

Galimberti, Jaqueson K.

Working Paper

An approximation of the distribution of learning estimates in macroeconomic models

KOF Working Papers, No. 453

Provided in Cooperation with:

KOF Swiss Economic Institute, ETH Zurich

Suggested Citation: Galimberti, Jaqueson K. (2019) : An approximation of the distribution of learning estimates in macroeconomic models, KOF Working Papers, No. 453, ETH Zurich, KOF Swiss Economic Institute, Zurich, <https://doi.org/10.3929/ethz-b-000332885>

This Version is available at:

<https://hdl.handle.net/10419/204790>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.



Working Paper

An approximation of the distribution of learning estimates in macroeconomic models

Author(s):

Kingeski Galimberti, Jaqueson

Publication Date:

2019-03

Permanent Link:

<https://doi.org/10.3929/ethz-b-000332885> →

Rights / License:

[In Copyright - Non-Commercial Use Permitted](#) →

This page was generated automatically upon download from the [ETH Zurich Research Collection](#). For more information please consult the [Terms of use](#).

KOF Swiss Economic Institute

An approximation of the distribution of learning estimates in macroeconomic models

Jaqueson K. Galimberti

KOF Working Papers, No. 453, March 2019

KOF

ETH Zurich
KOF Swiss Economic Institute
LEE G 116
Leonhardstrasse 21
8092 Zurich, Switzerland

Phone +41 44 632 42 39
Fax +41 44 632 12 18
www.kof.ethz.ch
kof@kof.ethz.ch

An approximation of the distribution of learning estimates in macroeconomic models*

Jaqueson K. Galimberti[†]
ETH Zurich

Abstract

Adaptive learning under constant-gain allows persistent deviations of beliefs from equilibrium so as to more realistically reflect agents' attempt of tracking the continuous evolution of the economy. A characterization of these beliefs is therefore paramount to a proper understanding of the role of expectations in the determination of macroeconomic outcomes. In this paper we propose a simple approximation of the first two moments (mean and variance) of the asymptotic distribution of learning estimates for a general class of dynamic macroeconomic models under constant-gain learning. Our approximation provides renewed convergence conditions that depend on the learning gain and the model's structural parameters. We validate the accuracy of our approximation with numerical simulations of a Cobweb model, a standard New-Keynesian model, and a model including a lagged endogenous variable. The relevance of our results is further evidenced by an analysis of learning stability and the effects of alternative specifications of interest rate policy rules on the distribution of agents' beliefs.

Keywords: expectations, adaptive learning, constant-gain, policy stability.

JEL codes: D84, E03, E37, C62, C63.

Declaration of interest: none.

*Paper accepted for publication at the **Journal of Economic Dynamics & Control**. This research was presented at the 24th International Conference Computing in Economics and Finance, held in Milan, 2018. I gratefully acknowledge the comments and evaluations provided by attendees of my presentation, an Associate Editor, and two anonymous reviewers. Any remaining errors are my own.

[†]Contact e-mail: galimberti@kof.ethz.ch.

Personal website: <https://sites.google.com/site/jkgeconoeng/>.

Correspondence address: KOF Swiss Economic Institute, ETH Zurich, LEE G 116, Leonhardstrasse 21, 8092 Zürich, Switzerland. Phone: +41 44 6328529. Fax: +41 44 6321218.

1 Introduction

The modeling of expectations is a central issue of dynamic macroeconomics. Adaptive learning proposes that agents form expectations as econometricians, replacing the assumption of an instantaneous adjustment of beliefs towards the model's implied rational expectations equilibrium (REE) by a more realistic recursive learning mechanism. According to such a learning algorithm, agents update their beliefs about the economy they inhabit as new information becomes available in real-time. These potentially out-of-equilibrium beliefs are then used by agents to make their economic decisions, generating a feedback loop in the determination of expectations and actual outcomes. One key question in this context is under which conditions this process would eventually converge to an equilibrium, and a considerable literature has been devoted to provide an answer to this question.

Importantly, the type of convergence of this learning process crucially depends on how agents are assumed to adjust their beliefs over time, a behavior that is regulated by the specification of a sequence of learning gains. In short, the learning gain determines how quickly a given information is incorporated into agents' beliefs. Whereas convergence to a non-stochastic equilibrium is possible when learning is assumed to cease in the long run, a situation reflected in the traditional assumption of decreasing-gains, such a scenario has restricted value for applied purposes. Realistically, adaptive learning should be allowed to persist over time so as to reflect agents' capability of tracking the continuous evolution of their economic environment. However, analysis of convergence in this latter situation poses a much more intricate challenge than in the former. Namely, under constant-gain learning beliefs never converge to a fixed equilibrium point, but may, instead, converge to a limiting probability distribution. In spite of its importance, few attempts have been made in the literature to provide ready-to-use analytical expressions describing this distribution of learning estimates, and we attempt to fill this gap in this paper.

Our main contribution is to propose a simple approximation of the first two moments of the asymptotic distribution of learning estimates for a general class of dynamic macroeconomic models under constant-gain learning. Particularly, we derive analytical expressions for the mean and the variance of such a distribution based entirely on standard time series statistics. Our approach is in contrast with the stochastic approximation approach of Benveniste et al. (1990), which has already been used by Evans and Honkapohja (2001, Section 7.4) to provide a characterization of the distribution of learning estimates under the assumption of a "small" gain; as we will show in an application, the restriction on the magnitude of the learning gain seems to harm the accuracy of their approximation. Moreover, our approximation provides renewed conditions for the convergence of the

learning process, and we show that these conditions are more stringent than the usual E-stability conditions derived under assumptions of arbitrarily small or decreasing gains. Importantly, our convergence conditions depend on the interaction between the learning gain and the model's structural parameters.

Our approximation relies on a simplifying assumption that reduces substantially the complexity of statistical interactions in the recursive learning expressions. In spite of being an approximation, we argue that the benefits of having such analytical expressions outweigh the drawbacks of their potential inaccuracies, which we show to be small with three numerical applications, namely: a Cobweb model of prices, a forward-looking New-Keynesian (NK) model of aggregate inflation and output, and a univariate model with a lagged endogenous variable. Overall, we find that our analytical expressions provide close approximations to the mean and the dispersion of simulated learning estimates. We also evaluate the relevance of our results for policy design by considering two alternative Taylor rules under the NK model: one responding to contemporaneous data, and the other responding to expectations. We find that the stability of the second, in spite of being more realistic, is more sensitive to the learning gain than the data-based rule. Our analysis of constant-gain learning stability parallels that of Evans and Honkapohja (2009), who derived similar stability conditions focusing on the case of steady-state learning, i.e., when the REE takes the form of a stochastic steady state value. Here, our approximation of the variance of the learning estimates provides additional insights on the long run dynamics of this model under persistent learning. Particularly, we find that the dispersion of the learning estimates is affected by the specification of the policy rules, providing an implicit learning channel through which policy can affect perceived uncertainty about the economy.

However, the accuracy of our approximation tends to deteriorate as the model's parameter values draw closer to the threshold values determining stability of the learning process, a result that can be associated to the stochastic nature inherent to the deterministically approximated learning process. Besides, in contrast to the more general stochastic approximation approach, our method is geared solely towards asymptotic analysis, hence disregarding the transient dynamic effects of learning over finite stretches of time, which, in turn, can lead to other interesting types of dynamics, such as escapes (see, e.g., Cho et al., 2002; Williams, 2018). Another important limitation of this paper is that our framework precludes non-linearities between agents' expectations, a feature commonly emerging in forward-looking macroeconomic models that assume a simultaneous determination of expectation and equilibrium variables (see Woodford, 2003, Chapter 3 for a discussion in the context New Keynesian models). Models with multiple equilibria are also ruled out in such a linear expectations framework.

Our results are relevant to other applied strands of the literature too. Particularly, the estimation of models with adaptive learning requires an *a priori* definition of a range of possible gain values; that is the case for both Bayesian (e.g., Milani, 2007) and classical (e.g., Chevillon et al., 2010) estimation approaches. Besides, the stability of the model depends on the interaction between the learning gain and other structural and policy parameter estimates. In that context our analytical expressions can improve the robustness of estimation to learning instabilities by conditioning the learning gain upper bound on the model's structural parameters. Similarly, recent approaches proposing the use of time-varying gains (e.g., Milani, 2014; Berardi and Galimberti, 2017) can benefit with our convergence-based upper bounds on the learning gain.

The remainder of this paper proceeds as follows. In section §2 we outline the general model under which we derive our approximation, presented in section §3. We then proceed with numerical applications in section §4, and conclude with some remarks in section §5. Details about derivations and supplementary statistics are provided in the Appendix.

2 Dynamic modeling framework

2.1 General environment

Let the model's actual law of motion (ALM), i.e., the equation(s) describing the determination of the model's endogenous variable(s), be given by

$$\mathbf{y}_t = \mathbf{z}_t' (\mathbf{A} + \Phi_{t-1} \mathbf{B}) + \mathbf{u}_t' \mathbf{C}, \quad (1)$$

where: $\mathbf{y}_t = [y_{1,t}, \dots, y_{n,t}]'$ contains the endogenous variables; $\mathbf{z}_t = [z_{1,t}, \dots, z_{p,t}]'$ contains the predetermined variables, assumed to be stationary, and which may include lags of the endogenous variables (this is further discussed in Remark 10 below); $\Phi_t = [\phi_{1,t}, \dots, \phi_{p,t}]'$ contains the coefficients estimates that form agents' perceived law of motion (PLM) of the variables for which expectations are required; $\mathbf{u}_t = [u_{1,t}, \dots, u_{n,t}]'$ contains i.i.d. zero mean random disturbances with variance given by $E[\mathbf{u}_t \mathbf{u}_t'] = \sigma_u^2 \mathbf{I}_{(n)}$; and, $\mathbf{A}_{(p \times n)}$, $\mathbf{B}_{(n \times n)}$ and $\mathbf{C}_{(n \times n)}$ are matrices of conformable sizes containing model parameters.

Agents are assumed to form expectations according to a PLM of the form

$$\mathbf{y}_t = \mathbf{z}_t' \Phi_{t-1} + \mathbf{e}_t', \quad (2)$$

where $\mathbf{e}_t = [e_{1,t}, \dots, e_{n,t}]'$ contains agents' forecasting errors.

We illustrate the usefulness of this framework with two standard economic models.

Example 1 (Cobweb model). The Cobweb model is a partial equilibrium model of demand and supply in which production is assumed to have a time lag. Hence, producers have to form expectations of next period's price in order to optimize their current period's production decisions. The demand and supply equations are usually formulated as

$$\begin{aligned} d_t &= m_1 - m_2 p_t + v_{1t}, \\ s_t &= r_1 + r_2 p_t^e + r_3 \omega_{t-1} + v_{2t}, \end{aligned}$$

respectively, where $m_2, r_2 > 0$, v_{1t} and v_{2t} are unobserved white noise shocks, and supply may also be affected by an observable shock, ω_{t-1} , e.g., an input price. Assuming market clearing, $d_t = s_t$, we obtain the reduced form

$$p_t = \mu + \alpha p_t^e + \delta \omega_{t-1} + \eta_t,$$

where $\mu = (m_1 - r_1) / m_2$, $\alpha = -r_2 / m_2 < 0$, $\delta = -r_3 / m_2$, and $\eta_t = (v_{1t} - v_{2t}) / m_2$. Assuming agents condition their expectations on a model containing a constant plus the observable exogenous shock, the PLM is given by

$$p_t^e = \phi_{1,t-1} + \phi_{2,t-1} \omega_{t-1},$$

and the corresponding ALM is given by equation (1) with $\mathbf{y}_t = p_t$, $\mathbf{z}_t = [1, \omega_{t-1}]'$, $\Phi_t = [\phi_{1,t}, \phi_{2,t}]'$, $\mathbf{u}_t = u_{1,t}$ with $\sigma_u^2 = 1$, and $\mathbf{A} = [\mu, \delta]'$, $\mathbf{B} = \alpha$, $\mathbf{C} = \sigma_\eta$.

Remark 1. Notice that autocorrelated shocks can be easily incorporated in this framework by augmenting the vectors of endogenous variables and random disturbances together with the appropriate adjustments to the matrices of parameters. For example, in the context of the Cobweb model above, if input prices are assumed to follow a first-order autocorrelated process given by

$$\omega_t = \varpi + \rho \omega_{t-1} + \varepsilon_t,$$

then the ALM would be given by equation (1) with $\mathbf{y}_t = [p_t, \omega_t]$, $\mathbf{z}_t = [1, \omega_{t-1}]'$, $\Phi_t = \begin{bmatrix} \phi_{1,t} & 0 \\ \phi_{2,t} & 0 \end{bmatrix}$, $\mathbf{u}_t = [u_{1,t}, u_{2,t}]'$, and $\mathbf{A} = \begin{bmatrix} \mu & \varpi \\ \delta & \rho \end{bmatrix}$, $\mathbf{B} = \begin{bmatrix} \alpha & 0 \\ 0 & 0 \end{bmatrix}$, $\mathbf{C} = \begin{bmatrix} \sigma_\eta & 0 \\ 0 & \sigma_\varepsilon \end{bmatrix}$.

Example 2 (New-Keynesian model). A standard New-Keynesian macroeconomic model can be derived from a simple intertemporal general equilibrium model with sticky prices (see Woodford, 2003), which in log-linearized reduced form is given by

$$x_t = -\varphi (i_t - \pi_{t+1}^e) + x_{t+1}^e + g_t, \quad (3)$$

$$\pi_t = \lambda x_t + \beta \pi_{t+1}^e + u_t, \quad (4)$$

where x_t stands for the output gap, π_t for the inflation rate, i_t for the nominal interest rate, all expressed in deviations from the steady state equilibrium; g_t and u_t are exogenous disturbances which can be interpreted as autonomous expenditures changes and a cost-push shock, respectively; (3) and (4) are commonly referred as the economy's intertemporal IS (investment/saving) and AS (aggregate supply) equations, respectively, where the parameters are related to primitive structural features: $\beta \in (0, 1)$ and $\varphi^{-1} > 0$ reflect the representative household's discount factor and elasticity of intertemporal substitution, respectively; and $\lambda > 0$ is related to the degree of price stickiness in firms' price setting behavior.

The model is closed with the definition of central bank's behavior, i.e., an interest rate setting rule, often referred as the Taylor rule. Different formulations have been advocated for that purpose and the analysis of determinacy and stability of such rules has attracted great interest in the literature (see Bullard and Mitra, 2002; Evans and Honkapohja, 2003). Two of the most common specifications are: (i) a *data-based policy rule*, as given by equation (5), where the central bank is assumed to observe and react to current observations of inflation and the output gap; and, (ii) an *expectations-based policy rule*, as given by equation (6), where the central bank responds instead to current expectations, formed under the same informational restrictions as those determining dynamics in the economy.

$$i_t = \chi_\pi \pi_t + \chi_x x_t, \quad (5)$$

$$i_t = \chi'_\pi \pi_t^e + \chi'_x x_t^e, \quad (6)$$

Solving the model for these alternative policy rules we obtain two different model specifications determining the dynamics of inflation and output gap as a function of forward looking expectations and unobserved disturbances. Hence, both policies imply a PLM where agents need to learn constants, which are equal to zero by the definition of the variables as deviations from the model's steady state equilibrium. More specifically, letting the superscripts (i) and (ii) denote the models resulting from the incorporation of the policy rules (5) and (6), respectively, into the reduced form equations (3)-(4), the New-Keynesian model can be translated into our general formulation given by equation (1) by setting: $\mathbf{y}_t = [x_t \ \pi_t]'$, $\mathbf{z}_t^{(i)} = \mathbf{z}_t^{(ii)} = 1$, $\mathbf{u}_t = [g_t \ u_t]'$, also adjusting the matrix of learning coefficients accordingly, i.e., $\Phi_t^\square = [\phi_{x,t}^\square \ \phi_{\pi,t}^\square]$ with $\square = \{(i), (ii)\}$; finally, the matrices of model parameters are given by

$$\begin{aligned} \mathbf{A}^{(i)} &= \mathbf{A}^{(ii)} = \begin{bmatrix} 0 & 0 \end{bmatrix}, \\ \mathbf{B}^{(i)} &= \psi \begin{bmatrix} 1 & \lambda \\ \varphi(1 - \beta\chi_\pi) & \beta + \lambda\varphi(1 - \beta\chi_\pi) \end{bmatrix}, \\ \mathbf{B}^{(ii)} &= \begin{bmatrix} 1 - \varphi\chi'_x & \lambda(1 - \varphi\chi'_x) \\ \varphi(1 - \chi'_\pi) & \beta + \lambda\varphi(1 - \chi'_\pi) \end{bmatrix}, \end{aligned}$$

$$\mathbf{C}^{(i)} = \begin{bmatrix} \psi & \lambda\psi \\ -\varphi\chi_\pi\psi & 1 - \lambda\varphi\chi_\pi\psi \end{bmatrix}, \quad \mathbf{C}^{(ii)} = \begin{bmatrix} 1 & \lambda \\ 0 & 1 \end{bmatrix},$$

where $\psi = (1 + \varphi(\chi_x + \chi_\pi\lambda))^{-1}$.

Remark 2. Although non-linearities between structural parameters can be clearly incorporated into the matrices of model parameters, the framework of equation (1) explicitly precludes non-linearities between agents' PLM estimates in the determination of the endogenous variables; particularly, models with contemporaneous expectations, where equilibrium values of the endogenous variables and expectations are determined simultaneously (see Evans and Honkapohja, 2001, section 10.3), are difficult to fit in the ALM representation of equation (1). Such non-linearities can also emerge in models with informational gaps longer than one period, i.e., when the lag-lead difference between the period of information on which expectations are based and the period for which the expectations are made is greater than one.

2.2 Expectations and learning

The implied dynamics of an economy represented by equation (1) is clearly influenced by the assumption of how expectations are determined. Under RE the PLM's coefficients are assumed to be constant and determined so as to make expectations consistent with observed actuals, i.e.,

$$\Phi^{RE} = \mathbf{A} (\mathbf{I}_{(n)} - \mathbf{B})^{-1} \quad (7)$$

and $\mathbf{e}'_t \sim \mathbf{u}'_t \mathbf{C}$, where $(\mathbf{I}_{(n)} - \mathbf{B})$ is assumed to be invertible. Hence, under RE the evolution of the endogenous variables is entirely determined by the model parameters and the behavior of the forcing variables and disturbance shocks included in the definitions of \mathbf{z}_t and \mathbf{u}_t , respectively.

Under adaptive learning, in contrast, agents' perceptions about the economy's law of motion may deviate from the model's implied self-referential equilibrium, and these deviations will also matter in the determination of the endogenous variables. More specifically, agents are assumed to hold estimates of their PLM's coefficients and to update these estimates using a learning algorithm. The most common choice for that purpose (see, e.g., Berardi and Galimberti, 2014, on merits and alternatives), is given by a Least Squares algorithm of the following form:

Algorithm 1 (LS). *For a PLM given by equation (2), the Least Squares estimates Φ_t , conditional on observations up to time t , are given by*

$$\Phi_t = \Phi_{t-1} + \gamma \mathbf{R}_t^{-1} \mathbf{z}_t (\mathbf{y}_t - \mathbf{z}'_t \Phi_{t-1}), \quad (8)$$

$$\mathbf{R}_t = \mathbf{R}_{t-1} + \gamma (\mathbf{z}_t \mathbf{z}'_t - \mathbf{R}_{t-1}), \quad (9)$$

where γ is the learning gain parameter, here assumed to be constant, and \mathbf{R}_t stands for an estimate of the regressors matrix of second moments.

One major question in the literature of adaptive learning is whether the learning estimates will converge to the model's corresponding REE, and the conditions required for that convergence to take place have long been established for the case of decreasing gains (see, e.g., Marcet and Sargent, 1989). The LS recursions, together with equation (1), constitute a self-referential system, where the evolution of \mathbf{y}_t both determines and is determined by the evolution of agents' beliefs according to their estimates of Φ_t . This feedback mechanism constitutes the basis for the definition of the E-stability principle, which states that the stability of equilibria under (decreasing-gain) LS learning is governed by the mapping between the PLM and the ALM (see Evans and Honkapohja, 2001). In our context, such mapping is given by

$$\mathbf{T}(\Phi) = \mathbf{A} + \Phi\mathbf{B}, \quad (10)$$

and E-stability requires that all eigenvalues of the Jacobian of \mathbf{T} (i.e., \mathbf{B}) have real parts less than one.

In order to derive model predictions under learning one has to characterize the dynamics of the learning estimates in the self-referential context described above. Under the assumption of constant-gain learning, nevertheless, convergence to a non-stochastic point, such as the REE, can be promptly ruled out: in a stochastic environment, random shocks will recurrently disturb the convergence of agents' beliefs to a fixed point. Yet, the model's asymptotic dynamics can still be useful if the learning estimates converge to a limiting probability distribution¹. More formally, we are interested in the notion of convergence in distribution, also known as weak convergence, i.e., letting \mathcal{F}_t and \mathcal{F} stand for the time t and the invariant distribution functions of the learning estimates, respectively, convergence is obtained if $\lim_{t \rightarrow \infty} \mathcal{F}_t(\Phi) = \mathcal{F}(\Phi)$. In fact, Evans and Honkapohja (2001, section 7.4) show that, under certain conditions, the asymptotic distribution of the learning estimates converges to a normal distribution, which has the attractive property of being fully characterized by a (vector of) mean(s) and a (co)variance (matrix, in case of PLMs with more than one coefficient). Nevertheless, although we are mostly concerned with models satisfying the conditions for asymptotic normality, it is important to note that constant gain learning models often violate these conditions, leading to the emergence of features such as escapes and fat tails (see, e.g., Cho et al., 2002). Hence, we proceed by attempting to characterize the asymptotic behavior of Φ_t , particularly focusing on its first and the second moments,

¹The focus on asymptotic dynamics can also be motivated by the fact that at any point in time, apart from periods following a structural break, learning can be thought as an ongoing process that has already settled in its long run operative state, i.e., a state where recursive updates are still responsive to perturbations but remaining in the vicinities of an equilibrium.

i.e., $\vec{E}[\Phi_t]$ and $\vec{E}[\Phi_t \Phi_t']$, where we let $\vec{E}[\bullet]$ stand for $\lim_{t \rightarrow \infty} E[\bullet]$, and $E[\bullet]$ stands for the unconditional expected value in the stochastic process sense of that operator of averaging across multiple realizations of sequences of learning estimates.

3 Approximating the distribution of learning estimates

3.1 Preliminaries

In order to characterize the first and second moments of the implied distribution of the learning estimates in our dynamic self-referential framework we introduce the ALM of (1) in equation (8) and solve it recursively. For that purpose, it is convenient to define $\theta_t = \text{vec}(\Phi_t)$, in which case we have that (see Appendix A.1)

$$\theta_t = \mathbf{F}_t(t) \theta_0 + \gamma \sum_{i=0}^{t-1} \mathbf{F}_t(i) (\mathbf{P}_{t-i} \text{vec}(\mathbf{A}) + \mathbf{Q}_{t-i} \text{vec}(\mathbf{C})), \quad (11)$$

where

$$\mathbf{F}_t(k) = \begin{cases} \prod_{j=0}^{k-1} \mathbf{H}_{t-j}, & \text{for } k > 0, \\ \mathbf{I}_{(np)}, & \text{otherwise.} \end{cases} \quad (12)$$

$$\mathbf{H}_t = \mathbf{I}_{(np)} - \gamma (\mathbf{I}_{(n)} - \mathbf{B}') \otimes (\mathbf{R}_t^{-1} \mathbf{z}_t \mathbf{z}_t'), \quad (13)$$

$$\mathbf{P}_t = \mathbf{I}_{(n)} \otimes (\mathbf{R}_t^{-1} \mathbf{z}_t \mathbf{z}_t'), \quad (14)$$

$$\mathbf{Q}_t = \mathbf{I}_{(n)} \otimes (\mathbf{R}_t^{-1} \mathbf{z}_t \mathbf{u}_t'). \quad (15)$$

As motivated above we are interested in expressions for the asymptotic first and second moments of the learning estimates. Before moving to our proposed approximation we illustrate how to obtain such expressions solely on the basis of time series statistics for the simplest case of a “guess-the-average” model, where an approximation is in fact unnecessary.

Example 3 (Guess-the-average model). Let the current value of a variable of interest, y_t , depend on a constant, α , plus a fraction, β , of the average value expected for that variable in the next period, y_{t+1}^e , plus a mean zero random shock, u_t , i.e.,

$$y_t = \alpha + \beta y_{t+1}^e + u_t.$$

Assuming agents condition their forecasts on a constant, ϕ_{t-1} , estimated using a constant-gain LS algorithm using past observations of y_t , this model’s ALM is

easily translated into the form of equation (1) by setting $\mathbf{y}_t = y_t$, $\mathbf{z}_t = 1$, $\Phi_t = \phi_t$, $\mathbf{u}_t = u_t$, and $\mathbf{A} = \alpha$, $\mathbf{B} = \beta$, $\mathbf{C} = 1$. Letting $|\delta| = |1 - \gamma(1 - \beta)| < 1$, in the form of equation (11) we then have that

$$\phi_t = \delta^t \phi_0 + \gamma \sum_{i=0}^{t-1} \delta^i (\alpha + u_{t-i}),$$

for which the asymptotic mean and variance can be easily evaluated as: $\vec{E}[\phi_t] = \alpha / (1 - \beta)$, which is also equal to this model's REE; and, $\vec{Var}[\phi_t] = \gamma \sigma_u^2 / (1 - \beta) (1 + \delta)$.

Unfortunately, for most other cases of interest, such as multivariate models and models where the PLM includes one or more variable regressors, a straightforward analytical solution is not available. Hence, an approximation is required.

3.2 An approximation

The main difficulty in evaluating the expected values of the learning estimates relates to the expansion of the geometric factor represented by $\mathbf{F}_t(k)$ in equation (11)². In the context of stochastic difference equations, this factor directly affects the transition function of the learning estimates, and has also been the focus of the stochastic approximation approach previously proposed in the literature. Namely, Evans and Honkapohja (2001, Section 7.4) approximate the distribution of the learning estimates by establishing the conditions under which the transitions of a corresponding continuous-time process imply convergence to a stationary distribution; one such a condition is that the learning gain must be arbitrarily small³.

Here, in contrast, we propose to obtain an approximation of the learning invariant distribution by following a time series statistical approach, as illustrated in example 3, yet also adding a simplifying assumption regarding the expected behavior of $\mathbf{F}_t(k)$ in order to cover the general case of multivariate models that can be translated into the ALM of equation (1). Particularly, we notice that, assuming \mathbf{z}_t is stationary, the first moment of \mathbf{R}_t converges asymptotically to $E[\mathbf{z}_t \mathbf{z}_t']$ (see Appendix A.2). Hence, the only time-varying component of $\mathbf{F}_t(k)$, namely, $\mathbf{R}_t^{-1} \mathbf{z}_t \mathbf{z}_t'$, can be seen to converge asymptotically, in expectation, to a constant matrix. That observation lead us to propose the following approximation.

²Also notice that this complication is particular to the case of constant-gain learning: under the traditional assumption of a decreasing gain, where $\gamma_t \rightarrow 0$ as $t \rightarrow \infty$, $\mathbf{F}_t(k)$ becomes asymptotically irrelevant, as $\mathbf{H}_t \xrightarrow{\infty} \mathbf{I}$.

³Specifically, convergence in continuous-time requires that $\gamma \rightarrow 0$, where the continuous-time approximation of the learning estimates is also tied to the gain by letting $\tau_t^\gamma = t\gamma$, and defining $\theta(\tau) = \theta_t$ if $\tau_t^\gamma \leq \tau < \tau_{t+1}^\gamma$.

Assumption 1. *The expected value of $\mathbf{F}_t(k)$ can be asymptotically approximated as*

$$\tilde{\mathbf{F}}(k) = \overrightarrow{E}[\mathbf{F}_t(k)] = (\mathbf{I}_{(np)} - \gamma(\mathbf{I}_{(n)} - \mathbf{B}') \otimes \mathbf{K})^k = \tilde{\mathbf{H}}^k, \quad (16)$$

where \mathbf{K} is a $p \times p$ positive semi-definite matrix.

Remark 3. Underlying the approximation proposed in equation (16) are the assumptions that $\overrightarrow{E}[\mathbf{R}_t^{-1} \mathbf{z}_t \mathbf{z}_t'] = \overrightarrow{E}[\mathbf{R}_t^{-1} \mathbf{z}_t \mathbf{z}_t' \mathbf{R}_{t-l}^{-1} \mathbf{z}_{t-l} \mathbf{z}_{t-l}'] = \mathbf{K}$, $\forall t, l$. Under the stationarity conditions guaranteeing the convergence of \mathbf{R}_t , it is quite clear that, indeed, these quantities converge to constants. To see that take, for instance, the first assumption for an univariate PLM; from the definition of covariances we have that

$$\overrightarrow{E}[R_t^{-1} z_t^2] = \overrightarrow{E}[R_t^{-1}] \overrightarrow{E}[z_t^2] + Cov(R_t^{-1}, z_t^2),$$

where both terms on the right hand side of this equality are constants. It is important to note, however, that the stationarity of \mathbf{z}_t and the convergence of \mathbf{R}_t can be interdependent for the case where lagged endogeneous variables are included in agents' PLM; see Remark 10 below for further discussion about this case.

Remark 4. More importantly, the accuracy of the approximation depends on the accuracy of \mathbf{K} , and an expression for the latter is not readily available. For applied purposes we argue that the identify matrix can be used as a first estimate of \mathbf{K} , although we notice that this estimate will tend to lose accuracy the higher the learning gain. Particularly, in the context of the univariate example above, notice that whereas Jensen's inequality implies that $\overrightarrow{E}[R_t^{-1}] \overrightarrow{E}[z_t^2] \geq 1$, irrespective of the learning gain, the covariance between the inverse of R_t and z_t^2 is negative and scaled by γ (see equation (27), in the Appendix, expanding the summation), attenuating the positive effect that convexity of the inverse has on the value of the first term.

Proposition 1. *Assuming the approximation given by equation (16), and provided that all eigenvalues of $\tilde{\mathbf{H}} = \overrightarrow{E}[\mathbf{H}_t]$ lie inside the unit circle, the mean and the (co)variance of the adaptive learning estimates associated to agents' expectations in the model given by equation (1), and updated according to the LS algorithm given by equations (8) and (9), converge asymptotically to*

$$\overrightarrow{E}[\Phi_t] = \mathbf{A}(\mathbf{I}_{(n)} - \mathbf{B})^{-1}, \quad (17)$$

and

$$Var(vec(\Phi_t)) = \gamma^2 \mathcal{B}, \quad (18)$$

respectively, where

$$vec(\mathcal{B}) = \sigma_u^2 \left(\mathbf{I}_{(n^2 p^2)} - \tilde{\mathbf{H}} \otimes \tilde{\mathbf{H}} \right)^{-1} vec\left((\mathbf{C}' \mathbf{C}) \otimes \overrightarrow{E}[\mathbf{R}_t^{-1}] \right). \quad (19)$$

Proof. The proof follows directly from the derivation of the first and second moments under the referred assumptions, which are detailed in the Appendices A.2 and A.3. \square

Remark 5. Notice that $\vec{E}[\Phi_t] = \Phi^{RE}$, i.e., convergence to RE under constant-gain LS learning in this model requires that the eigenvalues of $\tilde{\mathbf{H}}$ lie inside the unit circle, a condition different than the usual E-stability criterion for the case of a decreasing gain LS (see equation (10)). In fact, it is possible to show that whereas E-stability does not imply convergence under constant-gain learning, convergence according to proposition 1 requires E-stability.

Corollary 1. *E-stability conditions are necessary but not sufficient conditions for convergence under constant-gain learning.*

Proof. Letting $\{\lambda\}$ stand for the set of np eigenvalues of $\tilde{\mathbf{H}}$, we have that

$$\begin{aligned}\{\lambda\} &= \text{eigs}(\mathbf{I}_{(np)} - \gamma(\mathbf{I}_{(n)} - \mathbf{B}') \otimes \mathbf{K}), \\ &= \gamma \text{eigs}((\mathbf{B}' - \mathbf{I}_{(n)}) \otimes \mathbf{K}) + 1.\end{aligned}$$

Furthermore, let $\{\beta_i\}_{i=1}^n$ and $\{\kappa_j\}_{j=1}^p$ stand for the eigenvalues of \mathbf{B} and \mathbf{K} , respectively. From proposition 1, convergence under constant-gain learning requires that $|\{\lambda\}| < 1$, which is equivalent to $1 - 2/\gamma\kappa_j < \beta_i < 1, \forall i, j$. Notice the latter inequality is the same as that required for E-stability (necessity) whereas the former adds a lower bound restriction on the set of stable solutions (E-stability insufficiency). \square

Remark 6. Notice that as $\gamma \rightarrow 0$, $\tilde{\mathbf{H}} \rightarrow \mathbf{I}_{(np)}$ and $\mathcal{B} \rightarrow \mathbf{0}$, thence $\text{Var}(\boldsymbol{\theta}_t) \rightarrow \mathbf{0}$, which is consistent with the implied convergence of learning under decreasing gains.

Remark 7. Although there is some resemblance between our approximation and the generalized stochastic gradient learning rule proposed by Evans et al. (2010), these are not equal. Particularly, Evans et al. (2010) show that replacing \mathbf{R}_t^{-1} by a constant matrix in equation (8) renders an asymptotic approximation of the Kalman filter associated with a model with time-varying parameters. Our approximation, in contrast, is equivalent to replacing \mathbf{R}_t^{-1} by a time-varying matrix that, on average, yields a constant matrix after being post-multiplied by the second moment of the regressors, $\mathbf{z}_t \mathbf{z}_t'$.

Remark 8. An estimate of $\vec{E}[\mathbf{R}_t^{-1}]$ is required to calculate the (co)variance approximation. In principle, such statistic could be calculated using established convergence results for the Kalman filter specification corresponding to the LS algorithm (see Hamilton, 1994, section 13.5, for the convergence results, and Berardi and Galimberti, 2013, for the correspondence with the LS algorithm).

However, that approach quickly turns impracticable for non-trivial models. As a simpler alternative we suggest using $\vec{E} [\mathbf{z}_t \mathbf{z}_t']^{-1}$ as an approximation; in fact, notice that, again under a univariate setup for simplicity, Jensen's inequality implies that $\vec{E} [R_t^{-1}] \geq \vec{E} [R_t]^{-1} = \vec{E} [z_t^2]^{-1}$, i.e., this approximation provides a lower bound to the calculation of the (co)variance of the invariant distribution of the learning estimates.

Remark 9. The notion of stability implied by proposition 1 can be interpreted similarly as that obtained with the analysis of “mean dynamics” employed in the stochastic approximation approach (see, e.g., Williams, 2018). Namely, whereas the stability condition implies sure convergence in the deterministic approximation, it may often be the case that some stochastic realizations of the process, even though satisfying the stability conditions, will diverge. In the context of our proposition, this can be more clearly illustrated by considering a case where the eigenvalues of $\tilde{\mathbf{H}}$ are smaller but close to unity, while the actual realizations of \mathbf{H}_t are still affected by stochastic variation from the state variables (a numeric example is given in the next remark below). Hence, the occurrence of unstable dynamics can increase as the model parameters are drawn closer to the deterministic stability threshold, a result we in fact observe in our numerical simulations in the next section.

Remark 10. The inclusion of lagged endogenous variables in \mathbf{z}_t may pose another challenge to the use of the approximation (16) for the analysis of learning stability. Particularly, the accuracy of the approximation can deteriorate because the stationarity of \mathbf{z}_t , which is required for the convergence of \mathbf{R}_t , would then be affected by the learning estimates. To see that consider the example of a simple univariate model given by

$$y_t = \beta y_t^e + \lambda y_{t-1} + u_t, \quad (20)$$

which can be easily translated into the framework of the paper by assuming agents form their expectations according to a first-order autoregression and by setting $\mathbf{y}_t = y_t$, $\mathbf{z}_t = y_{t-1}$, $\Phi_t = \phi_t$, $\mathbf{u}_t = u_t$, $\mathbf{A} = \lambda$, $\mathbf{B} = \beta$, $\mathbf{C} = \sigma_u$; notice this implies that $H_t = 1 - \gamma(1 - \beta)(R_t^{-1} y_{t-1}^2)$. According to proposition 1, and assuming $K = 1$ for simplicity, the learning estimates would converge to a distribution around the REE mean, given by $\lambda/(1 - \beta)$, as long as $\tilde{H} = 1 - \gamma(1 - \beta) < 1$, which is satisfied for $\gamma > 0$ and $\beta < 1$. Notice, however, that the learning estimates also determine the stationarity of the endogenous variable, which, by being included in the agents PLM with a lag, will also affect the validity of approximation (16). For example, say $\beta = 0.5$ and $\lambda = 0.495$, implying that $\vec{E} [\phi_t] = 0.99$; for a learning gain of $\gamma = 0.02$, proposition 1 implies that the learning estimates will have a standard deviation equal to 0.02 (also assuming $\sigma_u = 1$), already

implying these estimates may spend some time outside the stationarity region, $|\phi_t| < 1$. Hence, additional boundedness restrictions are required in order to guarantee the validity of the approximation. One alternative that has been proposed in the earlier literature (e.g., Marcet and Sargent, 1989) comes with the introduction of a projection facility, which is a mechanism that is coupled to the learning algorithm in order to prevent the estimates from escaping a compact region of the parameter space surrounding the equilibrium point. Another alternative is the imposition of further primitive assumptions on the dynamics of the state variables, particularly constraining the support of \mathbf{u}_t and the (implied) transition of \mathbf{z}_t to be linear (see, e.g., assumptions B.1-B.3 in Evans and Honkapohja, 1998).

4 Numerical applications

We now validate the approximation proposed above with the numerical simulation of some economic models.

4.1 Cobweb model

The Cobweb model has been extensively analyzed in the literature. As described in example 1 on page 5, in this model prices are determined according to

$$p_t = \mu + \alpha p_t^e + \delta \omega_{t-1} + \eta_t, \quad (21)$$

with varying assumptions about the data generating process of ω_t . Particularly of our interest, Evans and Honkapohja (2001, section 14.2) derive an approximation to the invariant distribution of the learning estimates under the special case where ω_t is an iid process and the intercept in equation (21) is assumed to be zero. Whereas the cases with an intercept and more relaxed assumptions about the exogenous shock could be easily accommodated in our framework, for comparative purposes, we focus our numerical analysis on that simpler case too, i.e., with $\mu = 0$ and $\omega_t \sim N(0, \sigma_\omega^2)$.

Under these conditions proposition 1 implies expectations converge to an equilibrium in this model if: (i) $\alpha < 1$ (E-stability condition); and, (ii) $\gamma < 2/(1 - \alpha)$; in which case

$$\vec{E}[\phi_t] = \frac{\delta}{1 - \alpha} = \phi^{RE}, \quad (22)$$

and

$$Var(\phi_t) = \gamma \frac{\sigma_\eta^2}{\sigma_\omega^2} \frac{1}{(1 - \alpha)(2 - \gamma(1 - \alpha))}, \quad (23)$$

where we assume $\vec{E}[R_t^{-1}]$ can be approximated as $\vec{E}[R_t]^{-1} = \vec{E}[\omega_t^2]^{-1} = \sigma_\omega^{-2}$ and $K = 1$. Using a different approximation approach, Evans and Honkapohja

(2001, section 14.2) derive analytical expressions similar to ours, except for two main differences: (i) besides the E-stability condition, their convergence results require the learning gain to be “small”, whilst ours is less arbitrary in that respect; and, (ii) their approximation to the variance of the learning estimates is given by

$$\text{Var}(\phi_t)^{EH} = \gamma \frac{\sigma_\eta^2}{\sigma_\omega^2} \frac{1}{(1 - \alpha) 2}, \quad (24)$$

which will always imply a narrower distribution of the learning estimates than our approximation. This is easy to see by evaluating the ratio between these expressions, i.e., $\text{Var}(\phi_t)^{EH} / \text{Var}(\phi_t) = 1 - \gamma(1 - \alpha)/2 < 1$ for $\gamma > 0$ and $\alpha < 1$; also notice that as γ increases, the difference between the approximations increases, while it decreases as α increases. Another difference between these approximations is the expected effect of α on the dispersion of the learning estimates: whereas equation (24) implies a monotone positive relationship, according to our approximation the variance of the learning estimates can only be expected to increase with α if $\gamma < (1 - \alpha)^{-1}$, but decrease otherwise.

We now validate these approximations numerically by simulating the model. For that purpose we draw 10,000 random samples of $\{\omega_t, \eta_t\}_{t=1}^{1,000}$ from the standard normal distribution, i.e., assuming $\sigma_\omega^2 = \sigma_\eta^2 = 1$ and $\sigma_{\omega\eta} = 0$, and use the model’s ALM to construct Cobweb’s artificial price time series under a constant-gain LS learning mechanism. For the model’s parameters we fix $\delta = 1$ and conduct two main exercises with respect to α and γ : first, fixing $\gamma = 0.10$, we vary α from 0.5 to -18.5 (in -0.5 steps); and, second, fixing $\alpha = -2$, we vary γ from 0.025 to 0.65 (in 0.025 steps).

The mean and variance of the invariant distributions of the learning estimates obtained for these experiments are depicted in figure 1, where we take the 1,000th observation as representative of the long run. In some simulations, particularly as the combinations of parameter values draw closer to the stability threshold, we observe that the convergence of the learning estimates is compromised (see Remark 9 for an explanation). In order to attenuate the effect of such “outliers” on the averaged evaluation of the accuracy of the approximations, we also calculate trimmed statistics of the simulated distributions of learning estimates by discarding estimates that deviate by more than ± 5 standard deviations from the REE-implied mean. For most cases the fraction of discarded estimates remained well below 10% of the simulations, only rising above that figure for α values smaller than -11; see Appendix A.5 for more trimming statistics.

Overall, we find that our proposition provides a closer approximation to the invariant distribution of the learning estimates than that proposed by Evans and Honkapohja (2001, section 14.2), particularly for higher learning gains and lower α ’s. Interestingly, we observe that, as α decreases, the dispersion of the learning

estimates start to increase after a threshold α^* . This is consistent with our approximation, though the threshold predicted by proposition 1 is still not accurate: from equation (23) we obtain $\alpha^* = 1 - 1/\gamma = -9$, while in our simulation we find a threshold around -3.35 .


Another important difference in our approximation is that it implies that E-stability conditions are not restrictive enough to rule out learning instabilities. To evaluate this corollary in the context of the Cobweb model we construct the Mean Squared Deviation (MSD) learning curves associated to different combinations of parameters. The MSD learning curves are obtained by averaging the squared difference between the learning coefficients and the corresponding REE value across simulations of the model⁴. The results are presented in figure 2, varying α for a fixed γ in panel (a), and varying γ for a fixed α in panel (b); importantly, the E-stability condition, $\alpha < 1$, is satisfied in all cases. We observe that convergence is indeed governed by an additional restriction on the values of α and γ ; particularly, notice the area surrounding our approximation's restriction, corresponding to $\gamma < 2/(1 - \alpha)$ and plotted as a thicker black line: the further α and γ pass beyond their implied threshold, the likelier the occurrence of learning instabilities.

4.2 New-Keynesian model

In order to illustrate the generality of our results we now evaluate our approximation in the context of the standard New-Keynesian macroeconomic model described in example 2 on page 5. This model has served as the workhorse of an extensive literature on issues related with the design of monetary policy. Particularly of our interest is the analysis of Evans and Honkapohja (2009), who show that the connection between determinacy and E-stability of REE may not be robust under constant-gain learning. We reconsider this issue using our approximation and focusing on the two standard Taylor rule specifications described in equations (5) and (6). However, due to the interactions between the structural parameters introduced by these policy feedback rules, the expressions provided by proposition 1 become quite complicated. Hence, we proceed with a numerical analysis by fixing the model parameters according to standard calibrations in the literature (see Evans and Honkapohja, 2009, p.151): $\beta = 0.99$, $\varphi^{-1} = 0.157$, and $\lambda = 0.024$. That allows us to focus on the effects of the different policy assumptions and the learning gain on the dynamics of learning in this model.

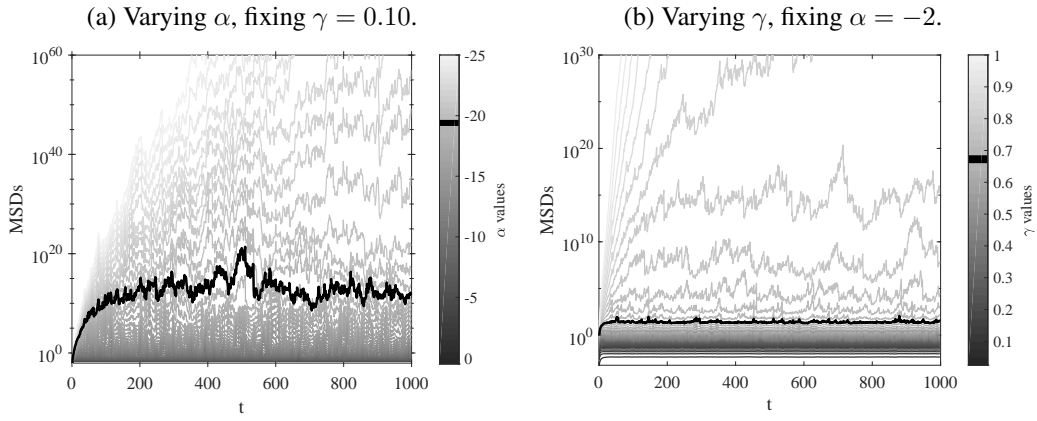
⁴Notice that, in contrast to the decreasing gain case, convergence under constant gain does not imply a zero MSD; in fact, when the constant gain learning estimates converge to a distribution around the REE, the mean deviation of these estimates in relation to the REE will tend to stabilize around a positive value that corresponds to the variance of the learning estimates.

Figure 1: Mean and variance of long run learning estimates in the Cobweb model.



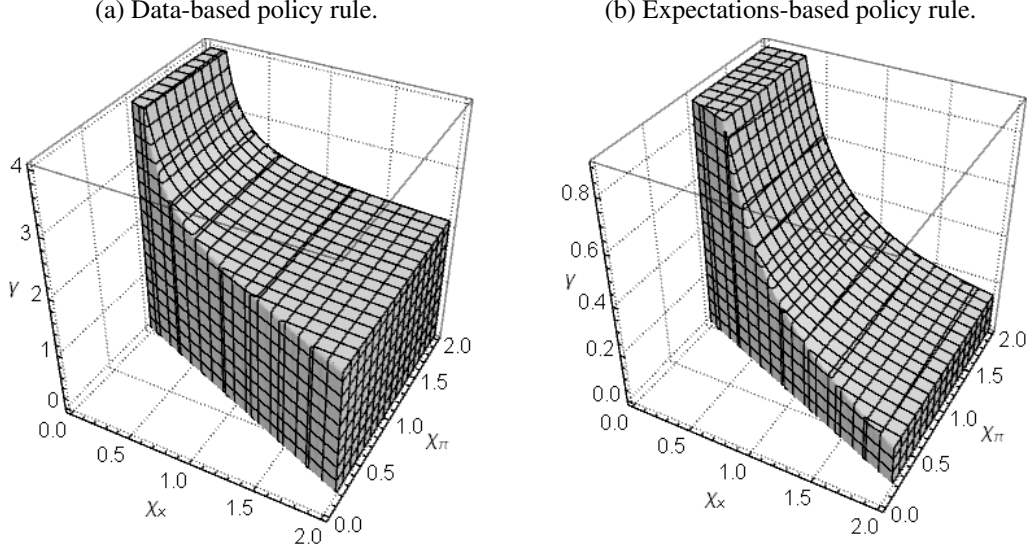
Notes: The plots are based on the period $t = 1,000$ distribution of learning coefficients obtained from 10,000 simulations of the Cobweb model with different parameter combinations: in (a) we fix $\gamma = 0.10$ and vary α from 0.5 to -18.5 (in -0.5 steps); in (b) we fix $\alpha = -2$ and vary γ from 0.025 to 0.65 (in 0.025 steps); these are indicated in the horizontal axis of each plot. The trimmed simulated statistics are obtained by discarding estimates that deviate by more than ± 5 standard deviations (using the variance implied by proposition 1) from the REE-implied mean; see Appendix A.5 for trimming statistics. The approximations (solid and dashed lines) are calculated according to equations (23) and (24).

Figure 2: Learning curves in the Cobweb model.



Notes: The learning curves are constructed by averaging the squared difference between the learning coefficients and the corresponding REE value across 10,000 simulations of the Cobweb model with different parameter combinations: in (a) we fix $\gamma = 0.10$ and vary α from 0.5 to -25 (in -0.5 steps); in (b) we fix $\alpha = -2$ and vary γ from 0.025 to 1 (in 0.025 steps); these are indicated in the colorbars next to each plot. The thicker black line refers to the learning curve obtained with $\alpha = 1 - 2/\gamma$, derived from proposition 1's convergence condition: in (a) this is given by $\alpha = -19$; in (b) it corresponds to $\gamma = 2/3$.

Figure 3: Learning stability in the New-Keynesian model.



Notes: The three-dimensional regions represent the combinations of the parameters γ , χ_π , and χ_x , for which the learning estimates converge to the corresponding REE values, according to proposition 1, assuming the other model parameters are fixed as: $\beta = 0.99$, $\varphi^{-1} = 0.157$, and $\lambda = 0.024$.

The conditions for stability of the learning process under the two alternative policy rules are presented in figure 3. Interestingly, we observe that the learning gain is less relevant under the data-based rule: for any value of γ below 2, learnability is determined mainly by the intensity of the policymaker's response to contemporaneous inflation and output gap; also, adjusting the interest rate by more than one-to-one changes in the inflation rate (Taylor principle) guarantees stability for any intensity of response to output gap changes. A different result is obtained under the more realistic expectations-based rule, where learning gains closer to the range of plausible values (between 0.01 and 0.20, see Berardi and Galimberti, 2017) can cause instabilities; that is particularly the case as the policymaker becomes more sensitive to output gap variations, i.e., the upper bound on γ decreases as χ_x increases. Although these results are qualitatively consistent with Evans and Honkapohja (2009, p. 151), our approximation results in looser constraints on the learning gain; for example, when $\chi_\pi = 1.5$ those authors calculate that the equilibrium would be unstable under learning for $\chi_x > 1.57$ and $\gamma \geq 0.10$, whereas our approximation points to a stability threshold of $\chi_x > 3.13$ for a $\gamma \geq 0.10$, or a $\gamma > 0.20$ if one fixes $\chi_x = 1.57$.

Furthermore, we evaluate the accuracy of our approximation by simulating the New-Keynesian model in an approach similar to our analysis of the Cobweb model above. Boxplots of the learning estimates are presented in figure 4, where the two different policies are distinguished by color: black for the data-based rule, gray for the expectations-based rule. Across the panels we also check the distribution of the learning estimates for varying values of the policy parameters and the learning gain. Our main conclusion is that our approximation captures the dispersion of the learning estimates pretty well, as evidenced by the tight connections between the approximation bands and the endpoints of the boxplots' whiskers. Also, the learning gain, in spite of its relevance for stability conditions as shown above, seems to have a similar effect on the dispersion of the learning estimates across the policy rules (see panels (a) and (d) in figure 4).

These results also show the effects that different policy settings can have on agents' perceived uncertainty about the economy. I.e., assuming a higher dispersion of learning estimates can be associated with a higher uncertainty about the underlying point estimates agents use to form their expectations, we observe that: (i) a policy reacting to contemporaneous data leads to higher (relative to the expectations-based policy) perceived uncertainty about the output gap if the policymaker is not reacting strongly to inflation (see the first two boxplots in panel (b) of figure 4), but (ii) a lower relative uncertainty on both inflation and output gap expectations if the policymaker is not reacting strongly to output gap (see the first two boxplots in panels (c) and (f) of figure 4); (iii) for both rules, the higher the reaction to inflation relative to the reaction to output gap, the lower (higher) the perceived uncertainty about inflation (output gap), and vice versa (this is evident by the narrowing/broadening of the dispersion bands in panels (c) and (e)/(b) and (f) of figure 4 as the policy reaction parameters increase).

4.3 Model with a lagged endogenous variable


Finally, we validate the accuracy of our approximation under a model including a lagged endogenous variable. As discussed in Remark 10, although our general framework can cope with this situation, the implied interdependence between the learning estimates and the stationarity of the state variables can affect the validity of our approximation for parameter values near the stability threshold. Focusing on the simple univariate model given by equation (20), and assuming agents form expectations according to a first-order autoregression, the ALM is given by

$$y_t = (\beta\phi_{t-1} + \lambda) y_{t-1} + u_t. \quad (25)$$


Using our approximation, and assuming $K = 1$, learning stability requires that $\beta < 1$, which is also the E-stability condition for this model. Besides, notice that

Figure 4: Boxplots of long run learning estimates in the New-Keynesian model.

(a) Output gap PLM, varying γ . (b) Output gap PLM, varying χ_π . (c) Output gap PLM, varying χ_x .



(d) Inflation PLM, varying γ . (e) Inflation PLM, varying χ_π . (f) Inflation PLM, varying χ_x .




-- REE mean — Data-based policy rule — Expectations-based policy rule

Notes:

The boxplots are based on the period $t = 1,000$ distribution of learning coefficients obtained from 10,000 simulations of the New-Keynesian model with the data-based policy rule (black) and the expectations-based policy rule (gray), and with different parameter combinations: in (a) and (d) we fix $\chi_\pi = 1.5$ and $\chi_x = 0.125$, and vary γ from 0.05 to 0.85 (in 0.20 steps); in (b) and (e) we fix $\gamma = 0.10$ and $\chi_x = 0.125$, and vary χ_π from 1.0 to 3.0 (in 0.5 steps); in (c) and (f) we fix $\gamma = 0.10$ and $\chi_\pi = 1.5$, and vary χ_x from 0.0 to 2.0 (in 0.5 steps); the other model parameters are fixed as: $\beta = 0.99$, $\varphi^{-1} = 0.157$, and $\lambda = 0.024$. The whiskers are constructed so as to correspond to 2 standard deviations from the mean. The approximated dispersions (black and gray lines) also refer to 2 standard deviations around the REE mean.

Figure 5: Variance of long run learning estimates in a model with a lagged endogenous variable.



Notes:

The plots are based on the period $t = 1,000$ distribution of learning coefficients obtained from 10,000 simulations of model (25) with different parameter combinations: in (a) we fix $\lambda = 0.25$, $\gamma = 0.10$, and vary β from -1.5 to 0.7 (in 0.1 steps); in (b) we fix $\beta = 0.5$, $\gamma = 0.10$, and vary λ from -0.45 to 0.45 (in 0.05 steps); in (c) we fix $\beta = 0.5$, $\lambda = 0.25$, and vary γ from 0.025 to 0.5 (in 0.025 steps); these are indicated in the horizontal axis of each plot.

stationarity of y_t requires that $|\beta\phi_{t-1} + \lambda| < 1$; although guaranteeing this condition under learning would require the imposition of further bounding assumptions, we adopt the implied REE in order to restrict the set of parameters for simulation, i.e., under REE this model is stationary if $|\lambda/(1 - \beta)| < 1$.

We simulate this model following a similar approach to that adopted in the previous examples, i.e., by drawing random samples of the stochastic disturbance and using the model's ALM together with a LS learning algorithm to obtain the implied learning estimates. Figure 5 presents the variances of the invariant distributions of these learning estimates for different parameter combinations, together with their approximated variances as implied by our proposition⁵. Interestingly, in spite of all reservations, our approximation is capturing the main effects of varying this model's parameters on the dispersion of the learning estimates, namely: (a) the increase of dispersion with β up to a critical value near the stationarity threshold, and decreasing dispersion beyond that point; (b) the concavity of dispersion with respect to λ ; and, (c) the standard effect of the learning gain in increasing the dispersion of the learning estimates.

⁵One difference in using (19) for the case with a lagged endogenous variable is that $\vec{E}[\mathbf{R}_t^{-1}]$ will depend on the learning estimates; here we take the REE implied value of $\vec{E}[\mathbf{z}_t\mathbf{z}_t']^{-1}$ as an approximation (see also Remark 8).

5 Concluding remarks

In this paper we proposed a simple approximation of the first two moments of the long run distribution of learning estimates for a general class of dynamic macroeconomic models under constant-gain learning. This approximation provides analytical expressions for the mean and the variance of the distribution of estimates that a constant-gain learning process, taking place within a self-referential model, converges in the long run. The simplicity of our approximation lies in a time series simplifying assumption that greatly reduces the complexity of statistical interactions in the recursive expressions. As a byproduct of this approximation, we also obtained new conditions for the convergence of the learning process, and showed that these conditions are more stringent than the usual E-stability conditions derived under assumptions of arbitrarily small or decreasing gains.

We showed the usefulness and evaluated the accuracy of our approximation numerically in the context of two standard economic models, namely, a Cobweb model of prices and a New-Keynesian (NK) model of aggregate inflation and output; we also considered the case of a model including a lagged endogenous variable. Overall, we found evidence in support of our approach. In the NK model, we also show how monetary policy settings can affect the dispersion of learning estimates and the implicit perceptions of uncertainty about the economy. In that context, we believe our analytical expressions can prove useful for further analysis of the role of learning and uncertainty for policy design.

A Appendix

A.1 Recursive solution

Defining $\boldsymbol{\theta}_t = \text{vec}(\boldsymbol{\Phi}_t)$, we find that the learning algorithm recursion from equation (8) is equivalent to

$$\boldsymbol{\theta}_t = (\mathbf{I}_{(np)} - \gamma (\mathbf{I}_{(n)} \otimes \mathbf{R}_t^{-1} \mathbf{z}_t \mathbf{z}_t')) \boldsymbol{\theta}_{t-1} + \gamma (\mathbf{I}_{(n)} \otimes \mathbf{R}_t^{-1} \mathbf{z}_t) \text{vec}(\mathbf{y}_t).$$

Similarly, the ALM, from equation (1), implies that

$$\text{vec}(\mathbf{y}_t) = (\mathbf{I}_{(n)} \otimes \mathbf{z}_t') \text{vec}(\mathbf{A}) + (\mathbf{B}' \otimes \mathbf{z}_t') \boldsymbol{\theta}_{t-1} + (\mathbf{I}_{(n)} \otimes \mathbf{u}_t') \text{vec}(\mathbf{C}),$$

hence

$$\begin{aligned}
\boldsymbol{\theta}_t &= (\mathbf{I}_{(np)} - \gamma (\mathbf{I}_{(n)} \otimes \mathbf{R}_t^{-1} \mathbf{z}_t \mathbf{z}_t')) + \gamma (\mathbf{I}_{(n)} \otimes \mathbf{R}_t^{-1} \mathbf{z}_t) (\mathbf{B}' \otimes \mathbf{z}_t') \boldsymbol{\theta}_{t-1} \\
&\quad + \gamma (\mathbf{I}_{(n)} \otimes \mathbf{R}_t^{-1} \mathbf{z}_t) ((\mathbf{I}_{(n)} \otimes \mathbf{z}_t') \text{vec}(\mathbf{A}) + (\mathbf{I}_{(n)} \otimes \mathbf{u}_t') \text{vec}(\mathbf{C})), \\
&= (\mathbf{I}_{(np)} - \gamma (\mathbf{I}_{(n)} - \mathbf{B}') \otimes (\mathbf{R}_t^{-1} \mathbf{z}_t \mathbf{z}_t')) \boldsymbol{\theta}_{t-1} \\
&\quad + \gamma \mathbf{I}_{(n)} \otimes (\mathbf{R}_t^{-1} \mathbf{z}_t \mathbf{z}_t') \text{vec}(\mathbf{A}) + \gamma \mathbf{I}_{(n)} \otimes (\mathbf{R}_t^{-1} \mathbf{z}_t \mathbf{u}_t') \text{vec}(\mathbf{C}), \\
\boldsymbol{\theta}_t &= \mathbf{H}_t \boldsymbol{\theta}_{t-1} + \gamma \mathbf{P}_t \text{vec}(\mathbf{A}) + \gamma \mathbf{Q}_t \text{vec}(\mathbf{C}), \tag{26}
\end{aligned}$$

where \mathbf{H}_t , \mathbf{P}_t , and \mathbf{Q}_t are defined in the main text in equations (13), (14), and (15), respectively. Solving equation (26) recursively we obtain

$$\boldsymbol{\theta}_t = \prod_{j=0}^{t-1} \mathbf{H}_{t-j} \boldsymbol{\theta}_0 + \gamma \sum_{i=0}^{t-1} \prod_{j=0}^{i-1} \mathbf{H}_{t-j} (\mathbf{P}_{t-i} \text{vec}(\mathbf{A}) + \mathbf{Q}_{t-i} \text{vec}(\mathbf{C})),$$

which is equivalent to equation (11) after defining $\mathbf{F}_t(k)$ according to equation (12).

A.2 Convergence of LS estimates of regressors' second moments

From equation (9) we have that

$$\mathbf{R}_t = (1 - \gamma) \mathbf{R}_{t-1} + \gamma \mathbf{z}_t \mathbf{z}_t',$$

which solved recursively is equivalent to

$$\mathbf{R}_t = (1 - \gamma)^t \mathbf{R}_0 + \gamma \sum_{i=0}^{t-1} (1 - \gamma)^i \mathbf{z}_{t-i} \mathbf{z}_{t-i}'. \tag{27}$$

Assuming \mathbf{z}_t is stationary, i.e., $E[\mathbf{z}_t \mathbf{z}_t'] = E[\mathbf{z}_i \mathbf{z}_i'] \forall t, i$, the expected value of equation (27) is given by

$$E[\mathbf{R}_t] = (1 - \gamma)^t \mathbf{R}_0 + E[\mathbf{z}_t \mathbf{z}_t'] (1 - (1 - \gamma)^t),$$

which in the long run simplifies to

$$\vec{E}[\mathbf{R}_t] = E[\mathbf{z}_t \mathbf{z}_t'],$$

for $0 < \gamma < 2$.

A.3 First moment of learning estimates

To obtain the first moment of the distribution of learning estimates we use the approximation in equation (16) to evaluate the expectation of equation (11) as

$$\begin{aligned} E[\boldsymbol{\theta}_t] &= \tilde{\mathbf{F}}(t) \boldsymbol{\theta}_0 + \gamma \sum_{i=0}^{t-1} \tilde{\mathbf{F}}(i) (\mathbf{I}_{(n)} \otimes \mathbf{K}) \text{vec}(\mathbf{A}), \\ &= \tilde{\mathbf{H}}^t \boldsymbol{\theta}_0 + ((\mathbf{I}_{(n)} - \mathbf{B}') \otimes \mathbf{K})^{-1} (\mathbf{I}_{(np)} - \tilde{\mathbf{H}}^t) (\mathbf{I}_{(n)} \otimes \mathbf{K}) \text{vec}(\mathbf{A}), \end{aligned}$$

where the solution of the summation only requires invertibility of $(\mathbf{I}_{(n)} - \mathbf{B}')$, which is equivalent to the RE solution requirement, and invertibility of \mathbf{K} . To obtain the long run expectation we evaluate this expression in the asymptotic limit, $t \rightarrow \infty$, for which case convergence is obtained if all the eigenvalues of $\tilde{\mathbf{H}}$ lie inside the unit circle, resulting in

$$\vec{E}[\boldsymbol{\theta}_t] = \left((\mathbf{I}_{(n)} - \mathbf{B}')^{-1} \otimes \mathbf{I}_{(p)} \right) \text{vec}(\mathbf{A}),$$

which, after re-arranging the coefficients back into their original form, is equivalent to equation (17).

A.4 Second moment of learning estimates


For the second moment we focus directly on the invariant distribution, i.e., dropping the terms that vanish in the limit we have

$$\begin{aligned} \vec{E}[\boldsymbol{\theta}_t \boldsymbol{\theta}_t'] &= \gamma^2 \underbrace{\vec{E} \left[\sum_{i=0}^{t-1} \mathbf{F}_t(i) \mathbf{P}_{t-i} \text{vec}(\mathbf{A}) \sum_{i=0}^{t-1} \text{vec}(\mathbf{A})' \mathbf{P}_{t-i}' \mathbf{F}_t(i)' \right]}_{\mathcal{A}} \\ &\quad + \gamma^2 \underbrace{\vec{E} \left[\sum_{i=0}^{t-1} \mathbf{F}_t(i) \mathbf{Q}_{t-i} \text{vec}(\mathbf{C}) \sum_{i=0}^{t-1} \text{vec}(\mathbf{C})' \mathbf{Q}_{t-i}' \mathbf{F}_t(i)' \right]}_{\mathcal{B}}. \quad (28) \end{aligned}$$

Expanding the summations, one can find that whereas the first is related to $\vec{E}[\boldsymbol{\theta}_t] \vec{E}[\boldsymbol{\theta}_t']$, i.e.,

$$\mathcal{A} = \gamma^{-2} \left((\mathbf{I}_{(n)} - \mathbf{B}')^{-1} \otimes \mathbf{I}_{(p)} \right) \text{vec}(\mathbf{A}) \text{vec}(\mathbf{A})' \left((\mathbf{I}_{(n)} - \mathbf{B})^{-1} \otimes \mathbf{I}_{(p)} \right), \quad (29)$$

Figure 6: Trimming statistics.



the second simplifies to the matrix equivalent of the sum of a geometric progression containing only even powers, namely,

$$\begin{aligned} \mathcal{B} &= \sigma_u^2 \sum_{i=0}^{t-1} \tilde{\mathbf{H}}^k \left((\mathbf{C}'\mathbf{C}) \otimes \overline{\mathbf{R}^{-1}} \right) \tilde{\mathbf{H}}'^k, \\ \text{vec}(\mathcal{B}) &= \sigma_u^2 \left(\mathbf{I}_{(n^2 p^2)} - \tilde{\mathbf{H}} \otimes \tilde{\mathbf{H}} \right)^{-1} \text{vec} \left((\mathbf{C}'\mathbf{C}) \otimes \overline{\mathbf{R}^{-1}} \right). \end{aligned} \quad (30)$$

where $\overline{\mathbf{R}^{-1}}$ stands for $E[\mathbf{R}_t^{-1}]$, which can be approximated as the inverse of $E[\mathbf{z}_t \mathbf{z}_t']$.

Finally, notice that the variance of this invariant distribution is determined by \mathcal{B} ,

$$\begin{aligned} \text{Var}(\boldsymbol{\theta}_t) &= \overrightarrow{E} \left[\left(\boldsymbol{\theta}_t - \overrightarrow{E}[\boldsymbol{\theta}_t] \right) \left(\boldsymbol{\theta}_t' - \overrightarrow{E}[\boldsymbol{\theta}_t'] \right) \right], \\ &= \overrightarrow{E}[\boldsymbol{\theta}_t \boldsymbol{\theta}_t'] - \overrightarrow{E}[\boldsymbol{\theta}_t] \overrightarrow{E}[\boldsymbol{\theta}_t'], \\ &= \gamma^2 (\mathcal{A} + \mathcal{B}) - \gamma^2 \mathcal{A}, \\ &= \gamma^2 \mathcal{B}, \end{aligned} \quad (31)$$

and that as $\gamma \rightarrow 0$, $\mathcal{B} \rightarrow \mathbf{0}$, thence $\text{Var}(\boldsymbol{\theta}_t) \rightarrow \mathbf{0}$, which is consistent with the implied convergence of learning under decreasing gains.

A.5 Trimming statistics

In order to reduce the effect of outliers in the numerical evaluation of the accuracy of the proposed approximations, the statistics reported in Figure 1 are based

on a trimmed sample of simulated learning estimates. The trimming is calculated on the basis of the approximation's implied variance, discarding estimates that deviate by more than ± 5 standard deviations from the REE-implied mean. The number of simulations discarded depended on the combination of parameters and are presented in Figure 6.

References

- Benveniste, A., M. Metivier, and P. Priouret (1990). *Adaptive Algorithms and Stochastic Approximations*. Springer-Verlag.
- Berardi, M. and J. K. Galimberti (2013). A note on exact correspondences between adaptive learning algorithms and the kalman filter. *Economics Letters* 118(1), 139–142.
- Berardi, M. and J. K. Galimberti (2014). A note on the representative adaptive learning algorithm. *Economics Letters* 124(1), 104 – 107.
- Berardi, M. and J. K. Galimberti (2017). Empirical calibration of adaptive learning. *Journal of Economic Behavior & Organization* 144, 219 – 237.
- Bullard, J. and K. Mitra (2002). Learning about monetary policy rules. *Journal of Monetary Economics* 49(6), 1105 – 1129.
- Chevillon, G., M. Massmann, and S. Mavroeidis (2010). Inference in models with adaptive learning. *Journal of Monetary Economics* 57(3), 341–351.
- Cho, I.-K., N. Williams, and T. J. Sargent (2002). Escaping nash inflation. *The Review of Economic Studies* 69(1), 1–40.
- Evans, G. W. and S. Honkapohja (1998). Economic dynamics with learning: New stability results. *Review of Economic Studies* 65(1), 23–44.
- Evans, G. W. and S. Honkapohja (2001). *Learning and expectations in macroeconomics*. Frontiers of Economic Research. Princeton, NJ: Princeton University Press.
- Evans, G. W. and S. Honkapohja (2003). Expectations and the stability problem for optimal monetary policies. *The Review of Economic Studies* 70(4), pp. 807–824.
- 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*, Volume 13 of *Central Banking, Analysis, and Economic Policies Book Series*, Chapter 5, pp. 145–170. Central Bank of Chile.
- Evans, G. W., S. Honkapohja, and N. Williams (2010). Generalized stochastic gradient learning. *International Economic Review* 51(1), 237–262.
- Hamilton, J. D. (1994). *Time Series Analysis*. Princeton University Press.
- Marcet, A. and T. J. Sargent (1989). Convergence of least squares learning mechanisms in self-referential linear stochastic models. *Journal of Economic Theory* 48(2), 337–368.
- Milani, F. (2007). Expectations, learning and macroeconomic persistence. *Journal of Monetary Economics* 54(7), 2065–2082.
- Milani, F. (2014). Learning and time-varying macroeconomic volatility. *Journal of Economic Dynamics and Control* 47, 94–114.
- Williams, N. (2018). Escape dynamics in learning models. *The Review of Economic Studies* (forthcoming).
- Woodford, M. (2003). *Interest and Prices*. Princeton, NJ: Princeton University Press.