On the determination of adaptive learning gains

Michele Berardi — *University of Manchester*

Jaqueson K. Galimberti* — *ETH Zurich*

Abstract

We investigate the empirical plausibility of different interpretations for the learning gains of adaptive algorithms adopted in the learning literature to depict the evolution of agents' beliefs over time. We distinguish between two possible rationales to its determination: as a *choice* of rationally optimizing agents, or as a *primitive* parameter of bounded rational agents. Our results provide strong evidence in favor of the latter, thus suggesting that agents adaptive behavior may represent a misspecification in the statistical adjustment of their learning algorithms. Our evidence also points to some heterogeneity in the time evolution of this behavior with respect to the variable forecasted and the algorithm adopted.

Keywords: adaptive learning, algorithms, gains calibration, bounded rationality.

JEL codes: D83, E03, E37.

1 Introduction

One recurrent issue in the application of adaptive learning algorithms in order to mimic the process through which agents form their expectations refers to the calibration of these

*Corresponding author. E-mail: galimberti@kof.ethz.ch.

algorithms. The computational operation of these algorithms requires the pre-specification of a sequence of learning gain values, or of a mechanism through which these gains are determined in real-time. This paper is devoted to investigate this issue empirically, thus producing renewed estimates of the learning gains. We hope these results will serve as guidance for the calibration of the learning algorithms in applied macroeconomics.

We introduce an insightful distinction between two assumptions on the rationale given to the learning gain. The first view is that the learning gain is determined as a *choice* by the agents adopting a given learning algorithm to update their expectations. The second view, in contrast, is to assume that the learning gain stands as a *primitive* parameter of agents learning-to-forecast behavior. In spite of their evident dissimilarity, an understanding of the effects of these different assumptions over the resulting gain calibrations seems to have been neglected in the previous literature.

In order to shed some light on this point we propose an evaluation framework that mimics the real-time process of expectation formation through what we define as learning-to-forecast exercises. We focus on the use of the two algorithms that have received most of the attention in the literature, namely, the Least Squares (LS) and the Stochastic Gradient (SG), with a particular interest in applications to time-varying estimation environments. Using such framework we evaluate empirically the performance of the distinct approaches to the calibration of the learning gains. We analyze the quality of the forecasts produced by each calibration along two dimensions: their *forecasting accuracy*, and their *resemblance to surveys.*

We carry out these exercises using real-time quarterly data on US inflation and output growth covering a broad post-WWII period of time, from 1947q2 to 2011q4. Our results provide strong evidence in favor of the gain as a *primitive* approach, hence favoring the use of surveys data for their calibration. Furthermore, the performances of particular calibrations of the learning algorithms, applied to forecast different variables, are found to be rather sensitive to the samples used for selection of the gains.

The remainder of this paper proceeds as follows. In section 2 we begin with a description of the methodological framework we adopt, encompassing the specifications of the standard learning algorithms and their implementation details, the data we use, and the comparative exercises we propose. In section 3 we discuss the issue of the determination of the learning gains, relating our work to previous debates in the literature. We also provide a precise specification of the different learning gain calibrations we evaluate, leaving the presentation and discussion of their comparative results to section 4. Finally, we conclude this paper with some remarks in section 5.

2 Methodological framework

Consider an estimation context faced by a real-time agent wishing to obtain inferences about the law of motion of a variable of interest, say y_t. From an economic perspective, these inferences can be thought of as the middle step agents undertake in a process of learning-to-forecast in order to form their expectations.

To narrow down our focus, we assume this agent attempts to construct such inferences assuming that y_t is statistically related to other observed variables, say a vector of (predetermined) variables $\mathbf{x}_t = (x_{1,t}, \ldots, x_{K,t})'$, through a linear regression of the form[1]

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

where $\boldsymbol{\theta}_t = (\theta_{1,t}, \ldots, \theta_{K,t})'$ stands for a vector of (possibly time-varying) coefficients, and ε_t denotes a white noise disturbance with variance given by σ_t^2. Both coefficients and disturbances are assumed not to be directly observable by the agent. Under this context, a technique for estimation of $\boldsymbol{\theta}_t$ is required to allow the agent to construct inferences for y_t on the basis of (1).

[1] This specification can be straightforwardly extended to a multivariate regressions context, an autoregressive context, or yet in both dimensions to a vector autoregression (VAR) specification.

2.1 Learning algorithms

In the literature of learning and expectations in macroeconomics (see Evans and Honkapohja, 2001) recursive algorithms have been proposed for this task. Two of the main forms adopted are the LS and the SG specifications.

Algorithm 1 (LS)**.** *Under the estimation context of (1), the* LS *algorithm assumes the form of*[2]

$$\begin{aligned} \hat{\boldsymbol{\theta}}_t^{LS} &= \hat{\boldsymbol{\theta}}_{t-1}^{LS} + \gamma_t \mathbf{R}_t^{-1} \mathbf{x}_t \left(y_t - \mathbf{x}_t' \hat{\boldsymbol{\theta}}_{t-1}^{LS} \right), \qquad (2) \\ \mathbf{R}_t &= \mathbf{R}_{t-1} + \gamma_t \left(\mathbf{x}_t \mathbf{x}_t' - \mathbf{R}_{t-1} \right), \qquad (3) \end{aligned}$$

where γ_t is a learning gain parameter, and $\mathbf{R}_t$ stands for an estimate of regressors matrix of second moments, $E\left[\mathbf{x}_t \mathbf{x}_t'\right]$.

Algorithm 2 (SG)**.** *Under the estimation context of (1), the* SG *algorithm is given by*[3]

$$\hat{\boldsymbol{\theta}}_t^{SG} = \hat{\boldsymbol{\theta}}_{t-1}^{SG} + \mu_t \mathbf{x}_t \left(y_t - \mathbf{x}_t' \hat{\boldsymbol{\theta}}_{t-1}^{SG} \right), \qquad (4)$$

with μ_t standing for the learning gain parameter.

Notice the hats in $\hat{\boldsymbol{\theta}}_t^{LS}$ and $\hat{\boldsymbol{\theta}}_t^{SG}$ indicate that they stand for estimates of $\boldsymbol{\theta}_t$ in (1), and these estimates are based on period t information.

Since the seminal works of Bray (1982); Marcet and Sargent (1989) the LS algorithm has been taken as the natural choice to represent agents mechanism of adaptive learning. This choice is in general attributed to the widespread knowledge of its non-recursive counterpart,

[2] This form is closer to that used in the adaptive learning literature under the name of Recursive Least Squares, for the case where the gain is decreasing with time, or Constant-Gain (Recursive) Least Squares, for the case of a time-invariant gain. In the engineering literature other variations in the nomenclature can be found and a computationally less demanding form is more common where the inversion of $\mathbf{R}_t$ in (3) is avoided by the use of the matrix inversion lemma (see Haykin, 2001).

[3] This form is common to both the adaptive learning (see also Evans et al., 2010) and the engineering literature. In the latter this algorithm is generally known as least mean squares, although commonly referring to the constant gain case, while stochastic gradient is often referred to the case of a time-decreasing gain. In some contexts, however, stochastic gradient is referred as a whole family of filters (see Macchi, 1995, p. 52).

the so-called Ordinary Least Squares (OLS) estimator, between econometricians. The SG algorithm, on the other hand, provides a computationally simpler alternative, a feature clearly apparent in (4) for the absence of the LS "normalization" step given by the inverse of the matrix of second moments. For this reason some authors have advocated for its use as a more plausible learning device from a bounded rationality standpoint (Barucci and Landi, 1997; Evans and Honkapohja, 1998).

2.2 Data and implementation details

Our interest is on the comparative evaluation of the forecasts obtained from model estimates provided by the different combinations of algorithms and gain calibrations applied to a sample of data. With respect to the model, our focus is on unrestricted VAR specifications applied to inflation and growth. We denote the time series for these variables by π_t and g_t, respectively. For robustness we estimate VARs with lag orders from 1 to 4.

We use quarterly data on the US real GNP/GDP and its price index from 1947q2 to 2011q4, which sums up to 259 observations for each variable. Our data on these series come from the Philadelphia's Fed Real-Time Data Research Center[4] and consists of vintages from 1966q1 to 2012q1, i.e., a total of 185 snapshots of what was known on these variables by a market participant in real-time (see Stark and Croushore, 2002). As our interest is on forecasts for output growth and inflation, we obtain these rates from the above data on levels computing their associated annual growth rates by compounding their simple quarterly growth factors.

For the purpose of comparing the algorithms' forecasts to those provided by survey respondents, we use data from the Survey of Professional Forecasters (SPF)[5], which are made available by the Philadelphia's Fed as well. Each quarter, this survey asks professional economists to give their forecasts for several macroeconomic variables, including those we

[4]See http://www.philadelphiafed.org/research-and-data/real-time-center/. We have done some specific adjustments to the original dataset, as detailed in Appendix A.1.

[5]See http://www.philadelphiafed.org/research-and-data/real-time-center/survey-of-professional-forecasters/.

Figure 1: Data series on US inflation, output growth, and survey forecasts.

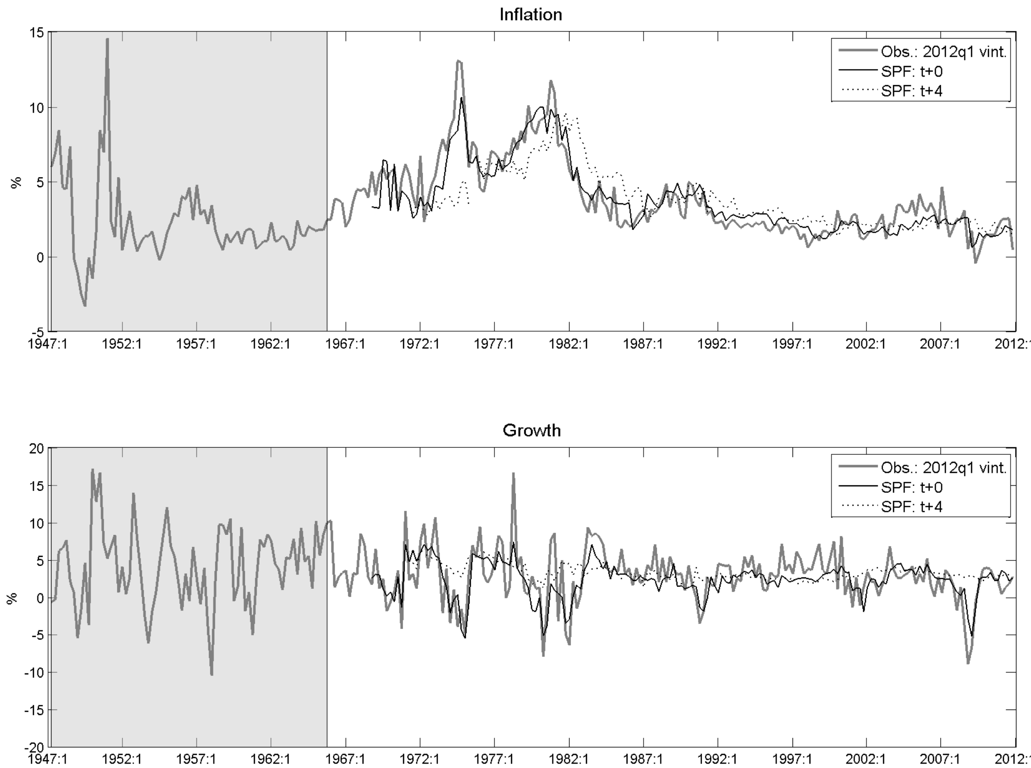

The series of actual observations refer to those we have available at the latest vintage in the real-time dataset we are using, i.e. 2012q1. The SPF's forecasts refer to those made for the pointed period from the information set of the corresponding forecasting horizon. The shaded area indicates the observations left aside for the initialization procedure of the algorithms.

indicated above, and also over different forecasting horizons. Here we use the median of the individual forecasts for output growth and inflation made for a total of five horizons, namely from t (nowcast) to $t+4$. The SPF data is available from 1968q4 onwards, and, consistent to our data on actuals, the last survey data we use is that of 2010q4, which contains forecasts up to 2011q4. An overview of these data series is presented in Figure 1.

The implementation of the LS and the SG learning algorithms involves two main stages. The first step in the process of obtaining forecasts from these algorithms is to set their initializations, i.e., estimates for $\hat{\boldsymbol{\theta}}_0$. Here we follow the smoothing approach of Berardi and Galimberti (2012a): first, set aside an initial portion of the available sample of data to serve as a training sample, for which we use the first 75 observations of our sample (1947q2-1965q4, shaded area in Figure 1); and second, repeat a smoothing routine within the training sample until the smoothed initial estimates converge to values consistent with the algorithm's steady state operation. The simulation evidence reported in Berardi and Galimberti (2012a) indicates that this method provides the most consistent estimates of initials within our estimation framework.

With the initials given, the second implementation stage is the estimation of the VAR model specifications with the series of data on inflation and growth. Estimation and forecasting are carried out by vintage as follows:

1. The recursions for each algorithm/gain calibration are computed departing from the vintage/algorithm/gain initials until exhaustion of the vintage sample.

2. The $t, \ldots, t+4$ forecasts for each vintage/algorithm/gain are computed using the last estimates of the model specification, where t stands for the vintage quarter.

We repeat these computations for each vintage of data from 1966q1 to 2010q4, which results in a total of 180 forecasts for each algorithm/gain/horizon both for inflation and growth.

2.3 Learning-to-forecast exercises

We propose two comparative exercises in order to assess the different approaches to the determination of the learning gains. In the first we take the point of view of an economic agent, who has to build forecasts of variables relevant for economic decisions. What matters for this agent is the accuracy of such forecasts, and so our first evaluation criterion is given by the *forecasting accuracy* of the algorithms/gain calibrations[6]. In the second exercise we assume the point of view of the researcher whose interest, in contrast, is in uncovering which mechanism better represents the learning-to-forecast behavior of the economic agents being

[6] We use the first-available observations in our real-time dataset to compute the average forecasting errors.

modeled. Taking survey forecasts as proxy measures for that behavior, our second evaluation criterion is given by the *resemblance to surveys* of the algorithms/gain calibrations forecasts.

To further substantiate our comparative analysis, we also make use of tests common to the literature on forecast evaluation. Namely, we adopt both the Diebold and Mariano (1995) (DM) test for equal (unconditional) predictive ability, and its more recently developed conditional counterpart test of Giacomini and White (2006) (GW)[7]. Other than for robustness purposes, our choice for these two tests can also be well justified: while the first stands as a classical test, whose properties have been long studied in the literature (see Diebold, 2013), the second clearly represents a more appropriate test for our purposes of comparison of different estimation methods[8].

3 Determination of the learning gains

Both the LS and the SG algorithms require the specification of a sequence of learning gains. The learning gain stands for a parameter determining how quickly a given information is incorporated into the algorithm's coefficients estimates. Three of the main alternatives for the specification of this learning gain are those of a time-decreasing, a time-constant, and a time-varying (not restricted to be decreasing) sequence of values, and their suitability depends on the time-varying nature of the environment.

A decreasing-gain LS was the seminal choice in the learning literature, so as to match the recursive form of the OLS estimator. For the case of linear models with time-invariant parameters, this estimator is known to possess some well desired properties, such as consistency and efficiency, though these properties do not extend to a time-varying context. This latter fact implies the intriguing observation that a decreasing-gain LS learning mechanism is appropriate only along the time-invariant path of a RE equilibrium, where learning itself

[7]We summarize the calculations involved in each of these tests in Appendix A.2.

[8]Although the GW test is not suitable for the comparison of recursively estimated model-based forecasts, our focus on constant gain specifications attach geometrically decaying weights to past observations, hence approximating a rolling window estimation scheme.

is indeed pointless (see Bray and Savin, 1986, and the discussion below).

Extensive evidence (see Stock and Watson, 2003; Cogley and Sargent, 2005; Sims and Zha, 2006; Sargent et al., 2006) favoring time-varying parameter models of the economy has, nevertheless, challenged this paradigm, and the departure from the parameter constancy assumption (see Margaritis, 1990; Bullard, 1992; McGough, 2003) has naturally led to the requirement of adjustments to the learning rules as well. The constant gain specification has been in the spotlight of most applied research since (Sargent, 1999), given its tracking capabilities and its suitability for time-varying environments. We will also assess the case of time-varying gains, determined according to an outer adaptive mechanism.

3.1 Theoretical background

Applied research requires going beyond the incipient debate between the use of a decreasing or a constant sequence of learning gains, and at this point the researcher is usually left with a decision between two calibration approaches. The first is to assume that the learning gain is determined as a *choice* by the agents adopting a given learning algorithm to update their expectations. The second, in contrast, is to assume that the learning gain stands as a *primitive* parameter of agents learning-to-forecast behavior.

Gain as a *choice*

One important debate related to the traditional assumption of decreasing gains was initiated by an intriguing observation made by Bray and Savin (1986). Adaptive learning attempts to capture agents expectations off the equilibrium path, which implies that until convergence is reached the agents' perceived law of motion (PLM) is changing. Thus, even if the structural model is assumed to be time-invariant, learning itself introduces time-variation in the economy's actual law of motion (ALM). Statistical theory, however, indicates that decreasing gains estimators are not appropriate for estimation of models with time-varying parameters. Therefore, it would be reasonable to expect that even if agents commence learning with decreasing gains, they would eventually have collected enough evidence to prompt a change towards tracking-consistent gain calibrations.

The natural follow up to this compelling argument by Bray and Savin (1986) was the incorporation of more sophisticated mechanisms into the analysis of learning convergence. Assuming agents learn using Kalman filtering techniques Margaritis (1990) was able to show convergence to RE in a model with time-varying parameters, though still requiring that the gains tend to zero. Relaxing this latter assumption, however, Bullard (1992) found that convergence results are lost if agents' PLMs are not restricted to become (asymptotically) time-invariant; hence there would be no reason, on rationality grounds, to expect vanishing learning gains. Finally, a reconcilement between these results was proposed by McGough (2003, p.120): "for convergence to a RE equilibrium to occur, the agents must believe that the conditional variance of the time-varying parameters decreases to zero." If the only source of time-variation comes from learning, convergence to RE will indeed lead the estimates of this conditional variance to zero[9].

This seemingly circular reasoning actually reflects much deeper implications on the internal consistency between adaptive learning and convergence to RE: either that agents initially believe in such a convergence, thus not abandoning the decreasing gain during out-of-equilibrium learning for its poor performance; or that agents are aware of this trade-off between tracking and convergence, and attempt to guarantee the latter at the same time that tuning their gain calibration to improve on the former during transient dynamics. Clearly, this last upshot from this debate provides a rationale for the view that the learning gains are determined as a *choice* by the agents.

Finally, we can also relate the approach of gain calibration as a *choice* with alternative approaches to bounded rationality. One popular alternative has been the adaptively rational expectations approach of Brock and Hommes (1997), where agents are assumed to select between a set of predictors to form their expectations. The selection is determined according

[9]As McGough (2003) appropriately recognizes, the conditional variance of the system's time-varying coefficients has to decrease at a sufficiently high rate for this argument to hold.

to a discrete-choice model, which is microfounded in a random utility framework, depending on fitness measures associated to each predictor's past performance. One appealing feature of this approach is that a diversity of predictors may be jointly in use at a given period depending on how the intensity of choice is regulated.

Another related approach is the "expectation calculation" approach of Evans and Ramey (1998), where agents are assumed to face a calculation decision pondering between the benefits and the costs of extra cognitive efforts to form expectations. Then, in a similar vein to the optimizing behavior we assume to impose discipline on agents *choice* of gains, their agents are assumed to optimize on their choice of calculation intensity. More in line with our above discussion about internal consistency, Marcet and Nicolini (2003) propose a threshold rule under which agents learn with decreasing gains during stable periods, but switch to a constant gain when some instability is detected through large prediction errors.

Gain as a *primitive*

From a theoretical point of view, the learning gain is determinant both of whether convergence to a RE equilibrium takes place and of its transitional dynamic properties. In order to obtain positive convergence results, the traditional analysis of learning usually places strong restrictions in the sequence of gains. Examples of these restrictions include the requirement of decreasing gains or that of a "small" constant gain to guarantee weaker (in distribution) convergence results (see Evans and Honkapohja, 2001, Ch.7). Furthermore, under the constant gain specification a phenomenon known as "escape dynamics," recurrently pushing the economy away from its equilibrium, has been found to have its frequency of occurrence associated to the value of the learning gain (Cho et al., 2002).

An understanding of the interplay between these features is provided by Sargent and Williams (2005), who establish a Bayesian interpretation to the LS algorithm, later extended to the SG case by Evans et al. (2010). Such derivations are obtained along similar lines to those explored in establishing an unified framework by Berardi and Galimberti (2013):

assuming data is generated by a model with time-varying parameters governed by a random walk, it is well known that the Kalman filter provides the Bayes optimal estimator; drawing correspondences between the learning algorithms and the Kalman filter is, thus, equivalent to finding the conditions under which the learning algorithms approximate the Bayesian estimator.

Under the Bayesian interpretation, however, these correspondences define particular priors imputed as agents beliefs about the statistical properties of the data. Besides, other than relating the learning algorithms to the Kalman filter, the priors also carry implications about the specific calibrations of the learning gains within each algorithm. Given that in the Bayesian tradition these priors are specified in advance to estimation, this interpretation provides an appealing rationale for the view that the learning gains stand as *primitive* parameters of agents learning-to-forecast behavior.

Another useful interpretation given to the learning gains is that of determining the memory of the learning mechanism (see, e.g., Barucci, 1999; Honkapohja and Mitra, 2006). The gains determine the weights the learning algorithms assign to each observation in updating their estimates (see Berardi and Galimberti, 2013, p. 140, for the precise correspondence). Namely, the higher the gain, the higher the emphasis given to more recent observations; hence, higher gains imply a lesser memory of the past observations of the data in the current estimates provided by the learning algorithms. Recognizing that it would be forcible to think of an agent's memory as a choice, this interpretation provides an alternative rationale to the calibration of gains as a *primitive*.

3.2 Gain calibrations

To approach empirically the gain determination issue, we associate to the alternative assumptions distinct measures of gain selection. Under the first alternative, additional discipline is required to model the decision problem by assuming that agents would be willing to optimize on their *choice* of a gain. Given that the ultimate purpose of the learning algorithms is to represent agents process of learning-to-forecast, a natural measure to select the gain calibrations under this assumption is the accuracy of their associated forecasts.

As for the second alternative, taking the learning gain as a *primitive* obviously does not require any additional behavioral assumption. Nevertheless, an appropriate measure of fit under this alternative requires the actual observation of agents behavior of learning-to-forecast. Consistent to our previous analysis, we once again use data of survey forecasts as a proxy. Therefore, in order to obtain a calibration of the learning gains as agents *primitive* parameters we can adopt a measure of resemblance of their associated forecasts to those forecasts observed from actual agents collected through surveys.

The main challenge to the calibration of the learning gains, however, is about how these measures are evaluated in order for a particular gain to be selected for the algorithm's recursive computations. This issue clearly represents an important obstacle for the view of the gains as an agent's *choice*; namely, one has to be careful about the conciliation of this choice with the real-time informational assumptions underlying the learning process.

Another related issue is about the specification of the range, or the options, of gain values of interest. We follow two computational approaches to obtain the gain values. In the first we begin by constructing a grid of admissible values, and then proceed imposing selection rules that will represent our different assumptions on the role of the gains. The second approach involves the use of an outer mechanism to adaptively adjust the gain in response to changes in the recent performance of each algorithm.

Grid-based

We construct a grid of gain values by setting upper bounds on their admissible values so as to ensure the algorithm's stability. We use a grid of 100 values for this purpose, meaning that our estimation routine is applied to each algorithm with 100 different constant gain values[10].

[10] We have also computed the algorithms with a decreasing gain on the form of $\gamma_t = \bar{\gamma}/t$ and $\mu_t = \bar{\mu}/t$ to benchmark our results.

When it comes to represent agents learning-to-forecast behavior, though, an unique gain value is required for each time an iteration on the recursive algorithms is performed. Under the gain as a *choice* assumption we are interested in finding the gain value that would be representative of agents' optimization of each algorithm's forecasting accuracy. In other words, we can uncover what would have been an agent's *choice* of a gain by minimizing the average (accuracy) loss associated to the gains available in our grid. Similarly, for the gain as a *primitive* assumption we can also pick the gain following such an average loss minimization, though here the forecasts losses need to be defined relative to their resemblance to the surveys.

It still remains to specify under which samples these losses are computed. We explore three alternatives on this aspect. The first, denoted as *full-sample*, is to pick the gain yielding the minimum average loss over the full sample of forecasts that we have computed. Clearly, under the gain as a *choice*, this selection sample violates the restrictions of a "fair" out-of-sample forecasting exercise that we are exploring in connection to the idea of real-time learning. Nevertheless, this alternative has been adopted in some of the previous calibration attempts in the literature (see, e.g., Milani, 2007, 2008, 2011). It does not present any conflict under the gain as *primitive*, where it would just imply that agents hold immutable beliefs about the system they are forecasting. This alternative has also found some applications in the previous literature (see, e.g., Orphanides and Williams, 2005; Pfajfar and Santoro, 2010)[11].

The second selection sample we adopt, denoted as *in-sample*, involves splitting the whole sample of forecasts in two parts: the first part is used as an in-sample period on which the minimization of the average loss is applied in order to pick a gain for each algorithm; the second part is then used for the evaluation, keeping fixed the gain calibration. This alternative goes in line with traditional exercises of forecast evaluation and has also been used by Branch and Evans (2006); Weber (2010) for purposes similar to ours.

The last selection sample we explore is to allow the choice of the gain to be *recursive*,

[11] We note that in all these papers the calibration has been made within a structural model and solely focusing on the LS.

through a minimization of the average loss over a rolling window sample of forecasts. From the real-time learning and forecasting perspective, the sequence of gains selected using this recursive approach is the extreme opposite to that proposed in the first alternative; particularly, it does not allow the use of information on the future quality of the forecasts for the calibration of the algorithms in real-time. Moreover, there is no reason to restrict the gains to be fixed throughout the whole forecasting sample.

To make our comparative between these different selection samples conformable we set the length of the rolling window equal to the amount of data set aside for the *in-sample* calibration. Namely, we will adopt an in-sample/window length of 60 observations (1966q1-1980q4)[12]. These selection samples determine what is left as our evaluation sample: from 1981q1 to 2010q4. Moreover, given our interest in multi-horizons forecasts, we also need to specify which horizon is used in the loss evaluations. Here we allow the gain selections to differ by forecasting horizon, hence evaluating the above selection criteria for each forecasting horizon.

Adaptive gain

An alternative approach to the selection of the learning gains in real-time is to turn the gain calibration itself into an adaptive estimation problem. The idea is to plug each learning algorithm with an outer mechanism to adjust the gain in response to changes in estimates of the algorithm's recent performance. Such an automatic tuning of the recursive algorithms was first suggested in Benveniste et al. (1990), and later analyzed by Kushner and Yang (1995) who presented evidence favoring this approach. In economics, a recent application has been presented in Kostyshyna (2012), evidencing some quantitative improvements to model hyperinflation episodes.

The derivation of the adaptive gain learning algorithms departs from the specification of an objective function, which in our context is related to the loss functions we associate with

[12] Recall the learning-to-forecasts exercises take the first 75 observations (1947q2-1965q4) of our sample for the initialization of the algorithms.

the two alternative roles given to the learning gains. Hence, the adaptive gain approach is also split in two, with the gain as a *choice* represented by a mechanism responding to the algorithm's accuracy, whereas under the gain as a *primitive* the adaptation is driven by the algorithm's resemblance. To be concise, we adopt the following notation in the derivation that follows: z_t stands for the forecast's target (either the actuals, y_t, or the survey forecasts, s_t); $\mathbf{x}_t$ continues to represent a vector of regressors, i.e., an unity constant and lags of inflation and growth under our VAR specification; and $\hat{\boldsymbol{\theta}}_t$ is a vector of estimates of our model's parameters[13]. Moreover, we also assume a squared loss function.

Let the objective be to select the gain so as to minimize the expected loss of forecast (comparison) errors

$$J_t = \frac{1}{2} E\left[\left(z_t - \mathbf{x}_t' \hat{\boldsymbol{\theta}}_{t-1}\right)^2\right]. \tag{5}$$

If the true properties of the data generating process (DGP) in (1) were known beforehand, one could ideally pick each algorithm's gain so as to optimize the above criterion. In the lack of this information, as it is the case in the learning-to-forecast situation, one alternative is to superimpose an outer adaptive scheme for the purpose of automatic tuning the gain parameter.

The general idea is to use a recursive updating scheme that corrects the gain parameter in the direction opposite to a stochastic approximation of the loss function gradient. Deriving[14] these gradients and plugging their stochastic approximation into each algorithm's recursion we obtain the LS and SG algorithms with adaptive gains (LSA and SGA, respectively).

Algorithm 3 (LSA)**.** *Under the estimation context of (1), the* LSA *algorithm assumes the*

[13]Without loss of generality, our presentation focus in a single equation.

[14]See appendix A.3.

form of

$$\begin{aligned}
\hat{\boldsymbol{\theta}}_t^{LSA} &= \hat{\boldsymbol{\theta}}_{t-1}^{LSA} + \gamma_t \mathbf{R}_t^{-1} \mathbf{x}_t \left(y_t - \mathbf{x}_t' \hat{\boldsymbol{\theta}}_{t-1}^{LSA} \right), && (6) \\
\mathbf{R}_t &= \mathbf{R}_{t-1} + \gamma_t \left(\mathbf{x}_t \mathbf{x}_t' - \mathbf{R}_{t-1} \right), && (7) \\
\gamma_t &= \left[\gamma_{t-1} + \alpha_\gamma \mathbf{x}_t' \hat{\boldsymbol{\Psi}}_{t-1}^{LSA} \left(z_t - \mathbf{x}_t' \hat{\boldsymbol{\theta}}_{t-1}^{LSA} \right) \right]_{\overline{\gamma}_{\min}}^{\overline{\gamma}_{\max}}, && (8) \\
\hat{\boldsymbol{\Psi}}_t^{LSA} &= \left(\mathbf{I} - \gamma_t \mathbf{R}_t^{-1} \mathbf{x}_t \mathbf{x}_t' \right) \hat{\boldsymbol{\Psi}}_{t-1}^{LSA} \ldots \\
&\quad + \left(\mathbf{I} - \gamma_t \mathbf{R}_t^{-1} \hat{\mathbf{S}}_t \right) \mathbf{R}_t^{-1} \mathbf{x}_t \left(y_t - \mathbf{x}_t' \hat{\boldsymbol{\theta}}_{t-1}^{LSA} \right), && (9) \\
\hat{\mathbf{S}}_t &= \left(1 - \gamma_t \right) \hat{\mathbf{S}}_{t-1} + \mathbf{x}_t \mathbf{x}_t' - \mathbf{R}_{t-1}, && (10)
\end{aligned}$$

where α_γ *is an adaptation constant,* $\hat{\boldsymbol{\Psi}}_t^{LSA}$ *stands for an estimate of* $\partial \hat{\boldsymbol{\theta}}_t^{LSA} / \partial \gamma$*,* $\hat{\mathbf{S}}_t$ *stands for an estimate of* $\partial \mathbf{R}_t / \partial \gamma$*, and* $[\bullet]_{\overline{\gamma}_{\min}}^{\overline{\gamma}_{\max}}$ *is a truncation operator setting* γ_t *to* $\overline{\gamma}_{\min}$ *if it falls below this value, or to* $\overline{\gamma}_{\max}$ *if it rises above this value. The remaining components follow from the definition of the LS algorithm.*

Algorithm 4 (SGA)**.** *Under the estimation context of (1), the* SGA *algorithm assumes the form of*

$$\begin{aligned}
\hat{\boldsymbol{\theta}}_t^{SGA} &= \hat{\boldsymbol{\theta}}_{t-1}^{SGA} + \mu_t \mathbf{x}_t \left(y_t - \mathbf{x}_t' \hat{\boldsymbol{\theta}}_{t-1}^{SGA} \right), && (11) \\
\mu_t &= \left[\mu_{t-1} + \alpha_\mu \mathbf{x}_t' \hat{\boldsymbol{\Psi}}_{t-1}^{SGA} \left(z_t - \mathbf{x}_t' \hat{\boldsymbol{\theta}}_{t-1}^{SGA} \right) \right]_{\overline{\mu}_{\min}}^{\overline{\mu}_{\max}}, && (12) \\
\hat{\boldsymbol{\Psi}}_t^{SGA} &= \left(\mathbf{I} - \mu_t \mathbf{x}_t \mathbf{x}_t' \right) \hat{\boldsymbol{\Psi}}_{t-1}^{SGA} + \mathbf{x}_t \left(y_t - \mathbf{x}_t' \hat{\boldsymbol{\theta}}_{t-1}^{SGA} \right), && (13)
\end{aligned}$$

where α_μ *is an adaptation constant,* $\hat{\boldsymbol{\Psi}}_t^{SGA}$ *stands for an estimate of* $\partial \hat{\boldsymbol{\theta}}_t^{SGA} / \partial \mu$*, and* $[\bullet]_{\overline{\mu}_{\min}}^{\overline{\mu}_{\max}}$ *is a truncation operator setting* μ_t *to* $\overline{\mu}_{\min}$ *if it falls below this value, or to* $\overline{\mu}_{\max}$ *if it rises above this value. The remaining components follow from the definition of the SG algorithm.*

An intuition for this gain adaptation mechanism follows directly from its interpretation as a numerical optimization method, browsing along the error-performance curve, (5), in search for an optimal gain. Here $\hat{\boldsymbol{\Psi}}_t$ keeps track of the algorithm's past estimation performance, accumulating its past gradients discounted according to each algorithm's forgetting factors. This synthetic measure of past performance is taken as a reference in the gain update equations, (8) and (12): when the latest gradient points towards the same (a different) direction as of $\hat{\boldsymbol{\Psi}}_t$, the adaptive mechanism interprets this as evidence of systematic (contradictory) mistakes; hence the gain is increased (decreased) to intensify (lessen) the algorithm's response to last period error[15].

Nevertheless, the above interpretation becomes less compelling under the gain as a *primitive* approach, where $z_t \equiv s_t$. In this case the gain adaptation mechanism has to bear with two distinct estimation objectives: the algorithm's accuracy performance, and its resemblance to the surveys. Following the above interpretation, that means the adaptation mechanism has to browse along two error-performance curves at the same time. Whereas the latest gradient estimates is drawn from the resemblance to surveys objective, the past gradients synthesized in $\hat{\boldsymbol{\Psi}}_t$ refer to the algorithm's accuracy performance. Hence, the gain is increased to intensify the algorithm's response to last period error if its recent performance indicates systematic mistakes on both dimensions, and vice versa.

Computation of the LSA and SGA algorithms still requires specification of their adaptation constants, α_γ and α_μ. According to the analysis of Kushner and Yang (1995), stability requires this parameter to be small, so as to satisfy $0 < \alpha_\mu \leq \overline{\gamma}_{\min}$ and $0 < \alpha_\mu \leq \overline{\mu}_{\min}$, together with an appropriate upper bound to the gains. For our purposes we take the extreme gain values in the grids defined in the previous section as the bounds for the adaptive gains. Most importantly, Kushner and Yang (1995) also present simulation evidence indicating that the algorithm's performance is not as sensitive to the calibration of α as it is to the learning gain. Therefore, our calibration of these adaptation constants will precede any attempt of performance optimization in favor of a calibration presenting a sequence of gains not too jumpy, but neither constant.

We end this section with a word of caution in the use of the adaptive gains approach

[15] A similar intuition is given in Kostyshyna (2012), though $\hat{\boldsymbol{\Psi}}_t$ is interpreted as the discounted past errors due to a simpler model specification where agents are learning only the value of a constant.

within a learning context. In the derivation of the gradient of (5) we assumed that the derivative of z_t with respect to the gain is null, which may not be a realistic assumption when $z_t \equiv y_t$ (gain as a *choice*): self-referentiality implies that the determination of the actuals are affected by agents expectations; if agents are learning, the gain would have an indirect effect over y_t by determining agents expectations, hence violating[16] the assumption above. However, relaxing this assumption requires the specification of a structural model, which goes beyond our scope.

4 Results and discussion

In this section we present the comparative results on the different approaches to the calibration of the learning gains as outlined above. We first present the numerical results of the calibrated gains. Our evaluation begins with an overview on the statistical properties of the forecasts associated to the combinations of learning algorithms and their calibrations. We then compare the different methods of gain selection, namely, the *full-sample*, the *in-sample*, and the *recursive* grid-based methods, and the *adaptive* gains approach. Finally, we shift our focus to our main issue on the determination of the learning gains, namely, we compare the gain as a *choice* and the gain as a *primitive* assumptions.

4.1 Gain calibrations

In our approach a sequence of gain values is associated to each application of the learning algorithms to forecast inflation and growth. These are determined according to the combinations of gain assumptions and selection methods. In Table 1 we present the resulting numerical calibrations we obtained for the fixed gain calibrations, namely, the *full-sample* and the *in-sample* methods.

We make three main observations regarding these fixed calibrations. First, it is clear that

[16] Agents would also be required to be aware of their expectation's effects for this relationship to be manifest in their loss function.

Table 1: Fixed gain calibrations.

Variable	Gains as a *choice*						Gain as a *primitive*					
Algorithm	Full-sample			In-sample			Full-sample			In-sample		
- Model	h=0	h=2	h=4	h=0	h=2	h=4	h=0	h=2	h=4	h=0	h=2	h=4
(a) Inflation:												
Least Squares												
- VAR(1)	0.0644	0.1287	0.0941	0.0545	0.1040	0.0891	0.0297	0.0347	0.0347	0.0347	0.0297	0.0347
- VAR(2)	0.0396	0.0347	0.0297	0.0446	0.0446	0.0297	0.0248	0.0248	0.0248	0.0297	0.0248	0.0248
- VAR(3)	0.0396	0.0545	0.0347	0.0396	0.0594	0.0446	0.0297	0.0248	0.0248	0.0297	0.0297	0.0248
- VAR(4)	0.0396	0.0594	0.0594	0.0446	0.0644	0.0693	0.0297	0.0297	0.0347	0.0347	0.0347	0.0347
Stochastic Gradient												
- VAR(1)	0.0023	0.0023	0.0023	0.0023	0.0023	0.0023	0.0023	0.0023	0.0023	0.0022	0.0023	0.0023
- VAR(2)	0.0012	0.0011	0.0009	0.0012	0.0011	0.0011	0.0009	0.0007	0.0007	0.0010	0.0009	0.0009
- VAR(3)	0.0008	0.0008	0.0007	0.0008	0.0008	0.0008	0.0008	0.0005	0.0004	0.0008	0.0006	0.0005
- VAR(4)	0.0006	0.0006	0.0006	0.0006	0.0006	0.0006	0.0006	0.0004	0.0003	0.0006	0.0005	0.0003
(b) Growth:												
Least Squares												
- VAR(1)	0.0495	0.0050	0.0050	0.0644	0.0050	0.0050	0.0248	0.0149	0.0099	0.0495	0.0149	0.0099
- VAR(2)	0.0149	0.0050	0.0050	0.0495	0.0099	0.0050	0.0050	0.0050	0.0050	0.0149	0.0050	0.0099
- VAR(3)	0.0149	0.0050	0.0050	0.0495	0.0099	0.0050	0.0149	0.0050	0.0050	0.0149	0.0050	0.0099
- VAR(4)	0.0198	0.0050	0.0050	0.0396	0.0099	0.0050	0.0198	0.0050	0.0149	0.0248	0.0099	0.0149
Stochastic Gradient												
- VAR(1)	0.0003	0.0004	0.0006	0.0002	0.0003	0.0005	0.0003	0.0006	0.0023	0.0002	0.0005	0.0008
- VAR(2)	0.0002	0.0002	0.0002	0.0002	0.0002	0.0002	0.0003	0.0002	0.0004	0.0002	0.0002	0.0004
- VAR(3)	0.0002	0.0001	0.0002	0.0002	0.0001	0.0002	0.0002	0.0002	0.0003	0.0002	0.0002	0.0003
- VAR(4)	0.0002	0.0001	0.0002	0.0001	0.0001	0.0002	0.0002	0.0002	0.0003	0.0002	0.0002	0.0003

The Full-sample calibrations reflect the gain values minimizing the average squared distance between the forecasts and (i) the actual observations for the gain as *choice*, or (ii) the survey forecasts for the gain as a *primitive*, over the entire sample of data available after initialization, i.e., from 1966q1 to 2010q4, whilst for the In-sample the gains are selected over an initial sample of 60 observations, from 1966q1 to 1980q4.

the data scales substantially pushed the gain calibrations for the SG algorithm downwards relative to those for the LS, even though there is no direct proportionality between them[17]. Second, the *primitive* gain calibrations tended to be smaller (bigger) than those as a *choice* for inflation (growth). Lastly, the gains displayed a decreasing trend in relation to the VAR lag orders, but no clear relationship with the forecasting horizons.

In contrast to these fixed calibrations, the *recursive* and *adaptive* methods introduce time-variation to the sequence of gains. Their evolution through time is presented in Figures 2 and 3, respectively. Overall, these figures point to an agreement between the *choice* and *primitive* time-varying calibrations. However, little harmony is found comparing the gains picked by the *recursive* and *adaptive* calibrations; one difference is observed in their scales, with the *adaptive* gains presenting a wider range of variation. Distinct behaviors are also observed on these calibrations with respect to the variable forecasted, and the algorithm. It remains to see whether these differences have been relevant for their performances.

4.2 Comparative overview

We analyze the quality of the forecasts produced by each combination of algorithm/calibration along the two evaluation dimensions of our learning-to-forecast exercises: their *forecasting accuracy*, and their *resemblance to surveys*. Although we take these two criteria as mutually desirable, it is not clear whether they are compatible with each other; recall that whilst *forecasting accuracy* represents the goal of optimizing agents, *resemblance to surveys* is indicative of actual agents behavior. In spite of this conflict, there is some overlap between these criteria and their homonym measures used in this paper's definitions of gain determination approaches. As such *forecasting accuracy* would tend to favor calibrations of the gain as a *choice*, whereas *resemblance to surveys* would have a similar bias towards the gain as a *primitive*. Therefore, their joint usage can be justified as an attempt to prevent favoring any of these approaches.

[17] Their relative ratio (LS/SG) vary from as least as 12x (4x) up to 120x (335x) for inflation (growth).

Figure 2: Evolution of recursive gain calibrations through time.

(a) LS on inflation.

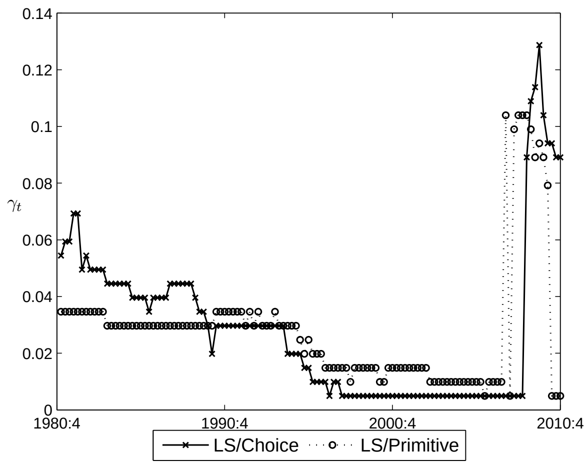

(b) LS on growth.

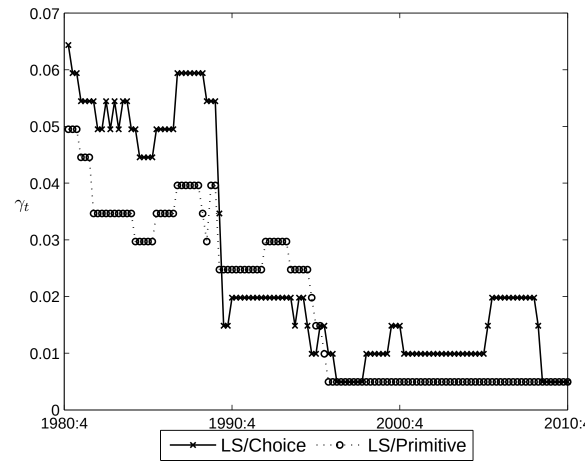

(c) SG on inflation.

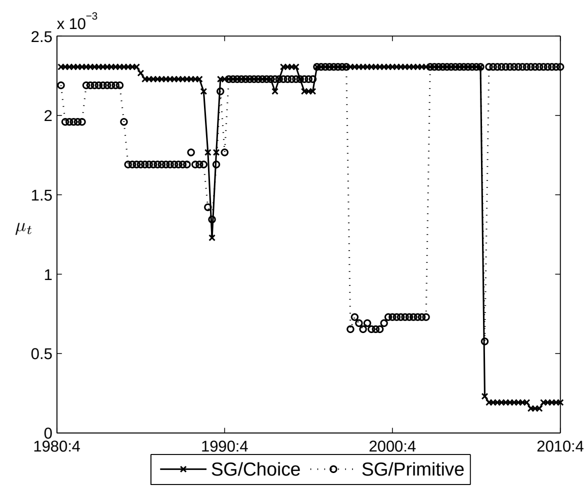

(d) SG on growth.

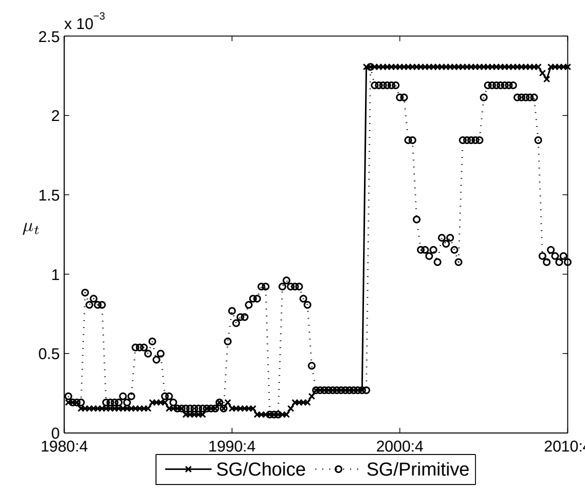

The recursive calibrations of the gains represent discrete values based on the grid of gains computed for each algorithm.

Figure 3: Evolution of adaptive gains through time.

(a) LS on inflation.

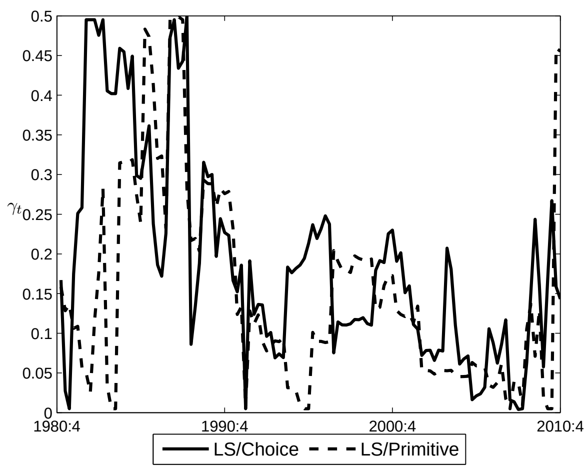

(b) LS on growth.

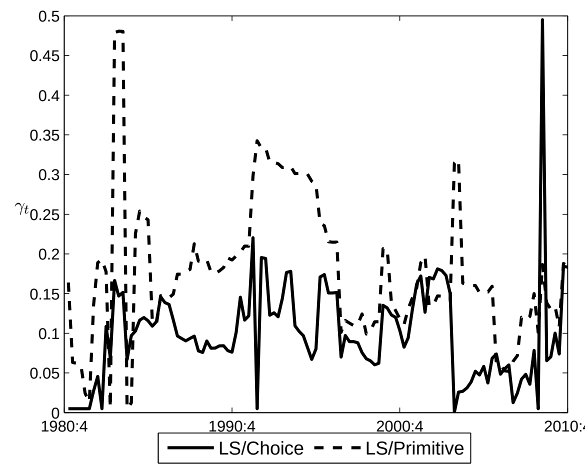

(c) SG on inflation.

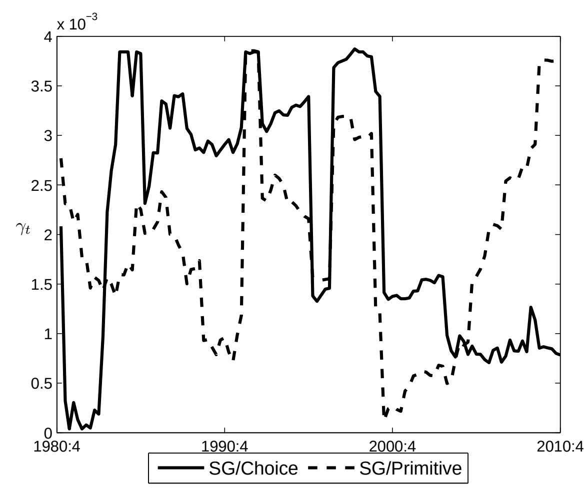

(d) SG on growth.

The series of adaptive gains were obtained using adaptation constants given by $\alpha_\gamma = 0.01$ and $\alpha_\gamma = 0.001$ for the LS on inflation and growth, respectively, and $\alpha_\mu = 1 \times 10^{-6}$ and $\alpha_\mu = 1 \times 10^{-7}$ for the SG on inflation and growth, respectively.

We start looking over the forecasts associated to each algorithm and gain value included in the grid computations. In Figure 4 we average the performance of each algorithm's forecasts for the different grid gain values, also presenting the algorithms' decreasing gain performances as a benchmark. The evidence is clearly favoring the constant gain SG algorithm, but not so remarkably for the LS case. This latter result appears to be at odds with that of Branch and Evans (2006), who found that the LS with constant gain tends to outperform its decreasing gain version.

An explanation for these results is due to our use of a different evaluation sample from that used by those authors, particularly for the inclusion of data covering the recent (2007-08) financial crisis. Intriguingly, the instabilities associated to this period seems to have more negatively affected the constant gain LS than its decreasing gain version[18]. Hence, the potential tracking benefits provided by the constant gain specifications seem to have been overcome by the higher level of noise affecting the economy during this period.

It is also evident in this figure how the performance of each algorithm is affected as the forecasting horizon varies. For the SG there is mainly a scale effect, whereas the LS performance is found to be rather sensitive to the gain values at longer horizons.

Finally, before getting into the comparative of the gain calibrations, we present in Tables 2 and 3 statistics for each individual series of forecasts. These statistics point towards the following observations: (i) the LS (SG) forecasts tend to be biased up(down)wards; (ii) the forecasts fail to replicate growth rates variability, whereas for inflation the SG forecasts presented variances closer to that of the actuals; (iii) between each algorithm's calibrations there is little variation in terms of their forecasts statistical properties.

The *adaptive* gain specifications seem to be the only escaping to these regularities. Forecasting growth, e.g., their variance were remarkably higher (lower) than the other calibrations for the LS (SG) algorithm. The *adaptive* gains' performances, here assessed through

[18] Inspection of results restricted to the Branch and Evans (2006) sample corroborate this point. Another difference is that we use a real-time dataset both for the computation of forecasts and for their evaluation; comparatively, this was found to make a bigger difference for the results on growth.

Figure 4: Algorithms' forecasting *accuracy* by gain aggregated over time.

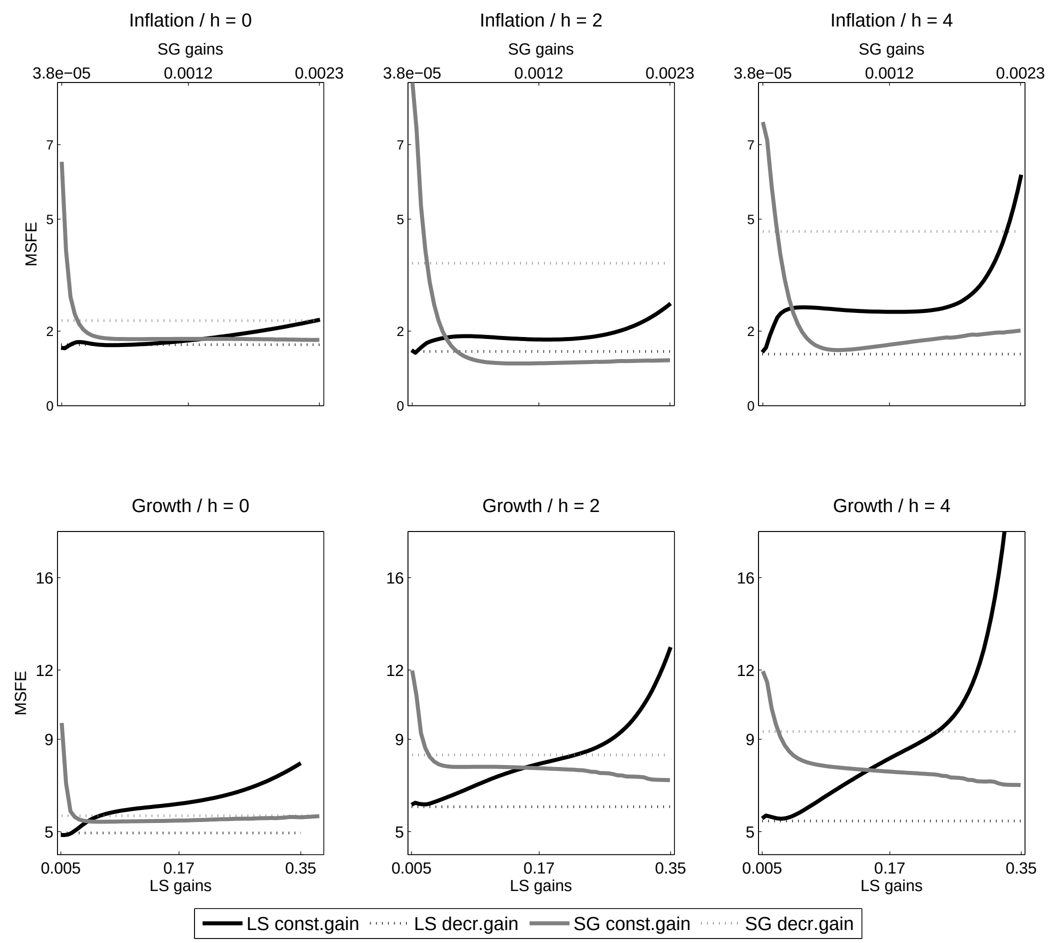

The MSFE plotted are computed for each gain over the evaluation sample of forecasts, 1981q1-2010q4. The gain calibration used for the LS/SG algorithms are indicated into the lower/upper horizontal axes, respectively, and correspond to those experimentally calibrated for algorithm's stability. Notice the minimums of these curves are not equivalent to the Full-sample gain calibrations presented in Table 1 due to the use of different sample periods. Similar figures corresponding to the Full-sample gain calibration can be found in Berardi and Galimberti (2012b, Fig. 2).

Table 2: Data and forecasts statistics by calibration - Inflation.

Series/Algorithm	Mean	Min	Max	Var	AR(1)	CorA	CorS	MSFE	MSFCE
Actuals	2.67	-0.33	9.39	2.68	0.63	1.00	0.81	0.00	1.03
Surveys	2.93	0.62	9.49	2.46	0.94	0.81	1.00	1.03	0.00
Least Squares									
-Full-smpl. /choice	2.99	-0.14	10.30	2.94	0.83	0.73	0.90	1.62	0.55
-In-smpl. /choice	2.98	-0.30	10.35	2.92	0.81	0.72	0.90	1.64	0.55
-Recur. /choice	3.07	0.13	10.35	2.80	0.81	0.72	0.90	1.69	0.56
-Adapt. /choice	2.99	0.71	10.80	3.30	0.85	0.71	0.88	1.86	0.74
-Full-smpl./primit.	2.97	-0.34	10.32	2.75	0.74	0.70	0.90	1.71	0.54
-In-smpl. /primit.	2.97	-0.40	10.39	2.83	0.76	0.71	0.90	1.70	0.55
-Recur. /primit.	3.06	0.01	10.39	2.57	0.78	0.72	0.89	1.62	0.55
-Adapt. /primit.	3.01	0.52	10.28	3.37	0.88	0.70	0.86	1.93	0.87
Stochastic Gradient									
-Full-smpl. /choice	2.77	0.12	11.89	2.65	0.69	0.67	0.88	1.76	0.62
-In-smpl. /choice	2.77	0.12	11.89	2.65	0.69	0.67	0.88	1.76	0.62
-Recur. /choice	2.69	-0.83	11.89	2.79	0.74	0.69	0.89	1.69	0.63
-Adapt. /choice	2.98	0.23	10.89	2.79	0.76	0.69	0.89	1.76	0.58
-Full-smpl./primit.	2.77	0.12	11.89	2.65	0.69	0.67	0.88	1.76	0.62
-In-smpl. /primit.	2.77	0.11	11.84	2.65	0.69	0.67	0.88	1.77	0.62
-Recur. /primit.	2.78	-0.08	11.84	2.62	0.68	0.67	0.88	1.76	0.63
-Adapt. /primit.	2.97	0.19	11.53	2.79	0.74	0.69	0.89	1.79	0.56

The forecasts statistics refer to those obtained for the first forecasting horizon, $h = 0$, and over the full evaluation sample from 1981q1 to 2010q4. The algorithms' forecasts refer to those from the VAR(1). Other than the usual descriptive statistics, AR(1) stands for the first order autocorrelation of each series, CorA and CorS stands for the correlation with the series of actual (real-time) and survey forecasts, respectively, and MSF(C)E stands for the mean squared forecast (comparison) errors.

Table 3: Data and forecasts statistics by calibration - Growth.

Series/Algorithm	Mean	Min	Max	Var	AR(1)	CorA	CorS	MSFE	MSFCE
Actuals	2.54	-6.14	8.67	6.17	0.50	1.00	0.78	0.00	2.56
Surveys	2.24	-5.19	7.01	3.14	0.72	0.78	1.00	2.56	0.00
Least Squares									
-Full-smpl. /choice	3.04	-3.66	6.58	1.58	0.60	0.39	0.58	5.55	2.74
-In-smpl. /choice	3.00	-4.15	6.87	1.90	0.59	0.37	0.57	5.72	2.85
-Recur. /choice	3.08	-0.95	6.70	1.18	0.53	0.38	0.63	5.55	2.61
-Adapt. /choice	2.80	-6.79	9.27	3.68	0.40	0.21	0.35	7.84	4.71
-Full-smpl./primit.	3.01	-2.02	5.46	1.07	0.58	0.47	0.64	5.03	2.46
-In-smpl. /primit.	3.04	-3.66	6.58	1.58	0.60	0.39	0.58	5.55	2.74
-Recur. /primit.	2.98	-0.95	6.00	1.10	0.52	0.42	0.66	5.24	2.34
-Adapt. /primit.	2.81	-5.48	7.43	3.47	0.57	0.30	0.50	6.90	3.60
Stochastic Gradient									
-Full-smpl. /choice	1.88	-2.13	4.66	1.32	0.48	0.43	0.59	5.44	2.16
-In-smpl. /choice	1.82	-2.10	4.57	1.24	0.47	0.43	0.59	5.52	2.20
-Recur. /choice	2.10	-1.24	4.75	1.32	0.53	0.40	0.55	5.38	2.23
-Adapt. /choice	2.13	-0.54	4.74	0.90	0.68	0.14	0.30	6.51	3.02
-Full-smpl./primit.	1.88	-2.13	4.66	1.32	0.48	0.43	0.59	5.44	2.16
-In-smpl. /primit.	1.86	-2.13	4.64	1.29	0.48	0.43	0.59	5.47	2.17
-Recur. /primit.	2.07	-1.85	4.68	1.46	0.53	0.41	0.56	5.37	2.20
-Adapt. /primit.	2.09	-0.56	4.35	0.77	0.64	0.20	0.38	6.23	2.73

See footnotes to Table 2.

the statistics in the last four columns, also stand out relative to the others, though negatively. These results, however, are representative of only the first horizon of forecasts obtained with the VAR(1) model specification. To achieve a more broad assessment of how the different gain calibrations compare, we introduce a key synthetic evaluation statistic in what follows.

4.3 Gain selection

A relative assessment of the learning mechanisms and their calibrations may be obtained by conducting paired comparisons of their forecasting performance. However, our coverage of multiple forecasting horizons and VAR lag order specifications, for robustness, requires performing a high quantity of such comparisons[19]. Hence, for the ease of exposition we now adopt *hit rate* measures to synthesize these evaluations. The *hit rate* stands for the frequency by which the forecasts associated to a given algorithm/calibration is found to outperform those associate to its competitor(s) with respect to one of our evaluation criteria.

To illustrate the above definition consider the hit rates presented in Table 4 comparing the different methods of gain selection. The first comparison focus on the LS algorithm applied in forecasting inflation with gains determined as a *choice*; here the hit rate of 45% associated to the Full-sample selection method, e.g., indicates that this method outperformed the others in 9 out of the 20 comparisons conducted for each combination of VAR lag order and forecasting horizon[20]. The DM-20% and GW-20% statistics associated to these hit rates then represent the frequency by which the losses associated to the outperforming method is found to be statistically different of the second-best method according to the corresponding test at a 20% level of significance. So, in our example case of the LS/choice calibrations we have that only 3 (15% of 20) of the 9 cases where the Full-sample was favored presented a GW test p-value below 20% when compared to the second-best method of gain selection.

The comparative evaluations of the different methods of gain selection are presented in

[19]To be specific, 40 for each pair of algorithms/calibrations: 5 horizons $\times$ 4 VARs $\times$ 2 evaluation criteria.

[20]Draws are handled by counting for both tied competitors. These may occur when, e.g., different methods of gain selection pick the same gain.

Table 4: Hit rates comparing gain selection methods.

(a) Forecast *accuracy*.

Variables	Full-sample wins			In-sample wins			Recursive wins			Adaptive wins		
- Alg./Gain	Hit rate	DM-20%	GW-20%	Hit rate	DM-20%	GW-20%	Hit rate	DM-20%	GW-20%	Hit rate	DM-20%	GW-20%
Inflation												
- LS/choice	45%	0%	15%	5%	0%	0%	**50%**	0%	10%	0%	0%	0%
- LS/primitive	22.5%	0%	0%	7.5%	0%	5%	**70%**	35%	30%	0%	0%	0%
- SG/choice	**45%**	0%	0%	20%	0%	0%	35%	5%	30%	0%	0%	0%
- SG/primitive	**55%**	0%	10%	15%	0%	10%	20%	15%	10%	10%	0%	0%
Growth												
- LS/choice	**70%**	20%	10%	20%	0%	0%	10%	0%	0%	0%	0%	0%
- LS/primitive	**47.5%**	5%	5%	7.5%	0%	0%	45%	0%	10%	0%	0%	0%
- SG/choice	**40%**	10%	15%	10%	0%	0%	10%	0%	0%	**40%**	0%	10%
- SG/primitive	10%	5%	0%	0%	0%	0%	35%	0%	10%	**55%**	0%	10%

(b) Forecast *resemblance*.

Variables	Full-sample wins			In-sample wins			Recursive wins			Adaptive wins		
- Alg./Gain	Hit rate	DM-20%	GW-20%	Hit rate	DM-20%	GW-20%	Hit rate	DM-20%	GW-20%	Hit rate	DM-20%	GW-20%
Inflation												
- LS/choice	0%	0%	0%	10%	5%	0%	**90%**	80%	50%	0%	0%	0%
- LS/primitive	15%	0%	0%	5%	0%	0%	**80%**	45%	35%	0%	0%	0%
- SG/choice	**47.5%**	15%	20%	22.5%	0%	0%	5%	0%	5%	25%	20%	10%
- SG/primitive	**45%**	25%	5%	0%	0%	0%	10%	0%	0%	**45%**	20%	15%
Growth												
- LS/choice	**70%**	15%	20%	20%	0%	0%	10%	5%	5%	0%	0%	0%
- LS/primitive	**50%**	0%	10%	10%	0%	0%	40%	0%	25%	0%	0%	0%
- SG/choice	30%	10%	15%	0%	0%	0%	**50%**	0%	45%	20%	0%	20%
- SG/primitive	15%	5%	10%	0%	0%	0%	**70%**	25%	25%	15%	0%	5%

The hit rates sum up forecasting horizons (5) and VAR lag orders (4). DM-20% and GW-20% represents the frequency within which the associated specification was found to have superior performance over the runner up specification with statistical significance below the 20% level using the Diebold and Mariano (1995) and the Giacomini and White (2006) test(s).

Table 4 according to the variable forecasted, the algorithm, and the gain determination approach. There is notable sensitivity on the results presented along these dimensions, though the combinations of variable/algorithm seem to be a main factor: in forecasting inflation, the *recursive* method is favored both in terms of *accuracy* as in terms of *resemblance* for the LS, whilst a similar result is found with the *full-sample* method for the SG algorithm; in forecasting growth these patterns change place, with the *full-sample* gain specification being favored for the LS, whilst under the SG the results favor time-varying gain specifications. Thus, where the evidence is favorable to a fixed gain calibration for one algorithm, it is favorable for time-varying gains for the other[21].

In spite of some evidence favoring fixed gains, the *in-sample* method has found scarce evidence in its favor. This indicates the relevance of the information contained in the post *in-sample* observations (after 1980q4) for the calibration of the algorithms, which is not surprising given the evidence of structural breaks around this period (see, e.g., Stock and Watson, 2003). Statistical significance was also scarce in most of the comparisons presented, with some exceptions under the *resemblance* criterion, e.g., the LS *recursive* wins in forecasting inflation. This result does not come at surprise neither; our previous inspection of the forecasts statistics already indicated these calibrations presented similar properties, which probably hindered the potential of the statistical tests to distinguish between these series.

Regarding the time-varying gain specifications the most remarkable result relates to the poorer performance of the *adaptive* method under the LS algorithm. An explanation for this result may be drawn looking back into how the *recursive* selections evolved compared to the *adaptive* gains, in the top panels of figures 2 and 3, respectively; clearly the *adaptive* gain values are generally higher than those picked by the *recursive* method. Hence, our results indicate that the upper bound for the LS *adaptive* gains may have been too loose, which is also consistent with the observation by Kushner and Yang (1995) on the relevance

[21]One might have mistakenly expected that the *recursive* method, for its flexibility, should have picked the same gain as the *full-sample*, giving the superiority of this latter; notice, however, that the *recursive* method does not have access to the same information set as that given to the *full-sample* method.

of this parameter for the algorithm's performance. Clearly, this was not a problem for the SG algorithm.

4.4 Gain determination and implications

We now get to the issue on the determination of the learning gains. The *hit rates* comparing the gain as a *choice* and the gain as a *primitive* assumptions are presented in Table 5. The calculations adopted to obtain these hit rates follow in similar lines as those for Table 4, as explained above. The only difference is that here there are only two "competitors". Take the case of the LS applied to forecast inflation using the Full-sample gain selection method: the 90% hit rate associated to the gain determined as a *primitive* indicates that this assumption outperformed the alternative determination of the gain as a *choice* in 18 out of the 20 comparisons combining the VAR lag orders and forecasting horizons we consider. Furthermore, we can see that in 5 (9) of these victories the DM (GW) test indicated that the series of losses associated the two methods of gain determination are statistically different at a significance level below 20%.

In contrast to our previous results, here the evidence is crystal clear: the gain as a *primitive* is overwhelmingly favored on both evaluation criteria. The few threats to the *primitive* gain dominance can be easily disqualified: (i) the LS with *adaptive* gains, we recall, was found to have a poor performance relative to its alternatives; (ii) under the other few calibrations where the gain as *choice* prevailed, i.e., the LS fixed gains for growth, statistical significance is still favoring the *primitive* gain.

From an applied perspective, our findings provide some important guidance on how the learning gain can be appropriately calibrated. The gain as a *primitive* is associated to a minimization of the distance between the algorithm's forecasts and their survey's counterparts. To select a particular sequence of gains on the basis of this measure, one needs to take into account the algorithm and variable of interest. For inflation, flexibility in the time-variations of the gain is required by the LS algorithm; the SG, on the other hand, is more appropriately

Table 5: Hit rates comparing gain determination assumptions.

(a) Forecast *accuracy*.

Variables	*Gain as a choice* wins			*Gain as a primitive* wins		
- Alg./Gain Select.	Hit rate	DM-20%	GW-20%	Hit rate	DM-20%	GW-20%
Inflation						
- LS/Full-sample	10%	0%	0%	**90%**	25%	45%
- LS/In-sample	5%	0%	0%	**95%**	35%	45%
- LS/Recursive	5%	0%	0%	**95%**	50%	35%
- LS/Adaptive	40%	5%	15%	**60%**	15%	5%
- SG/Full-sample	17.5%	0%	0%	**82.5%**	0%	20%
- SG/In-sample	22.5%	0%	0%	**77.5%**	0%	15%
- SG/Recursive	40%	5%	25%	**60%**	0%	15%
- SG/Adaptive	15%	0%	5%	**85%**	40%	20%
Growth						
- LS/Full-sample	**50%**	0%	5%	**50%**	10%	10%
- LS/In-sample	**55%**	10%	15%	45%	25%	0%
- LS/Recursive	15%	5%	5%	85%	35%	10%
- LS/Adaptive	**75%**	30%	15%	25%	0%	5%
- SG/Full-sample	27.5%	0%	5%	**72.5%**	5%	10%
- SG/In-sample	25%	0%	0%	**75%**	15%	15%
- SG/Recursive	0%	0%	0%	**100%**	35%	55%
- SG/Adaptive	20%	0%	0%	**80%**	40%	10%

(b) Forecast *resemblance*.

Variables	*Gain as a choice* wins			*Gain as a primitive* wins		
- Alg./Gain Select.	Hit rate	DM-20%	GW-20%	Hit rate	DM-20%	GW-20%
Inflation						
- LS/Full-sample	15%	0%	0%	**85%**	50%	30%
- LS/In-sample	15%	0%	0%	**85%**	55%	35%
- LS/Recursive	40%	25%	15%	**60%**	20%	15%
- LS/Adaptive	35%	10%	10%	**65%**	15%	5%
- SG/Full-sample	17.5%	0%	0%	**82.5%**	65%	20%
- SG/In-sample	22.5%	0%	0%	**77.5%**	60%	25%
- SG/Recursive	5%	0%	5%	**95%**	65%	15%
- SG/Adaptive	20%	0%	5%	**80%**	75%	45%
Growth						
- LS/Full-sample	45%	5%	20%	**55%**	20%	20%
- LS/In-sample	**55%**	25%	45%	45%	30%	25%
- LS/Recursive	10%	5%	5%	**90%**	30%	30%
- LS/Adaptive	**65%**	10%	5%	35%	0%	10%
- SG/Full-sample	2.5%	0%	0%	**97.5%**	60%	85%
- SG/In-sample	5%	0%	0%	**95%**	70%	80%
- SG/Recursive	0%	0%	0%	**100%**	75%	80%
- SG/Adaptive	5%	0%	5%	**95%**	50%	45%

See footnotes to Table 4.

calibrated with a fixed gain value. For growth, these specifications change place: the SG is the one requiring time-varying gains, whilst a fixed gain does the job with the LS.

Finally, our evidence favoring the interpretation of the learning gains as *primitive* parameters of agents learning-to-forecast behavior also has important implications within the context of the theoretical literature, as we outlined in section 3. Our results contribute to the debate on the internal (in)consistency between the adaptive learning approach and the rationality of mechanisms of expectations formation. Namely, by favoring the gain as a *primitive*, we provide support to the view that agents are essentially characterized as bounded rational; when their expectations formation process is represented by adaptive learning algorithms, even if this process is likely to converge to a RE equilibrium, it does not imply that these agents will rationally optimize on their calibration. Agents simply do not seem to take it as a *choice*.

5 Concluding remarks

In this paper we have studied empirically the issue of what is the most appropriate interpretation of the learning gains in adaptive algorithms, and how can a calibration strategy be developed so as to reflect it. Our main insight relates to a distinction on the rationale given to the determination of the learning gains: as a *choice* of rationally optimizing agents, or as a *primitive* parameter of bounded rational agents.

We have also produced some renewed numerical calibrations of the learning gains for applications on real-time US data of inflation and output growth. Consistent to our analysis, our gain calibrations are segmented by the different methods and assumptions in the determination of these gains, and according to the lag orders of VAR models and the forecasting horizons. Significant heterogeneity was found with respect to these dimensions as well.

Our results provide strong evidence in favor of the interpretation of the learning gains as *primitive* parameters of agents learning-to-forecast behavior. Furthermore, our evidence also points to some heterogeneity in the time evolution of this behavior with respect to the variable forecasted and the algorithm adopted.

Therefore, the main implication of our results is that agents' adaptive behavior in the adjustment of their learning algorithms is better represented from a bounded rationality point of view rather than from a rationally optimizing interpretation. This finding turns out to be relevant to a long established debate on the internal (in)consistency between the adaptive learning approach and the rationality of expectations. Namely, there is no reason to expect, or question, a rational agent to optimize on the *choice* of a learning gain, given that our evidence suggest that it stands as a *primitive* parameter of agents bounded rationality.

A Appendix

A.1 Details on data

Short time series history: some vintages lack of earlier observations due to delays into BEA revisions (see Philadelphia's Fed documentations). This was the case of the vintages of 1992q1-1992q4 (missing data from 1947-1958), 1996q1-1997q1 (missing data from 1947-1959q2), and 1999q4-2000q1 (missing data from 1947-1958). We circumvent this problem (to turn the dataset vintages-balanced) by reproducing observations from the last available vintage while rescaling in accordance to the ratio between the first observation available in the missing observation vintage and the value observed for the same period in the vintage being used as source for the missing observations.

Missing observation for 1995q4 in vintage 1996q1: as a result of the US federal government shutdown in late 1995, the observation for 1995q4 was missing in the 1996q1 vintage. Fortunately, this is the only point in this dataset that this happens. We fulfill this gap by using the observation available in the March 1996 monthly vintage for the same series. Incidentally, the SPF 1996q1 median backcast for 1995q4 is identical to the value later observed in March 1996, thence, our simplifying procedure is not favoring any method.

Caveat on SPF's forecasts for Real GDP: forecasts for real GDP were not asked in the surveys prior to 1981q3. To extend this series of forecast back to 1968q4, real GDP prior to 1981q3 is computed by using the formula (nominal GDP / GDP prices) * 100.

A.2 Review of statistical tests for equal predictive ability

We want to determine whether two series of forecasts are statistically different from each other. Let $f_{1,t,h}$ and $f_{2,t,h}$ stand for these forecasts, where h (going from 0 to 4 in our case) denotes the horizon at which these forecasts were made, and y_t stand for the series of targets of these forecasts. Let the losses associated to each of these forecasts be given by $\mathcal{L}(f_{1,t,h}, y_t)$ and $\mathcal{L}(f_{2,t,h}, y_t)$. Letting $d_{t,h} = \mathcal{L}(f_{1,t,h}, y_t) - \mathcal{L}(f_{2,t,h}, y_t)$ denote the series of loss differentials between the two forecasts at horizon h, the Diebold and Mariano (1995) test evaluates whether their average loss differences,

$$\overline{d}_h = \frac{1}{T}\sum_{i=1}^{T} d_{i,h}, \tag{14}$$

is significantly different from zero. Under the null hypothesis of equal predictive ability the DM statistic,

$$DM_h = \frac{\overline{d}_h}{\sqrt{\hat{\sigma}_d^2/T}}, \tag{15}$$

has a t-distribution with $T-1$ degrees of freedom, where $\hat{\sigma}_d^2$ is an estimate of the long-run variance of $d_{t,h}$. For the estimation of $\hat{\sigma}_d^2$ we adopt the heteroskedasticity and autocorrelation consistent (HAC) estimator proposed by Newey and West (1987).

The Giacomini and White (2006) test, in contrast, evaluates the null hypothesis of equal conditional predictive ability. The main caveat on this test relates to the specification of a test function, $\mathbf{q}_{t,h}$ containing q instruments, which attempts to control for the informational conditioning required by the null hypothesis. To test the conditional moment restriction $E\left[\mathbf{q}_{t,h}d_{t,h}\right]=\mathbf{0}$, a Wald-type test statistic is proposed having the form of

$$GW_h = T\left(T^{-1}\sum_{i=1}^{T}\mathbf{q}_{i,h}d_{i,h}\right)'\hat{\mathbf{\Omega}}_h^{-1}\left(T^{-1}\sum_{i=1}^{T}\mathbf{q}_{i,h}d_{i,h}\right), \tag{16}$$

where $\hat{\mathbf{\Omega}}_h$ is a $q\times q$ consistent estimate of the covariance matrix of $\mathbf{q}_{t,h}d_{t,h}$. Under the null hypothesis of equal conditional predictive ability GW_h has a χ^2_q distribution.

Apart from the first horizon, $\hat{\mathbf{\Omega}}_h$ is again estimated using the HAC estimator of Newey and West (1987), with h determining the truncated kernel bandwidth. For the case of the first horizon, Giacomini and White (2006) simplify the computation of (16) to be given by TR^2, where R^2 is the uncentered squared multiple correlation coefficient obtained by regressing a constant unity on $\mathbf{q}_{t,h}d_{t,h}$. Finally, regarding the specification of $\mathbf{q}_{t,h}$, in the lack of better alternatives, the recommendation is for the use of h-lagged loss differentials. Thus, in our calculations we set $\mathbf{q}_{t,h}=d_{t-h,h}$.

A.3 Derivation of adaptive gain algorithms

To derive the LSA and the SGA expressions in (6)-(10) and (11)-(13), we start by defining the gains adaptation recursions, as given by

$$\gamma_t = \gamma_{t-1}-\alpha_\gamma\hat{\nabla}_t^\gamma, \tag{17}$$

$$\mu_t = \mu_{t-1}-\alpha_\mu\hat{\nabla}_t^\mu, \tag{18}$$

where α_γ and α_μ represent small adaptation constants, and $\hat{\nabla}_t^\gamma$ and $\hat{\nabla}_t^\mu$ stand for estimates of the gradients $\nabla_t^\gamma=\frac{\partial J_t}{\partial\gamma}$ and $\nabla_t^\mu=\frac{\partial J_t}{\partial\mu}$, for the LS and the SG, respectively. The key step then is to find the relevant gradients and plug their stochastic approximation in the above recursions.

We first derive the LSA. Taking the first derivative of J_t with respect to γ we obtain the gradient

$$\nabla_t^\gamma = \frac{\partial J_t^{LSA}}{\partial \gamma} = -E\left[\mathbf{x}_t' \frac{\partial \hat{\boldsymbol{\theta}}_{t-1}^{LSA}}{\partial \gamma}\left(z_t - \mathbf{x}_t'\hat{\boldsymbol{\theta}}_{t-1}^{LSA}\right)\right], \tag{19}$$

which is stochastically approximated as

$$\hat{\nabla}_t^\gamma = -\mathbf{x}_t'\hat{\boldsymbol{\Psi}}_{t-1}^{LSA}\left(z_t - \mathbf{x}_t'\hat{\boldsymbol{\theta}}_{t-1}^{LSA}\right), \tag{20}$$

where $\hat{\boldsymbol{\Psi}}_t^{LSA}$ stands for a recursive estimate of $\partial\hat{\boldsymbol{\theta}}_t^{LSA}/\partial\gamma$. Differentiating (2) and (3) we obtain

$$\begin{aligned}\frac{\partial \hat{\boldsymbol{\theta}}_t^{LSA}}{\partial \gamma} \quad = \quad & \frac{\partial \hat{\boldsymbol{\theta}}_{t-1}^{LSA}}{\partial \gamma} + \mathbf{R}_t^{-1}\mathbf{x}_t\left(y_t - \mathbf{x}_t'\hat{\boldsymbol{\theta}}_{t-1}^{LSA}\right)\ldots \\ & -\gamma_t\mathbf{R}_t^{-1}\frac{\partial \mathbf{R}_t}{\partial \gamma}\mathbf{R}_t^{-1}\mathbf{x}_t\left(y_t - \mathbf{x}_t'\hat{\boldsymbol{\theta}}_{t-1}^{LSA}\right) - \gamma_t\mathbf{R}_t^{-1}\mathbf{x}_t\mathbf{x}_t'\frac{\partial \hat{\boldsymbol{\theta}}_{t-1}^{LSA}}{\partial \gamma},\end{aligned} \tag{21}$$

$$\frac{\partial \mathbf{R}_t}{\partial \gamma} = \frac{\partial \mathbf{R}_{t-1}}{\partial \gamma} + \mathbf{x}_t\mathbf{x}_t' - \mathbf{R}_{t-1} - \gamma_t\frac{\partial \mathbf{R}_{t-1}}{\partial \gamma}. \tag{22}$$

Letting $\hat{\mathbf{S}}_t$ stand for the recursive estimate of $\partial\mathbf{R}_t/\partial\gamma$, and subsituting (20), (21), and (22) into (17) we obtain the LSA algorithm.

We follow similar steps for the SGA. Particularly, as estimate of the loss gradient is given by

$$\hat{\nabla}_t^\mu = -\mathbf{x}_t'\hat{\boldsymbol{\Psi}}_{t-1}^{SGA}\left(z_t - \mathbf{x}_t'\hat{\boldsymbol{\theta}}_{t-1}^{SGA}\right), \tag{23}$$

with $\hat{\boldsymbol{\Psi}}_t^{SGA} = \partial\hat{\boldsymbol{\theta}}_t^{SGA}/\partial\mu$ obtained differentiating (4),

$$\frac{\partial \hat{\boldsymbol{\theta}}_t^{SGA}}{\partial \mu} = \frac{\partial \hat{\boldsymbol{\theta}}_{t-1}^{SGA}}{\partial \mu} + \mathbf{x}_t\left(y_t - \mathbf{x}_t'\hat{\boldsymbol{\theta}}_{t-1}^{SGA}\right) - \mu_t\mathbf{x}_t\mathbf{x}_t'\frac{\partial \hat{\boldsymbol{\theta}}_{t-1}^{SGA}}{\partial \mu}, \tag{24}$$

which together with (18), reduces to the SGA algorithm.

References

Barucci, E. (1999). Heterogeneous beliefs and learning in forward looking economic models. *Journal of Evolutionary Economics 9*(4), 453–464.

Barucci, E. and L. Landi (1997). Least mean squares learning in self-referential linear stochastic models. *Economics Letters 57*(3), 313–317.

Benveniste, A., M. Metivier, and P. Priouret (1990). *Adaptive Algorithms and Stochastic Approximations*. Springer-Verlag.

Berardi, M. and J. K. Galimberti (2012a). On the initialization of adaptive learning algorithms: A review of methods and a new smoothing-based routine. Centre for Growth and Business Cycle Research Discussion Paper Series 175, Economics, The Univeristy of Manchester.

Berardi, M. and J. K. Galimberti (2012b). On the plausibility of adaptive learning in macroeconomics: A puzzling conflict in the choice of the representative algorithm. Centre for Growth and Business Cycle Research Discussion Paper Series 177, Economics, The Univeristy of Manchester.

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.

Branch, W. A. and G. W. Evans (2006). A simple recursive forecasting model. *Economics Letters 91*(2), 158–166.

Bray, M. (1982). Learning, estimation, and the stability of rational expectations. *Journal of Economic Theory 26*(2), 318–339.

Bray, M. M. and N. E. Savin (1986). Rational expectations equilibria, learning, and model specification. *Econometrica 54*(5), 1129–1160.

Brock, W. A. and C. H. Hommes (1997). A rational route to randomness. *Econometrica 65*(5), 1059–1095.

Bullard, J. (1992). Time-varying parameters and nonconvergence to rational expectations under least squares learning. *Economics Letters 40*(2), 159 – 166.

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

Cogley, T. and T. J. Sargent (2005). Drifts and volatilities: monetary policies and outcomes in the post wwii us. *Review of Economic Dynamics 8*(2), 262 – 302.

Diebold, F. X. (2013, November). Comparing predictive accuracy, twenty years later: a personal perspective on the use and abuse of diebold-mariano tests. Mimeo.

Diebold, F. X. and R. S. Mariano (1995). Comparing predictive accuracy. *Journal of Business & Economic Statistics 13*(3), 253–63.

Evans, G. W. and S. Honkapohja (1998). Stochastic gradient learning in the cobweb model. *Economics Letters 61*(3), 333–337.

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.

Evans, G. W. and G. Ramey (1998). Calculation, adaptation and rational expectations. *Macroeconomic Dynamics 2*(02), 156–182.

Giacomini, R. and H. White (2006). Tests of conditional predictive ability. *Econometrica 74*(6), 1545–1578.

Haykin, S. S. (2001). *Adaptive Filter Theory* (4th ed.). Prentice Hall Information and System Sciences Series. New Jersey, USA: Prentice Hall.

Honkapohja, S. and K. Mitra (2006). Learning stability in economies with heterogeneous agents. *Review of Economic Dynamics 9*(2), 284–309.

Kostyshyna, O. (2012, 11). 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.

Macchi, O. (1995). *Adaptive Processing: the least mean squares approach with applications in transmission*. John Wiley & Sons.

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.

Margaritis, D. (1990). A time-varying model of rational learning. *Economics Letters 33*(4), 309 – 314.

McGough, B. (2003). Statistical learning with time-varying parameters. *Macroeconomic Dynamics 7*(01), 119–139.

Milani, F. (2007, October). Expectations, learning and macroeconomic persistence. *Journal of Monetary Economics 54*(7), 2065–2082.

Milani, F. (2008). Learning, monetary policy rules, and macroeconomic stability. *Journal of Economic Dynamics and Control 32*(10), 3148 – 3165.

Milani, F. (2011). Expectation shocks and learning as drivers of the business cycle. *The Economic Journal 121*(552), 379–401.

Newey, W. K. and K. D. West (1987, May). A simple, positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix. *Econometrica 55*(3), 703–08.

Orphanides, A. and J. C. Williams (2005, November). The decline of activist stabilization policy: Natural rate misperceptions, learning, and expectations. *Journal of Economic Dynamics and Control 29*(11), 1927–1950.

Pfajfar, D. and E. Santoro (2010). Heterogeneity, learning and information stickiness in inflation expectations. *Journal of Economic Behavior & Organization 75*(3), 426–444.

Sargent, T., N. Williams, and T. Zha (2006). Shocks and government beliefs: The rise and fall of american inflation. *American Economic Review 96*(4), 1193–1224.

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

Sargent, T. J. and N. Williams (2005). Impacts of priors on convergence and escapes from nash inflation. *Review of Economic Dynamics 8*(2), 360 – 391.

Sims, C. A. and T. Zha (2006). Were there regime switches in u.s. monetary policy? *American Economic Review 96*(1), 54–81.

Stark, T. and D. Croushore (2002). Forecasting with a real-time data set for macroeconomists. *Journal of Macroeconomics 24*(4), 507 – 531.

Stock, J. H. and M. W. Watson (2003, December). Has the business cycle changed and why? In *NBER Macroeconomics Annual 2002, Volume 17*, NBER Chapters, pp. 159–230. National Bureau of Economic Research, Inc.

Weber, A. (2010). *Heterogeneous expectations, learning and European inflation dynamics*, Chapter 12, pp. 261–305. Cambridge University Press.