

Berardi, Michele; Galimberti, Jaqueson K.

Working Paper

Empirical calibration of adaptive learning

KOF Working Papers, No. 392

Provided in Cooperation with:

KOF Swiss Economic Institute, ETH Zurich

Suggested Citation: Berardi, Michele; Galimberti, Jaqueson K. (2015) : Empirical calibration of adaptive learning, KOF Working Papers, No. 392, ETH Zurich, KOF Swiss Economic Institute, Zurich, <https://doi.org/10.3929/ethz-a-010513113>

This Version is available at:

<https://hdl.handle.net/10419/126239>

Standard-Nutzungsbedingungen:

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

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

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

Terms of use:

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

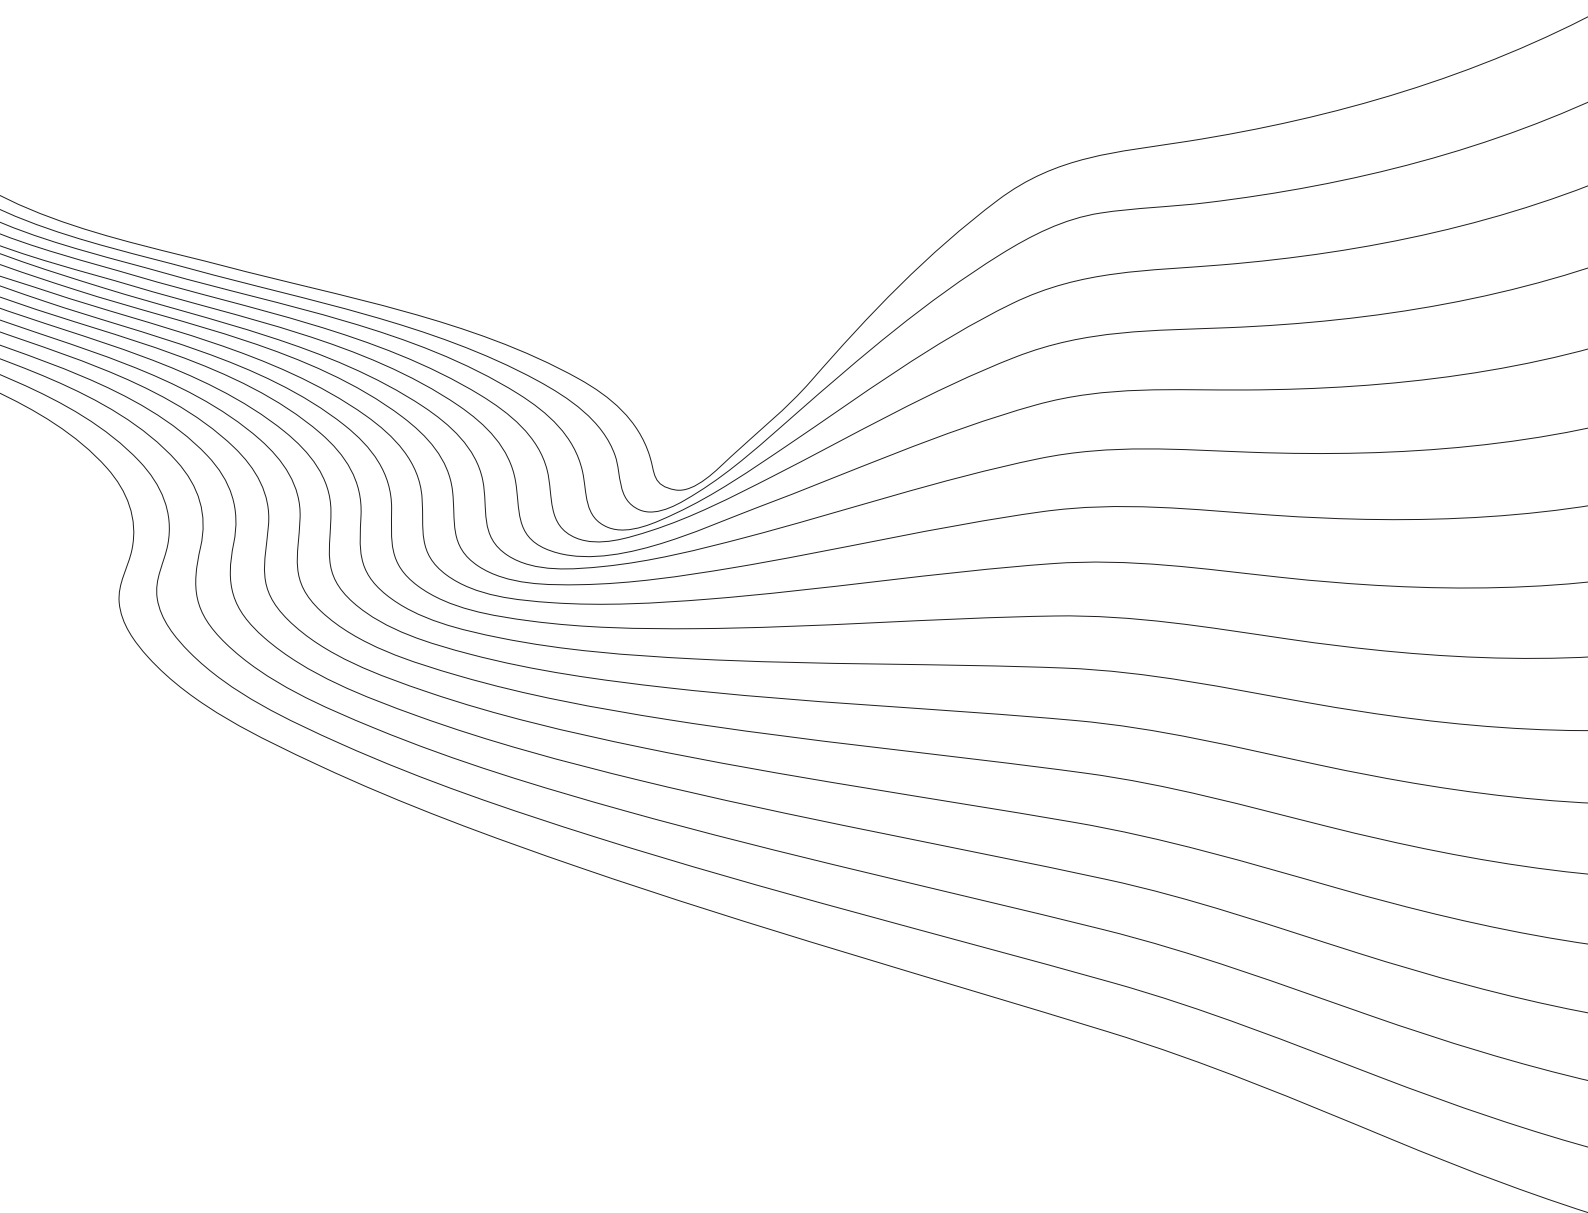
You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

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

KOF Working Papers

Empirical Calibration of Adaptive Learning

Michele Berardi and Jaqueson K. Galimberti



KOF

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

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

Empirical Calibration of Adaptive Learning*

MICHELE BERARDI

University of Manchester

JAQUESON K. GALIMBERTI[†]

ETH Zurich

Abstract

Adaptive learning introduces persistence in the evolution of agents' beliefs over time. For applied purposes this is a convenient feature to help explain why economies present sluggish adjustments towards equilibrium. The pace of learning is directly determined by the gain parameter, which regulates how quickly new information is incorporated into agents' beliefs.

We document renewed empirical calibrations of plausible gain values for adaptive learning applications to macroeconomic data. We cover a broad range of model specifications of applied interest. Our analysis also includes innovative approaches to the endogenous determination of time-varying gains in real-time, and a thorough discussion of the different theoretical interpretations of the learning gain. We also evaluate the merits of different approaches to the gain calibration according to their performance in forecasting macroeconomic variables and in matching survey forecasts.

Our results indicate a great degree of heterogeneity in the gain calibrations according to the variable forecasted and the lag length of the model specifications. Calibrations to match survey forecasts are found to be lower than those derived according to the forecasting performance, suggesting some degree of bounded rationality in the speed with which agents update their beliefs.

Keywords: expectations, forecasting, bounded rationality, real-time, recursive estimation.

JEL codes: D83, E03, E37.

1 Introduction

The practice of modeling expectations through the use of adaptive learning in macroeconomic models has become increasingly popular in the recent applied literature¹. By allowing persis-

*An earlier version of this paper was presented at the conference on Expectations in Dynamic Macroeconomic Models, held in Helsinki, 2014. We thank to the organizers of that conference for the opportunity, to the Bank of Finland's staff for the hospitality, and to the participants for the comments that helped us shape the current presentation.

[†]Corresponding author. E-mail: galimberti@kof.ethz.ch.

¹For some time the literature on adaptive learning developed with a focus on the theoretical debate about the convergence of learning to rational expectations equilibria (see Evans and Honkapohja, 2001). The seminal

tence in the evolution of agents' beliefs over time, adaptive learning provides an alternative² explanation for the common observation of inertia in the dynamics of macroeconomic variables. Importantly, the persistence introduced by learning is directly related to the calibration of the learning algorithm assumed to represent the process through which agents update their beliefs. The implementation of this recursive algorithm 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, particularly aiming to provide a guide for applied researchers on plausible calibration values of the learning gain.

From a theoretical point of view, the gain determines both whether convergence to a rational expectations (RE) equilibrium takes place and the dynamic properties of the transition towards that equilibrium. In order to obtain positive convergence results, the analysis of learning places strong restrictions on the sequence of gains. Examples of these restrictions include the requirement of a "sufficiently small" gain to guarantee weaker (in distribution) convergence results (see Evans and Honkapohja, 2001, Ch.7). Furthermore, 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).

For applied purposes, the importance of the gain is related to its role in determining the statistical properties of the evolution of agents' beliefs. While adaptive learning establishes a new channel for the dynamic dissipation of structural (e.g., Huang et al., 2009; Eusepi and Preston, 2011) and expectational shocks (Milani, 2011) throughout the economy, the intensity of these effects is directly determined by the value of the learning gain. On one hand, higher gains can be associated with higher degrees of variability in agents beliefs (see, e.g., Carceles-Poveda and Giannitsarou, 2007, pp. 2691-4), which would ultimately improve the explanatory power of models with learning; on the other, it is under smaller gain values that learning becomes more sluggish, hence increasing the persistence of deviations from the equilibrium path³.

In this context, our main contribution in this paper is the documentation of plausible gain calibrations according to their quality in forecasting macroeconomic variables and in matching survey forecasts. Moreover, we explore different elements of interest for applied research, such

contribution in the applied literature is Sargent (1999), with a study on the role of learning by the monetary authority in the evolution of US inflation rates, followed by Marcat and Nicolini (2003), who showed how learning can help explain the recurrence of hyperinflation episodes observed in several South-American countries during the 1980's. The literature that followed then unfolded into two main strands: policy-oriented studies of particular episodes of macroeconomic turmoil (mainly the US Great Inflation period during the 1970's; see, e.g., Bullard and Eusepi, 2005; Orphanides and Williams, 2005; Primiceri, 2006; Sargent et al., 2006; Milani, 2008); and business cycles modeling studies (Williams, 2003; Huang et al., 2009; Chevillon et al., 2010; Eusepi and Preston, 2011; Slobodyan and Wouters, 2012; Milani, 2007, 2011, 2014; Ormeño and Molnár, 2015).

²Common extensions to add persistence in the standard framework of dynamic stochastic general equilibrium (DSGE) models include consumption habits, price indexation and adjustment costs (Christiano et al., 2005), or more directly through the introduction of persistent structural shocks (Smets and Wouters, 2007).

³This is consistent with a well known trade-off between speed and accuracy of recursive estimation from the engineering literature: on one extreme, tracking can be slower than the system actual time variations, but with less noisy estimates; on the other extreme, tracking can be made as rapid as the time-varying context, but with estimates much more contaminated by noise (see Benveniste et al., 1990, Part I, Chapters 1 and 4).

as: several model specifications; a proper account of real-time data restrictions; alternative foundations for an endogenous determination of the gains; and different roles for the gains. Although our results indicate that the gain calibrations are sensitive to all of these features, one of our most important findings is that the gains should be allowed to differ for different variables requiring forecasting.

1.1 Related literature

Theoretical foundations

There are different interpretations for the role played by the learning gains. Sargent and Williams (2005); Evans et al. (2010) establish a Bayesian interpretation of learning. Other than relating the learning algorithm to the Kalman filter, the Bayesian optimal estimator of a random walk time-varying parameters model, the priors also carry implications about the specific calibrations of the learning gains. Hence, in a Bayesian framework the gains can be interpreted as representing agents' prior beliefs about the evolution of the economy.

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). From the estimation problem originating the learning algorithm, it is possible to infer how the gains determine the weights assigned to each observation included in the estimation (the precise relation will be presented below, in Eqs. 10-11). Namely, the higher the gain, the higher the emphasis given to more recent observations; hence, higher gains imply a smaller memory of past observations in the current estimates provided by the learning algorithm.

Finally, a broader distinction on the possible foundations of the learning gains is obtained considering how these are assumed to be determined from the point of view of the learning agents (see Galimberti, 2013). First, one can assume that the learning gain stands as a primitive parameter of agents learning-to-forecast behavior. An immediate example is given by the interpretation of the gain as a memory parameter, as discussed above.

The second option, in contrast, is to assume that the learning gain is determined as a choice by the agents adopting a given learning algorithm to update their expectations. Approaches following along these lines often require additional behavioral assumptions to discipline such a choice. Examples include: the adaptively rational expectations approach of Brock and Hommes (1997), where agents are assumed to select between a set of predictors according to a fitness measure; and, 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.

Previous calibration attempts

Three of the main alternatives for the specification of the 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 was the seminal choice in the learning literature, so as to match the recursive form of the well known Ordinary Least Squares (OLS) estimator from basic econometrics. 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 learning mechanism is appropriate only along the time-invariant path of a RE equilibrium, where learning itself is indeed pointless (see Bray and Savin, 1986).

Extensive evidence (see, e.g., Stock and Watson, 2003; Cogley and Sargent, 2005) favoring time-varying parameter models of the economy has, nevertheless, challenged this paradigm, and the departure from the parameter constancy assumption has naturally led to the requirement of adjustments to the learning rules as well (see Margaritis, 1990; Bullard, 1992; McGough, 2003). 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. More recently, alternative approaches of time-varying gains also have been proposed to deal with relatively stable systems that occasionally undergo structural breaks (see Marcet and Nicolini, 2003; Kostyshyna, 2012; Milani, 2014); one key issue in the introduction of time-varying gains is to find an appropriate statistical foundation for the gain adjustment mechanism.

Finally, different approaches have been proposed in the recent applied literature to obtain empirical estimates of the learning gain. The main distinctive feature in these approaches refers to the joint estimation of a constant gain with other model's structural parameters; both Bayesian (Milani, 2007, 2008) and classical (Chevillon et al., 2010) estimation approaches have been successfully applied, though inference is strongly affected by weak identification issues.

More closely related to our approach (detailed below), are calibration procedures attempting to optimize the fit of learning rules in forecasting macroeconomic data and in matching forecasts from surveys. One earlier application of this approach is provided by Orphanides and Williams (2005) in a study of US monetary policy during the 1970s; Branch and Evans (2006) also provide gain calibrations to US data on inflation and output growth, which have been recently extended to other variables by Markiewicz and Pick (2014), and to other countries focusing on inflation by Molnár and Reppa (2010); Weber (2010) provides a similar application to European data; and Ormeño and Molnár (2015) show the usefulness of information contained in survey data for the estimation of structural models with adaptive learning.

1.2 Approach and main results

In order to shed some light on the gain calibration issue we develop an empirical framework that mimics a real-time learning-to-forecast process. We document renewed numerical calibrations of the gains for empirical applications with US quarterly data on inflation, output growth, and interest rates. One key feature in our analysis is our coverage of a broad range of model specifications to represent agents beliefs: we explore all possible combinations of the variables above in vector autoregressive (VAR) forecasting models with up to four lags. This is motivated in section §2 by the pervasive possibilities of misspecification introduced by the adaptive learning approach, particularly with respect to agents perceived laws of motion (PLM). Regarding the learning algorithm, our focus is on the least squares (LS) algorithm, which has received most of the attention in the literature.

Using such framework, detailed in section §3, we conduct several gain calibration exercises covering data over the period from 1981 to 2012; data from earlier periods are used to initialize the learning algorithm and the calibrations. We segment the calibrations according to different assumptions in the determination of the learning gains, particularly with respect to the measure used for their selection and their variation over time. Regarding the selection measure, we distinguish between two alternatives depending on the reference data: actual-based calibrations are selected by maximizing the accuracy of the forecasts; survey-based calibration, instead, are selected by maximizing the resemblance of the learning-based forecasts to those obtained from survey data.

These calibrations are documented in section §4, where we find that the actual-based approach leads to higher gain calibrations than the survey-based. We also find evidence of a great degree of heterogeneity in the gain calibrations, depending mainly on the variable forecasted and the lag length of the forecasting model. As a guide for applied research we reported averaged gain values along these dimensions. Inflation presented the highest gain calibrations, followed by output growth and interest rates, with maximum averaged values on VAR(1) models of about 0.11, 0.02, and 0.005, respectively. These calibrations then tended to decrease as the VAR lag order increased. Interestingly, the calibrations on the forecasting models of interest rates were found to be very small, suggesting that the real-time tracking provided by learning adds little predictive content to the estimates obtained in the pre-evaluation sample period.

In order to assess the different behavioral assumptions on the determination of the learning gains, in section §5 we conduct a comparative analysis on the forecasting performance of the gain calibrations. Our results point to mixed evidence, providing support to the choice of different gain calibrations, depending mainly on the variable forecasted and the model specification. But the effects of these different gain assumptions on the forecasting performance of the learning models are found to be of limited statistical relevance, particularly relative to the role played by the model specifications.

Finally, in section §6 we evaluate the effects of further dimensions of applied interest over the gain calibrations, such as: the use of revised instead of real-time data; alternative data definitions for the variables requiring the formation of expectations; and calibrations representative of the evolution of policymakers beliefs. In section §7 we conclude this paper with some remarks.

2 Adaptive Learning Context

Macroeconomic models imply relationships between aggregate variables such as inflation, output growth and interest rates, and often involve also past and future expected values of these same variables. In the latter case, one therefore has to specify how expectations are formed, that is, one has to specify a forecasting model (or PLM) for agents. Perhaps one of the greatest virtues of the RE paradigm is the containment of the set of such admissible models by the disciplinary assumption that expectations must be consistent with the implied actual outcomes. The introduction of adaptive learning, in contrast, opens the way for a plethora of (mis)specification possibilities, particularly in multivariate contexts. Which variables should enter agents PLMs? How does the calibration of the learning mechanism depend on the statistical properties of these variables?

Consider for example a reduced form model given by

$$y_t = \alpha + \beta E_{t-1}^* [y_t] + \delta x_{t-1} + u_t, \quad (1)$$

where y_t is an endogenous variable, x_t is a zero mean exogenous variable, and u_t is an unobservable white noise disturbance. With $\beta < 0$, this is a standard specification of Muth's (1961) partial equilibrium cobweb model of supply and demand, taking y_t as the market clearing price and x_t as a supply shock. With $\beta > 0$, it represents instead a Lucas (1973) aggregate supply model. Multivariate extensions and alternative timing assumptions are possible to fit more complex general equilibrium models. The minimum state variable (MSV) RE equilibrium associated to (1) can be obtained by assuming agents include a constant term and the lagged exogenous variable in their PLM, i.e.,

$$E_{t-1}^{RE} [y_t] = a + bx_{t-1}, \quad (2)$$

where RE dictates that $a = (1 - \beta)^{-1} \alpha$ and $b = (1 - \beta)^{-1} \delta$. Thus, the assumption of RE establishes an unique trajectory of y_t that satisfies (1) for given parameters and a sequence of stochastic shocks u_t .

Under adaptive learning, instead, one relaxes the rationality requirement and assumes that agents act like econometricians. There are two important features in the shift from the RE to the adaptive learning approach. First, under adaptive learning agents are not certain about

their PLM's parameter values, and hence follow a real-time process to obtain estimates relevant for their expectations formation at the same time that such expectations are determining actual outcomes. Second, bounded rationality features may restrict agents perceptions about the relevant variables to be included in their PLMs. The first feature is commonly accomplished by assuming agents update their PLM parameter estimates using a recursive LS algorithm of the form

$$\hat{\boldsymbol{\theta}}_t = \hat{\boldsymbol{\theta}}_{t-1} + \gamma_t \mathbf{R}_t^{-1} \mathbf{x}_t \left(y_t - \mathbf{x}_t' \hat{\boldsymbol{\theta}}_{t-1} \right), \quad (3)$$

$$\mathbf{R}_t = \mathbf{R}_{t-1} + \gamma_t (\mathbf{x}_t \mathbf{x}_t' - \mathbf{R}_{t-1}), \quad (4)$$

where γ_t is a learning gain parameter, \mathbf{R}_t stands for an estimate of regressors matrix of second moments, and \mathbf{x}_t and $\hat{\boldsymbol{\theta}}_t$ collect the variables included in the PLM and their associated parameter estimates, respectively; in our example model, under a PLM of the form of (2), $\mathbf{x}_t = \begin{pmatrix} 1 & x_{t-1} \end{pmatrix}'$ and $\hat{\boldsymbol{\theta}}_t = \begin{pmatrix} a_t & b_t \end{pmatrix}'$. Substituting this into the reduced form model one finds that the actual law of motion (ALM) of y_t under adaptive learning,

$$y_t = \alpha + \beta a_{t-1} + (\beta b_{t-1} + \delta) x_{t-1} + u_t, \quad (5)$$

now depends on the evolution of the PLM parameter estimates, which are jointly determined by the learning algorithm recursions for given initials, $\hat{\boldsymbol{\theta}}_0$ and \mathbf{R}_0 , and a sequence of learning gains, γ_t . For a sufficiently small constant gain, and provided stability conditions on the model parameters and the exogenous variable are satisfied, it is well known that the PLM parameter estimates will converge to a stationary distribution centered on the RE equilibrium values (Evans and Honkapohja, 2001). Since the statistical properties of this distribution are directly affected by the learning gain, it is important to have some guidance on plausible calibration values for applied purposes.

The convergence of adaptive learning towards the RE equilibrium is nevertheless conditional on the assumption of a correctly specified PLM. Consider for instance what would happen in our example if agents did not realize that y_t depends on the lagged values of x_t , assuming instead a naive PLM containing only a constant term. Clearly, such a PLM automatically rules out a RE solution when $\delta \neq 0$. Under learning, however, one can still find support for this PLM misspecification as a restricted perceptions (RP) equilibrium. Substituting the implied ALM,

$$y_t = \alpha + \beta a_{t-1} + \delta x_{t-1} + u_t, \quad (6)$$

into the learning algorithm, (3), one finds that

$$a_t = a_{t-1} + \gamma_t (\alpha + \beta a_{t-1} + \delta x_{t-1} + u_t - a_{t-1}) \quad (7)$$

is a stochastic recursive algorithm governing the evolution of agents expectations. For sufficiently small learning gains, the mean dynamics for a_t , represented by

$$\dot{a} = \lim_{t \rightarrow \infty} E(\alpha + (\beta - 1) a_{t-1} + \delta x_{t-1} + u_t) \quad (8)$$

converges to a distribution centered on $a = (1 - \beta)^{-1} \alpha$, thus defining a RP equilibrium.

A dimension of particular interest for the calibration of adaptive learning algorithms therefore relates to the choice of variables to include in agents information set. Such choice is particularly important in applied works, where it is not necessarily superior to assume that agents are aware of the correct specification for their PLM. For this reason, our calibration analysis will focus on the performance of several PLMs that may be of common use for different macroeconomic models.

3 Empirical Framework

3.1 Model and estimation

Consider an estimation context faced by a real-time agent wishing to obtain inferences about the law of motion of a set of variables of interest, say $\mathbf{y}_t = (y_{1,t}, \dots, y_{n,t})$. 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 believes that those variables are statistically related to their own lagged values through a linear regression of the form

$$\mathbf{y}_t = \mathbf{x}_t' \boldsymbol{\Theta}_t + \boldsymbol{\epsilon}_t, \quad (9)$$

where $\mathbf{x}_t = (1, \mathbf{y}_{t-1}, \dots, \mathbf{y}_{t-p})'$ is a vector collecting a constant term plus the np lagged values of the endogenous variables, $\boldsymbol{\Theta}_t = (\boldsymbol{\theta}_{0,t}, \boldsymbol{\theta}_{1,t}, \dots, \boldsymbol{\theta}_{np,t})'$ stands for a $(np + 1) \times n$ matrix of (possibly time-varying) coefficients associated to each regressor in \mathbf{x}_t , and $\boldsymbol{\epsilon}_t$ denotes a vector of uncorrelated white noise disturbances. For $n > 1$, (9) is commonly referred as an n -variables vector autoregression of order p , or n -VAR(p) for short. Both coefficients and disturbances are assumed not to be directly observable by the agent, hence a technique for the estimation of $\boldsymbol{\theta}_t$ is required to allow the construction of the inferences of interest.

As it turns out, the LS algorithm of (3)-(4) provides a recursive form of the equation-by-equation⁴ estimator obtained by the minimization of the sum of weighted error squares,

$$\hat{\boldsymbol{\theta}}_t = \arg \min \sum_{i=1}^t \beta(t, i) \left(y_i - \mathbf{x}_i' \hat{\boldsymbol{\theta}}_t \right)^2, \quad (10)$$

⁴For simplicity we abuse notation and let $\hat{\boldsymbol{\theta}}_t$ stand for column j in $\boldsymbol{\Theta}_t$, i.e., the vector of coefficients associated to the j th equation in the VAR system; similarly, y_t stands for $y_{j,t}$.

where the weights given to each observation are determined by the sequence of gains according to

$$\beta(t, i) = \begin{cases} \frac{\gamma_i}{\gamma^t} \prod_{k=i+1}^t (1 - \gamma_k) & \text{for } i < t, \\ 1 & \text{for } i = t. \end{cases} \quad (11)$$

Under a constant gain specification, $\beta(t, i) = (1 - \bar{\gamma})^{t-i}$, so that past observations are given geometrically decaying weights, whereas a decreasing gain leads to the famous OLS estimator of basic econometrics (see Berardi and Galimberti, 2013, for derivations). These properties may provide an explanation for the prominence of the LS algorithm in the adaptive learning literature as the choice to represent agents mechanism of adaptive learning; here we follow such practice and focus our calibration analysis on the LS case⁵.

3.2 Data

We consider the inclusion of three key macroeconomic variables in agents learning-to-forecast problem: inflation, output growth, and interest rate. Following our previous discussion, we evaluate all possible combinations of these variables in the specification of the learning PLMs; namely, one univariate $AR(p)$ specification for each variable, three bivariate 2- $VAR(p)$ specifications, and one trivariate 3- $VAR(p)$ specification. For further robustness we also vary the VAR lag orders from $p = 1, \dots, 4$, hence totaling 16 PLM specifications for each variable.

We use US quarterly data on real GNP/GDP and its price index to obtain measures of output growth and inflation, respectively, and the 3-month Treasury Bill rate as a measure of the interest rate. The sample covers the period from 1947q2 to 2013q4, and is organized in a real-time data format with vintages from 1966q1 to 2014q1; i.e., a total of 193 snapshots of what was known on the above variables by a market participant in real-time.

For the purpose of comparing the learning-based forecasts to those provided by survey respondents, we use data from the Survey of Professional Forecasters (SPF); each quarter, this survey asks professional economists to give their forecasts for several macroeconomic variables and also over different forecasting horizons. Here we use the median of the individual forecasts made for a total of five horizons, namely from t (nowcast) to $t + 4$. The SPF data is available from 1968q4 onwards for inflation and output growth, but only from 1981q3 for the interest rates. To close this gap we follow Milani (2011, p. 385) and derive expectations about 1-period-ahead interest rates according to the expectations theory of the term structure⁶.

⁵A less frequently adopted alternative is provided by the stochastic gradient (SG) algorithm, which, in spite of being computationally simpler than the LS, raises rather challenging issues in its calibration due to a sensitiveness to the data scales. See Berardi and Galimberti (2014) for a comparison between these algorithms, and Galimberti (2013) for details on the calibration of the SG algorithm.

⁶Further details about the data sources and adjustments are provided in Appendix A.1.

3.3 Further implementation details

In order to implement the LS recursions one needs to first specify how the initial estimates are obtained. Here we adopt a training sample-based method, where an initial portion of the available data is set aside to the application of the algorithm departing from diffuse initials, i.e., $\hat{\theta}_0 = \mathbf{0}_{np \times 1}$ and $\mathbf{R}_0 = \mathbf{0}_{np \times np}$. Berardi and Galimberti (2015) recently surveyed different methods adopted in the literature and reported evidence in favor of the convergence properties provided by the use of a training sample, in comparison with other more model-dependent initials such as equilibrium-related and estimation-based methods. Also in line with the previous applied literature, we use 75 observations for the training sample, which corresponds to the period from 1947q2 to 1965q4.

Estimation and forecasting are then carried out by vintage as follows: first, recursions for each gain calibration are computed departing from the vintage/gain initials until exhaustion of the vintage sample; then, the $t, \dots, t + 4$ forecasts for each vintage/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 2012q4, which results in a total of 188 forecasts for each combination of gain calibration, model specification, variable, and forecasting horizon. It remains to be specified how the gain calibrations are determined.

4 Gain Calibrations

In order to approach empirically the gain calibration issue, we adopt two distinct measures of gain selection. Under the first alternative, we assume gains are determined as a choice by agents aiming to optimize their learning performance. A natural measure to select gain calibrations under this assumption is the accuracy of their associated forecasts. We denote the gains selected according to this criterion as *actual-based* calibrations. As for the second selection measure, we take the learning gain as a primitive parameter of agents' learning-to-forecast behavior. An appropriate measure of fit under this alternative requires the observation of agents' actual expectations. For that purpose we use data of survey forecasts as a proxy, calling this gain selection method a *survey-based* calibration.

Different assumptions are possible regarding the use of real-time data, and about the specification of the range of gain values of interest. We consider two computational approaches to obtain these gain values. In the first, we begin by constructing a grid of admissible values, and then proceed by imposing selection rules that represent our different assumptions on the evolution 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 the algorithm.

4.1 Grid-based gains

We construct a grid of 100 gain values by setting an upper bound on admissible values so as to ensure the algorithm’s stability. Hence, our estimation routine is applied to each model specification with 100 different constant gain values. When it comes to represent agents learning-to-forecast behavior, nevertheless, an unique gain value is required for each time an iteration on the learning algorithm is performed. Under our two selection criteria, i.e., actual-based and survey-based, we pick the period specific gains by minimizing averages of forecast errors and forecast comparison errors, defined as the difference between the forecasts associated to each gain and the actual⁷ and survey observations, respectively. Moreover, since we have multiple forecasting horizons available from the surveys, our gain selections are also based on the average errors over these horizons. Letting z_{t+h} stand for the selection measure (either the actuals, y_{t+h} , or the surveys, s_{t+h}), the gain is obtained according to

$$\gamma_z = \arg \min_{0 < \gamma \leq 0.5} \left\{ T_{\mathcal{S}}^{-1} \sum_{t \in \mathcal{S}} H_t^{-1} \sum_{h \in \mathcal{H}_t} (z_{t+h} - \hat{y}_{t+h}(\gamma))^2 \right\}, \quad (12)$$

where $\hat{y}_{t+h}(\gamma)$ is the forecast of y_{t+h} associated with the model estimates obtained with the gain γ and using data in vintage t , h is the forecasting horizon (usually $\mathcal{H}_t = \{0, \dots, 4\}$ and $H_t = 5$, except when particular horizons are missing in the survey data), and \mathcal{S} is a selection sample of length $T_{\mathcal{S}}$. The definition of this selection sample depends on whether the gains are assumed to remain fixed throughout the sample, or are allowed to vary over time.

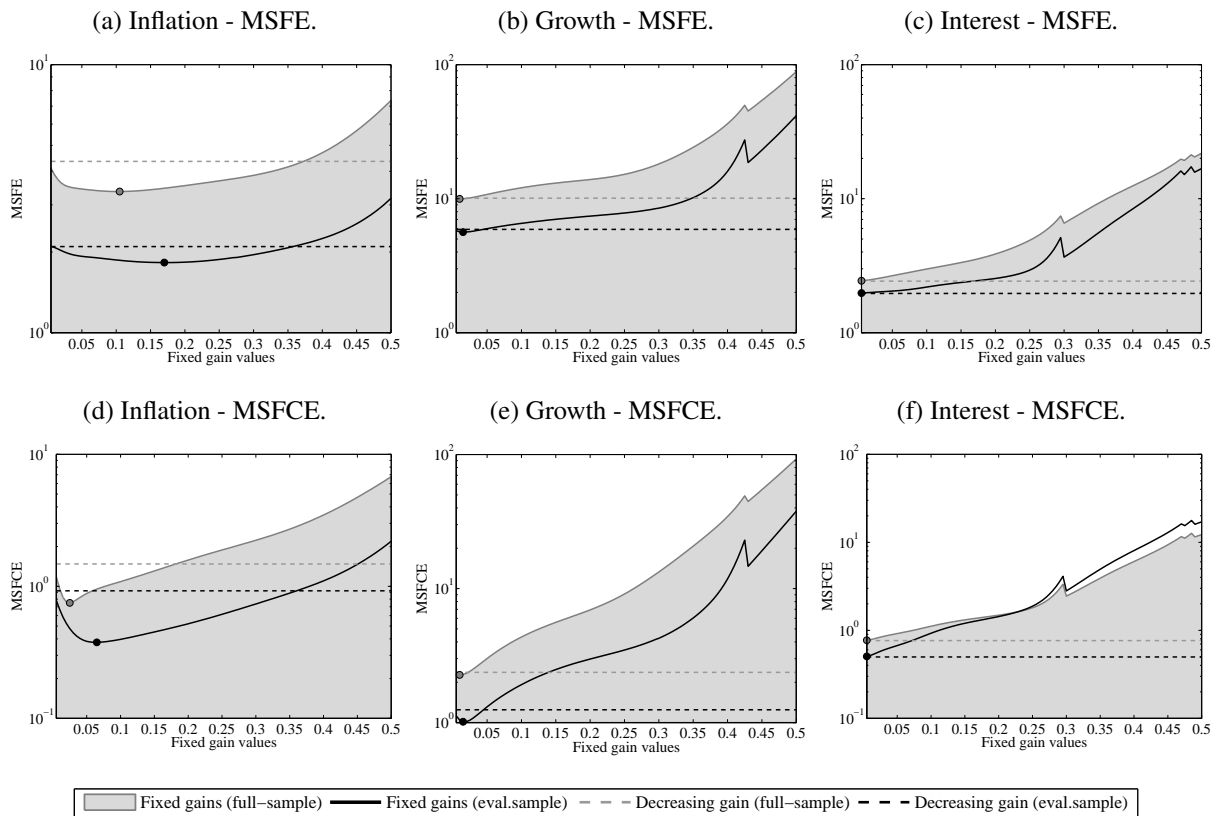
Fixed gains

Fixed gains are selected by minimizing the average forecast (comparison) errors over the full sample of forecasts that we compute, i.e., 188 forecasts covering the period from 1966q1 to 2012q4. Figure 1 illustrates this selection for the case of univariate specifications on each variable forecasted. The performance of the fixed gains are also compared to the traditional case of a decreasing gain sequence, which, overall, is found to provide inferior accuracy and resemblance to the surveys than the best fixed gains. Selections based on an alternative evaluation sample, from 1981q1 to 2012q4, are also presented; these suggest some potential misspecification in the use of a fixed gain for inflation, since the optimal gain for this variable depends on the selection sample.

Some “spikes” are observed in the forecasting performance curves of output growth and interest rates. These are caused by the activation of a projection facility (PF), which is a mechanism assumed to be coupled to the learning algorithms such that whenever the estimates leave a bounded region in the parameters space the device is activated in order to contain the escape. Naturally, as the gain increases there is more instability and the estimates get closer to the PF

⁷Following the real-time spirit of learning, the actual-based forecast errors are computed in relation to the first data measures available for each variable.

Figure 1: Forecasting performance curves.



The forecasts are based on AR(1) univariate specifications for each variable. Performance is measured by the sample average of the squared forecast (comparison) errors averaged over the forecasting horizons. The bullet points indicate the fixed gains that minimize the performance criterion over the corresponding samples (full-sample: 1966q1-2012q4; evaluation sample: 1981q1-2012q4). The major “spikes” in the forecasting performance curves of output growth and interest rates are caused by activations of the projection facility at particular periods (2006q1 for the growth spikes; 1992q1-q2 for the interest rates) for high gain values, as explained in the text.

activation threshold⁸. Before reaching that threshold, nevertheless, the forecasts become more “explosive”, which deteriorates the performance.

Clearly, under the actual-based calibration, the selection using the full-sample of forecasts violates the restrictions of a “fair” out-of-sample forecasting exercise, since it uses information not available for a real-time learner; e.g., an agent forecasting future inflation rates in real-time from 1990q1 does not know how each of the available gain values will perform over 1990q1, 1990q2, and so on. Branch and Evans (2006) circumvent this issue by defining the calibration within a pre-sample. In contrast, the use of the full-sample of forecast comparison errors under the survey-based gain calibration does not present any conflict, as 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).

Recursive gains

Another possibility is to allow the selection of the gain to be recursive, through a minimization of the average forecast (comparison) errors over a rolling window of forecasts. Here we adopt a window length of 60 observations, the first of which covering the period from 1966q1 to 1980q4⁹. Hence, the recursive gains are obtained through a period-by-period application of (12) with $\mathcal{S}_t = \{t - 1, \dots, t - 60\}$.

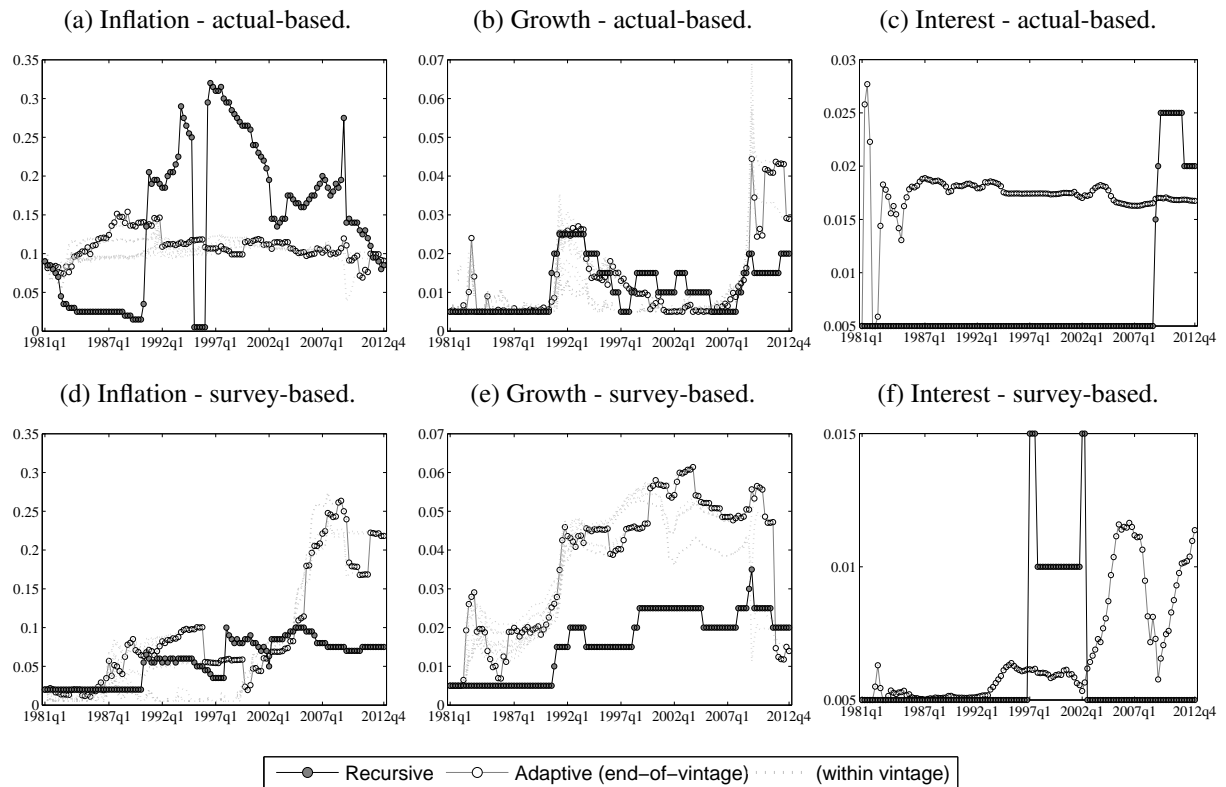
Figure 2 presents the evolution of these gain calibrations for the illustrative case of univariate forecasting model specifications for each variable (the adaptive gains will be detailed below). Clearly, the flexibility provided by the recursive optimization of the gain selection leads to the emergence of interesting patterns; particularly, the gains for inflation and output growth tended to increase over the 1990s, a period characterized by reduced volatility in these variables compared to the preceding decades (during the 1970s and early 80s the US went through the Great Inflation period). Also notice there are more differences between the actual-based and the survey-based recursive calibrations than observed under the fixed gain assumption.

The sequence of gains selected using this recursive approach provides a more realistic approximation to the informational environment faced by an agent learning and forecasting in real-time. Moreover, there is no reason to restrict the gains to be fixed throughout the whole forecasting sample, particularly considering the evidence of misspecification for the case of inflation reported above. Notice, however, that the gain is still kept fixed within each vintage of data used in the calculation of past forecasting performance. E.g., the recursive gain selected to estimate the forecasting model from 1990q1 vintage of data is obtained by minimizing the

⁸The projection facility is activated only if the (9)’s companion matrix eigenvalue with highest modulus gets bigger than a critical threshold. Whereas a value of one is the standard requirement for stationarity (see Lütkepohl, 2005, pp. 13-18), we adopt a critical value of 2 so as to allow the algorithm some space to return to the stability region on its own.

⁹Recall we take the first 75 observations (1947q2-1965q4) of our sample for the initialization of the algorithm.

Figure 2: Evolution of time-varying gain calibrations.



The gains are selected in real-time under AR(1) univariate specifications for each variable. The recursive gains are selected on the basis of the forecasting performance (actual or survey-based) of fixed gains over a window of past 60 observations. Adaptive gains are time-varying within each vintage, but with gain values calculated by an outer algorithm instead of a selection based on a grid; the adaptation constants are 0.001, for inflation, 0.0001, for growth, and 0.0005, for interest.

average forecast (comparison) errors based on the end-of-vintage-sample forecasts of 60 different vintages from 1975q1 to 1989q4. The next approach relaxes this assumption by allowing within vintage time-variation of the gain.

4.2 Adaptive gains

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 complement the learning algorithm with an outer mechanism to adjust the gain in response to 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 based on a PLM including only a constant.

The derivation of the adaptive gain mechanism starts with the specification of an objective function, which in our context is related to the loss functions we associate to the two alternative interpretations given to the learning gains. Hence, the adaptive gain approach is also split in two, with the actual-based gains determined by a mechanism responding to the algorithm's accuracy, while the survey-based gains are driven by the algorithm's resemblance to the survey forecasts. Assuming a squared loss function, the gain is selected by minimizing the expected loss of forecast (comparison) errors

$$J_t = \frac{1}{2} E \left[\left(z_t - \mathbf{x}'_t \hat{\boldsymbol{\theta}}_{t-1} \right)^2 \right]. \quad (13)$$

If the true data generating process of z_t 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 this gradient and plugging in its stochastic approximation we obtain the adaptive gains recursions¹⁰,

$$\gamma_t = \left[\gamma_{t-1} + \alpha \mathbf{x}'_t \hat{\boldsymbol{\Psi}}_{t-1} \left(z_t - \mathbf{x}'_t \hat{\boldsymbol{\theta}}_{t-1} \right) \right]_{\bar{\gamma}_{\min}}^{\bar{\gamma}_{\max}}, \quad (14)$$

$$\begin{aligned} \hat{\boldsymbol{\Psi}}_t &= (\mathbf{I} - \gamma_t \mathbf{R}_t^{-1} \mathbf{x}_t \mathbf{x}'_t) \hat{\boldsymbol{\Psi}}_{t-1} \dots \\ &\quad + (\mathbf{I} - \gamma_t \mathbf{R}_t^{-1} \hat{\mathbf{S}}_t) \mathbf{R}_t^{-1} \mathbf{x}_t \left(y_t - \mathbf{x}'_t \hat{\boldsymbol{\theta}}_{t-1} \right), \end{aligned} \quad (15)$$

$$\hat{\mathbf{S}}_t = (1 - \gamma_t) \hat{\mathbf{S}}_{t-1} + \mathbf{x}_t \mathbf{x}'_t - \mathbf{R}_{t-1}, \quad (16)$$

¹⁰See appendix A.2 for these derivations. It is important to note that we assume that the derivative of y_t with respect to the gain is null, which neglects the effect of the gain, through the update of expectations, over the determination of y_t . Relaxing this assumption goes beyond the scope of this paper.

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

An intuition for this gain adaptation mechanism follows directly from its interpretation as a numerical optimization method, browsing along the error-performance curve in search for an optimal gain. Here $\hat{\Psi}_t$ keeps track of the algorithm's past estimation performance, accumulating its past gradients discounted according to forgetting factors implied by the evolution of the gains. This synthetic measure of past performance is taken as a reference in the gain update equation, (14): when the latest gradient points towards the same (a different) direction as of $\hat{\Psi}_t$, 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.

Nevertheless, the above interpretation becomes less compelling under the survey-based gains, 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 estimate is drawn from the resemblance to surveys objective, the past gradients synthesized in $\hat{\Psi}_t$ 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 adaptive gains still requires specification of the adaptation constant, α . According to the analysis of Kushner and Yang (1995), stability requires this parameter to be small, so as to satisfy $0 < \alpha \leq \bar{\gamma}_{\min}$, together with an appropriate upper bound to the gains. For our purposes we take the extreme gain values in the grids defined above 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 trades off performance optimization in favor of a calibration presenting a sequence of gains not too jumpy, but neither constant¹¹.

The resulting sequences of adaptive gains are presented in Figure 2. Comparison with the previous approach shows that the adaptive gains tend to move in tandem with the recursive gains, though we also observe some extra time variation within each vintage in the adaptive gains approach.

¹¹We also found that as the number of estimated parameters increased, the adaptive gains turned more unstable for a fixed adaptation constant. To deal with that we set $\alpha = \bar{\alpha}/(np + 1)$, where $np + 1$ is the number of parameters in the model specification.

Table 1: Averaged gain calibrations.

Calibration	Actual-based calibrations			Survey-based calibrations		
	Inflation	Growth	Interest	Inflation	Growth	Interest
(a) Fixed gains:						
- VAR(1)	10.8	1.8	0.5	1.6	1.4	0.5
- VAR(2)	2.4	1.3	0.5	0.9	1.1	0.5
- VAR(3)	2.3	0.9	0.5	0.8	0.9	0.5
- VAR(4)	3.0	1.3	0.5	1.4	1.1	0.5
- Overall	4.6	1.3	0.5	1.2	1.1	0.5
(b) Recursive gains:						
- VAR(1)	12.8	1.7	0.8	4.0	2.5	0.7
- VAR(2)	5.0	1.3	1.2	2.0	1.7	1.0
- VAR(3)	3.1	1.4	1.6	1.0	1.7	1.0
- VAR(4)	2.9	1.5	1.0	1.3	1.7	0.8
- Overall	6.0	1.5	1.1	2.1	1.9	0.9
(c) Adaptive gains:						
- VAR(1)	7.6	1.2	1.4	4.4	2.3	0.8
- VAR(2)	3.7	1.3	0.8	3.8	1.7	1.3
- VAR(3)	5.1	1.1	1.7	5.1	1.3	0.8
- VAR(4)	3.4	1.1	0.9	4.8	1.4	1.8
- Overall	4.9	1.2	1.2	4.5	1.7	1.1

The presented gain values are scaled by 100, e.g., 10.8 is equivalent to $\gamma = 0.108$. Averages are based on fixed and averaged time-varying gains selected under different model specifications (univariate, bivariate, and trivariate).

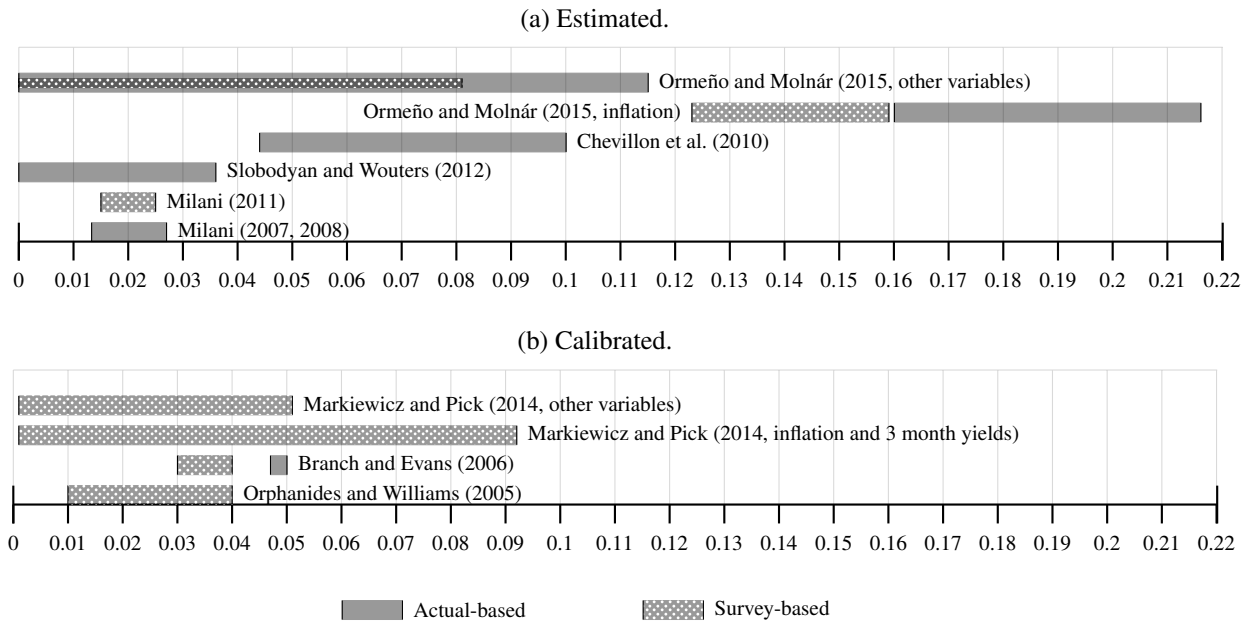
4.3 Summary and related literature

Other than the variables forecasted and the selection measures, we obtain gain calibrations over two additional dimensions related to the model specification: the explanatory variables and the number of lags included. Averaged gain calibrations over the first of these dimensions are reported in Table 1, showing non-negligible heterogeneity depending on all remaining dimensions. We draw several interesting observations from these results:

1. The gains selected to forecast inflation are overall higher than those used on output growth and the interest rate.
2. The gains tend to decrease as the number of lags included in the forecasting model specification increases.
3. Overall, the actual-based gains are higher than the survey-based ones.
4. The averaged time-varying gains also tend to be higher than the averaged fixed gains.

Most of these findings are consistent with the gain values adopted in the previous literature, as summarized in Figure 3. There are two main approaches to the determination of the learning

Figure 3: Gain values in the literature.



All gains are based on US data, but for diverse samples and model specifications. The intervals reported for estimated gains refer to 90% credible (Bayesian estimates) and confidence (GMM estimates) intervals, except for Milani (2007, 2008, 2011), who reports 95% credible intervals, and Ormeño and Molnár (2015), who report median and standard deviations of the estimates (intervals are ± 2 st.devs. around medians). Most studies adopt a common gain for the different forward variables requiring expectations; those that estimate/calibrate different gains for each variable are indicated.

gains in applied works: first, the gain can be jointly estimated with other parameters within a structural model; second, the gain can be calibrated according to the forecasting performance. Both instances allow the use of actual observations or survey forecasts as a reference, the latter often yielding lower gain values. Moreover, the use of different gains for each variable has been increasingly popular¹².

Given all this heterogeneity, an important question is which of these calibrations should be favored in applied research. We now turn to this issue.

5 Forecasting Performance

We compare the performance of the different gain calibration assumptions under two evaluation criteria defined according to the measure used for the gain selection. Namely, we evaluate the performance of the actual-based calibrations in terms of their forecasting accuracy, measured by mean squared forecast errors (MSFE), whereas for the survey-based calibrations we focus

¹²Calibrations for other countries also present a great diversity of gain values (all on inflation): for European economies, Weber (2010) documents actual-based gains between $[0.067, 0.3]$, experts-survey-based gains between $[0.013, 0.178]$, and households-survey-based gains between $[0.001, 0.093]$; Molnár and Reppa (2010) find survey-based gains between $[0, 0.08]$, for developed economies, and between $[0.04, 0.26]$ for Eastern European and Latin American economies.

on the forecasts resemblance to the surveys, measured by the mean squared forecast comparison errors (MSFCE). The motivation for these criteria naturally arise from the purpose of each calibration. 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. In the second criterion 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 modeled.

There is some overlap between these evaluation criteria and the measures we use for the gain selections, particularly for the case of the fixed gains. To circumvent this issue we attempt to reduce the informational advantage provided to the fixed gain calibration, relative to the backward-looking time-varying calibrations, by restricting our evaluation sample. Namely, our evaluation covers the period from 1981q1 to 2012q4, which also aligns with the forecasts obtained under the recursive calibration after the initial window of 60 observations (it also corresponds to the samples presented in Figure 1). Hence, because the calibrations are based on the gains performance in a past sample, they are not necessarily the optimal gains in the evaluation sample. Under this design, our evaluation can uncover evidence of model misspecification and time-varying features in the underlying data generating process.

5.1 Overview

Figure 4 compares the forecasting performance of the gain calibrations over the several dimensions we are considering. Starting with the evidence in forecasting inflation, panel (a) shows that the univariate model often outperforms the other specifications including additional variables; the optimal lag order also depends on the model specification. Finally, the fixed gain values reported in the middle panel confirm our earlier observation that the selections are mainly determined by the lag orders and the selection measure.

The results on output growth, panel (b) in Figure 4, reveal a slightly different pattern regarding the optimal lag order, overall favoring the use of only one lag. The univariate model specification is again favored, particularly in terms of resemblance to the surveys; in terms of forecasting accuracy, the trivariate model specification with only one lag is the best, but the performance of this model deteriorates substantially as the lag order increases.

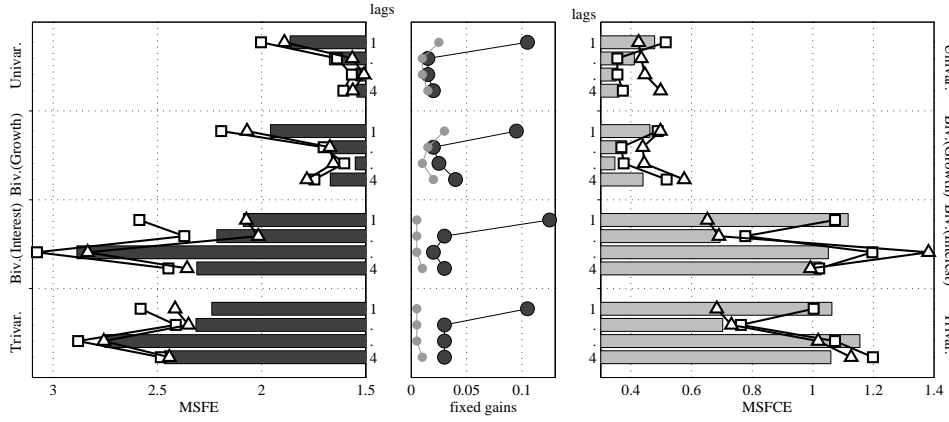
Perhaps more surprising are the results on forecasting the interest rate, presented in panel (c) of Figure 4, where learning seems to be less relevant: the selected gain calibrations equal the smallest value available in our grid, suggesting little advantages in the use of recursively estimated models to forecast this variable that is more directly affected by central bank policies.

5.2 Comparative exercises

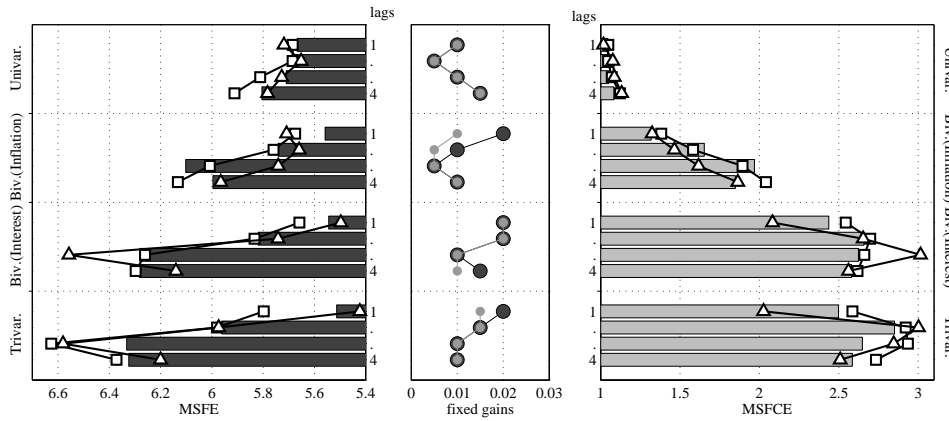
Our results above provide clear evidence that the relative performance of the gain calibrations is sensitive to dimensions involving the PLM model specification. In order to obtain a more

Figure 4: Forecasting performance of gain calibrations.

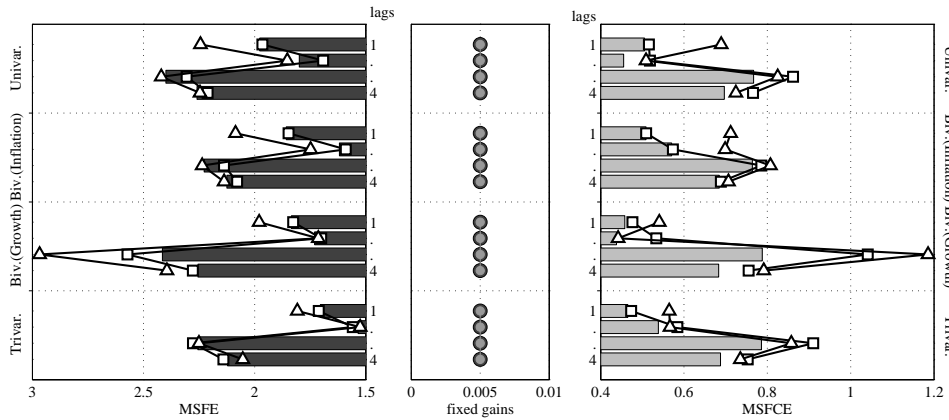
(a) Inflation.



(b) Growth.



(c) Interest.



Fixed actual-based
 Fixed survey-based
 Recursive
 Adaptive

The fixed gains are based on the full-sample, 1966q1-2012q4, whereas the recursive gains are selected over a moving window of past 60 observations. The adaptive gains are based on an outer algorithm with adaptation constants adjusted according to the variable and model specification (see footnote 11). The reported performance measures are then calculated over the evaluation sample, 1981q1-2012q4.

general picture, we now go beyond those dimensions and synthesize our comparisons in the form of *hit rate* measures. The hit rate stands for the frequency by which the forecasts associated to a given calibration is found to outperform those associated to a competitor with respect to one of our evaluation criteria. Here we conduct paired comparisons between 80 sequences of forecasts associated to each gain calibration, covering combinations of 4 model specifications per variable, 4 lag orders per specification, and 5 forecasting horizons per combination of the former two dimensions. The results are presented in Table 2.

The evidence is quite mixed, depending on the variable forecasted and the evaluation criterion for most of the comparisons. Whereas the actual-based fixed gains tend to outperform the recursive ones in forecasting inflation and output growth, the comparison is leveled against the adaptive gains. To forecast interest rates, the recursive gains seem to be the best choice. In terms of resemblance to the surveys, there is little support in favor of the adaptive gains. The comparison between the survey-based fixed gains and the recursive ones is balanced on inflation and output growth, but dominated by the former in forecasting interest rates.

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)¹³. 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, the second clearly represents a more appropriate test for our purposes of comparison of different estimation methods¹⁴.

The results on these statistical tests are again synthesized in hit rate measures, and are presented in Table 2. For most cases the tests for equal predictive ability suggest there is little difference between the forecasts associated to the alternative gain calibrations; the frequency of statistically significant rejections of the null of equal predictive ability observed with the GW tests, for example, average out to about only 5.6% of the comparisons over all variables and evaluation criteria.

5.3 Discussion

The results presented in this section provided mixed evidence to support the choice of a gain calibration. There was a great degree of heterogeneity on the forecasting performance of the gain calibrations, particularly with respect to the variable forecasted and the model specification. Importantly, most of the comparisons were not statistically conclusive, meaning that the predictive ability of the different gain calibrations was hard to distinguish.

¹³We summarize the calculations involved in each of these tests in Appendix A.3.

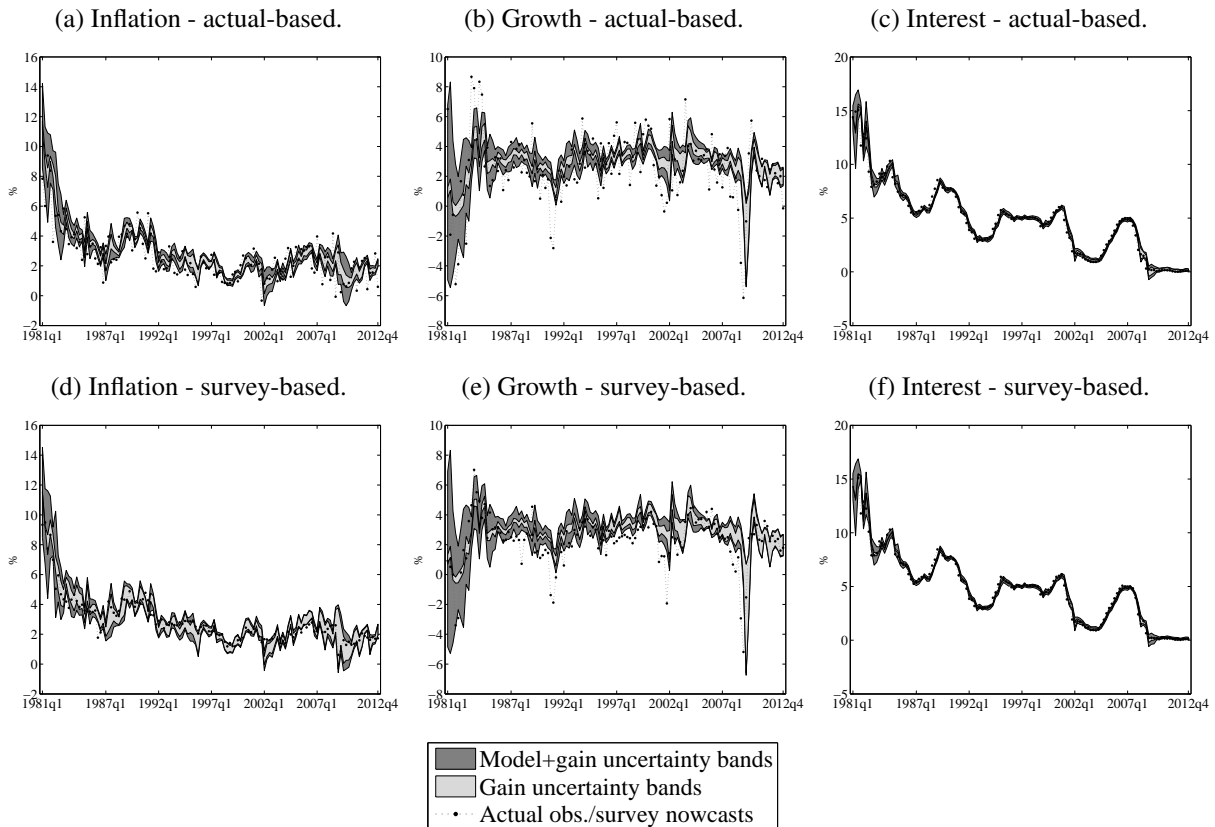
¹⁴Although the GW test is not suitable for the comparison of recursively estimated model-based forecasts, our focus on constant gain specifications attach geometrically decaying weights to past observations, hence approximating a rolling window estimation scheme.

Table 2: Hit rates on comparisons between gain calibrations.

(a) Actual-based calibrations - MSFE criterion.						
Statistic	Inflation		Growth		Interest	
	A wins	B wins	A wins	B wins	A wins	B wins
- Fixed (A) x Recursive (B):						
Hit-rates	85.0%	15.0%	70.0%	30.0%	38.8%	61.2%
DM 10% signif.	(15.0%)	(3.8%)	(2.5%)	(1.3%)	(0.0%)	(13.8%)
GW 10% signif.	[2.5%]	[1.3%]	[1.3%]	[1.3%]	[0.0%]	[8.8%]
- Fixed (A) x Adaptive (B):						
Hit-rates	58.8%	41.2%	42.5%	57.5%	71.2%	28.8%
DM 10% signif.	(10.0%)	(6.3%)	(5.0%)	(3.8%)	(5.0%)	(1.3%)
GW 10% signif.	[7.5%]	[7.5%]	[0.0%]	[3.8%]	[3.8%]	[2.5%]
- Recursive (A) x Adaptive (B):						
Hit-rates	26.2%	73.8%	26.2%	73.8%	81.2%	18.8%
DM 10% signif.	(1.3%)	(10.0%)	(2.5%)	(3.8%)	(10.0%)	(0.0%)
GW 10% signif.	[2.5%]	[8.8%]	[2.5%]	[13.8%]	[2.5%]	[0.0%]
(b) Survey-based calibrations - MSFCE criterion.						
Statistic	Inflation		Growth		Interest	
	A wins	B wins	A wins	B wins	A wins	B wins
- Fixed (A) x Recursive (B):						
Hit-rates	56.2%	43.8%	48.8%	51.2%	68.6%	31.4%
DM 10% signif.	(25.0%)	(13.8%)	(0.0%)	(6.3%)	(0.0%)	(4.3%)
GW 10% signif.	[6.3%]	[10.0%]	[5.0%]	[2.5%]	[0.0%]	[0.0%]
- Fixed (A) x Adaptive (B):						
Hit-rates	78.8%	21.2%	65.0%	35.0%	67.5%	32.5%
DM 10% signif.	(36.3%)	(11.3%)	(20.0%)	(3.8%)	(33.8%)	(6.3%)
GW 10% signif.	[18.8%]	[6.3%]	[10.0%]	[1.3%]	[1.3%]	[3.8%]
- Recursive (A) x Adaptive (B):						
Hit-rates	78.8%	21.2%	75.0%	25.0%	62.5%	37.5%
DM 10% signif.	(33.8%)	(10.0%)	(42.5%)	(1.3%)	(25.0%)	(7.5%)
GW 10% signif.	[17.5%]	[2.5%]	[32.5%]	[12.5%]	[1.3%]	[1.3%]

The hit rates sum up the paired comparisons of forecasting performance between the gain calibrations over the model specifications (4 per variable), the (V)AR lag orders (4 per specification), and the forecasting horizons (5 per spec./lag combination). The DM and GW hit rates refer to the frequency with which the hypothesis of equal predictive ability between the forecasts being compared was rejected at a 10% level of significance.

Figure 5: Evolution of nowcasts variability between models and calibrations.



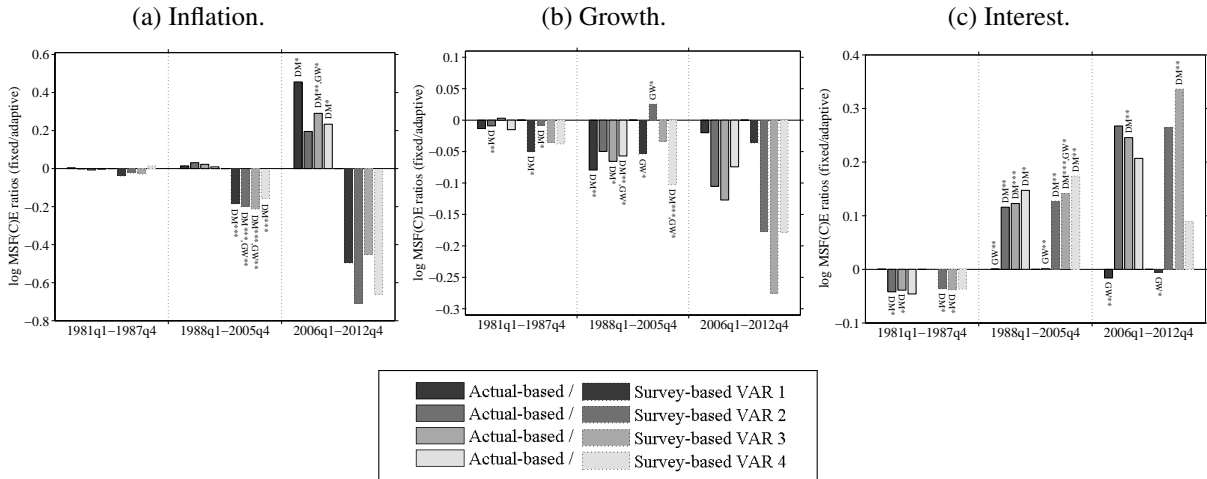
The uncertainty bands are ± 2 standard deviations around each period's nowcast averaged over the model specifications and the gain calibrations. The uncertainty due only to the gain calibrations are obtained by averaging the standard deviations between the calibrations over the model specifications. The actual observations refer to the first measures observed in real-time.

One implication of these results is that changes in the learning gains, within the reasonable limits we have established, have little effects over the forecasting performance of recursively estimated models involving macroeconomic variables. In contrast, our analysis highlights a greater impact of the model specification on such forecasting performance, both regarding the choice of variables and of number of lags to include in the forecasting models.

To corroborate this idea, in Figure 5 we draw uncertainty bands around the averaged forecasts obtained with the different calibrations and model specifications (see the Figure's footnote for details about their construction). Clearly, the variability of forecasts over gain calibrations is in general smaller than that observed across model specifications. The fact that the gain uncertainty bands are thin, compared for instance to the variability in the forecasting targets, helps explain the lack of power presented by the statistical tests in distinguishing between the predictive ability associated to the gain calibrations.

These results should not, nevertheless, diminish the relevance of properly designed gain calibration mechanisms. Particularly, the stability of the learning recursions is highly sensitive to the gains, hence the importance of setting appropriate limits to their values. Besides, we have found cases favoring the use of the time-varying calibrations that we have proposed in this paper

Figure 6: Comparison of nowcasting performance over sub-samples.



The nowcasts are based on univariate specifications for each variable. Actual-based (survey-based) calibrations are compared in terms of logs of MSF(C)E ratios between fixed and adaptive gains: positive values indicate the adaptive gains outperformed the fixed gain, and vice versa. Statistical rejections of the hypothesis of equal predictive ability using both the DM and GW tests are indicated by one, two, or three stars, for three corresponding significance levels: 1%, 5%, and 10%.

instead of the standard approach of a fixed gain. In fact, Figure 6 presents some evidence that the relevance of time-varying gains also depends on the overall state of the economy. We split our evaluation sample in three sub-samples so as to cover three distinct periods in the evolution of the US economy: the final years of the Great Inflation, the Great Moderation period, and the recent financial crisis, respectively.

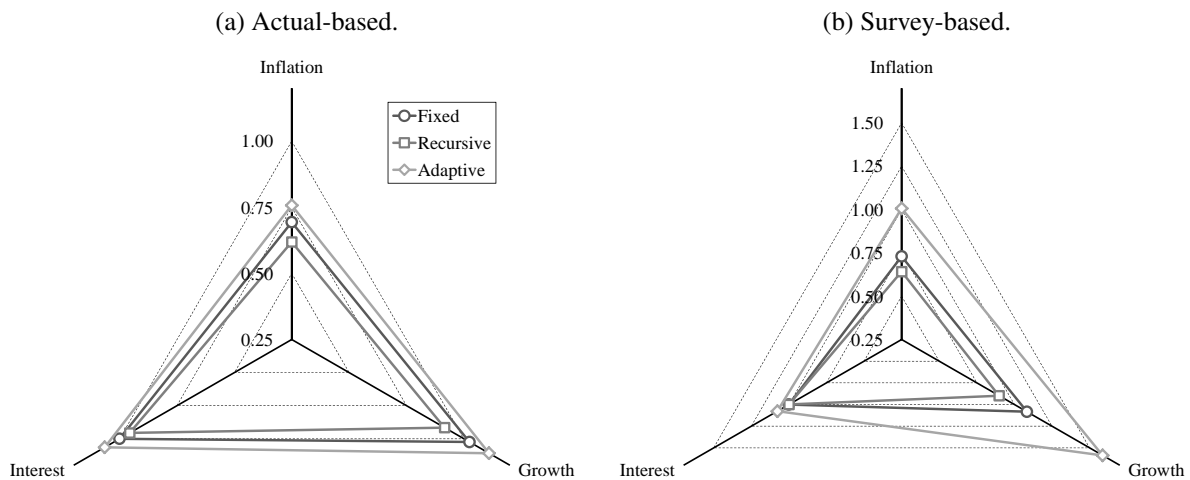
The results in Figure 6 provide some interesting observations. First, for inflation, the difference between the fixed and the adaptive gains performances gets more pronounced in the period surrounding the recent financial crisis; whereas the adaptive gains provided the best response to the instabilities of this period in terms of forecasting accuracy, the fixed gains provided better approximations to the survey forecasts. For growth, the results are rather favorable to the use of a fixed gain over the three sub-samples, and both in terms of forecasting accuracy and resemblance to the surveys. Finally, for interest the adaptive gains have an advantage on both criteria over the last two sub-samples. Altogether, we interpret these results as evidence of the relevance of the gain calibrations for the performance of learning during transitional periods.

6 Further Calibration Dimensions

6.1 Using revised data

Although the availability and use of real-time data has been increasingly widespread in the applied literature, the traditional approach to the calibration or estimation of macroeconomic models involves the use of only one vintage of data. Whereas we believe the focus on revised

Figure 7: Averaged gain values using revised data relative to real-time calibrations.



The plotted values represent ratios between the averaged gain calibrations obtained using revised data and those obtained in real-time. Values below 1 indicate cases where the calibrations for revised data are, on average, smaller than the real-time ones, and vice versa. The fixed gains are averaged over the model specifications and the lag orders, whereas the time-varying specifications are also averaged over the time dimension.

data restricts the analysis of learning, by violating the informational restrictions faced by agents when forming their expectations in real-time, it is important to assess its effects over the gain calibrations we have hitherto obtained using real-time data.

Figure 7 compares the averaged gain values obtained using revised data, from the 2014q1 vintage, to those obtained in real-time¹⁵. The effects of using revised data are similar across the gain calibrations, depending mainly on the variable forecasted; inflation appears to be the most affected variable, with revised calibrations about 25% smaller than those obtained in real-time; the survey-based adaptive gains also showed some sensitivity to the timing of the data, particularly on growth, where the average gain was more than 50% higher than that obtained in real-time.

It is important to mention that the use of survey forecasts jointly with revised data to calibrate the gains represents a critical violation to the timing of information implied by a realistic learning process. Namely, the forecasts based on learning models estimated with revised data incorporate information that was not available at the time the survey panelists submit their forecasts. Thus, we recommend that the survey-based calibration should be avoided in applications involving revised data.

6.2 Alternative variables definitions

Most macroeconomic models are based on relationships for the determination of the conceptual measures of aggregate economic activity that we have explored so far. For applied purposes,

¹⁵The averaged gain values calibrated to revised data are presented in Appendix A.4.

Table 3: Fixed gain calibrations for alternative variables.

Lag order	Actual-based calibrations			Survey-based calibrations		
	CPI inflation	Unemployment rate	10-year T-bond	CPI inflation	Unemployment rate	10-year T-bond
- AR(1)	3.0	0.5	0.5	4.0	0.5	0.5
- AR(2)	2.5	1.0	0.5	4.0	1.5	0.5
- AR(3)	1.5	1.0	0.5	0.5	1.0	0.5
- AR(4)	3.0	1.0	0.5	2.5	1.0	0.5
Average	2.5	0.875	0.5	2.75	1.0	0.5

The presented gain values are based on univariate specifications for each variable, and are scaled by 100.

nevertheless, there is no unanimous definition of how these concepts should be measured, hence the importance of evaluating how alternative data definitions may affect the calibrations of the learning gains.

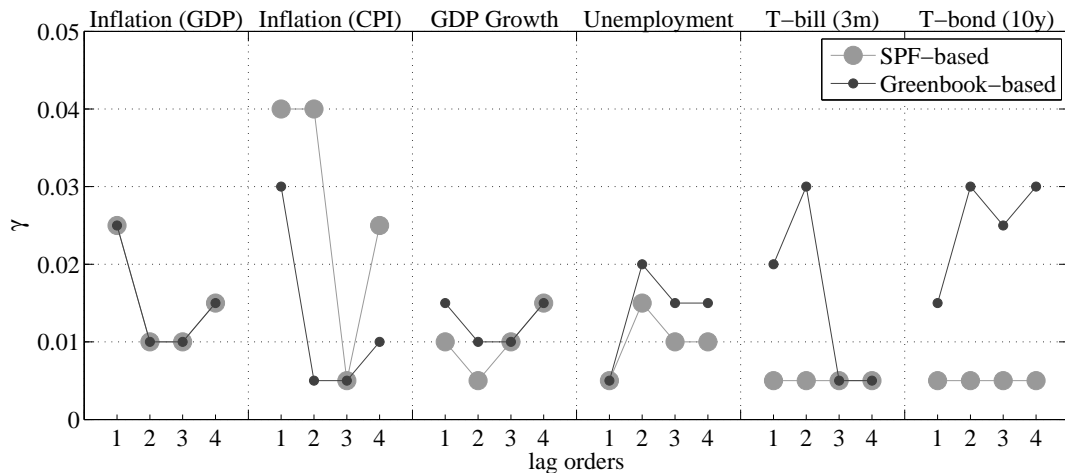
In Table 3 we report fixed gain calibrations obtained with alternative data measures that are also commonly employed in the literature, namely, CPI inflation, the unemployment rate, and the 10-year Treasury Bond rate. These are based on univariate specifications of the forecasting model for each variable, hence they are not directly comparable to the averages presented in Table 1, but see the middle panels of Figure 4. Overall, we observe that only the calibrations for inflation are affected by the use of the CPI instead of the GNP/GDP price index; whereas the actual-based gains for this variable decrease, the survey-based gains increase, attenuating the differences between these two calibration approaches.

6.3 Matching policymakers beliefs

Finally, another dimension of interest for researchers applying adaptive learning to policy issues relates to calibrations aimed to represent policymakers beliefs (see, e.g., Carboni and Ellison, 2009). For that purpose we re-evaluate our survey-based fixed gain calibrations using Greenbook forecasts. The Greenbooks are reports produced by the research staff of the Federal Reserve Board of Governors before each meeting of the Federal Open Market Committee. They contain projections and assumptions about the future evolution of many economic variables, including those we have analyzed here.

Figure 8 presents the resulting fixed gain calibrations matching the Greenbook forecasts. For comparative purposes, we also present our earlier survey-based calibrations using the SPF. Interestingly, there are only small deviations from the calibrations obtained to represent the professional forecasters. Except for inflation, the Greenbook-based gains tend to be slightly higher than those based on the SPF; this result suggests that the policymakers adjust their beliefs more quickly than the professionals to new observations on output-related and interest-related variables, but more slowly on inflation.

Figure 8: Policymakers fixed gain calibrations.



The presented gain values are based on univariate specifications for each variable and are selected over different sample periods, depending on the availability of the corresponding survey data.

7 Concluding Remarks

In this paper we have studied empirically the issue of how to calibrate adaptive learning mechanisms to mimic a real-time process of expectations formation in macroeconomic models. We have produced and documented renewed numerical calibrations of the learning gains for applications with US data on inflation, output growth, unemployment and interest rates. We have considered both actuals and survey forecasts data on these variables as reference for our calibrations. We have also explored other dimensions of applied interest, including several combinations of the above variables in the model specifications, and different assumptions regarding the determination of the learning gains.

We have found evidence of a great degree of heterogeneity in the gain calibrations, depending mainly on the variable forecasted and the model specification. The measure used as reference for the selection of the gains also mattered: survey-based calibrations were found to yield lower gain values than those based on actuals. In terms of forecasting performance, the effects of different assumptions regarding the time-variation of the learning gains were found to be of limited statistical relevance. In contrast, the model specification is found to play a greater role in the determination the learning models predictive ability.

Our main conclusion from these results is that an appropriate approach for the determination of the learning gains should be defined in accordance with each model's assumptions: whereas actual-based calibrations should be favored in models aiming to represent rationally optimizing agents, the survey-based calibration should be preferred on bounded rationality grounds; similarly, models allowing for abrupt structural breaks in the evolution of the economic system should consider the possibility of time-varying gains. We hope to have provided some guidance on the design of such mechanisms too.

A Appendix

A.1 Details on data

Our data on the US real GNP/GDP and its price index, unemployment, and CPI come from the Philadelphia Fed Real-Time Data Research Center. We obtain output growth and inflation rates from the data on levels computing their associated annual growth rates by compounding simple quarterly growth factors. For CPI we obtain annualized quarter-over-quarter inflation rates based on the quarterly-average price index level, using always the data from the first vintage month of each quarter, so as to make this variable consistent with the SPF forecasts. We obtain the quarterly averaged 3-month Treasury Bill rates and the 10-year Treasury Bond rates from the FRED database of the St. Louis Fed. The SPF data is also provided by the Philadelphia Fed. The following adjustments are necessary:

Short vintage time series history: some vintages lack earlier observations due to delays in BEA revisions (see Philadelphia's Fed documentations). This was the case for 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 $(\text{nominal GDP} / \text{GDP prices}) * 100$.

T-bill expectations data before 1981: the SPF only started asking for the Treasury Bill forecasts from 1981q3, whereas the GDP-related forecasts are available since 1968q4. To extend the sample of interest rate expectations, we exploit data on different maturities of the Treasury bill to obtain estimates of expected future interest rates as

implied by the expectations theory of the term structure, according to which

$$i_t^{6m} = \left(\frac{i_t^{3m} + E_t i_{t+1}^{3m}}{2} \right) + \varpi, \quad (17)$$

where ϖ is a constant term premium obtained as the average of the term premium implied by (17) over the period 1981q3-2013q4, replacing $E_t i_{t+1}^{3m}$ by the first horizon of T-bill SPF forecasts. The same equation is then solved for $E_t i_{t+1}^{3m}$ over the sample for which the SPF data for the T-bill is missing (see also Milani, 2011, p. 385).

The Greenbook forecasts were also obtained from the Philadelphia Fed. Since there are usually more than one Greenbook report per quarter, we chose to use those published after the first month of each quarter (only the first in case there are more), so as to align these forecasts with those we have from the SPF. Besides we use only up to 5 forecast horizons, including a first nowcast. The available samples of forecasts depend on the variable, and may include some missing horizons: for real GDP/GNP growth, its implicit price changes, and unemployment rates, the sample goes from 1967q1 to 2009q4; for CPI inflation it is 1979q4-2009q4; for the 3-month Treasury Bill it is 1981q1-2008q3; and for the 10-year Treasury Bond it is 2000q1-2008q3.

A.2 Derivation of adaptive gain algorithm

To derive the LS with adaptive gains, we start by defining the gain adaptation recursion, as given by

$$\gamma_t = \gamma_{t-1} - \alpha \hat{\nabla}_t, \quad (18)$$

where α represents a small adaptation constant, and $\hat{\nabla}_t$ stands for an estimate of the gradient $\nabla_t = \frac{\partial J_t}{\partial \gamma}$. The key step then is to find the relevant gradient and plug its stochastic approximation in the above recursions.

Taking the first derivative of J_t with respect to γ we obtain the gradient

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

which is stochastically approximated as

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

where $\hat{\Psi}_t$ stands for a recursive estimate of $\partial\hat{\theta}_t/\partial\gamma$. Differentiating (3) and (4) we obtain

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

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

Letting $\hat{\mathbf{S}}_t$ stand for the recursive estimate of $\partial\mathbf{R}_t/\partial\gamma$, and substituting (20), (21), and (22) into (18) we obtain the adaptive gain recursions presented in the main text. A similar derivation is also possible for the case of the SG algorithm (see footnote (5)), which can be found in Galimberti (2013).

A.3 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,

$$\bar{d}_h = \frac{1}{T} \sum_{i=1}^T d_{i,h}, \quad (23)$$

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

$$DM_h = \frac{\bar{d}_h}{\sqrt{\hat{\sigma}_d^2/T}}, \quad (24)$$

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 $E[\mathbf{q}_{t,h}d_{t,h}] = \mathbf{0}$, a Wald-type test statistic is proposed having the form of

$$GW_h = T \left(T^{-1} \sum_{i=1}^T \mathbf{q}_{i,h} d_{i,h} \right)' \hat{\Omega}_h^{-1} \left(T^{-1} \sum_{i=1}^T \mathbf{q}_{i,h} d_{i,h} \right), \quad (25)$$

Table 4: Averaged gain calibrations using revised data.

Calibration	Actual-based calibrations			Survey-based calibrations		
	Inflation	Growth	Interest	Inflation	Growth	Interest
(a) Fixed gains:						
- VAR(1)	6.0	1.9	0.5	1.3	1.4	0.5
- VAR(2)	2.3	1.4	0.5	0.8	1.3	0.5
- VAR(3)	2.1	1.0	0.5	0.6	1.1	0.5
- VAR(4)	2.4	1.0	0.5	0.8	1.1	0.5
- Overall	3.2	1.3	0.5	0.8	1.2	0.5
(b) Recursive gains:						
- VAR(1)	6.4	1.9	0.8	2.5	1.9	0.6
- VAR(2)	3.2	1.3	1.1	1.1	1.7	1.1
- VAR(3)	2.7	1.1	1.5	0.8	1.7	1.0
- VAR(4)	2.4	1.1	1.0	0.9	1.6	0.8
- Overall	3.7	1.4	1.1	1.3	1.7	0.9
(c) Adaptive gains:						
- VAR(1)	5.6	1.4	1.5	4.3	3.5	0.9
- VAR(2)	4.0	1.4	0.8	3.4	2.8	1.4
- VAR(3)	3.2	1.2	1.9	5.4	2.3	0.7
- VAR(4)	2.2	1.2	0.9	5.2	2.0	1.9
- Overall	3.7	1.3	1.3	4.6	2.7	1.2

The presented gain values are scaled by 100, e.g., 6.0 is equivalent to $\gamma = 0.06$. Averages are based on fixed and averaged time-varying gains selected under different model specifications (univariate, bivariate, and trivariate).

where $\hat{\Omega}_h$ 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_q^2 distribution.

Apart from the first horizon, $\hat{\Omega}_h$ 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 (25) 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.4 Supplementary results

References

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

- 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 (2015, April). On the initialization of adaptive learning in macroeconomic models. Mimeo.
- Branch, W. A. and G. W. Evans (2006). A simple recursive forecasting model. *Economics Letters* 91(2), 158–166.
- 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.
- Bullard, J. and S. Eusepi (2005). Did the great inflation occur despite policymaker commitment to a taylor rule? *Review of Economic Dynamics* 8(2), 324 – 359.
- Carboni, G. and M. Ellison (2009). The great inflation and the greenbook. *Journal of Monetary Economics* 56(6), 831 – 841.
- Carceles-Poveda, E. and C. Giannitsarou (2007). Adaptive learning in practice. *Journal of Economic Dynamics and Control* 31(8), 2659–2697.
- Chevillon, G., M. Massmann, and S. Mavroeidis (2010). Inference in models with adaptive learning. *Journal of Monetary Economics* 57(3), 341–351.
- Cho, I.-K., N. Williams, and T. J. Sargent (2002). Escaping nash inflation. *Review of Economic Studies* 69(1), 1–40.
- Christiano, L. J., M. Eichenbaum, and C. L. Evans (2005). Nominal rigidities and the dynamic effects of a shock to monetary policy. *Journal of Political Economy* 113(1), 1–45.
- 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. and R. S. Mariano (1995). Comparing predictive accuracy. *Journal of Business & Economic Statistics* 13(3), 253–63.
- Eusepi, S. and B. Preston (2011, October). Expectations, learning, and business cycle fluctuations. *American Economic Review* 101(6), 2844–72.
- 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.
- Galimberti, J. K. (2013). *Adaptive Learning for Applied Macroeconomics*. Ph. D. thesis, The University of Manchester.
- Giacomini, R. and H. White (2006). Tests of conditional predictive ability. *Econometrica* 74(6), 1545–1578.
- Honkapohja, S. and K. Mitra (2006). Learning stability in economies with heterogeneous agents. *Review of Economic Dynamics* 9(2), 284–309.
- Huang, K. X., Z. Liu, and T. Zha (2009). Learning, adaptive expectations and technology shocks. *The Economic Journal* 119(536), 377–405.
- 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.
- Lütkepohl, H. (2005). *New Introduction to Multiple Time Series Analysis*. Springer.
- Lucas, R. E. J. (1973, June). Some international evidence on output-inflation tradeoffs. *American Economic Review* 63(3), 326–34.
- Marcet, A. and J. P. Nicolini (2003). Recurrent hyperinflations and learning. *American Economic Review* 93(5), 1476–1498.
- Margaritis, D. (1990). A time-varying model of rational learning. *Economics Letters* 33(4), 309 – 314.
- Markiewicz, A. and A. Pick (2014). Adaptive learning and survey data. *Journal of Economic Behavior & Organization* 107, Part B(0), 685 – 707. Empirical Behavioral Finance.

- 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.
- Milani, F. (2014). Learning and time-varying macroeconomic volatility. *Journal of Economic Dynamics and Control* 47(0), 94 – 114.
- Molnár, K. and Z. Reppa (2010, October). Economic stability and the responsiveness of inflation expectations. Mimeo.
- Muth, J. F. (1961). Rational expectations and the theory of price movements. *Econometrica* 29(3), 315–335.
- 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.
- Ormeño, A. and K. Molnár (2015, June). Using survey data of inflation expectations in the estimation of learning and rational expectations models. *Journal of Money, Credit and Banking* 47(4), 673–699.
- 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.
- Primiceri, G. E. (2006). Why inflation rose and fell: Policy-makers’ beliefs and u. s. postwar stabilization policy. *The Quarterly Journal of Economics* 121(3), 867–901.
- 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.

- Slobodyan, S. and R. Wouters (2012). Learning in an estimated medium-scale dsge model. *Journal of Economic Dynamics and Control* 36(1), 26 – 46.
- Smets, F. and R. Wouters (2007). Shocks and frictions in us business cycles: A bayesian dsge approach. *American Economic Review* 97(3), 586–606.
- 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.
- Williams, N. (2003, January). Adaptive learning and business cycles. Mimeo.