϵ as low as 10^{-5} . For $\epsilon \lesssim 10^{-5}$ we observe typical fluctuations of beliefs around the SCE. For very high values of ϵ (above 0.027), one of the eigenvalues of the matrix $\epsilon J + (1 - \epsilon)I$ is located outside a unit circle, which also leads to divergence of simulations. Because the mechanism of divergence for large ϵ was described previously by Evans and Honkapohja (2009), we do not discuss it here.

There is thus a significant discrepancy between the mean dynamics and the simulations of real-time learning for $\epsilon \gtrsim 10^{-5}$. To explain this discrepancy, we utilize a stochastic Lyapunov function approach, described in detail in the Appendix.⁷ As a Lyapunov function, we choose a quadratic form $V(\gamma) = (\gamma - \bar{\gamma})^T P(\gamma - \bar{\gamma})$. Under the (linearized) mean dynamics, the time derivative of $V(\gamma)$ along the solution path of linearized equation (8),

$$\dot{\gamma} = J(\gamma - \bar{\gamma}), \tag{11}$$

equals $\dot{V}(\gamma) = (\gamma - \bar{\gamma})^T \cdot (PJ + J^T P) \cdot (\gamma - \bar{\gamma})$. Given that J is stable, for any symmetric positive definite matrix Q we can find P such that

$$PJ + J^T P + Q = 0.$$

The derivative of the Lyapunov function along the solution path, $\dot{V}(\gamma)$, equals $-(\gamma - \bar{\gamma})^T Q(\gamma - \bar{\gamma})$ and is always negative. To construct a suitable matrix Q , we take the orthonormal basis of the subspace spanned by the “slow” eigenvectors

of J plus the complement of this subspace, collected into matrix W , and retain the eigenvalues of $-J$. Thus, $Q = W \cdot \text{diag}[\text{eig}(-J)] \cdot W^T$. By construction, Q is positive definite, and therefore P exists. Also by construction, the Lyapunov function $V(\gamma)$ decreases very gradually in the subspace spanned by the “slow” eigenvectors of J and falls very fast in the complement of this subspace, resembling the spatial behavior of the linearized mean dynamics (11) to a significant degree.⁸ For a linear system such as (11), $\dot{V}(\gamma)$ will be negative everywhere but at the stationary point.

The Lyapunov function constructed in the preceding delivers stability results for the mean dynamics linearized around the SCE. For the nonlinear mean dynamics ODE (8), the expression for $\dot{V}(\gamma)$ becomes

$$\dot{V}(\gamma) = 2h(\gamma)^T P(\gamma - \bar{\gamma}).$$

There is no guarantee that the same Lyapunov function would work for nonlinear as well as linearized systems; in practice, it serves surprisingly well. Observe the dashed line in Figure 1, right panel. After an initial increase $V(\gamma)$ falls very slowly but monotonically, meaning that $\dot{V}(\gamma)$ becomes negative, even though we are simulating the nonlinear mean dynamics (8). However, along the simulation path of (6) and (7) (black dots), the situation is very different: $V(\gamma)$ increases steadily instead. Under the influence of stochastic shocks, the very slow convergence along the “slow” eigenvectors of J is replaced by divergence.

To get a deeper insight into the properties of (6) and (7), we utilize its nonlinear diffusion approximation (10), use the same Lyapunov function $V(\gamma)$, and derive a stochastic analog of $\dot{V}(\gamma)$. The concept of stochastic Lyapunov function is described in the Appendix:

$$\mathcal{L}V(\gamma) = \dot{V}(\gamma) + \epsilon \cdot \text{trace}(\Sigma P) = 2h(\gamma)^T P(\gamma - \bar{\gamma}) + \epsilon \cdot \text{trace}(\Sigma P).$$

We then calculate $\mathcal{L}V$ at every twentieth point of a typical simulation run of (6) and (7) with $\epsilon = 0.006$, and the increment of $V(\gamma)$ within the next twenty iterations. Recall that $\mathcal{L}V > 0$ means that V is expected to increase. The simulation results are presented in Figure 2. First, there indeed is a strong positive correlation of ~ 0.7 between $\mathcal{L}V(\gamma)$ and the subsequent increment of $V(\gamma)$. Second, stochastic shocks often push the beliefs into areas where $\mathcal{L}V$ is positive and large; this is explained by the fact that away from the SCE the mean dynamics is often unstable. Even when it is stable, the positive $\mathcal{L}V$ term, $\epsilon \cdot \text{trace}(\Sigma P)$, can be larger than the negative term, $2h(\gamma)^T P(\gamma - \bar{\gamma})$, given the very slow convergence of the mean dynamics. As a result, the average value of $\mathcal{L}V$ turns positive after several iterations of (6) and (7).⁹

Although it is impossible to derive a Lyapunov function for the nonlinear diffusion approximation (10) and formally prove instability of the SCE, this discussion is useful in illuminating the potential factors, which could lead to the discrepancy between the stability of mean dynamics trajectories and of real-time learning process (6) and (7). We identify three such factors.




FIGURE 2. Value of the Lyapunov derivative and consecutive increments of the Lyapunov function, SG learning.

First, the linearized mean dynamics has two extremely “slow” eigenvalues that are very close to zero. Once the beliefs are disturbed away from the SCE, they linger for a long time in the subspace spanned by these eigenvectors.¹⁰ Second, the symmetric part of J is not stable, which implies that random deviation of beliefs in a certain direction will be initially driven even further away from the SCE.¹¹ Finally, and most importantly, after such a deviation away from the SCE, the linear approximation (9) is no longer valid and (10) often has an unstable drift term: we observed a largest eigenvalue of locally linearized $h(\gamma)$ as high as $+4.5$. A locally unstable drift term leads to the Lyapunov derivative along the simulation run being positive on the average: Figure 3 plots the running average value of \mathcal{LV} , evaluated at the same points as in Figure 2. For almost all time intervals, the average value of \mathcal{LV} is positive and increasing, which means that the value $V(\gamma)$ is expected to grow with ever higher speed. This is exactly the behavior observed in the right panel of Figure 1, black dots.

The first two factors mean that because of the slow convergence of the mean dynamics, stochastic noise of the real-time learning has plenty of time to derail convergence of beliefs to the SCE, once an initial deviation has occurred. The third factor implies that the noise could be taking the beliefs into regions where




FIGURE 3. Cumulative Lyapunov derivative along a typical simulation path, SG learning.

the mean dynamics is divergent, not convergent, making further expansion of the deviation away from the SCE inevitable. Such regions are rather close to the SCE, because it might take a very small change in beliefs to move an eigenvalue from a stable -10^{-4} to an unstable 10^{-4} .

To further stress the importance of nonlinearities and locally unstable mean dynamics, we perform the following exercise. After calculating the invariant distribution of beliefs implied by (9),¹² we evaluate the largest eigenvalue of the locally linearized mean dynamics at 10,000 points, randomly drawn from this distribution. Averaging over the draws produces a stark result: even for $\epsilon = 6 \times 10^{-6}$, the average equals 0.16. Mean dynamics is locally unstable more often than stable, even in a very small neighborhood of the SCE.

Our conclusion is that the formal proof of the mean dynamics stability may matter little when eigenvalues close to the imaginary axis are present, as the actual real-time learning may still diverge. In itself, this is not surprising, as both the mean dynamics approximation (8) and the nonlinear diffusion approximation (10) are valid only as $\epsilon \rightarrow 0$. The novelty of our approach lies in showing that we were able to use the nonlinear diffusion approximation to ascertain the divergence of the real-time learning process for values of the constant gain ϵ close to those commonly used in the literature, using a stochastic Lyapunov function approach.

We have to comment on the stochastic stabilization and destabilization literature, in particular Mao (1994). Theorems 4.1 and 4.2 and Corollary 5.2 in that paper state that any nonlinear system of ODE could be destabilized by a Brownian motion with sufficiently high variance, provided the dimension of the system was greater than one. Destabilizing a given ODE requires a Brownian motion with lower variance when the norm of the linearized diffusion term (mean dynamics in our notation) is smaller. In our case this norm is essentially zero in some directions,

which makes destabilization very easy indeed. The variance of the noise term in (10) is proportional to ϵ , and therefore destabilization becomes impossible for extremely low values of ϵ , which is exactly what we observe.

As noted earlier, Evans et al. (2010) study SG learning in the “static” version of the Phelps model. The discrepancy between the mean dynamics convergence and real-time learning divergence is not observed in that paper. We note that in the “static” version of the model the Jacobian J is two-dimensional, with one eigenvalue being very “fast” and another “slow.” Therefore, the “slow” subspace has dimension equal to one. As shown in Mao (1994), it is impossible to destabilize one-dimensional stable ODE with stochastic noise.

SG real-time learning is divergent for gains greater than 10^{-5} . Does this range include values of ϵ that are “reasonable” for a macro model? Consider two rules of thumb for selecting ϵ . When CG adaptive learning is used in environments with time-varying parameters or structural breaks, for instance, in regime-switching model of Branch et al. (2013), the characteristic time $1/\epsilon$ is related to the typical time until a break, or to the time during which the variation becomes “large.” The second rule imagines agents who could use CG with an infinite amount of data or run OLS regression with T data points. In OLS regression, the mean age of a data point is $T/2$. A CG learning algorithm is equivalent to a weighed LS with weights proportional to points’ age, giving $\frac{1-\epsilon}{\epsilon}$ as the mean age. Therefore, $\epsilon \approx 2/T$ produces a CG algorithm approximately equivalent to the OLS regressions with sample length T ; see Orphanides and Williams (2009). For the Phelps problem at best monthly data are available, and $\epsilon \sim 0.002 \div 0.01$ seems to be empirically justified.

Recently, several papers have presented empirical estimates of ϵ . For example, Orphanides and Williams (2007) find ϵ between 0.01 and 0.03 fitting the data in a model with CG RLS learning. They also use $\epsilon = 0.005$ for CG SG learning of a natural real rate and a natural unemployment rate. Branch and Evans (2006) find that CG RLS with gains of 0.007 (GDP growth) and 0.062 (inflation) produces the smallest RMSE of out-of-sample forecasts. Milani (2008) estimates a simple New Keynesian model under a variety of policy and CG RLS specifications and finds gains between 0.01 and 0.03. Finally, Slobodyan and Wouters (2012) estimate a medium-scale DSGE model under several specifications of CG RLS and find gains from 0.001 to 0.02.

Could it be that the divergence that we document is a fragile property caused by very specific parameter values? To answer this question, we consider alternative values of the average unemployment rate $u \in [1; 4]$. The SCE remains stable under mean dynamics for all values of u . When u decreases from 5 to 1, the slowest eigenvalue of J gets faster, moving from -1.4×10^{-4} to -2.9×10^{-3} . As the slowest eigenvalue increases its magnitude, two effects happen. First, the deterministic contribution to $\mathcal{L}V$, $2h(\gamma)^T P(\gamma - \bar{\gamma})$, becomes more negative. Second, larger deviations of beliefs from the SCE are required to make the locally linearized mean dynamics unstable by turning one of its eigenvalues positive. Both effects make the divergence harder, and thus we expect that lowering u will lead

to slower divergence of the real-time learning out of any given neighborhood of the SCE.

To test the conjecture, we simulated (6) and (7) until the beliefs exit from the ball of radius 0.4 around the SCE. We use the same value of $\epsilon = 0.006$ as previously. For $u \geq 2$ we still observe a time path of Lyapunov function V that looks almost deterministic, as in Figure 1. In full accordance with our intuition, as u decreases from 5 to 2, average number of periods needed for a deviation increases by a factor of 4.¹³ On the other hand, divergence is observed for the model with $u = 1$ only for ϵ as high as 0.02. The result supports our message: it is closeness of the slowest eigenvalues of J to the imaginary axis that matters for the observed divergence. At $u = 1$, the mean dynamics is too fast to be destabilized by the real-time learning noise. In addition, nonlinearities, which are much milder in this case, ensure that the largest positive eigenvalue of the locally linearized mean dynamics never gets beyond approximately 0.05, which also prevents divergence.

As shown by Evans et al. (2010), SG learning is not scale-invariant, and thus the estimated gain could depend on the measurement units, whereas for RLS this problem does not appear. To test the dependence of our results on scaling, we switch to measuring unemployment and inflation as decimal proportions rather than percentages. Correspondingly, we set u equal to 0.05 and $\sigma_1 = \sigma_2$ to 0.003, which leads to the same problem as before.¹⁴ The “fast” eigenvalue becomes much less extreme at ~ -0.5 , whereas the “slowest” ones get even smaller, of order 10^{-5} . However, as the positive constant trace(ΣP) decreases in absolute value by much more than $\dot{V}(\gamma)$, the negative component of the stochastic Lyapunov function, the divergence result is now observed only for significantly larger values of the gain ϵ , as high as 1×10^{-2} .

4. CONCLUSION

We studied the performance of constant-gain stochastic gradient learning in a Phelps model of monetary policy that has been extensively studied previously for the RLS learning case. We showed that the model dynamics is divergent for all values of a constant gain greater than a very small threshold. This behavior occurs despite the fact that the mean dynamics is stable in a large neighborhood of the SCE. The divergence is caused by the mean dynamics being locally *almost* unstable at the SCE and unstable at many points close to the SCE. Shocks associated with constant-gain learning then take the beliefs into unstable regions often enough to lead to the divergence, which is reversed only for very small gain values. We implemented a Lyapunov function approach to confirm the intuition described in the preceding and showed that it is possible to infer the behavior of the SRA from the expected value of the derivative of the Lyapunov function $\mathcal{L}V$ along the SRA trajectory. We believe this to be the first application of the stochastic Lyapunov function approach to the analysis of destabilization of asymptotically stable adaptive learning algorithms under real-time learning dynamics in the economic literature.

The behavior of SG learning in real time leads us to express a warning. For constant gain values that are not very small, checking that the mean dynamics is asymptotically stable is not enough to guarantee that the learning algorithm in real time is “well behaved”; moreover, checking that the mean dynamics trajectories are stable in a large region is not enough either. If many eigenvalues of the mean dynamics are close to the imaginary axis and thus the mean dynamics is very slow, the SRA might exhibit *divergent* behavior despite *convergent* mean dynamics. This discrepancy disappears for “sufficiently small” values of the gain, as implied by the theory developed in EH, but these values might prove to be empirically irrelevant.

The mechanism that leads to destabilization of the convergent mean dynamics by noise related to real-time learning under realistic values of the constant gain is rather general and might manifest itself in RLS learning cases as well, when some of the eigenvalues of the E-stability ODE are close to the imaginary axis. We also believe that this mechanism is not specific to constant-gain learning and would appear in other perpetual learning algorithms such as the Kalman filter, provided the mean dynamics possesses several stable yet very small eigenvalues.

Bullard and Mitra (2002) argued that a good monetary policy rule should deliver equilibrium that is E-stable and thus learnable under RLS. Unfavorable outcomes of simple learning behavior in the Phelps model, both under RLS (studied by CWS and others) and under SG, suggest that a desirable policy rule should also be “well-behaved” in real time for a broad class of learning algorithms. Finding rules satisfying these additional restrictions is a topic for further research.

NOTES

1. Evans et al. (2010) report simulations of the “static” version of this model under CG SG learning.
2. Note that at the SCE, $E[g(\bar{\gamma}, \xi)] = 0$ because of (5).
3. In fact, EH show that approximation (9) holds for any initial beliefs $\tilde{\gamma}$, not just $\bar{\gamma}$.
4. The intuition for the existence of slow eigenvalues is that if the ratio σ^2/u is low, the rows of the mean dynamics Jacobian J become almost linearly dependent. Alternatively, the noise is weak relative to the average value of unemployment and inflation at the SCE, and the regressors in the government’s Phillips curve (4) are almost collinear. This feature is independent of inclusion of the lags of inflation and unemployment in the government’s misspecified Phillips curve.
5. The same trend is observed for simulation runs that are driven by other sequences of shocks. If we observe the simulations even longer, the beliefs γ eventually reach values at which the stochastic process for X_n loses stationarity and the government’s control problem becomes nonstabilizable, making derivation of an optimal policy impossible.
6. To emphasize the difference between the mean dynamics and real-time learning, we initiate the mean dynamics simulations of (8) at the last point of simulation of (6) and (7).
7. Informally, if there is a Lyapunov function V that is zero at the stationary point O of an ODE and strictly positive outside of it, and the full derivative of V along the ODE solutions, \dot{V} , is always negative, then the ODE solutions approach O as $t \rightarrow \infty$, and thus O is asymptotically stable. The stochastic analog of \dot{V} is an expectation of this derivative, usually denoted $\mathcal{L}V$. If $\mathcal{L}V$ is always negative, then V is expected to decrease over time, which for V positive everywhere but at O proves that as $t \rightarrow \infty$, the solution converges to O with probability one.
8. For a symmetric matrix J , this construction results in P equal to the identity matrix, and the Lyapunov function being the squared Euclidean distance from $\bar{\gamma}$.

9. This effect is not observed if the linear approximation (9) remains valid everywhere. In this case, the negative contribution of the drift term to $\mathcal{L}V$, $(\gamma - \bar{\gamma})^T \cdot (PJ + J^T P) \cdot (\gamma - \bar{\gamma})$, overcomes a positive but constant diffusion contribution $\epsilon \cdot \text{trace}(\Sigma P)$ for large values of $\gamma - \bar{\gamma}$.

10. A variance–covariance matrix of the beliefs along a typical simulation trajectory has a dominant component that explains more than 90% of total variance for sufficiently small values of the constant gain. The eigenvector associated with this component is essentially collinear with the subspace spanned by the two slowest eigenvalues of the linearized mean dynamics: the angle between them is approximately 3° . Therefore, over time, most of the movement in the simulated trajectory occurs along the directions associated with the two slowest eigenvalues of J .

11. For a solution of the ODE $\dot{\gamma} = J(\gamma - \bar{\gamma})$, the angle between the current deviation $\gamma - \bar{\gamma}$ and its speed of change is determined by

$$(\gamma - \bar{\gamma})^T J (\gamma - \bar{\gamma}) = (\gamma - \bar{\gamma})^T \frac{J + J^T}{2} (\gamma - \bar{\gamma}) = (\gamma - \bar{\gamma})^T J^{\text{SYM}} (\gamma - \bar{\gamma}).$$

The symmetric part of J is unstable: its largest eigenvalue equals +10.35. Thus, any deviation in the direction of the corresponding eigenvector, w_1^{SYM} , is initially strongly amplified. w_1^{SYM} is essentially collinear to the subspace spanned by the “slow” eigenvectors of J in which the mean dynamics trajectory spends most of the time; therefore such amplification is very likely to occur.

12. This is a multivariate Gaussian distribution $N(\bar{\gamma}, \epsilon C)$, with C given by

$$C = \int_0^\infty e^{sD_\gamma h(\bar{\gamma})} \Sigma e^{sD_\gamma^T h(\bar{\gamma})} ds.$$

For details, see EH, Theorem 7.9.

13. Decreasing χ also makes the “slowest” eigenvalues faster, thus raising the number of periods needed for a deviation. Increasing $\sigma_{\{1,2\}}$ makes mean dynamics faster but at the same time increases variance of shocks hitting (6) and (7), thus making destabilization easier.

14. We thank an anonymous referee of this paper for suggesting this exercise.

REFERENCES

Berardi, M. (2013) Escape dynamics and policy specification. *Macroeconomic Dynamics* 17, 123–142.
 Branch, W.A., T. Davig, and B. McGough (2013) Adaptive learning in regime-switching models. *Macroeconomic Dynamics* 17, 998–1022.
 Branch, W. and G. Evans (2006) A simple recursive forecasting model. *Economics Letters* 91, 158–166.
 Bullard, J., G.W. Evans and S. Honkapohja (2010) A model of near-rational exuberance. *Macroeconomic Dynamics* 14, 166–188.
 Bullard, J. and K. Mitra (2002) Learning about monetary policy rules. *Journal of Monetary Economics* 49(6), 1105–1129.
 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, NJ: Princeton University Press.
 Evans, G.W. and S. Honkapohja (2009) Robust learning stability with operational monetary policy rules. In K. Schmidt-Hebbel and C. E. Walsh (eds.), *Monetary Policy under Uncertainty and Learning*, pp. 145–170. Santiago: Central Bank of Chile.
 Evans, G.W., S. Honkapohja, and N. Williams (2010) Generalized stochastic gradient learning. *International Economic Review* 51, 237–262.
 Hommes, C. and T. Lux (2013) Individual expectations and aggregate behavior in learning-to-forecast experiments. *Macroeconomic Dynamics* 17, 373–401.
 Loparo, K.A. and X. Feng (1999) Stability of stochastic systems. In W.S. Levine (ed.), *The Control Handbook*, pp. 1105–1126. CRC Press with IEEE Press.
 Mao, X. (1994) Stochastic stabilization and destabilization. *Systems and Control Letters* 23, 279–290.

Milani, F. (2008) Learning, monetary policy rules, and macroeconomic stability. *Journal of Economic Dynamics and Control* 32, 3148–3165.

Orphanides, A. and J.C. Williams (2007) Inflation targeting under imperfect knowledge. *FRBSF Economic Review* 2007, 1–23.

Orphanides, A. and J.C. Williams (2009) Imperfect knowledge and the pitfalls of optimal control monetary policy In K. Schmidt-Hebbel and C.E. Walsh (ed.), *Monetary Policy under Uncertainty and Learning*, pp. 115–144. Santiago: Central Bank of Chile.

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

Slobodyan, S. and R. Wouters (2012) Learning in an estimated medium-scale DSGE model. *Journal of Economic Dynamics and Control* 36, 26–46.

Williams, N. (2009) *Escape Dynamics in Learning Models*. Mimeo, University of Wisconsin–Madison.

APPENDIX: LYAPUNOV FUNCTION IN DETERMINISTIC AND STOCHASTIC SYSTEMS

The following short description of Lyapunov’s second (direct) method, both deterministic and stochastic, is based on Loparo and Feng (1999). For a deterministic ODE such as (8), consider a continuous function $V(\gamma)$ such that $V(\bar{\gamma}) = 0$ and $V(\gamma) > 0$ for $\gamma \neq \bar{\gamma}$. Suppose that for some c the set $\Omega_c = \{\gamma : V(\gamma) < c\}$ is bounded and V has continuous first partial derivatives in this set. As $c \rightarrow 0$, the set Ω_c becomes $\{\bar{\gamma}\}$. If the total derivative of V along a solution of (8), defined as

$$\dot{V}(\gamma) = \frac{dV(\gamma)}{dt} = h^T(\gamma) \frac{\partial V(\gamma)}{\partial \gamma} = -\kappa(\gamma),$$

satisfies $-\kappa(\gamma) < 0$ for all $\gamma \in \Omega_c \setminus \{\bar{\gamma}\}$, then V is a monotonically decreasing function of time. This implies that $V(\gamma_t) \rightarrow 0$ as $t \rightarrow \infty$, and thus $\gamma_t \rightarrow \bar{\gamma}$ as $t \rightarrow \infty$. A useful relation is given by

$$0 < V(\gamma_0) - V(\gamma_t) = \int_0^t \kappa(\gamma_s) ds.$$

The stochastic analog of the Lyapunov function is introduced in the following way. Suppose that the system, such as (10), is Markov so that the solution process γ_t is a strong, time-homogeneous Markov process. Then the infinitesimal generator of the process γ_t is defined by

$$\mathcal{L}V(\gamma_0) = \lim_{t \rightarrow 0} \frac{E_{\gamma_0} [V(\gamma_{\Delta t}) - V(\gamma_0)]}{\Delta t}.$$

Thus, $\mathcal{L}V$ is a natural analog of the derivative of V along the solution path, \dot{V} , in the deterministic case. Suppose that for some Lyapunov function V we have $\mathcal{L}V(\gamma) = -\kappa(\gamma) < 0$. Then the following relation holds:

$$0 < V(\gamma_0) - E_{\gamma_0} [V(\gamma_t)] = E_{\gamma_0} \int_0^t \kappa(\gamma_s) ds = -E_{\gamma_0} \int_0^t \mathcal{L}V(\gamma_s) ds < +\infty,$$

and for $t, s > 0$,

$$E_{\gamma_s} [V(\gamma_{t+s})] - V(\gamma_s) < 0 \text{ almost surely.}$$

The last expression means that V is a supermartingale, which could be used to infer that γ_t converges to $\bar{\gamma}$. A similar relation (with changed sign) holds when $-\kappa(\gamma) > 0$, in which case the Lyapunov method could be used to establish divergence of solution paths a.s.

We will not be able to prove stability or instability of the SCE, because the diffusion term in (10) does not go to zero as $\gamma \rightarrow \bar{\gamma}$. However, we utilize the relation between the integral of $\mathcal{L}V$ and the expected increment of V in our discussion of the likely behavior of the simulation paths of (6) and (7).