



Australian
National
University

Crawford School of Public Policy

CAMA

Centre for Applied Macroeconomic Analysis

Information weighting under least squares learning

CAMA Working Paper 46/2020
May 2020

Jaqueson K. Galimberti

Auckland University of Technology

Centre for Applied Macroeconomic Analysis, ANU

Abstract

This paper evaluates how adaptive learning agents weight different pieces of information when forming expectations with a recursive least squares algorithm. The analysis is based on a renewed and more general non-recursive representation of the learning algorithm, namely, a penalized weighted least squares estimator, where a penalty term accounts for the effects of the learning initials. The paper then draws behavioral implications of alternative specifications of the learning mechanism, such as the cases with decreasing, constant, regime-switching, adaptive, and age-dependent gains, as well as practical recommendations on their computation. One key new finding is that without a proper account for the uncertainty about the learning initial, a constant-gain can generate a time-varying profile of weights given to past observations, particularly distorting the estimation and behavioral interpretation of this mechanism in small samples of data. In fact, simulations and empirical estimation of a Phillips curve model with learning indicate that this particular misspecification of the initials can lead to estimates where inflation rates are less responsive to expectations and output gaps than in reality, or “flatter” Phillips curves.

Keywords

bounded rationality, expectations, adaptive learning, memory

JEL Classification

E70, D83, D84, D90, E37, C32, C63

Address for correspondence:

(E) cama.admin@anu.edu.au

ISSN 2206-0332

[The Centre for Applied Macroeconomic Analysis](#) in the Crawford School of Public Policy has been established to build strong links between professional macroeconomists. It provides a forum for quality macroeconomic research and discussion of policy issues between academia, government and the private sector.

The Crawford School of Public Policy is the Australian National University's public policy school, serving and influencing Australia, Asia and the Pacific through advanced policy research, graduate and executive education, and policy impact.

Information weighting under least squares learning

Jaqueson K. Galimberti*

March 16, 2020

Abstract

This paper evaluates how adaptive learning agents weight different pieces of information when forming expectations with a recursive least squares algorithm. The analysis is based on a renewed and more general non-recursive representation of the learning algorithm, namely, a penalized weighted least squares estimator, where a penalty term accounts for the effects of the learning initials. The paper then draws behavioral implications of alternative specifications of the learning mechanism, such as the cases with decreasing, constant, regime-switching, adaptive, and age-dependent gains, as well as practical recommendations on their computation. One key new finding is that without a proper account for the uncertainty about the learning initial, a constant-gain can generate a time-varying profile of weights given to past observations, particularly distorting the estimation and behavioral interpretation of this mechanism in small samples of data. In fact, simulations and empirical estimation of a Phillips curve model with learning indicate that this particular misspecification of the initials can lead to estimates where inflation rates are less responsive to expectations and output gaps than in reality, or “flatter” Phillips curves.

Keywords: bounded rationality, expectations, adaptive learning, memory.

JEL codes: E70, D83, D84, D90, E37, C32, C63.

“The longer you can look back,
the farther you can look forward.”

–Winston Churchill

1 Introduction

Adaptive learning can generate out-of-equilibrium expectations that help explain deviations from rational expectations and an economy’s transitional dynamics towards equilibrium. Under adaptive learning, agents’ beliefs are modeled through the assumption of a recursive learning

*Auckland University of Technology, New Zealand; and Centre for Applied Macroeconomic Analysis, Australian National University. Corresponding author: jaqueson.galimberti@aut.ac.nz, <https://sites.google.com/site/jkgeconoeng/>.

mechanism that updates agents' perceptions about the economy as new data observations become available. The weight given to these observations is a key determinant of the degree of persistence introduced by adaptive learning in the evolution of expectations, and, hence, it is an important factor in the explanation of deviations from rational expectations predictions. In this paper we study the behavioral implications of alternative specifications of the learning mechanism, showing how these assumptions can affect the way agents weight information when forming expectations according to a recursive least squares algorithm.

The first contribution of this paper is the proposal of a new and more general non-recursive representation of the learning algorithm. Namely, we show that the recursive least squares (RLS), a commonly assumed mechanism to represent agents econometric learning, is more properly represented by a penalized weighted least squares (WLS) estimator, where a penalty term accounts for the effects of the learning initial estimates.¹ The pace of RLS learning is regulated by a sequence of learning gains, which also determines how different pieces of information are weighted in the implied estimates of agents' perceived law of motion. The framework proposed in this paper provides flexible analytical expressions for the calculation of these weights under alternative assumptions on the behavior of the learning gains, including the traditional decreasing-gain (Marcet and Sargent, 1989) and constant-gain (Sargent, 1999) specifications, as well as more sophisticated mechanisms such as endogenous gain-switching (Marcet and Nicolini, 2003) and age-dependent (Malmendier and Nagel, 2016) gain specifications.

Several instructive results are derived on the basis of this new framework. First, the influence of the learning initials can become non-negligible for applied purposes, especially under a constant-gain setup, where the weights given to past observations decrease geometrically. For example, for calibrations and samples of standard macroeconomic applications, it can take almost 5 decades of quarterly data for the latest observation to receive a weight as high as the weight given to the initial estimate—although the relevance of initials is well known in the literature (see, e.g., Carceles-Poveda and Giannitsarou, 2007; Slobodyan and Wouters, 2012b; Berardi and Galimberti, 2017b), the framework developed in this paper allows a more precise account of these effects. Second, although the effect of the initials are less persistent under a decreasing-gain setup, the often assumed equivalence between this formulation and the traditional ordinary least squares estimator from econometrics does not hold for arbitrary learning initials, but only under the assumption of a diffuse initial, which is equivalent to a diffuse prior in a Bayesian context—this is another puzzling observation that applied researchers quickly learn with the computation of learning estimates and that the framework of this paper clarifies.

One key new finding of this paper is that the assumption of a diffuse initial also implies that the profile of weights given to past observations under constant-gain learning is time-varying,

¹A similar correspondence between the RLS and the standard (i.e., without the penalty term) WLS was presented in Berardi and Galimberti (2013), except that under the assumption of a diffuse initial prior, in which case the effects of the learning initials are not taken into account; also note that most recursive formulations of statistical estimators adopted in the adaptive learning literature are drawn from the system identification literature in engineering (e.g., Ljung and Soderstrom, 1983).

causing a geometric decay of weighting stronger than would have been expected from asymptotic analysis of the learning algorithm. This result can have important implications for the estimation of models with learning, because it can lead to an overweighting of the initial sample of observations used for model estimation. In fact, simulating the estimation of a new Keynesian Phillips curve model we find that diffuse initials lead to stronger small sample distortions in the model estimates. Particularly, the misspecified initials result in “flatter” Phillips curve estimates in the sense that inflation rates seem less responsive to changes in real marginal costs than they actually are. A back-of-the-envelope calculation based on the simulation results with a sample of 50 observations indicates the diffuse initial’s distortion can lead to Phillips curve slope estimates about half its true value. Empirical estimates with US data also indicate that the diffuse initials lead to systematically lower estimates of the relevance of expectations for the determination of inflation, and imply flatter Phillips curve estimates than those obtained with initials consistent with the constant-gain learning mechanism.

The new framework is also useful for the analysis of more recent specifications of the learning gains. Particularly, the learning from experience approach of Malmendier and Nagel (2016) stands, in the framework of this paper, as a generalization of decreasing-gain learning applied at cohort-level data. Consistent with their analysis, we find that this mechanism implies a profile of weights given to lifetime observations that is varying with the agent’s age. Nevertheless, another interesting new observation facilitated by our framework is the finding that this setup can lead to hump-shaped profiles of weights given to lagged observations throughout an agent’s lifetime. For example, under an usual calibration of this mechanism, an agent would be giving an increasing weight to 20-years lagged observations all the way till reaching the age of 60 years old. Finally, the case of time-varying learning gains is also considered to show how the framework of this paper can be informative for the design of mechanisms consistent with a behavioral concern with ongoing structural changes in the economy.

The remainder of this paper is split into four other sections: Section 2 outlines the general framework of recursive learning and the new non-recursive representation of the learning algorithm; Section 3 shows how to use this framework for drawing the information weighting and behavioral implications of different specifications of the learning gains; Section 4 presents simulation and empirical analysis of the distortionary effects of diffuse initials in the estimation of models with constant-gain learning; Section 5 concludes with some remarks.

2 Framework

In models with adaptive learning a perceived law of motion (PLM) is specified relating the variables agents are assumed to observe and those variables they care and need to form expectations about. Focusing on a univariate case a typical PLM specification is given by a linear regression model of the form

$$y_t = \mathbf{x}_t' \phi_t + \varepsilon_t, \tag{1}$$

where y_t is assumed to be related to a vector of (pre-determined) variables, $\mathbf{x}_t = (x_{1,t}, \dots, x_{k,t})'$, through the vector of (possibly time-varying) coefficients $\phi_t = (\phi_{1,t}, \dots, \phi_{k,t})'$, and ε_t denotes a white noise disturbance term. This specification can be easily extended to a multivariate context, by augmenting y_t , ε_t , and ϕ_t with extra columns, and to different specifications of lag/lead in the timing of expectations, by adjusting the timing of \mathbf{x}_t .

In a typical economic modeling context, the observations of y_t needed to estimate (1), as well as some or all of the regressors in \mathbf{x}_t , are endogenously determined within a hypothetical structural model. These observations are the result of market equilibrium and the interaction between the economic decisions by the many different actors that compose the economy, such as households, firms, and policymakers. Hence, due to the relevance of expectations for these agents' economic decisions, the same macroeconomic outcomes that are relevant for the formation of expectations are themselves determined by expectations, a feature often called self-referentiality. Notwithstanding, for the purposes of deriving weighting expressions we will abstract from one side of this self-referential nature of expectations, and focus instead on the modeling of agents' PLM. I.e., we will assume agents form expectations according to equation (1) without accounting for the endogeneity of the involved variables. Notice, as is usually the case in applications of adaptive learning, this implies some degree of bounded rationality in the way agents form expectations. The effects of self-referentiality are more fully taken into account in the simulation and empirical estimation exercises presented in Section 4.

Recursive learning A recursive estimator is assumed to represent how agents update their PLM estimates as new observations become available. One popular algorithm (see Berardi and Galimberti, 2014, for rationale and a comparative of alternatives) is given by the Recursive Least Squares (RLS),

$$\hat{\phi}_t = \hat{\phi}_{t-1} + \gamma_t \mathbf{R}_t^{-1} \mathbf{x}_t (y_t - \mathbf{x}_t' \hat{\phi}_{t-1}), \quad (2)$$

$$\mathbf{R}_t = \mathbf{R}_{t-1} + \gamma_t (\mathbf{x}_t \mathbf{x}_t' - \mathbf{R}_{t-1}), \quad (3)$$

where γ_t is a learning gain parameter, \mathbf{R}_t stands for an estimate of regressors' matrix of second moments, $E[\mathbf{x}_t \mathbf{x}_t']$, and the $\{\hat{\phi}_0, \mathbf{R}_0\}$ initial estimates are set to be consistent with plausible agents' beliefs at the beginning of the modeled sample (see Berardi and Galimberti, 2017b, for further discussion and methods). The learning gain is an important parameter of this learning mechanism because it determines how quickly new information is incorporated into the recursive estimates, and hence, how quickly agents react to different pieces of information. More precisely, the sequence of learning gains can be related with the relative weights given to sample observations in the estimation process. In order to draw this relationship it is useful to consider the non-recursive formulation corresponding to this estimation problem.

Non-recursive form One important contribution of this paper is to provide a new and more general non-recursive representation of the RLS learning algorithm accounting for the initials. When initialized from arbitrary initials, $\hat{\phi}_0$ and \mathbf{R}_0 , the RLS has a non-recursive form given by

$$\hat{\phi}_t = \arg \min \sum_{i=1}^t \omega_{t,i} (y_i - \mathbf{x}'_i \hat{\phi}_t)^2 + \omega_{t,0} (\hat{\phi}'_0 - \hat{\phi}'_t) \mathbf{R}_0 (\hat{\phi}_0 - \hat{\phi}_t), \quad (4)$$

$$= \left[\sum_{i=1}^t \omega_{t,i} \mathbf{x}_i \mathbf{x}'_i + \omega_{t,0} \mathbf{R}_0 \right]^{-1} \left[\sum_{i=1}^t \omega_{t,i} \mathbf{x}_i y_i + \omega_{t,0} \mathbf{R}_0 \hat{\phi}_0 \right], \quad (5)$$

where the weights are related to the sequence of learning gains according to

$$\omega_{t,i} = \begin{cases} \prod_{j=1}^t (1 - \gamma_j) & \text{for } i = 0 \text{ (initial),} \\ \gamma_i \prod_{j=i+1}^t (1 - \gamma_j) & \text{for } 0 < i < t, \\ \gamma_t & \text{for } i = t, \end{cases} \quad (6)$$

Interestingly, when the initial is taken into account, the RLS is equivalent to a Weighted Least Squares (WLS) estimation problem augmented with a penalty on squared deviations between estimates and initials.

Correspondence between penalized WLS and RLS To see how the RLS of (2)-(3) can be derived from the penalized WLS formulation of (5) and (6), first notice that iterating (3) recursively from \mathbf{R}_0 we have that

$$\mathbf{R}_t = \sum_{i=1}^t \omega_{t,i} \mathbf{x}_i \mathbf{x}'_i + \omega_{t,0} \mathbf{R}_0,$$

which is the inverse of the first term in (5), leading to

$$\hat{\phi}_t = \mathbf{R}_t^{-1} \left[\sum_{i=1}^t \omega_{t,i} \mathbf{x}_i y_i + \omega_{t,0} \mathbf{R}_0 \hat{\phi}_0 \right]. \quad (7)$$

For the second term notice that

$$\begin{aligned} \sum_{i=1}^t \omega_{t,i} \mathbf{x}_i y_i &= \sum_{i=1}^{t-1} \omega_{t,i} \mathbf{x}_i y_i + \gamma_t \mathbf{x}_t y_t, \\ &= (1 - \gamma_t) \sum_{i=1}^{t-1} \omega_{t-1,i} \mathbf{x}_i y_i + \gamma_t \mathbf{x}_t y_t, \end{aligned}$$

and

$$\omega_{t,0} \mathbf{R}_0 \hat{\phi}_0 = (1 - \gamma_t) \omega_{t-1,0} \mathbf{R}_0 \hat{\phi}_0,$$

where we use

$$\omega_{t,i} = (1 - \gamma_t) \omega_{t-1,i},$$

which follows from (6). Hence, (7) is equivalent to

$$\hat{\phi}_t = \mathbf{R}_t^{-1} \left[\gamma_t \mathbf{x}_t y_t + (1 - \gamma_t) \left(\sum_{i=1}^{t-1} \omega_{t-1,i} \mathbf{x}_i y_i + \omega_{t-1,0} \mathbf{R}_0 \hat{\phi}_0 \right) \right]. \quad (8)$$

Lagging (7) one period we find that

$$\mathbf{R}_{t-1} \hat{\phi}_{t-1} = \sum_{i=1}^{t-1} \omega_{t-1,i} \mathbf{x}_i y_i + \omega_{t-1,0} \mathbf{R}_0 \hat{\phi}_0,$$

which can be substituted into (8) to yield

$$\hat{\phi}_t = \mathbf{R}_t^{-1} \left[\gamma_t \mathbf{x}_t y_t + (1 - \gamma_t) \mathbf{R}_{t-1} \hat{\phi}_{t-1} \right]. \quad (9)$$

From (3) notice that

$$(1 - \gamma_t) \mathbf{R}_{t-1} = \mathbf{R}_t - \gamma_t \mathbf{x}_t \mathbf{x}_t',$$

which substituted into (9) and after rearranging leads to

$$\begin{aligned} \hat{\phi}_t &= \mathbf{R}_t^{-1} \left[\gamma_t \mathbf{x}_t y_t + (\mathbf{R}_t - \gamma_t \mathbf{x}_t \mathbf{x}_t') \hat{\phi}_{t-1} \right], \\ &= \gamma_t \mathbf{R}_t^{-1} \mathbf{x}_t y_t + \hat{\phi}_{t-1} - \gamma_t \mathbf{R}_t^{-1} \mathbf{x}_t \mathbf{x}_t' \hat{\phi}_{t-1}, \\ &= \hat{\phi}_{t-1} + \gamma_t \mathbf{R}_t^{-1} \mathbf{x}_t (y_t - \mathbf{x}_t' \hat{\phi}_{t-1}), \end{aligned}$$

establishing the correspondence between the penalized WLS solution of (5) and the RLS of (2).

Penalized WLS in matrix form It is sometimes useful to have estimators in matrix form. For the case with arbitrary initials, $\hat{\phi}_0$ and \mathbf{R}_0 , the RLS is equivalent to the solution of a penalized regression problem given by

$$\hat{\phi}_t = \arg \min \left\{ \left(\mathbf{y}_t - \hat{\phi}_t' \mathbf{X}_t \right) \Omega_t \left(\mathbf{y}_t' - \mathbf{X}_t' \hat{\phi}_t \right) + \omega_{t,0} \left(\hat{\phi}_0' - \hat{\phi}_t' \right) \mathbf{R}_0 \left(\hat{\phi}_0 - \hat{\phi}_t \right) \right\},$$

where $\mathbf{X}_t = (\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_t)$ is a $(k \times t)$ matrix collecting the regressors' time series, $\mathbf{y}_t = (y_1, y_2, \dots, y_t)$, $\Omega_t = \text{diag}(\omega_{t,1}, \omega_{t,2}, \dots, \omega_{t,t})$ is a weighting matrix with the sequence of weights in the main diagonal and zeros otherwise. The solution to this minimization problem is then given by

$$\hat{\phi}_t = (\mathbf{X}_t \Omega_t \mathbf{X}_t' + \omega_{t,0} \mathbf{R}_0)^{-1} (\mathbf{X}_t \Omega_t \mathbf{y}_t' + \omega_{t,0} \mathbf{R}_0 \hat{\phi}_0).$$

Relation to literature To the best of the author's knowledge this non-recursive formulation of the RLS for arbitrary initials has never been outlined in the previous literature. Berardi

and Galimberti (2013), for example, have drawn a correspondence between the RLS and the standard (without penalty) WLS, except that under the assumption of a diffuse initial prior, when $\mathbf{R}_0 \rightarrow \mathbf{0}$. From a Bayesian point of view, \mathbf{R}_t is inversely related to the uncertainty in the corresponding Kalman filter estimates of ϕ_t (see Evans et al., 2010). Hence, $\mathbf{R}_0 \rightarrow \mathbf{0}$ can be interpreted as increasing the uncertainty about the initial estimates, and it is only under this diffuse prior assumption that the standard WLS is equivalent to the RLS. Also, notice that if $\mathbf{R}_0 = \mathbf{0}$ (exactly rather than as a limit), (2)-(3) implies that $\hat{\phi}_1 = (\mathbf{x}_1 \mathbf{x}'_1)^{-1} \mathbf{x}_1 y_1$, which will be indeterminate² for $k > 1$.

For applied purposes, the learning process implied by the RLS of (2)-(3) will require the specification of initials. Hence, an appropriate account of the implied relative weights given to sample observations can only be obtained by including the penalty on deviations of the estimates from the learning initials, as in (5). The weight given to this penalty, nevertheless, is adjusted according to the weights given to the sample observations and the initial estimate of the regressors' matrix of second moments. We discuss the relevance of these initials and the observation weights in the next section.

3 Information weighting

The weight given to a sample observation determines the amount of information from that particular observation that is incorporated into the parameters' estimates. The non-recursive estimator above allows the calculation of such weights for any arbitrary sequence of learning gains. Also notice that the weights, $\omega_{t,i}$, defined in equation (6), are already in relative terms, though this equivalence does not hold under diffuse initials.

Absolute and relative weights What in fact matters for estimation are the relative weights instead of their absolute values. To see that, consider the effects of multiplying the sequence of weights in the solution to the penalized WLS problem, $\{\omega_{t,0}, \omega_{t,1}, \dots, \omega_{t,t}\}$ in equation (5), by a constant κ ; clearly, such a re-scaling of the weights will have no effect over the resulting estimates because the constant factor entering in the numerator will cancel out with that entering in the denominator of the estimator. In order to calculate the relative weights we need to divide the absolute weights by the sum of weights given to sample observations and the initial. It turns out that the relationship defined between the weights and the learning gains implies that this sum will always be equal to unity. Letting W_t^n stand for the sum of weights starting from weight n up to weight t , from the definition of the absolute weights, (6), this sum of weights

²More precisely, $\mathbf{x}_1 \mathbf{x}'_1$ must have full rank for a non-degenerate inverse to exist; to guarantee the existence of this inverse, $\gamma_t = 1$ must also be ruled out.

can be expanded according to

$$\begin{aligned}
W_t^0 &= \sum_{i=0}^t \omega_{t,i}, \\
&= \prod_{j=1}^t (1 - \gamma_j) + \sum_{i=1}^{t-1} \gamma_i \prod_{j=i+1}^t (1 - \gamma_j) + \gamma_t.
\end{aligned} \tag{10}$$

Expanding the first term of (10) we have that

$$\begin{aligned}
\omega_{t,0} &= (1 - \gamma_1)(1 - \gamma_2) \dots (1 - \gamma_{t-1})(1 - \gamma_t), \\
&= (1 - \gamma_2) \dots (1 - \gamma_{t-1})(1 - \gamma_t) - \gamma_1 \prod_{j=2}^t (1 - \gamma_j), \\
&= 1 - \gamma_t - \sum_{i=1}^{t-1} \gamma_i \prod_{j=i+1}^t (1 - \gamma_j).
\end{aligned} \tag{11}$$

Returning to (10) we then have

$$\begin{aligned}
W_t^0 &= 1 - \gamma_t - \sum_{i=1}^{t-1} \gamma_i \prod_{j=i+1}^t (1 - \gamma_j) + \sum_{i=1}^{t-1} \gamma_i \prod_{j=i+1}^t (1 - \gamma_j) + \gamma_t, \\
&= 1.
\end{aligned}$$

Hence, in the context of the correspondence between the RLS and the penalized WLS outlined in this paper, the relative weights will be equal to their corresponding absolute weights, and we can focus directly on the weights defined by equation (6). Moreover, it is often more interesting to evaluate how the observation weights evolve relative to the last observation in the sample, i.e., redefining equation (6) in terms of lags relative to the end of the sample, $\bar{\omega}_{t,l} = \omega_{t,t-l}$, in which case we look at

$$\bar{\omega}_{t,l} = \begin{cases} \gamma_{t-l} \prod_{j=0}^{l-1} (1 - \gamma_{t-j}) & \text{for } 0 < l < t, \\ \gamma_t & \text{for } l = 0. \end{cases} \tag{12}$$

For clarity of notation, in what follows we also add superscripts to the different definitions of learning gains and their implied weights.

3.1 Decreasing-gain

The decreasing-gain (DG) specification has been prominent in the adaptive learning literature since the seminal contributions of Bray (1982); Marcet and Sargent (1989). Under a decreasing

gain, traditionally modeled as $\gamma_t^{dg} = 1/(t+1)$, the weights are given by

$$\omega_{t,i}^{dg} = \begin{cases} \prod_{j=1}^t \left(\frac{j}{j+1}\right) = \frac{1}{2} \frac{2}{3} \cdots \frac{t-1}{t} \frac{t}{t+1} = \frac{1}{t+1} & \text{for } i = 0, \\ \frac{1}{i+1} \prod_{j=i+1}^t \left(\frac{j}{j+1}\right) = \frac{1}{i+1} \frac{i+1}{i+2} \cdots \frac{t}{t+1} = \frac{1}{t+1} & \text{for } 0 < i < t, \\ \frac{1}{t+1} & \text{for } i = t, \end{cases}$$

and every observation receives an equal weight that is decreasing with the sample size, which makes the DG-RLS particularly interesting for the analysis of learning convergence towards equilibrium. Particularly, notice that asymptotically, as $t \rightarrow \infty$, the weight given by the DG-RLS to any observation, and the initial, decreases towards zero.

OLS equivalence breakup The DG-RLS is often motivated as representative of real-time econometricians due to its resemblance with the Ordinary Least Squares (OLS) estimator from econometrics,

$$\hat{\phi}_t^{ols} = \left[\sum_{i=1}^t \mathbf{x}_i \mathbf{x}_i' \right]^{-1} \left[\sum_{i=1}^t \mathbf{x}_i y_i \right].$$

Considering the prominence of the OLS estimator, an equivalence with the RLS may provide a rationale for adaptive learning as a device representative of, say, real-time econometricians. However, as equation (5) now makes clear, such an equivalence would require disregarding the learning initials (or assuming a diffuse prior as discussed above).

3.2 Generalized decreasing-gain

A recent strand of the literature has revived the decreasing-gain specification with the proposal of so-called learning from experience, where agents expectations are modeled at cohort-level (Malmendier and Nagel, 2016, MN). Under this type of learning individuals' beliefs are more strongly affected by the data observed during their lifetimes, hence leading to a dispersion in the cross-section of expectations depending on the age structure of the population of agents; MN find some empirical support for this result from consumer expectations data collected through surveys. Their modeling of learning is based on a generalization of the decreasing-gain setup

(denoted by GDG-RLS in what follows) by assuming that³

$$\gamma_t^{gdg} = \frac{\theta}{t + \theta}, \quad (13)$$

with $\theta > 0$; in fact, notice that the original DG-RLS is obtained when $\theta = 1$. Otherwise, it is interesting to cast the GDG-RLS weights in lag recursive form, starting from $\omega_{t,0}^{gdg} = \theta/t + \theta$ and expanding to

$$\begin{aligned} \omega_{t,1}^{gdg} &= \omega_{t,0}^{gdg} \left(\frac{t}{t-1+\theta} \right), \\ \omega_{t,2}^{gdg} &= \omega_{t,1}^{gdg} \left(\frac{t-1}{t-2+\theta} \right), \\ &\vdots \\ \omega_{t,l}^{gdg} &= \omega_{t,l-1}^{gdg} \left(\frac{t-l+1}{t-l+\theta} \right), \end{aligned}$$

which makes clear that the weight given to lagged observations within a sample of data decreases with the lag if $\theta > 1$, and increases if $\theta < 1$. Also note that the higher (lower) θ is above (below) 1, the quicker the decay (ascent) in the relative weights given by the GDG-RLS to observations farther into the past.

Age-dependent experiences The GDG specification also implies a time-varying profile of weights as the sample size grows. In behavioral terms this is what MN (p. 59) refer as “experiences earlier and later in life to have a different influence” on expectations. To see this consider how the weight given to a fixed observation lag evolves as the sample increases by one observation; the ratio between the two corresponding weights is given by

$$\begin{aligned} \frac{\omega_{t,l}^{gdg}}{\omega_{t-1,l}^{gdg}} &= \frac{\frac{\theta}{t-l+\theta} \prod_{j=0}^{l-1} \left(\frac{t-j}{t-j+\theta} \right)}{\frac{\theta}{t-1-l+\theta} \prod_{j=0}^{l-1} \left(\frac{t-1-j}{t-1-j+\theta} \right)}, \\ &= \frac{t}{t-l} \left(\frac{t-1-l+\theta}{t+\theta} \right), \end{aligned} \quad (14)$$

which is smaller than 1 if $\theta < t/l$; this latter condition is always true for the latest observations ($l = 0$)—younger individuals are more strongly influenced by the latest observation than older


³The exact gain specification in MN is actually given by

$$\gamma_{t,s}^{mn} = \begin{cases} \frac{\theta}{t-s} & \text{if } t-s \geq \theta, \\ 1 & \text{if } t-s < \theta, \end{cases}$$

where $t-s$ stands for the age of the agent; however, from a computational point of view, this specification is clearly problematic because it can lead to degenerate inverses when $t-s \leq \theta$, as discussed in footnote (2), and indeterminate initials. The alternative proposed in equation (13) is not affected by this issues at the same time that it captures the intended weighting profiles.

ones—and for $\theta \leq 1$. For lagged observations, the weight decreases with the sample size as long as θ is small enough relative to the sample size/lag ratio. These effects are depicted in Figures 1a and 1b.

Figure 1: Generalized decreasing-gain weights.



Notes: Weights calculated for quarterly data using eqs. (6)-(13) and $\theta = 3$. The circles in panel (b) indicate the turning points of the corresponding curves.

Recurring memories When $\theta > 1$, the weight given to some lagged observations may increase as the sample size grows—although an individual’s memories of past experiences tend to fade away, “recurring memories” may have a comeback as the individual ages. Here, by recurring memories we mean the memories recalled over a moving window of lagged observations⁴. This effect is more clearly illustrated in Figure 1b: notice how the weights given to 10 and 20-years lagged observations increase during an individual’s lifetime at the same time that the weights given to more recent observations are decreasing; after a threshold age, here of 30 years for the 10-years lag and 60 years for the 20-years lag, the lagged weights start to decrease as well, giving rise to a hump-shaped profile of lagged weights. Also notice, from equation (14), that the higher θ the earlier the effects of recurring memories will start to show up within the agent’s lifetime. Although this property seems to have gone unnoticed so far in the learning from experience literature, its effects are likely to be of secondary importance in the explanation of cohort level expectations data. Particularly, notice that the increase in the lagged weights tend to be small relative to the weights given to more recent observations.

3.3 Constant-gain

The constant-gain (CG) learning specification was introduced in the applied learning literature by Evans and Honkapohja (1993) and became popular after Sargent (1999) for its improved

⁴An earlier version of this paper denoted these as “dormant memories,” which would be more appropriate if a weight $\omega_{t,i}^{dgs}$ increased over time for a given i .

capability of tracking the evolution of time-varying environments. This specification has also been under the spotlight of the most recent research on the dynamic modeling of expectations for its potential of generating escape dynamics over finite stretches of time (see Williams, 2019) and asymptotically stable distributions of beliefs (see Galimberti, 2019). Under CG-RLS learning, $\gamma_t^{cg} = \bar{\gamma}$, and the weights are given by

$$\omega_{t,i}^{cg} = \begin{cases} \prod_{j=1}^t (1 - \bar{\gamma}) = (1 - \bar{\gamma})^t & \text{for } i = 0, \\ \bar{\gamma} \prod_{j=i+1}^t (1 - \bar{\gamma}) = \bar{\gamma} (1 - \bar{\gamma})^{t-i} & \text{for } 0 < i \leq t. \end{cases}$$

i.e., the sequence of sample observation weights is related to the learning gain by a geometric progression given by $\omega_{t,i}^{cg} = \bar{\gamma} (1 - \bar{\gamma})^{t-i}$. Hence, the relative weights given to past information by the CG-RLS decrease with the observation lag ($l = t - i$); moreover, notice that this result also holds asymptotically: although the CG-RLS also assigns a vanishing weight to any given sample observation i , the information weight given to observations at a fixed lag l do not change with t . This property makes the CG-RLS particularly well suited for modeling the behavioral assumption that agents give a higher emphasis to the more recent observations than to those collected farther into the past.

Persistently influential initials In contrast with the decreasing-gain specification, under the CG-RLS the influence of the learning initials can become non-negligible for applied purposes. Particularly, whereas $\omega_{t,0}^{cg}$ also decreases with the sample size, the duration of the effects of the initials within finite samples will depend on the learning gain. In fact, it is possible to calculate how far, within a sample of data, the initial will have a greater weight than the data observation by solving:

$$\begin{aligned} \omega_{t,i^*}^{cg} &= \omega_{t,0}^{cg}, \\ \bar{\gamma} (1 - \bar{\gamma})^{t-i^*} &= (1 - \bar{\gamma})^t, \\ i^* &= \frac{\log(\bar{\gamma})}{\log(1 - \bar{\gamma})}, \end{aligned}$$

which is decreasing with $\bar{\gamma}$, i.e., the lower the learning gain the longer it takes for a data point to receive a weight higher than the initial. Strikingly, for a gain value of $\bar{\gamma} = 0.02$, which is typically found in applications with quarterly macroeconomic data (see Berardi and Galimberti, 2017a, for a review), it takes over 193 observations (or over 48 years of quarterly data) for the CG-RLS to assign a weight to the latest observation as high as the weight given to the learning initial. Alternatively, one may be interested in the cumulative weight given to the sample of

observations relative to the initial:

$$\sum_{j=1}^{i^*} \omega_{t,j}^{cg} = \omega_{t,0}^{cg},$$

$$i^* = \frac{\log(1/2)}{\log(1 - \bar{\gamma})},$$

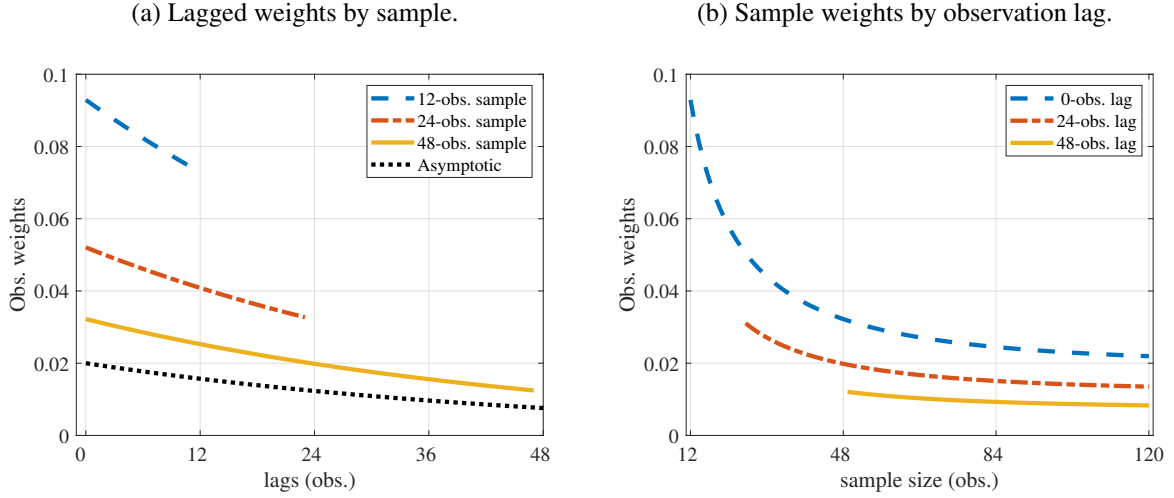
which is again decreasing with $\bar{\gamma}$, but leading to less startling accounts of the initials effect: for $\bar{\gamma} = 0.02$, $i^* \simeq 34$, i.e., it takes about $8^{1/2}$ years of quarterly data for the CG-RLS to assign a higher weight to the sample of observations than the learning initials in the PLM estimates.

Finite sample distortion under diffuse initials Without a proper initialization of R_0 , the weights given to lagged observations by the CG-RLS can decay faster than the profile of weights expected from its asymptotic operation. Particularly, under diffuse initials, $R_0 = 0$, the CG-RLS sample weights are given by

$$\omega_{t,l}^{dcg} = \frac{\bar{\gamma}(1 - \bar{\gamma})^l}{1 - (1 - \bar{\gamma})^t}, \quad (15)$$

which are declining not only with the observation lag, but also with the size of the sample. These effects are illustrated in Figures 2a and 2b; notice the asymptotic weights depicted in Figure 2a are in fact equivalent to $\omega_{t,l}^{cg}$. Hence, although the relative sample weights also decrease as the observation becomes outdated, the actual profile of relative sample weights is not time-invariant, which can distort the behavioral interpretation of a constant gain in small samples. Interestingly, in a parallel with the age-dependent gain specification discussed above, letting t stand for age, equation (15) implies that a younger agent would assign a higher weight to any given observation than an older one experiencing the same observation. We will return to a quantification of this effect in the simulation analysis of Section 4 below.

Figure 2: Constant-gain weights under diffuse initials.



Notes: Weights calculated using equation (15) with $\bar{\gamma} = 0.02$.

3.4 Time-varying gains

Time-varying gains offer an alternative to relax the determination of information weighting under learning. In the literature, this has been achieved either by merging the decreasing- and constant-gain specifications (Marcet and Nicolini, 2003; Milani, 2014), or by turning the determination of the gains endogenous with an additional adaptation mechanism (Kostyshyna, 2012; Berardi and Galimberti, 2017a). In both cases the learning gain is adjusted according to the recursive forecasting performance of the implied expectations, increasing/decreasing the gain (or switching from decreasing- to constant-gain and vice-versa) during periods of elevated/low forecasting errors. These time-varying approaches allow the modeling of behavioral shifts of attention that agents give to incoming data, which could be motivated as a concern with structural changes.

Adaptive gains revisited Benveniste et al. (1990) and Kushner and Yang (1995) propose the use of an outer mechanism to adjust the learning gain for tracking slowly time-varying systems. In its most general form the gains are adjusted according to

$$\gamma_t = [\gamma_{t-1} + \alpha \mathbf{x}'_t \hat{\Psi}_{t-1} (y_t - \mathbf{x}'_t \hat{\phi}_{t-1})]_{\gamma_-}^{\gamma_+}, \quad (16)$$

where $\hat{\Psi}_t$ keeps track of the algorithm's past estimation performance, α is an adaptation constant, and $[\bullet]_{\gamma_-}^{\gamma_+}$ is a truncation operator. In this approach, the gain is increased (decreased) following evidence of systematic (contradictory) estimation errors (see also Berardi and Galimberti, 2017a, p.226-8, for more details and intuition). However, one interesting implication of our results is that the learning gain does not need to increase for the weight given to the latest

observation to rise; more formally, $\bar{\omega}_{t,0} > \bar{\omega}_{t,1}$ only requires that

$$\begin{aligned}\gamma_t &> \gamma_{t-1}(1 - \gamma_t), \\ \gamma_t &> \frac{\gamma_{t-1}}{1 + \gamma_{t-1}}.\end{aligned}\tag{17}$$

Hence, our results imply that the symmetric adaptation of learning gains, according to (16), leads to an asymmetric over/under-weighting of new observations following to evidence of systematic/contradictory estimation errors. In fact, Kushner and Yang (1995, p.1405) suggest adding a correction term in (16) to stabilize γ_t around values that are slightly smaller than the optimal value implied by the stochastic approximation used to derive this mechanism.

Gain-switching revisited Marcet and Nicolini (2003) propose an endogenous gain mechanism that can be represented in our framework as⁵

$$\gamma_t^{gs} = \begin{cases} \frac{\bar{\gamma}}{1+s\bar{\gamma}} & \text{if } \mathcal{C} \\ \bar{\gamma} & \text{otherwise,} \end{cases}\tag{18}$$

where $s = t - t_s$ stands for the number of periods since the last time the constant-gain was used, and \mathcal{C} for a regime-switching condition. Using equation (17) we find that a switch from the decreasing-gain, $\gamma_{t-1}^{gs} = \bar{\gamma}/(1 + (s-1)\bar{\gamma})$, to the constant-gain regime, $\gamma_t^{gs} = \bar{\gamma}$, would always imply that the latest observation will be given a higher weight than the previous one—which is consistent with the structural change/tracking rationale given to this learning mechanism. When the switch goes on the other direction, from the constant-gain, $\gamma_{t-1}^{gs} = \bar{\gamma}$, to a decreasing-gain, $\gamma_t^{gs} = \bar{\gamma}/(1 + \bar{\gamma})$, equation (17) turns into an equality, which means the new observation is given a weight that is equal to that given to the previous observation. Table 1 summarizes the effects of all possible regime transitions on the weight given to the new observation, both within the latest sample (1a) and relative to the previous sample (1b). It is interesting to note that the magnitude of the effect of some of these regime switches depends on $\bar{\gamma}$ and s . Particularly, a higher constant learning gain leads to a greater weight to the new observation not only within the CG learning regime, but also in the transition from DG ($t-1$) to CG (t), and the latter effect will be greater the longer learning has operated under DG before the switch.

⁵The decreasing-gain specification in Marcet and Nicolini (2003) is actually equivalent to $\gamma_t = \frac{\gamma_{t-1}}{1+\gamma_{t-1}}$, which implies, for an initial gain given by $\gamma_0 = \bar{\gamma}$, the formulation presented in equation (18) with $t_s = 0$.

Table 1: Gain-switching weighting dynamics.

(a) Last ($\bar{\omega}_{t,0}$) vs. lagged ($\bar{\omega}_{t,1}$) weight:		(b) Last ($\bar{\omega}_{t,0}$) vs. previous ($\bar{\omega}_{t-1,0}$) weight:			
	DG ($t-1$)	CG ($t-1$)		DG ($t-1$)	CG ($t-1$)
DG (t)	Equal	Equal	DG (t)	Decrease (more for $\uparrow \bar{\gamma}$ and $\downarrow s$)	Decrease (more for $\uparrow \bar{\gamma}$)
CG (t)	Increase (more for $\uparrow \bar{\gamma}$ and $\uparrow s$)	Increase (more for $\uparrow \bar{\gamma}$)	CG (t)	Increase (more for $\uparrow \bar{\gamma}$ and $\uparrow s$)	Equal

4 Simulation and empirical application

One key finding of our analysis of information weighting under least squares learning is that the assumption of diffuse initials can distort the profile of weights given to sample observations by a constant-gain mechanism. Given the popularity of this learning mechanism, we now turn to a quantification of these effects with a simulation of the estimation of a macroeconomic model and an empirical application. Particularly, we focus on a standard new Keynesian Phillips curve model, given by

$$\begin{aligned}\pi_t &= \beta \pi_{t+1}^e + \lambda x_t + \alpha + u_t, \\ x_t &= \rho x_{t-1} + v_t,\end{aligned}\tag{19}$$

where π_t is inflation, π_{t+1}^e represents agents' expectations for next period inflation, x_t is a proxy for real marginal cost, often measured as output gaps, and u_t is a disturbance. Under adaptive learning agents form expectations according to a PLM given by

$$\pi_t = \phi_{t-1} x_t + \varphi_{t-1} + z_t,$$

where ϕ_t and φ_t are parameters estimated with the LS algorithm of (2)-(3), and are expected to converge to the rational expectations equilibrium (REE), $\phi^* = \lambda / (1 - \beta\rho)$ and $\varphi^* = \alpha / (1 - \beta)$, as long as $\beta < 1$ and $\beta\rho < 1$ (E-stability condition, see Evans and Honkapohja, 2001, pp. 198-200).

4.1 Simulation analysis

We generate 10,000 samples of artificial series of π_t and x_t assuming that $\gamma = 0.03$, $\beta = 0.99$, $\lambda = 0.2$, $\alpha = 0$, $\rho = 0.7$, $u_t \sim N(0, 3)$, and $v_t \sim N(0, 1)$. For each sample we simulate the model for 2,000 observations, discarding the first 1,000 and using the remaining data for estimation of $\hat{\beta}$, under varying sample sizes and \hat{R}_0 prior assumptions. We focus on the estimation of $\hat{\beta}$ because this is a key parameter determining the relevance of forward-looking expectations for inflation in this model. All other parameters, including the learning initial ϕ_0 , are set to their actual values—because we set $\alpha = 0$, we drop the intercept from the PLM.

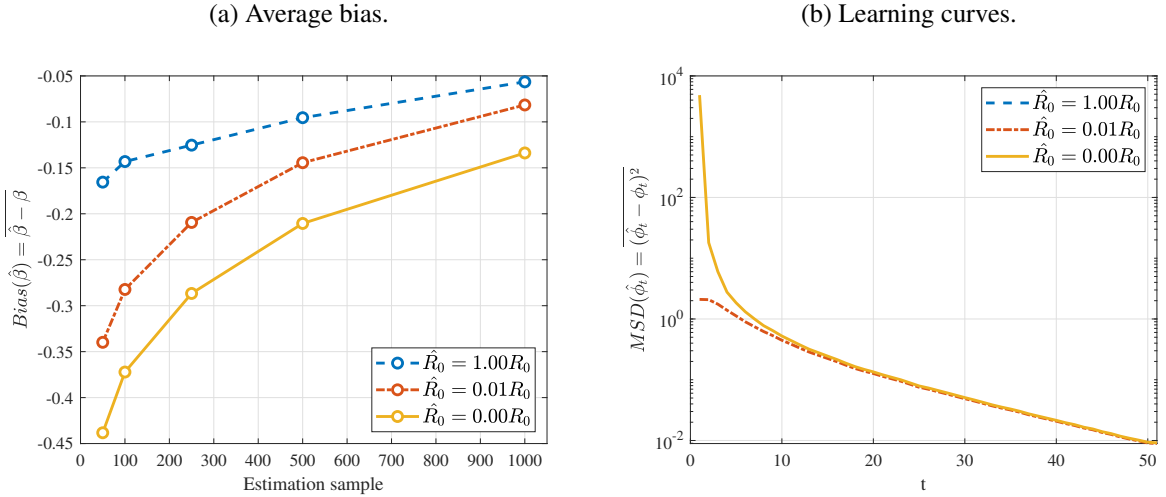
Under learning, the actual law of motion is given by

$$\pi_t = (\beta\rho\phi_{t-1} + \lambda)x_t + u_t.$$

Since we are fixing γ , ρ , and λ , estimation of $\hat{\beta}$ is linear on the interaction between ϕ_{t-1} and x_t . Hence, we estimate $\hat{\beta}$ by regressing $\pi_t - \lambda x_t$ on $\rho\phi_{t-1}x_t$ and a constant using OLS—our results on $\hat{\beta}$ are qualitatively similar when adding $\hat{\lambda}$ or $\hat{\gamma}$ for joint estimation, although the latter requires nonlinear estimation methods. To avoid effects of unstable dynamics we discard simulations where $\hat{\beta} < 0$ or $\hat{\beta} > 1/\rho$, the latter following from the E-stability condition.

We evaluate the effect of a diffuse initial by considering how alternative assumptions of \hat{R}_0 affect the average estimates of $\hat{\beta}$. Figure 3a presents the average estimation results over different estimation samples—here we focus on mean deviations (bias) between the estimates and their simulated values, but their mean squared deviations (variance) lead to similar results on the positive domain. The $\hat{\beta}$ estimates are, on average, below their actual value. As expected, these estimates tend to converge to their actual values for bigger estimation samples. However, the use of diffuse initials, $\hat{R}_0 = 0$, leads to greater biases, which also tend to be stronger for the smaller estimation samples. The learning curves presented in Figure 3b offer an explanation for these effects: the diffuse initials make the learning estimates oversensitive to the first observations in the sample, hence confirming our prediction that this approach can add up to finite sample distortions in the estimation of models with learning.

Figure 3: Constant-gain simulated estimation exercise.



Notes: Statistics are averaged across repeated simulations of model (19), discarding boundary estimates as described in the text. The blue dashed line, corresponding to $\hat{R}_0 = 1.00R_0$, does not show up in 3b because its corresponding MSD is always equal to 0 (by definition) and the MSDs are depicted in logarithmic scale.

Quantitatively, we find that the use of diffuse initials lead to $\hat{\beta}$ estimates between 0.44 and 0.13 below the true value, and 0.27 and 0.08 below the estimates obtained with the correct

initial, for the samples with 50 and 1,000 observations, respectively. This parameter is among the determinants of the slope of the Phillips curve in this model, or how responsive inflation is to changes in real marginal costs. Under learning this slope will be varying with ϕ_t , but taking this model’s REE as a reference, a $\hat{\beta}$ deviation of 0.44 from its true value implies a Phillips curve slope estimate about half the magnitude (or “flatter”) of its true value, 0.33 versus 0.65, respectively—notice part of this effect comes from the small sample distortion of estimates as the use of the correct initial still leads to a downward biased slope of 0.47, on average. This is equivalent to saying that estimation of this model with diffuse initials, and a sample of 50 observations, would lead a researcher to infer that when this economy is producing 10% above its full-employment capacity, prices would be rising at a 3.3% inflation rate, when they will actually be rising 6.5% (plus a random shock). Hence, misspecified initials, and their associated distortions to how sample information is weighted under constant-gain learning, can lead to flatter Phillips curve estimates.

4.2 Empirical estimates

We now turn to an empirical evaluation of the effects of diffuse initials on the estimation of the standard new Keynesian Phillips curve model under constant-gain learning, as outlined above. For this purpose we obtain quarterly US data on CPI and real GDP covering the period from 1947 to 2019. We then measure inflation, π_t , by the log change in CPI relative to the same quarter in the previous year, and derive output gap measures using the Hodrick-Prescott filter on the log real GDP series to stand for the forcing variable, x_t .

We again focus on the estimation of $\hat{\beta}$, while pre-fixing the other parameters to plausible values: $\gamma = 0.03$ is in the range of calibrations reported by Berardi and Galimberti (2017a, Fig.8) to match survey forecasts of US CPI inflation from professionals, consumers, and policymakers; $\lambda = 0.05$ is close to previous empirical estimates of hybrid forms of this model in the literature, which tend to be positive and borderline significant (see Mavroeidis et al., 2014, for a survey); and, $\rho = 0.85$ is directly obtained by estimating the first-order autocorrelation in the output gap series. For the learning initials we use pre-sample data to estimate $\{\phi_0, \varphi_0\}$ using WLS in order to obtain initials consistent with the constant-gain learning adopted in the estimation sample (see Berardi and Galimberti, 2017b). We then consider two variants for initial matrix of second moments: (i) the estimate obtained with the pre-sample training, $\hat{R}_0 = R_0^{wls}$; (ii) a re-scaled version to represent the diffuse initial approach, $\hat{R}_0 = 10^{-6} R_0^{wls}$.

In order to avail the effects of different estimation samples, we fit the model both with the full-sample and across sub-samples of the data. Table 2 presents the estimation results. The $\hat{\beta}$ estimates are generally positive and statistically significant across the different samples and \hat{R}_0 assumptions. We also find that the implied estimates of the slope of the Phillips curve have decreased during the so-called Great Moderation period (1986–2006), which is consistent with previous evidence in the literature (Slobodyan and Wouters, 2012a), but rebounded since

the financial crisis of 2007-09. The latter is also consistent with the broader view on the role of deviations from rational expectations as an explanation for the apparent continuation of a flattening of the US Phillips curve in recent periods (see Coibion et al., 2018). Namely, we find no evidence of a flattening following the financial crisis when expectations are modeled according to adaptive learning.

Table 2: Empirical estimates of a Phillips curve model with constant-gain learning.

	1965-2019 (220 quarters)	1965-1985 (84 quarters)	1986-2006 (84 quarters)	2007-2019 (52 quarters)
- Under pre-initialized \hat{R}_0 :				
$\hat{\beta}$	0.938*** (0.258)	0.750** (0.323)	0.486** (0.240)	1.040** (0.464)
implied $\tilde{\phi}$	0.247	0.138	0.085	0.433
- Under diffuse \hat{R}_0 :				
$\hat{\beta}$	0.890*** (0.213)	0.668*** (0.195)	0.387* (0.227)	0.481** (0.195)
implied $\tilde{\phi}$	0.206	0.116	0.074	0.085

Notes: Standard errors in parentheses are heteroscedasticity and autocorrelation robust (Newey and West, 1994), and statistical significance at the 1%, 5%, and 10% levels of significance are indicated by ***, **, *, respectively. The implied slope of the Phillips curve is obtained according to the REE.

We find that the use of diffuse learning initials leads to systematically lower estimates of $\hat{\beta}$, confirming our simulation evidence. This means that the misspecification of that crucial learning parameter can lead to misleading conclusions about the sensitiveness of inflation to expectations and its effects on output gaps. In fact, using a pre-initialized \hat{R}_0 , we find a sharp increase in the relevance of expectations in the latest sample, whereas under diffuse initials the increase is less evident—the standard errors of the point estimates, nevertheless, suggest these changes across sub-samples are not statistically significant at standard levels of significance.

5 Concluding remarks

This paper proposed a more general non-recursive representation of the recursive least squares algorithm that is used in the adaptive learning literature to represent how agents form their expectations in economic settings. According to this new formulation, the recursive learning mechanism is more properly represented by a penalized weighted least squares estimator, where a penalty term accounts for the effects of the learning initial estimates. The non-recursive formulation also allowed a renewed analysis of how information is weighted in the implied estimates of agents' perceived law of motion. Such weights are directly determined by the sequence of learning gains used in the recursive least squares algorithm, and the framework

proposed in this paper provides flexible analytical expressions for their calculation under different assumptions on the evolution of the learning gains.

Drawing on this refreshed framework this paper provided renewed behavioral implications of alternative assumptions about the learning gains. One interesting finding is that the application of decreasing-gains to cohort-level data, an approach that has found empirical support from recent research on consumers' expectations survey data, can lead to the emergence of a hump-shaped profile of weights to lagged observations. Another key finding is that, without a proper account for the learning initial, the estimation of models under the assumption of a constant gain over increasing samples of data would imply agents give a decreasing weight to more recent observations, distorting the real-time tracking interpretation of this mechanism. The relevance of this distortion was evidenced by simulation and empirical exercises, where the misspecified initials led to a systematic underestimation of the slope of a Phillips curve model. Hence, a proper account of how information is weighted under alternative learning mechanisms is an important aspect for the estimation of models of imperfect information such as adaptive learning. This paper attempted to provide useful guidelines on the design and choice of such learning mechanisms for applied economics research.

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 (2017a). Empirical calibration of adaptive learning. *Journal of Economic Behavior & Organization* 144, 219 – 237.
- Berardi, M. and J. K. Galimberti (2017b). On the initialization of adaptive learning in macroeconomic models. *Journal of Economic Dynamics and Control* 78, 26 – 53.
- Bray, M. (1982). Learning, estimation, and the stability of rational expectations. *Journal of Economic Theory* 26(2), 318–339.
- Carceles-Poveda, E. and C. Giannitsarou (2007). Adaptive learning in practice. *Journal of Economic Dynamics and Control* 31(8), 2659–2697.
- Coibion, O., Y. Gorodnichenko, and R. Kamdar (2018, December). The formation of expectations, inflation, and the phillips curve. *Journal of Economic Literature* 56(4), 1447–91.

- Evans, G. W. and S. Honkapohja (1993). Adaptive forecasts, hysteresis, and endogenous fluctuations. *Economic Review*, 3–13.
- Evans, G. W. and S. Honkapohja (2001). *Learning and expectations in macroeconomics*. Frontiers of Economic Research. Princeton, NJ: Princeton University Press.
- Evans, G. W., S. Honkapohja, and N. Williams (2010). Generalized stochastic gradient learning. *International Economic Review* 51(1), 237–262.
- Galimberti, J. K. (2019). An approximation of the distribution of learning estimates in macroeconomic models. *Journal of Economic Dynamics & Control* 102, 29–43.
- Kostyshyna, O. (2012). Application of an adaptive step-size algorithm in models of hyperinflation. *Macroeconomic Dynamics* 16, 355–375.
- Kushner, H. and J. Yang (1995, aug). Analysis of adaptive step-size sa algorithms for parameter tracking. *Automatic Control, IEEE Transactions on* 40(8), 1403 –1410.
- Ljung, L. and T. Soderstrom (1983). *Theory and Practice of Recursive Identification*. The MIT Press.
- Malmendier, U. and S. Nagel (2016). Learning from inflation experiences. *The Quarterly Journal of Economics* 131(1), 53–87.
- Marcet, A. and J. P. Nicolini (2003). Recurrent hyperinflations and learning. *American Economic Review* 93(5), 1476–1498.
- 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.
- Mavroeidis, S., M. Plagborg-Møller, and J. H. Stock (2014). Empirical evidence on inflation expectations in the new keynesian phillips curve. *Journal of Economic Literature* 52(1), 124–88.
- Milani, F. (2014). Learning and time-varying macroeconomic volatility. *Journal of Economic Dynamics and Control* 47, 94–114.
- Newey, W. K. and K. D. West (1994). Automatic lag selection in covariance matrix estimation. *The Review of Economic Studies* 61(4), 631–653.
- Sargent, T. J. (1999). *The Conquest of American Inflation*. Princeton, NJ: Princeton University Press.
- Slobodyan, S. and R. Wouters (2012a). Learning in a medium-scale dsge model with expectations based on small forecasting models. *American Economic Journal: Macroeconomics* 4(2), 65–101.

Slobodyan, S. and R. Wouters (2012b). Learning in an estimated medium-scale dsge model. *Journal of Economic Dynamics and Control* 36(1), 26 – 46.

Williams, N. (2019). Escape dynamics in learning models. *The Review of Economic Studies* 86(2), 882–912.