

The background of the entire page is a photograph of the De Nederlandsche Bank (DNB) building. The building is a modern, multi-story structure with a prominent curved glass facade on the left and a more angular, brick-clad section on the right. A flag flies from the top of the building. In the foreground, a canal with a stone quay and several people walking is visible.

DNB Working Paper

The formation of European
inflation expectations: One
learning rule does not fit all

Christina Strobach and
Carin van der Crujsen

No. 472 / April 2015

DeNederlandscheBank

EUROSYSTEEM

The formation of European inflation expectations: One learning rule does not fit all

Christina Strobach ^a and Carin van der Cruijssen ^{*}

^{*} Views expressed are those of the authors and do not necessarily reflect official positions of De Nederlandsche Bank.

Working Paper No. 472

April 2015

De Nederlandsche Bank NV
P.O. Box 98
1000 AB AMSTERDAM
The Netherlands

The formation of European inflation expectations: One learning rule does not fit all ^{*}

Christina Strobach ^a and Carin van der Cruijssen ^a

^a *De Nederlandsche Bank, the Netherlands*

8 April 2015

Abstract

We empirically investigate how well different learning rules manage to explain the formation of household inflation expectations in six key member countries of the euro area. Our findings reveal a pronounced heterogeneity in the learning rules employed on the country level. While the expectation formation process in some countries can be best explained by rules that incorporate forward-looking elements (Germany, Italy, the Netherlands), households in other countries employ information on energy prices (France) or form their expectations by means of more traditional learning rules (Belgium, Spain). Moreover, our findings suggest that least squares based algorithms significantly outperform their stochastic gradient counterparts, not only in replicating inflation expectation data but also in forecasting actual inflation rates.

Keywords: Inflation expectations, adaptive learning algorithms, household survey.

JEL classifications: E31, E37, D84, C53.

^{*} Christina Strobach is the corresponding author, e-mail: c.a.b.strobach@dnb.nl. We are grateful for discussions with Falk Bräuning, Michael Massmann, Damjan Pfajfar, Andreas Pick, and seminar participants at DNB. The views expressed in this paper do not necessarily reflect the views of De Nederlandsche Bank.

1 Introduction

Central banks around the world have a particular interest in obtaining more insights on how economic agents form their expectations about future inflation as this information supports them in their aim to achieve price stability. Bernanke (2007) underlines the importance of understanding this formation process by pointing to its relevance for the determination of optimal policy decisions, for the process of forecasting inflation as well as for the judgment of the central bank's own credibility. So far however, there is no consensus among policy makers and academics on how to best model this expectation formation process. In this paper, we pursue a particular expectation formation scheme, namely adaptive learning and empirically investigate how well different adaptive learning rules manage to explain inflation expectations of households in six member countries of the euro area: Belgium, France, Germany, Italy, the Netherlands, and Spain. Given the fact that the ECB has declared an area-wide objective for price stability, adopting a comparative approach provides useful insights. It reveals heterogeneities in the formation processes on the national level that carry implications for the area as a whole.

In this paper, we present evidence that the adaptive learning approach provides a good description of the expectation formation process in all six European countries. We compare different learning rules which have been suggested by Pfajfar and Santoro (2010) and evaluate their ability to replicate household inflation forecasts. Our findings reveal some heterogeneity in the combination of information sets and algorithms which best describe the expectation formation on the individual country level, whereas no regional pattern can be identified. More specifically, our results suggest that households in Germany, Italy, and the Netherlands employ some forward-looking type of information in their expectation formation process, while those in France, Belgium, and Spain use information on energy prices or even more traditional learning rules focusing on the current rate of inflation. To benchmark our findings on household survey data, we also fit the learning rules to the actual inflation rates and thereby obtain information which algorithms and parameter values would have been optimal in forecasting changes in national CPIs. Our results imply that learning models in fact depict a better description of the survey data than they serve as a forecasting rule for the actual inflation rates. We also check for time variation in the gain parameters as well as the underlying information sets. While some variation becomes apparent, we do not see that the recent financial turmoil has had a systematic effect on the learning behavior of European consumers. In addition, we shed some light on the question whether European expectations are better modeled using least squares (LS) or stochastic gradient (SG) based learning algorithms. Statistical tests reveal that LS algorithms, which are commonly used in economic research, outperform SG based learning rules, both when tracking household survey data but also when forecasting actual inflation.

Modeling inflation expectations by adaptive learning has become widespread practice in the macroeconomic literature (see, e.g., Orphanides and Williams, 2004, 2005 or Milani, 2007, 2008, 2011). Compared to the concept of rational expectations, which is also widely used in economic modeling, this approach has the attractive feature of relaxing the informational requirements on the expectation forming agent. That is, agents are not required to know the true parameterization of the underlying economic model, as they learn those parameters over time by drawing inference based on new incoming data.¹ In fact, empirical evidence supporting the assumption that European households are strictly rational when forming their inflation expectations is rather weak, see Forsells and Kenny (2002) or Dias et al. (2010). Considering alternative expectation formation models such as adaptive learning therefore comes in at hand.

¹For a systematic discussion of adaptive learning see Evans and Honkapohja (2001).

The analysis presented in this paper relates to several contributions investigating the empirical support for different adaptive learning models in their ability to replicate survey data, in particular inflation expectations. Branch and Evans (2006) for example focus on the ability of different recursive forecasting models which are commonly used in the adaptive learning literature to track inflation and GDP growth expectations of professional forecasters as well as to forecast the variables themselves. Based on data for the US, they find that constant gain learning yields the most accurate description in both cases. Markiewicz and Pick (2014) extend the analysis as they replicate not only forecasts of macroeconomic but also of financial variables, allowing for several different learning specifications. Their results indicate that an adaptive expectation model performs best in tracking US inflation expectations of professional forecasters. Malmendier and Nagel (2015) furthermore use micro data on consumer inflation expectations to fit a model which reflects the idea of “learning from experience”. To do so, the authors employ a constant gain specification where the gain depends on the age of the respective agent and thereby show that young individuals assign a higher weight to more recent observations than older individuals. With respect to the country coverage, our study is closest to the work of Weber (2010) who investigates the formation of inflation expectations for several European countries by means of adaptive learning techniques. She finds that constant gain learning describes the expectation formation mechanism of European households as well as professional forecasters better than a decreasing gain specification.

While our contribution builds on the existing evidence for household learning in European countries, it extends the literature by testing new learning mechanisms which are motivated by the advancement of Pfajfar and Santoro (2010). In particular, we consider different laws of motion for the evolution of actual inflation and introduce more flexibility in the updating process of the estimated parameters than commonly assumed in traditional learning algorithms. The latter extension is specifically designed to detect forward-looking elements in the updating process of consumers. One updating mechanism considered allows for the combination of sticky information and adaptive learning in the expectation formation process and reflects the idea of the epidemiological model suggested by Carroll (2003). Sticky information models have been proposed as another alternative to rational expectations since they perform well in explaining inflation dispersion on the micro level, see Cukierman and Wachtel (1979) or Mankiw et al. (2004). Based on the observation that only a certain fraction of all agents update their expectations each period, Carroll (2003) derives a micro-founded model which accounts for this particular feature of the data. Building on Carroll’s epidemiological approach, Doepke et al. (2008) show that European households update their expectations on average every 18 months.

Besides testing the empirical evidence of forward-looking updating rules, we furthermore extend the learning algorithms of Pfajfar and Santoro (2010) by introducing the possibility that consumers rely on information about frequently purchased items such as energy prices. This idea is based on the observation that households often think about energy prices when asked to provide their opinion on future price changes as documented by Bruine de Bruin et al. (2010). Overall, our results imply that the actual models used to form expectations on the national level are quite heterogeneous as we find evidence in favor of every learning rule that we consider. Jointly for all countries however, the findings suggest that households rely on rather simple information sets when conducting their inflation forecasts.

Finally, this paper also relates to the contribution of Berardi and Galimberti (2014) who compare the ability of LS and SG learning algorithms to replicate survey data as well as to forecast different macroeconomic variables. Based on data for the US, the authors find that LS algorithms better replicate observations of professional inflation forecasts, while the actual inflation rate can

be projected more accurately by means of SG algorithms. We conduct a similar exercise for the European countries, providing statistical evidence that LS based algorithms outperform their SG counterparts in both cases.

The remainder of this paper is organized as follows. In the next section, we introduce the household survey data and discuss certain features of it. Section 3 then outlines a general version of the adaptive learning model. Issues related to the estimation of the time-varying parameters, the initialization as well as the underlying information sets are discussed in Section 4. The actual learning rules that we assess are presented in Section 5, followed by the corresponding estimation results, whereas we separately discuss the role of different information sets. Section 6 presents evidence on a time-varying gain estimation, while Section 7 concludes.

2 Household inflation expectations

We derive consumer inflation expectation figures from qualitative data on inflation perceptions and expectations. The data is obtained from the monthly European Commission’s consumer survey. More specifically, information on current inflation perceptions is obtained by asking households *“How do you think that consumer prices have developed over the last 12 months”*. The question on inflation expectations furthermore is phrased as *“By comparison with the past 12 months, how do you expect that consumer prices will develop in the next 12 months”*. The survey respondents can choose their answer on a discrete scale of five categories, ranging between *“prices have fallen/ will fall”* to *“prices have risen a lot/ will increase more rapidly”*, depending on the respective question asked. They can also indicate that they do not know the answer to the question. The European Commission then publishes aggregate outcomes by country and response share on a monthly basis. As a result, we have time series over all response categories for each of the six countries that we consider.² We use data from January 1990 until December 2012 and the country selection has been conducted based on data availability not only with respect to the survey data but with respect to all variables we use in this study. Data is available for six key member countries of the euro area: Belgium, France, Germany, Italy, the Netherlands, and Spain. Even though central banks usually stress the importance of medium to long-term inflation expectations, investigating the formation and properties of expectations with a shorter horizon as in our case is also highly relevant. Bernanke (2007) for example points to the necessity of having reliable short-term forecasts as they provide a good “jumping-off point” for more accurate longer-term projections.³

To transform the qualitative data into quantitative figures of inflation expectations, we follow the literature such as Weber (2010) or Doepke et al. (2008) and apply the probability method suggested by Carlson and Parkin (1975) in its augmented form (see Batchelor and Orr, 1988) to accommodate cases with five response categories. Since the survey question on inflation expectations asks households to provide their outlook about future price developments relative to their perception of price changes over the last year, we use the qualitative information on inflation perceptions to compute a scaling variable as suggested by Berk (1999). This approach has the advantage that expectations are not assumed to be unbiased which is important when investigating their formation

²A description of the survey and the data can be found in European Commission (2007).

³For a discussion of short-term inflation and short-term inflation expectations in monetary policy making see, e.g., Giannone et al. (2014) or Clark and Davig (2008).

process.⁴ However, Nardo (2003) and Weber (2010) note that one has to keep in mind that the expectation figures display only an approximation of the unknown economic agent's expectations. This is the necessary result of the fact that consumer inflation expectations for European countries are only available in qualitative form.

Quantified inflation expectations and annual inflation rates calculated based on the respective national consumer price index (CPI) are shown for each country in Figure 2.1. In addition, Table 2.1 provides summary statistics of the two variables as well as the associated mean squared forecast errors (MSFE) by country. For each variable, we show the mean outcome over the entire period, which ranges from the beginning of 1991 until the end of 2013 as well as for three subperiods.⁵ The first subsample encompasses the period before the introduction of the single monetary policy (SMP) and includes data from 1991 until the end of 1998. Our second subsample covers the implementation of the SMP in 1999 as well as the euro currency changeover in 2002. It ranges until the declaration of bankruptcy by Lehman Brothers in September 2008. The third period covers the European sovereign debt crisis and ranges from September 2008 up until December 2013.

Table 2.1: Descriptive statistics for inflation, inflation expectations, and forecast errors


The table reports the mean CPI inflation rate together with the mean household forecast and the associated mean squared forecast error (MSFE) for each country. The means are calculated over the entire period of investigation (1991m1-2013m12) as well as three consecutive subperiods. "Pre-SMP" ranges from 1991m1 to 1998m12, "SMP I" from 1999m1 to 2008m8, and "SMP II" from 2008m9 to 2013m12.

	Belgium	France	Germany	Italy	Netherlands	Spain
Inflation						
Entire period	2.13	1.71	1.98	2.89	2.25	3.22
Pre-SMP	2.11	1.87	2.79	4.18	2.47	4.15
SMP I	2.18	1.76	1.61	2.34	2.22	3.23
SMP II	2.06	1.36	1.43	1.94	1.99	1.79
Forecast						
Entire period	1.66	1.30	1.75	2.43	1.85	2.44
Pre-SMP	1.67	0.81	2.31	3.76	1.84	2.42
SMP I	1.33	1.26	1.56	1.74	2.97	2.70
SMP II	2.24	2.12	1.24	1.70	1.67	2.00
MSFE						
Entire period	2.17	1.47	1.52	1.61	1.01	4.29
Pre-SMP	0.58	1.51	1.93	1.86	0.90	4.49
SMP I	1.74	0.76	1.30	0.79	0.89	2.08
SMP II	5.33	2.71	1.28	2.71	1.38	7.98


The results reported in the upper panel of Table 2.1 show that over those 23 years, inflation was on average somewhat below 2% in France and Germany and slightly above in Belgium and the

⁴In our context we are faced in very few cases with the problem that the quantification method breaks down whenever the response category indicating that prices have fallen or will fall is zero. To avoid that, we apply the correction suggested in Heinzl and Wollmershäuser (2005) based on the approximate number of participating household in each country as released by the European Commission.


⁵To make the comparison between CPI inflation and its forecast more comprehensive, we report each forecast at the date of its realization. For example, a forecast made in January 1990 is reported in January 1991 when the actual CPI materializes. We keep that notation throughout the entire paper unless stated otherwise.




(a) Belgium




(b) France




(c) Germany



(d) Italy



(e) Netherlands



(f) Spain

Figure 2.1: Inflation and inflation expectations

Household inflation expectations (blue solid line) are displayed 12 months after the forecast has been conducted together with CPI inflation (green dashed line).

Netherlands. The average inflation rate in Italy and Spain was instead close to 3%. Remarkably, Spain and Italy experienced quite elevated levels of inflation before the introduction of the SMP with an average of above 4%. While the forecast error for Spain is also high during that period as consumers expected a lower inflation rate, the level of expected price changes by Italian consumers is quite in line with the level of actual inflation. The forecast errors furthermore suggests that Dutch household were on average most successful in forecasting inflation while Spanish consumers, which also experienced the highest overall inflation rate, were least successful in their forecasting exercise.

A visual inspection of the period including the euro currency changeover suggests that the introduction of the new currency raised inflation expectations especially in Germany, the Netherlands, and Spain (see Figure 2.1). Table 2.1 however indicates that during that particular period, inflation was on average lower than before the introduction of the SMP, at least in five out of six countries. This observation is not only pronounced for Italy and Spain but also for Germany. Moreover, inflation was lowest in the last subperiod which covers the European sovereign debt crisis. The low level however does not imply that inflation forecasts became more accurate. In fact, for all countries but Germany, average forecast errors were highest during this episode which was marked by elevated economic uncertainty.

3 A general model

In this section, we briefly outline how the traditionally used adaptive learning rules can be derived from a general state space model. For a thorough treatment of this correspondence see Berardi and Galimberti (2013). First, consider a household which intends to form an inflation forecast. The household assumes that inflation evolves in the following way

$$\pi_t = \mathbf{x}_t' \boldsymbol{\theta}_t + \varepsilon_t, \quad (1)$$

where π_t is the inflation rate, \mathbf{x}_t is a $k \times 1$ vector of variables that the agent expects to be related to inflation, $\boldsymbol{\theta}_t$ denotes a $k \times 1$ vector of (potentially) time-varying parameters and $\varepsilon_t \stackrel{iid}{\sim} N(0, \sigma^2)$ is a disturbance term. The aspect of learning is added to the model through the development of the vector $\boldsymbol{\theta}_t$ which we allow to evolve either according to a recursive LS or an SG algorithm. Both learning algorithms can be derived under certain parameter restrictions as special cases of the Kalman filter. To do so, we follow Berardi and Galimberti (2013), who establish an exact correspondence between both learning algorithms and the Kalman filter, and begin from a state space representation of the learning exercise. In this context, Equation (1) is the observation equation and the state equation can be depicted as

$$\boldsymbol{\theta}_t = \boldsymbol{\theta}_{t-1} + \boldsymbol{\omega}_t, \quad (2)$$

with $\boldsymbol{\omega}_t \stackrel{iid}{\sim} N(\mathbf{0}, \boldsymbol{\Omega})$. The Kalman recursion corresponding to Equations (1) and (2) can be expressed in the following way

$$\hat{\boldsymbol{\theta}}_t = \hat{\boldsymbol{\theta}}_{t-1} + \mathbf{K}_t(\pi_t - \mathbf{x}_t' \hat{\boldsymbol{\theta}}_{t-1}), \quad (3)$$

where $\hat{\boldsymbol{\theta}}_t$ denotes the estimated coefficient parameter. \mathbf{K}_t , the Kalman gain, and \mathbf{P}_t , the covariance matrix of the estimated coefficients, are respectively defined as

$$\mathbf{K}_t = \frac{\mathbf{P}_{t-1} \mathbf{x}_t}{\mathbf{x}_t' \mathbf{P}_{t-1} \mathbf{x}_t + \sigma^2}$$

and

$$\mathbf{P}_t = \left(\mathbf{I} - \frac{\mathbf{P}_{t-1} \mathbf{x}_t \mathbf{x}_t'}{\mathbf{x}_t' \mathbf{P}_{t-1} \mathbf{x}_t + \sigma^2} \right) \mathbf{P}_{t-1} + \mathbf{\Omega}_t.$$

Berardi and Galimberti (2013) now show that setting

$$\sigma^2 = \frac{\gamma_{t-1}}{\gamma_t} (1 - \gamma_t),$$

while defining the parameter covariance matrix as

$$\mathbf{\Omega}_t = \left(\frac{1 - \sigma_t^2}{\sigma_t^2} \right) \left(\mathbf{I} - \frac{\mathbf{P}_{t-1} \mathbf{x}_t \mathbf{x}_t'}{\mathbf{x}_t' \mathbf{P}_{t-1} \mathbf{x}_t + \sigma_t^2} \right) \mathbf{P}_{t-1},$$

yields an equivalence to the LS learning algorithm which is characterized as

$$\begin{aligned} \hat{\boldsymbol{\theta}}_t &= \hat{\boldsymbol{\theta}}_{t-1} + \gamma_t \mathbf{R}_t^{-1} \mathbf{x}_t (\pi_t - \mathbf{x}_t' \hat{\boldsymbol{\theta}}_{t-1}), \\ \mathbf{R}_t &= \mathbf{R}_{t-1} + \gamma_t (\mathbf{x}_t \mathbf{x}_t' - \mathbf{R}_{t-1}), \end{aligned} \tag{4}$$

where γ_t is a gain parameter and \mathbf{R}_t is the matrix of second moments of the regressors \mathbf{x}_t . In addition, they establish an exact correspondence between the Kalman filter and the SG algorithm

$$\hat{\boldsymbol{\theta}}_t = \hat{\boldsymbol{\theta}}_{t-1} + \gamma_t \mathbf{x}_t (\pi_t - \mathbf{x}_t' \hat{\boldsymbol{\theta}}_{t-1}), \tag{5}$$

by imposing the following restrictions on the covariance of the measurement equation

$$\sigma^2 = \gamma_t^{-1} - \mathbf{x}_t' \mathbf{x}_t$$

as well as on the covariance of the state equation

$$\mathbf{\Omega}_t = \mathbf{I} - \left(\mathbf{I} - \frac{\mathbf{P}_{t-1} \mathbf{x}_t \mathbf{x}_t'}{\mathbf{x}_t' \mathbf{P}_{t-1} \mathbf{x}_t + \sigma_t^2} \right) \mathbf{P}_{t-1}.$$

A comparison between Equations (4) and (5) suggests that the LS and the SG algorithms are quite similar as the second algorithm resembles the first one neglecting only the normalization by the matrix of second moments. For a systematic discussion of the algorithms see Evans and Honkapohja (2001) and Evans et al. (2010).

While LS learning has received a lot of attention in the literature on macroeconomic modeling (see, e.g., Branch and Evans, 2006; Orphanides and Williams, 2005; Milani, 2007, 2008, 2011), only few studies in the applied literature such as Pfajfar and Santoro (2010) or Berardi and Galimberti (2014) consider the use of SG learning. Berardi and Galimberti (2014) argue that the superiority of neither algorithm can be established a priori, yet that intuitively one would expect agents to prefer the SG algorithm as it is computationally much less demanding than the LS one. Beyond that, the LS algorithm may be just better known among researchers because it converges under certain conditions to a well-known OLS estimator.

To further specify the model, we need to impose assumptions on the form of the gain parameter γ_t . As shown in Equations (4) and (5), the gain reflects the weight that an agent puts on the most recent forecast error. In the applied learning literature it has usually been defined to take one of two forms. The first specification commonly employed is the one of a time-decreasing gain such that $\gamma_t = t^{-1}$. Adapted to our context this means that households weigh each forecast error equally but that the value attached to each error declines over time. For the LS algorithm this implies that

as $t \rightarrow \infty$, the system converges to an OLS estimator. A popular alternative to the decreasing gain is to assume that the parameter remains constant over time and does not depend on the sample length so that $\gamma_t = \gamma, \forall t$. In this context, past observations receive geometrically decaying weights, that is, observations which are further back in time are weighted less than more recent ones.⁶ Since this specification is more robust to account for structural changes which are likely to occur given that our sample period covers the creation of the EMU as well as the financial and sovereign debt crises, we solely focus on a constant gain specification for γ in the remainder of this study.

4 Estimation of the time-varying parameters

4.1 Estimation of the gain parameter

To determine the time-varying parameter vector θ_t for the LS and the SG learning algorithms outlined in Equations (4) and (5) respectively, we have to estimate the gain parameter γ . We simulate⁷ different laws of motion which will be specified in Section 5. We then evaluate the simulated forecasts $\hat{\pi}_{t+12|t}$ against the true observable forecasts conducted by the representative household $\pi_{t+12|t}^{\text{HF}}$ by computing the associated mean squared error (MSE). The selected gain parameter is then the one that yields the smallest MSE over the sample period considered. Thus, it is the solution to the following minimization problem

$$\gamma^{HF} = \arg \min_{\gamma \in [0,1]} \left\{ \frac{1}{T} \sum_{t=1}^T (\hat{\pi}_t - \pi_{t|t-12}^{\text{HF}})^2 \right\}, \quad (6)$$

where γ^{HF} is the optimal gain parameter to replicate the household forecast and $t = 1, 2, \dots, T$ is the time period considered, that is, the time period after the initialization of the parameters. In choosing the gain and thereby also the optimal model over the entire sample of observations, we follow Pfajfar and Santoro (2010) and Orphanides and Williams (2005).⁸ In a second step however, we will check the robustness of our optimal gain parameters over time by computing the optimal γ^{HF} parameters for a much shorter rolling window. While the first approach uses all available data over the entire time period, the second one only makes use of data that is available to agents when they form their expectations and therefore, as noted by Berardi and Galimberti (2012), it is more in line with a “fair” out of sample forecast. We select the gain by means of a grid search procedure, where we consider 10001 equally spaced grids between zero and one. By explicitly including zero in the set of potential parameter values, we do not foreclose the option that learning does not take place. Therefore, the gain selection is at the same time also a general check of whether or not learning is an appropriate description of the households’ expectation formation process.

While the aim of this paper is to determine the algorithms which best replicate the expectation formation of households in different European countries, we also conduct a second exercise intending to identify which algorithms would have been most plausible in forecasting the national inflation rates. Computing this benchmark allows us to check if households use optimal algorithms as well as if they employ all relevant information when predicting future changes in the price level. We

⁶Berardi and Galimberti (2013) show that the way observation are discounted by the SG algorithm does not solely depend on the gain parameter as in the LS case but also on \mathbf{x}_t . The finding corresponds with the result of Evans et al. (2010) who point to the fact that the gain in SG learning is not scale invariant. This makes it harder to draw clear inference on the weight assigned to past observations when SG learning is applied.

⁷We use the term “simulate” interchangeably with “replicate” or “compute”.

⁸Alternative approaches pursuit in the literature often estimate the gain over a much shorter period of time, whereas the remaining observations are used to engage in a forecasting exercise, see, e.g., Markiewicz and Pick (2014) or Weber (2010).

conduct the gain selection for the forecasting models analogous to the one outlined in Equation (6) with the exception that $\pi_{t|t-12}^{\text{HF}}$ is replaced by π_t . Thus, the optimal gain parameter in this context, γ^{AI} , is the one that delivers the lowest MSE between our simulated forecasts ($\hat{\pi}_t$) and the actual inflation rate (π_t) which is

$$\gamma^{AI} = \arg \min_{\gamma \in [0,1]} \left\{ \frac{1}{T} \sum_{t=1}^T (\hat{\pi}_t - \pi_t)^2 \right\}. \quad (7)$$

The distinction between γ^{HF} and γ^{AI} allows to conceptually understand the different intentions of the minimization exercise, that is, either finding the optimal γ to replicate household forecast or to replicate actual inflation. In the following however, we will introduce the algorithms and just refer to the gain parameter as γ . We will only use the superscripts if we intend to specifically point to the difference between household forecasts and actual inflation. Note that we interpret the gain parameter not as resulting from the agent’s conscious choice, but see it more as a primitive parameter, which reflects the weight of the past error in the household’s memory process. For a thorough discussion of the difference between the interpretation as choice or primitive parameter, see Galimberti (2013). The latter exercise of fitting the algorithms to actual inflation therefore constitutes a benchmark that allows to infer how close the gain parameter underlying the expectation formation process is to the parameter that results as a choice from an optimization process when forecasting inflation.

4.2 Initialization of the learning algorithms

To set up the learning algorithm in Equations (4) and (5), we have to find appropriate initial values for θ_0 and in case of the LS algorithm also for the matrix of second moments \mathbf{R}_0 . Carceles-Poveda and Giannitsarou (2007) discuss different initialization methods and their implication for various learning algorithms, while Berardi and Galimberti (2012) provide a comprehensive overview of different initialization methods used in the applied learning literature. In this paper, we follow an approach referred to as “Diffuse-track” initialization by Berardi and Galimberti (2012). That is, we start the algorithm by setting $\theta = \mathbf{0}$ and $\mathbf{R} = \mathbf{I}$. We then allow the algorithm to update the subsequent parameters which should, after a certain period of time, settle at values that are consistent with the respective learning algorithm. The last observation of this initialization sequence is then used as our θ_0 as well as \mathbf{R}_0 , while previous observations are discarded. Similar to Berardi and Galimberti (2012) who use 75 periods, we set the length of the initialization period to 72 observations which corresponds to six years. By choosing this time window for the initialization, we keep 17 years of data to determine the optimal gain parameter.

4.3 Information sets

We start our analysis by assuming that households employ simple forecasting models relying only on information about the most recent change in the price level as a predictor for future inflation. This is in line with the results presented by Pfajfar and Santoro (2010) for US consumers. Thus, the information set of the representative household is $\mathbf{x}_t = (1, \pi_{t-1})'$, which implies that the household in period t has access to data on inflation only up to period $t - 1$. In addition, the vector is augmented by a constant term. Based on this information set, we introduce different laws of motion as well as different updating algorithms to determine first of all the learning schemes that best track the expectation formation processes of the consumers in different countries and second the optimal forecasting algorithms for the national inflation rates.

In the next step, we relax the assumption of homogeneous information sets across countries and allow the representative household to choose from a wider range of information sets, which

also allows us to check the robustness of previous findings. In our selection of additional variables we follow Weber (2010). Alternative to the first information set specified above, we consider a set of regressors which in addition to the inflation rate and the constant term also includes the annual growth rate of industrial production y_{t-1} , yielding $\mathbf{x}_t = (1, \pi_{t-1}, y_{t-1})'$. While we would prefer to use GDP as a measure for economic activity, those figures are not available at a monthly frequency and we therefore use industrial production as a proxy. Since real-time data of industrial production can't be accessed for the entire period of observation, we employ revised data.⁹ We furthermore consider the option that households employ the annual change in the nominal short-term interest rate i_{t-1} instead of economic activity so that $\mathbf{x}_t = (1, \pi_{t-1}, i_{t-1})'$. Finally, the fourth set of regressors consists of the industrial production growth rate, the annual change in the interest rate as well as the inflation variable, so that $\mathbf{x}_t = (1, \pi_{t-1}, y_{t-1}, i_{t-1})'$. All variables are taken from the OECD data base and growth rates of industrial production are calculated from the industrial production index.

4.4 Comparison between LS and SG algorithms

To evaluate whether the LS or the SG learning algorithm depicts a better description of the households' expectation formation process as well as which of the two algorithms delivers a more accurate forecast of the national CPIs, we consider two statistics. Primarily, we focus on the lowest MSE as our main indicator. In addition however, we also evaluate the predictive accuracy according to the Diebold and Mariano (2002) test statistic incorporating the small sample correction suggested by Harvey et al. (1997). In the following, we will refer to the latter statistic as "modified Diebold-Mariano test" (DM). The calculation of the test statistic is based on the differential d_t , which is defined as the difference between the squared errors of the forecasts based on the LS algorithm, $(e_t^{LS})^2$ and the SG algorithm $(e_t^{SG})^2$, so that

$$d_t = (e_t^{LS})^2 - (e_t^{SG})^2.$$

The forecast errors are computed as the difference between the constructed series and i) the survey data or ii) the actual inflation rate, depending on the purpose of the exercise.

5 Different learning mechanisms

To identify the empirical relevance of adaptive learning for households in different European countries, we test several alternative learning rules as proposed in Pfajfar and Santoro (2010). For each type of learning rule, we present a simple conceptual model illustrating the main sequence of events in Figure B.3 in the Appendix. We start by considering two alternative reduced form specifications for inflation which differ in their degree of backward-lookingness. Then, in the next step, we allow the updating process to be conducted based on variables other than the actual inflation rate, while we keep the economic law of motion constant. Even though this degree of freedom withholds the possibility of deriving the learning algorithms as nested models of the Kalman filter, it allows to incorporate assumptions that have been established in other models on the formation of inflation expectations. Having determined the optimal learning rule for each country, we test whether the choice of the optimal algorithm on the national level is robust to the inclusion of additional explanatory variables. Besides serving as a robustness check, this also allows us to draw conclusions on whether households use rather simple information sets in their forecasting exercise or if the forecasts are based on a larger set of variables. Finally, we evaluate based on all possible combinations of learning rules and information sets, whether LS or SG learning statistically outperforms

⁹Weber (2010) also considers revised instead of real-time data to test the validity of adaptive learning models in replicating inflation forecasts.

its counterpart in modeling survey inflation expectations as well as in forecasting future changes in the price level.

5.1 How do households perceive inflation to evolve?

The literature so far has mostly considered learning algorithms which are based on a specification suggesting that agents perceive inflation in period $t + 1$ to depend on variables available in period t . Thus, agents update their beliefs and use their updated beliefs to conduct (several) one period ahead forecasts. In this section, we extend the set of potential laws of motion by incorporating an alternative which is more backward-looking than the ones usually employed in the literature. Notably, for all algorithms that we discuss, agents form their forecasts for the upcoming 12 months based on information (data as well as parameter estimates) from the previous period.

5.1.1 Learning based on current forecast errors

In this section, we pursue the idea that households consider inflation to evolve according to a simple AR(1) process so that the economic law of motion can be expressed as

$$\pi_t = \mathbf{x}_t' \boldsymbol{\theta}_t + \varepsilon_t, \quad (8)$$

which is equivalent to Equation (1). This reduced form expression for inflation is in line with the ones usually employed in the applied learning literature, see Branch and Evans (2006) or Berardi and Galimberti (2014). The associated algorithms are for SG learning

$$\hat{\boldsymbol{\theta}}_t = \hat{\boldsymbol{\theta}}_{t-1} + \gamma \mathbf{x}_t (\pi_t - \mathbf{x}_t' \hat{\boldsymbol{\theta}}_{t-1}) \quad (9)$$

and for LS learning

$$\begin{aligned} \hat{\boldsymbol{\theta}}_t &= \hat{\boldsymbol{\theta}}_{t-1} + \gamma \mathbf{R}_t^{-1} \mathbf{x}_t (\pi_t - \mathbf{x}_t' \hat{\boldsymbol{\theta}}_{t-1}), \\ \mathbf{R}_t &= \mathbf{R}_{t-1} + \gamma (\mathbf{x}_t \mathbf{x}_t' - \mathbf{R}_{t-1}). \end{aligned} \quad (10)$$

The sequence of events implied by this learning rule is the following: At the beginning of the period, the household has expectations about the current rate of inflation in mind that have been formed in the previous period. It is then asked to provide its inflation forecast for the upcoming year. As the current rate of inflation has not been observed yet, it uses the information employed in calculating the forecast for period t to engage in 12 one-step ahead forecasts (since the data is at a monthly frequency). This yields the following conditional forecast for the univariate case

$$\begin{aligned} \pi_{t+1|t}^{\text{HF}} &= \mathbf{x}_{t+1}' \hat{\boldsymbol{\theta}}_{t-1}, \\ &\dots \\ \pi_{t+12|t}^{\text{HF}} &= \mathbf{x}_{t+12}' \hat{\boldsymbol{\theta}}_{t-1}, \end{aligned} \quad (11)$$

with $\mathbf{x}_{t+1} = (1, \pi_{t|t-1}^{\text{HF}})$, $\mathbf{x}_{t+2} = (1, \pi_{t+1|t}^{\text{HF}})$, ..., $\mathbf{x}_{t+12} = (1, \pi_{t+11|t}^{\text{HF}})$. In a multivariate setting, the algorithm turns itself into a VAR process, where forecasts of all variables included in the set of regressors evolve as outlined in Equation (11). After that, households observe the current inflation rate π_t which they use to update the coefficient estimate $\hat{\boldsymbol{\theta}}_t$ and, based on that, form expectations for the next period $t + 1$. The corresponding conceptual model is presented in the upper section of Figure B.3 in the Appendix.

5.1.2 Learning based on previous forecast errors

Next, we assume that the representative household is more backward-looking than commonly hypothesized. We define the economic law of motion that the agent considers to be

$$\pi_t = \mathbf{x}'_{t-12} \boldsymbol{\theta}_{t-13} + \varepsilon_t, \quad (12)$$

which yields the following SG learning algorithm

$$\hat{\boldsymbol{\theta}}_t = \hat{\boldsymbol{\theta}}_{t-1} + \gamma \mathbf{x}_{t-12} (\pi_t - \mathbf{x}'_{t-12} \hat{\boldsymbol{\theta}}_{t-13}), \quad (13)$$

while the LS learning in this case is defined as

$$\begin{aligned} \hat{\boldsymbol{\theta}}_t &= \hat{\boldsymbol{\theta}}_{t-1} + \gamma \mathbf{R}_t^{-1} \mathbf{x}_{t-12} (\pi_t - \mathbf{x}'_{t-12} \hat{\boldsymbol{\theta}}_{t-13}), \\ \mathbf{R}_t &= \mathbf{R}_{t-1} + \gamma (\mathbf{x}_{t-12} \mathbf{x}'_{t-12} - \mathbf{R}_{t-1}). \end{aligned} \quad (14)$$

The economic law of motion outlined in Equation (12) suggests that the household assumes inflation to respond with a lag of more than one year to different macroeconomic variables which makes the learning rule particularly backward-looking. The sequence of events implied by the learning algorithms therefore is slightly different from the one based on current forecast errors. At the beginning of each period t , the household has a forecast for the current inflation rate in mind which it conducted one year ago. In addition, it knows the data up to $t-1$ as well as the previous period's coefficient estimate. Before current macroeconomic variables are released, the household is then asked to provide its (conditional) forecast for inflation in $t+12$ which is $\pi_{t+12|t}^{\text{HF}} = \mathbf{x}'_t \hat{\boldsymbol{\theta}}_{t-1}$. After having forecast future inflation, the household observes the current change in the price level which it employs to update the coefficient estimate of $\boldsymbol{\theta}_t$ either based on the LS or the SG algorithm. The updated parameter estimate is then used in the following period to compute a new inflation forecast for the upcoming year $\pi_{t+13|t+1}^{\text{HF}}$. The corresponding conceptual model depicting the idea of learning based on previous errors is outlined in the middle section of Figure B.3 in the Appendix.

5.1.3 Estimation results for learning based on different laws of motion

Table 5.1 displays the optimal gain parameters γ^{HF} , together with the MSEs for household forecasts in the left panel and γ^{AI} for CPI inflation in the right panel. SG and LS learning based on current forecast errors are indicated by CE while algorithms based on previous forecast errors are indicated by PE. By comparing the ability of the LS and SG learning algorithm to track household forecasts, one can infer from the left panel of the table that LS learning outperforms its SG counterpart in four out of six cases. This suggests that when households are assumed to update their forecasts to the actual inflation variable, their learning is better described by an LS algorithm. The effect however turns out to be statistically significant only for the Netherlands. In addition, we find the optimal gain to be greater than zero for every single country. This result supports the assumption that learning is an adequate description of the households' inflation expectation formation processes throughout the euro area.

The gain parameters for the optimal learning rules of households vary between 0.0007 in the case of Italy and 0.0134 for Germany. This implies that an average Italian consumer attaches a relatively low weight to the most recent forecast error and thus also includes a longer history of observations into its forecasts. A representative German household on the other hand downweights previous observations more rapidly, which makes it more responsive to structural changes in its environment. More specifically, the gain parameters suggest that German consumers weigh a ten year old observation only by 0.2 ($= (1 - 0.0134)^{10 \cdot 12}$) while in Italy, the observation still receives

Table 5.1: Estimates for learning based on current and previous errors

The table reports optimal gain parameters together with the associated MSEs in parentheses for SG and LS learning. Optimal gains to replicate household forecasts are displayed in the left panel, while the right panel depicts optimal gains with respect to actual inflation. Gain parameters and MSEs are shown for each learning algorithm assuming learning is based on current forecasts errors (CE) or alternatively on previous forecast errors (PE). For each country, the combination of gain parameter and forecasting algorithm that yields the lowest MSE is highlighted. * indicates the algorithm that is more accurate than its counterpart (either LS or SG) according to the modified Diebold-Mariano two-sided test at the 20% significance level.

	Household forecasts				CPI inflation			
	SG-CE	LS-CE	SG-PE	LS-PE	SG-CE	LS-CE	SG-PE	LS-PE
Belgium	0.0031 (0.3408)	0.0098 (0.3568)	0.0012 (0.4130)	0.0021 (0.3840)	0.0029 (2.1211)	0.0075 (1.9514)	0.0011 (3.0129)	0.0156* (2.2355)
France	0.0038 (0.4094)	0.0127 (0.4539)	0.0018 (0.4136)	0.0035 (0.4086)	0.0040 (0.7002)	0.0094 (0.6791)	0.0016 (1.0645)	0.0144* (0.9272)
Germany	0.0170 (0.4229)	0.0134 (0.3935)	0.0067 (0.4166)	0.0104 (0.4183)	0.0024 (0.5595)	0.0073 (0.5574)	0.0013 (0.9096)	0.0176* (0.6946)
Italy	0.0011 (0.4483)	0.0045 (0.4514)	0.0004 (0.3172)	0.0007 (0.3137)	0.0013 (0.7932)	0.0057 (0.7255)	0.0005 (1.2482)	0.0009* (1.2175)
Netherlands	0.0087 (0.3906)	0.0197 (0.3201)	0.0018 (0.3587)	0.0034* (0.3145)	0.0073 (0.7948)	0.0161 (0.8036)	0.0022 (0.9786)	0.0116 (0.8354)
Spain	0.0013 (1.1286)	0.0081 (1.1973)	0.0018 (0.9498)	0.0010 (1.1043)	0.0013 (2.1940)	0.0065 (1.9421)	0.0005 (2.7557)	0.0155 (2.4294)

a pretty high weight of 0.9 ($= (1 - 0.0007)^{10 \cdot 12}$) in the household's memory process. Comparing the range of γ^{HF} parameters presented in the table to the outcomes reported in other studies, in particular to the ones of Weber (2010) who also considered European countries, we find that they are quite similar. Her optimal gains range between 0.0002 and 0.0530, however they are estimated solely based on an LS algorithm in combination with a slightly different learning rule for the time period from 1990 until mid 1998. Distinctively however, our results do not support the conclusion that households in southern European countries, who experienced on average higher levels of inflation, use higher constant gains than northern European consumers.¹⁰

The right panel of Table 5.1 provides information on the optimal constant gains in the forecasting exercise of the actual inflation rate. Analogous to the previous case, the LS algorithm turns out to be superior to SG learning. The finding however is significant only for the algorithm incorporating previous errors for four out of six countries. Furthermore, the range of optimal parameters obtained for predicting actual inflation is smaller than the one reported for tracking the survey data as the parameters lie in between 0.0057 and 0.0094. In addition, one can see that the optimal algorithm for the household forecast corresponds only for Germany with the one for actual inflation. The high MSEs furthermore suggest that adaptive learning, at least based on the algorithms and information sets considered so far, represents a better description for household learning than for the evolution of inflation.

5.2 Which information do households update their forecasts to?

Having discussed learning based on different reduced form specifications for inflation, this section proceeds to introduce more flexibility in the updating process of each learning algorithm. Follow-

¹⁰Since the gains for SG learning are sensitive to data scale, we draw this conclusion focusing solely on gains obtained from LS learning.

ing Pfajfar and Santoro (2010), we relax the assumption that the household necessarily updates its forecast to the variable that it intends to predict. This allows us to incorporate assumptions that have been established in other models on the formation of inflation expectations, such as in Carroll (2003). An alternative interpretation of the fact that households may use other information than actual inflation to update their forecasts to, is that they do not have the same concept of “inflation” in mind as for example economists might have when asked to provide their projections. That is, they may not forecast the change in the overall price level but some other variable that they define as inflation, such as price changes of frequently bought items. This then implies that households also update their forecasts to their perceived concept of inflation.¹¹

While we introduce flexibility in the updating process, we will, in contrast to the previous sections, keep the underlying economic law of motion for the inflation rate constant. In particular, we employ the concept of learning based on previous errors as outlined in Equation (12) because this law of motion provides the lowest MSEs for four out of six countries. However, it is also necessary to allow for this time lag in the reduced form, particularly when we intend to incorporate the idea of the epidemiological model. The sequence of events for all three following updating rules is similar and conceptually represented by the model presented in the lower section of Figure B.3 in the Appendix.

5.2.1 Learning based on professional forecasts

In this section, we introduce a learning mechanism that combines the assumption of adaptive learning with sticky information as suggested by Pfajfar and Santoro (2010). The SG algorithm in this context takes the following form

$$\hat{\theta}_t = \hat{\theta}_{t-1} + \gamma \mathbf{x}_t (\pi_{t+12|t}^F - \mathbf{x}_t' \hat{\theta}_{t-1}), \quad (15)$$

while LS learning is characterized by

$$\begin{aligned} \hat{\theta}_t &= \hat{\theta}_{t-1} + \gamma \mathbf{R}_t^{-1} \mathbf{x}_t (\pi_{t+12|t}^F - \mathbf{x}_t' \hat{\theta}_{t-1}), \\ \mathbf{R}_t &= \mathbf{R}_{t-1} + \gamma (\mathbf{x}_t \mathbf{x}_t' - \mathbf{R}_{t-1}), \end{aligned} \quad (16)$$

where $\pi_{t+12|t}^F$ is the professional inflation forecast made in period t for period $t + 12$. This type of updating corresponds with the idea reflected in the epidemiological approach brought forward by Carroll (2003). The approach is based on the assumption that households receive new information in the form of professional forecasts from various sources such as newspapers and TV. This information slowly disseminates among households, who then update their forecasts according to the experts’ outlook. Equations (15) and (16) imply exactly this idea, namely that the household updates the coefficient parameter based on the forecast error between its previous period’s forecast and the experts’ outlook on future price developments, which is subsequently used to form the conditional forecast $\pi_{t+12|t} = \mathbf{x}_t' \hat{\theta}_{t-1}$. As a measure for professional forecasts, we use monthly figures of inflation expectations published by Consensus Economics. These are provided for all countries of interest with a forecast horizon directed to the end of the current as well as the end of the following year. To transform them into forecasts with a constant one year horizon, we follow the methodology outlined in Siklos (2013) which provides us with weighted averages of the released forecasts. Even though it seems unlikely that the expectation figures are communicated to the public in this specific form, it is reasonable to assume that they capture the overall notion of different outlooks that are reported in the news coming from various forecasting institutions.

¹¹Following this interpretation implies that households are indifferent about the forecast horizon, that is one or 12 periods ahead, when they do their one step prediction.

5.2.2 Learning based on future information about inflation

In line with the previous assumption that households are more forward-looking in their updating process than commonly presumed, we now test the empirical validity of the idea that households possess information about future inflation that is better than the one obtained from professional forecasters. This information, as it is hard to exactly measure, is proxied by the actual future inflation rate π_{t+12} , that is, by the rate which the household intends to forecast. Pfajfar and Santoro (2010) argue that the well documented phenomenon of “herding” among professional forecasters (see, e.g., Bewley and Fiebig, 2002 or Pons-Novell, 2003) may create a situation where the forecasters’ outlook does not really reflect the true information content that can be extracted from the macroeconomic variables but rather the professionals’ notion not to deviate from their peers. Thus, households may potentially receive some other more precise information, in this case approximated by π_{t+12} , which they use to update their forecasts to.

To picture the situation of learning based on future information about inflation, we consider the following forward-looking SG learning algorithm

$$\hat{\theta}_t = \hat{\theta}_{t-1} + \gamma \mathbf{x}_t (\pi_{t+12} - \mathbf{x}_t' \hat{\theta}_{t-1}), \quad (17)$$

while LS learning is represented as

$$\begin{aligned} \hat{\theta}_t &= \hat{\theta}_{t-1} + \gamma \mathbf{R}_t^{-1} \mathbf{x}_t (\pi_{t+12} - \mathbf{x}_t' \hat{\theta}_{t-1}), \\ \mathbf{R}_t &= \mathbf{R}_{t-1} + \gamma (\mathbf{x}_t \mathbf{x}_t' - \mathbf{R}_{t-1}). \end{aligned} \quad (18)$$

The conditional forecasts associated with both learning algorithms are the same as the one outlined in Section 5.2.1.

5.2.3 Learning based on energy prices

The final algorithm we introduce constitutes an alternative to the flexible updating mechanisms suggested by Pfajfar and Santoro (2010). It is based on the idea that households, when asked to report their inflation forecasts, do not entirely think of the economic definition of inflation, but rather look at the price changes of products that they are mostly confronted with, such as energy prices. This may be due to the fact that changes in those prices are more obvious to consumers and therefore also easier for them to access. In fact, Lynch and Trehan (2013) point out that household inflation expectations in the US, UK, and Japan appear to be highly sensitive to changes in the level of the oil price. Bruine de Bruin et al. (2010) furthermore provide evidence that consumers, when asked about expected changes of “prices in general”, often think about prices for food and gasoline. In the following, we test the empirical validity of this assumption, using the energy price component of the national CPIs to represent frequently purchased items.¹² In doing so we assume that households, after having conducted their inflation forecast, receive information on current energy prices which they perceive as a good indicator for future inflation and therefore incorporate this information in their own forecasts.¹³

Based on the same law of motion as in Sections 5.2.1 and 5.2.2, the learning algorithms are

$$\hat{\theta}_t = \hat{\theta}_{t-1} + \gamma \mathbf{x}_t (\pi_t^E - \mathbf{x}_t' \hat{\theta}_{t-1}), \quad (19)$$

¹²We also regard food items as an important CPI component for households. As those time series however are not available for Belgium and Spain over the entire sample period, we refrain from reporting the results.

¹³We furthermore test whether using frequently purchased items as explanatory variable improves the fit of the learning models. This however is not the case.

in case of SG learning and

$$\begin{aligned}\hat{\boldsymbol{\theta}}_t &= \hat{\boldsymbol{\theta}}_{t-1} + \gamma \mathbf{R}_t^{-1} \mathbf{x}_t (\pi_t^E - \mathbf{x}_t' \hat{\boldsymbol{\theta}}_{t-1}), \\ \mathbf{R}_t &= \mathbf{R}_{t-1} + \gamma (\mathbf{x}_t \mathbf{x}_t' - \mathbf{R}_{t-1}),\end{aligned}\tag{20}$$

for LS learning. Since energy prices are quite volatile, we define π_t^E as the six-months moving average of the monthly growth rates of the CPI's energy price component.¹⁴ Conditional forecasts again are formed in the same way as outlined in Section 5.2.1. This particular updating and forecasting mechanism corresponds with the one used by Weber (2010) under the assumption that households have a law of motion in mind equivalent to Equation (8) with the dependent variable being π_t^E .

5.2.4 Estimation results for learning based on different updating schemes

Table 5.2 extends those findings presented in Section 5.1.3 on different laws of motion by evidence obtained for various updating schemes. The upper panel shows the optimal gain parameters and MSEs for tracking observed household forecasts, while results presented in the lower panel are based on the true inflation rate. Beyond the definitions introduced in Table 5.1, we indicate updating with respect to professional forecasts by PF, with respect to future information proxied by the actual future inflation rate by FI, and finally with respect to energy prices by EN.

Focusing first on the upper panel, we find that the choice of the optimal updating scheme turns out to be very dependent on the country. That is, there is no single algorithm which resembles the expectation formation process best throughout the euro area. While Italian households seem to possess information on future price changes that is best proxied by the actual future rate itself, Dutch and German consumers update their beliefs to news on inflation that they receive from professional forecasters. Only French consumers appear to focus on price changes of frequently purchased items in the form of energy prices. The latter finding is rather surprising given the fact, that households quite often report that this corresponds with their definition of inflation. However, this may be a particular characteristic of US households and different for consumers of the euro area.

The optimal gain parameters range between 0.0018 for Spain up to 0.0654 for the Netherlands, both for SG learning. Here again, no clear distinction can be made between northern and southern European countries in terms of the height of the gain parameters. The lowest MSEs are found in four out of six cases for SG based learning mechanisms, which seems plausible as the algorithm is computationally simpler and thus more appealing to associate it with household learning. However, we do not obtain any statistical evidence that SG algorithms outperform LS based one, in fact we rather find the opposite.

To assess which algorithm tracks the evolution of the actual inflation rates best, we turn to the results presented in the lower panel of Table 5.2. Before we proceed with the interpretation of the results, note that proxying better information by future inflation is appealing if one aims to infer whether households in fact possess knowledge that is superior to the one of forecasters. However, when judging which algorithm is best to forecast actual inflation, it might be misleading to consider this updating mechanism as it has to yield the lowest MSE by definition. Therefore, we allow two types of algorithms to be considered optimal: First, we identify the FI algorithm that yields the lowest MSE, that is either LS or SG. Second, we select the algorithm which turns out to perform

¹⁴We also tried the annual change or different moving average horizons but found this to provide the lowest MSE on average.

Table 5.2: Estimates based on all learning rules

The table reports optimal gain parameters together with the associated MSEs in parentheses for SG and LS learning. Gain parameters and MSEs are depicted for each learning algorithm assuming that households consider an economic law of motion based on current errors (CE) or based on previous errors (PE) as well as that they update their forecasts to professional forecasts (PF), to future inflation (FI) as well as energy prices (EN). For each country, the combination of gain parameter and forecasting algorithm that yields the lowest MSE is highlighted. * indicates the algorithm that is more accurate than its counterpart (either LS or SG) according to the modified Diebold-Mariano two-sided test at the 20% significance level.

	SG-CE	LS-CE	SG-PE	LS-PE	SG-PF	LS-PF	SG-FI	LS-FI	SG-EN	LS-EN
Household forecast										
Belgium	0.0031 (0.3408)	0.0098 (0.3568)	0.0012 (0.4130)	0.0021* (0.3840)	0.0443 (0.6237)	0.0023 (0.3778)	0.0017 (0.3474)	0.0036 (0.3434)	0.0003 (0.7410)	0.0004 (0.6692)
France	0.0038 (0.4094)	0.0127 (0.4539)	0.0018 (0.4136)	0.0035 (0.4086)	0.0018 (0.3978)	0.0035 (0.4130)	0.0026 (0.3839)	0.0063 (0.3762)	0.0008 (0.3671)	0.0012* (0.3211)
Germany	0.0170 (0.4229)	0.0134* (0.3935)	0.0067 (0.4166)	0.0104 (0.4183)	0.0823 (0.3928)	0.0305 (0.3665)	0.0030 (0.4839)	0.0071 (0.4433)	0.0020 (0.9014)	0.0066* (0.7936)
Italy	0.0011 (0.4483)	0.0045 (0.4514)	0.0004 (0.3173)	0.0007 (0.3137)	0.0005 (0.3042)	0.0010 (0.3042)	0.0063 (0.2891)	0.0013 (0.3018)	0.0003 (0.4147)	0.0004* (0.3611)
Netherlands	0.0087 (0.3906)	0.0197* (0.3201)	0.0018 (0.3587)	0.0034* (0.3145)	0.0654 (0.2629)	0.0049 (0.3011)	0.0025 (0.3251)	0.0060 (0.3001)	0.0005 (0.8876)	0.0008* (0.7449)
Spain	0.0013 (1.1286)	0.0081 (1.1973)	0.0018 (0.9498)	0.0010 (1.1043)	0.0006 (1.1112)	0.0018 (1.1362)	0.0009 (0.9916)	0.0028 (1.0739)	0.0004 (1.2719)	0.0009 (1.2014)
	SG-CE	LS-CE	SG-PE	LS-PE	SG-PF	LS-PF	SG-FI	LS-FI	SG-EN	LS-EN
CPI inflation										
Belgium	0.0029 (2.1211)	0.0075 (1.9514)	0.0011 (3.0129)	0.0156* (2.2355)	0.0543 (2.2105)	0.0300* (1.9325)	0.0630 (0.6526)	0.4252 (0.2801)	0.0002 (3.7727)	0.0003 (3.6849)
France	0.0040 (0.7002)	0.0094 (0.6791)	0.0016 (1.0645)	0.0144* (0.9272)	0.1513 (0.8951)	0.2795* (0.8755)	0.1501 (0.1646)	0.3916* (0.0911)	0.0005 (1.4917)	0.0008 (1.4043)
Germany	0.0024 (0.5595)	0.0073 (0.5574)	0.0013 (0.9096)	0.0176* (0.6946)	0.0796 (0.6751)	0.1847 (0.6910)	0.0759 (0.1730)	0.4880 (0.1348)	0.0018 (1.2481)	0.0047 (1.1918)
Italy	0.0013 (0.7932)	0.0057 (0.7255)	0.0005 (1.2482)	0.0009* (1.2175)	0.0525 (1.0366)	0.3313 (1.0153)	0.0665 (0.1277)	0.5006 (0.0592)	0.0003 (1.4285)	0.0005* (1.3217)
Netherlands	0.0073 (0.7948)	0.0161 (0.8036)	0.0022 (0.9786)	0.0116 (0.8354)	0.0968 (0.7118)	0.2843 (0.7016)	0.1031 (0.1588)	0.3605 (0.1286)	0.0005 (1.6747)	0.0009 (1.4779)
Spain	0.0013 (2.1940)	0.0065 (1.9421)	0.0005 (2.7557)	0.0155 (2.4294)	0.0520 (2.0093)	0.0560 (1.9669)	0.0540 (0.3294)	0.4757* (0.1722)	0.0004 (3.0762)	0.0009* (2.7723)

best neglecting the ones that update to the actual inflation rate.

Results presented in Table 5.2 are quite clear in that they highlight the superiority of LS learning in forecasting future inflation regardless of the updating scheme employed. Abstracting from FI based learning, we see that the LS algorithm in association with current forecast errors provides the lowest MSE for France, Germany, Italy, and Spain, while inflation in Belgium and the Netherlands is better forecast when professional predictions on the inflation rate are incorporated. In addition, the optimal gains for algorithms that update to future information in the form of professional forecasts or the actual future inflation rate are a lot higher than those that do not incorporate future information. As they, on average, display lower MSEs than the other algorithms, this implies that the optimal weights put on most recent observations are rather low (which is the case for a high γ) and thus relatively short periods of observations should be included when forecasting changes in the actual price level.

5.3 Which information is relevant for households when forecasting inflation?

Based on the assumption that past inflation is the only explanatory variable, our findings so far suggest that the learning algorithms considered by European households are quite heterogeneous, while no clear support could be found for either SG or LS learning. Even though the idea that households use simple univariate models to forecast inflation seems appealing given the associated costs of collecting data, we next pursue the question whether models that incorporate additional explanatory variables deliver a more accurate description of household inflation expectations. To do so, we re-estimate previous algorithms for each country, this time however based on a wider range of information sets as outlined in Section 4.3. More specifically, we augment the vector of explanatory variables consisting of actual CPI inflation and a constant (set 1) by the annual growth in industrial production (set 2) or instead by the annual change in the short-term interest rate (set 3) or, alternatively, by both, industrial production growth as well as interest rate changes (set 4). We then test which one of those four information sets performs best in replicating household forecasts. This exercise also represents a robustness check of whether the chosen model for each country remains optimal when additional information is included. The results are summarized in the left panels of Table 5.3. The table displays the optimal gain for each country together with the respective learning algorithm, information set, and MSE. More detailed results providing information on all algorithms and information sets are presented in Tables A.1–A.6 in the Appendix.

A comparison of the findings shown in Table 5.3 and 5.2 reveals that the previously selected algorithms remain the best choice for all countries but the Netherlands. From the type of updating that turns out to be optimal, we can infer that in four out of six countries, consumers incorporate some sort of forward-looking information in their expectation formation process. Moreover, we see that French, German, Italian, and Spanish households appear to conduct their forecasts based on simple univariate models consisting only of the past inflation rate. Including an additional interest rate term however improves the resemblance of our simulated series with the forecasts conducted by Belgium and Dutch households. The magnitude of the γ^{HF} parameter estimates shown in Table 5.3 furthermore suggests that French consumers assign the lowest weight to the previous forecast error, while German households are most responsive to new incoming data. To avoid that this result is driven by data sensitivity of the gain parameter for SG learning, we compute the average optimal γ^{HF} over all LS based algorithms for each country. The result is presented in the right column of Table 5.3. We still observe the highest gain parameter for Germany, while the lowest average LS based gain is found for Italy. Again, this result somewhat contradicts the evidence provided by Weber (2010).

Table 5.3: Summary of the optimal gain, model, and information set for each country

Column 2 - 5 of the table report the optimal gain parameters to track inflation expectations together with the associated MSEs, learning algorithms and information sets between 1991 and 2013 for different euro area countries. The sixth column displays the average optimal gain parameters calculated over all optimal LS based gains.

	γ^{HF}	MSE	Learning algorithm	Information set	\emptyset opt. γ^{HF} (LS)
Belgium	0.0030	0.2828	SG-CE	$\mathbf{x}_t = (1, \pi_{t-1}, i_{t-1})'$	0.0035
France	0.0012	0.3211	LS-EN	$\mathbf{x}_t = (1, \pi_{t-1})'$	0.0055
Germany	0.0305	0.3665	LS-PF	$\mathbf{x}_t = (1, \pi_{t-1})'$	0.0125
Italy	0.0063	0.2891	SG-FI	$\mathbf{x}_t = (1, \pi_{t-1})'$	0.0016
Netherlands	0.0062	0.2406	LS-FI	$\mathbf{x}_t = (1, \pi_{t-1}, i_{t-1})'$	0.0135
Spain	0.0018	0.9498	SG-PE	$\mathbf{x}_t = (1, \pi_{t-1})'$	0.0029


Turning to the individual results on the country level presented in Tables A.1 - A.6, we see that Dutch consumers are the only ones to employ the algorithm which in fact also yields the best prediction of the actual inflation rate. The finding corresponds with our initial observation derived from Table 2.1 that Dutch households are most successful in forecasting inflation compared to consumers in other countries. In addition, we can infer from the MSEs that adaptive learning describes the expectation formation of Dutch consumers best, while it is least suited to track the forecasts of Spanish households.

Results reported in the lower panel of the individual country tables suggest that with respect to the true forecasting exercise, adaptive learning techniques perform best for Germany and again worst for Spain. Inference here is drawn neglecting the updating mechanisms that incorporate the actual future inflation rates because this (precise) information, as outlined above, is not available at the time of the forecast. By comparing the magnitude of the optimal gain parameters for tracking household expectations with the ones obtained from the actual inflation forecast, we do not find that the latter are systematically higher, a result which has been brought forward by Weber (2010). It turns out that this only holds when we focus on learning based on future information of inflation and partially also for learning based on professional forecasts. Overall, our results suggest that the tracking ability of our learning algorithms for household expectations are much better than their forecasting performance.


Figure 5.1 displays individually for each country the simulated household forecasts based on the optimal model, gain, and information set together with the survey expectations and the actual inflation rates. Due to the fact that we need six years for the initialization of the algorithms, the time series are plotted from the beginning of 1997 onward. The graph underlines the good tracking ability of the algorithm for consumer inflation expectations in the Netherlands and its comparably poor fit for Spanish expectations. In addition, the graphical depiction suggests that there is no necessary connection between the height of the gain parameter as displayed in Table 5.3 and the level or volatility of the actual inflation rate. While the optimal γ^{HF} is the highest for Germany, which consequently renders the parameter update most responsive to new incoming data, we find that the inflation level has been relatively low over the entire sample period. The opposite however is found for Spain, where we observe a rather volatile and high inflation rate yet a comparably low gain parameter. The same holds if we compare the highest gain parameter solely for LS based algorithms, thus excluding potential biases arising from the sensitivity of the SG based gain to data scale.

5.4 Which type of algorithm performs better?


To draw a systematic comparison between the performance of SG and LS learning in forecasting inflation as well as in tracking the formation of household expectations, we summarize the comparative outcomes for each country over all algorithms and information sets available in Table 5.4 and furthermore report how often the finding has been statistically significant. While previous results suggested that SG learning provides the optimal learning mechanism to describe Belgian, Italian, and Spanish consumer forecasts for our sample of observations, evidence from a systematic comparison of all learning algorithms clearly indicates that LS learning (statistically) outperforms its SG counterpart. The result is even more pronounced for the actual forecasts of CPI inflation. Here, we do not obtain any statistical support for the assumption that SG learning is superior to LS based algorithms, while the opposite is found in 20% to 40% of all cases depending on the country. In that, our results are more straightforward than evidence provided by Berardi and Galimberti (2014) for the US. The authors conclude that the LS algorithm is superior to SG learning when




(a) Belgium




(b) France




(c) Germany



(d) Italy



(e) Netherlands



(f) Spain

Figure 5.1: Simulated household expectations based on optimal models and gain parameters
The figure displays the simulated household forecasts based on the optimal models and gain parameters (black solid line) together with the actual household inflation expectations (blue dashed line) both depicted 12 months after the forecasts together with CPI inflation (green dotted line).

Table 5.4: Comparison between LS and SG based learning algorithms

The table reports how often LS learning outperforms SG learning (and vice versa) in tracking household forecasts as well as in forecasting inflation for each country. In addition, the table reports how often the result has been found with statistical significance based on the modified Diebold-Mariano (DM) two-sided test at the 20% level.

	LS outperforms SG	DM 20% level	SG outperforms LS	DM 20% level
Household forecasts				
Belgium	80%	15%	20%	0%
France	75%	5%	25%	5%
Germany	85%	25%	15%	0%
Italy	50%	20%	50%	0%
Netherlands	90%	40%	10%	0%
Spain	60%	5%	40%	10%
	LS outperforms SG	DM 20% level	SG outperforms LS	DM 20% level
CPI inflation				
Belgium	100%	30%	0%	0%
France	95%	35%	5%	0%
Germany	85%	30%	15%	0%
Italy	95%	25%	5%	0%
Netherlands	85%	20%	15%	0%
Spain	100%	40%	0%	0%

evaluated with respect to its resemblance to observations from the survey of professional forecasters but not when intending to forecast actual inflation.¹⁵

6 Optimal gains and information sets over time

The simulated series in Figure 5.1 are based on the optimal gains which are estimated over the entire observation period. While those gain parameters provide a good impression of the *average* difference between countries, it forecloses the option that households change their information sets or gain parameters when learning over time.¹⁶ However this particular property is highly relevant in the context of our study. During major events such as crises or the introduction of a new currency, consumers might become more aware of changes in their environment and thus put higher weight on new data or start to consider different information sets in their forecasts. In this section, we therefore intend to further explore this particular idea by allowing for changes in the gain parameter and information set over the 22 years of observation. At the same time, the exercise also constitutes a robustness check whether our optimal gains and therefore also the associated models are sensitive to changes in the estimation period.

¹⁵When we evaluate forecast accuracy, we compare our results to the findings in Berardi and Galimberti (2014) where the gain is taken as a consciously chosen parameter. When we test the resemblance to observable forecasts, we compare our results to the ones in Berardi and Galimberti (2014) where the gain is interpreted as a primitive parameter.

¹⁶Galimberti (2013) points to the fact that when interpreting the gain as a primitive parameter, it is still possible to allow for changes in its level over time.

6.1 Time variation in the gain parameter


In contrast to our previous analysis, we estimate the gain parameter γ over a rolling window of 60 months, which is equivalent to the number of periods chosen in Berardi and Galimberti (2014). Given the fact that we need 72 months for the initialization, 145 periods of observations are left to determine a potential time variation in γ^{HF} , from 12/2000 onward until 12/2012. In contrast to previous expositions, survey expectations are now shown at the time of the forecast to facilitate their interpretation. Figure 6.1 displays for each country individually the corresponding time-varying gains (solid line for gains computed based on information set 1 and cross line for those computed relying on information set 3) together with the associated 95% confidence bands (dashed line).¹⁷ Gains pictured in t have been estimated from $t - 59$ until t . For the purpose of comparison, we also display the time-constant γ^{HF} computed over the entire observation period as shown in Table 5.3. The optimal algorithms and information sets underlying the gain estimation in this exercise are exactly the ones summarized in Section 5.3.

The variation in the scale of the gain parameter is rather low over the entire sample period in particular for France (0.001 - 0.0031), Spain (0.0003 - 0.0049), and Belgium (0.0025 - 0.0077), while the highest fluctuation can be observed in Germany (0.0049 - 0.1303). For German households, we identify a remarkable increase in the gain parameter when γ^{HF} is estimated over the period beginning with the establishment of the ECB up until right after the euro cash changeover. This supports the idea that households were aware of the institutional changes and thus more alert to new incoming data. In contrast, we observe a drop in the parameters for Spain and Italy when the estimation period encompasses solely the time after the establishment of the ECB. However, given the scale of the parameters, the drop is relatively low compared to the variation observed for Germany.


Due to the scale sensitivity of γ in the case of SG learning as emphasized by Evans et al. (2010) and Berardi and Galimberti (2014), the graphical expositions in Figure 6.1 do not allow for a meaningful comparison of the gain parameters among countries as they are computed for LS and SG algorithms depending on the optimal type of algorithm of each euro area member state. To obtain more insights into the time-varying updating behavior of households based on a data invariant parameter and to draw a comparison of our gains to the ones presented in Weber (2010), we redo the exercise, this time however focusing on the time variation solely of the most optimal LS algorithms. Therefore, we compute new series of gain parameters for the optimal LS based learning rules in the case of Belgium, Italy, and Spain. These are the countries that originally displayed the lowest MSEs for SG learning.¹⁸ The results are depicted in Figure B.1 in the Appendix. The graphical exposition allows to benchmark our results to the ones presented in Weber (2010). She selects the optimal gains associated with LS learning algorithms over the period beginning in 1990 until mid 1998, which corresponds most closely with the beginning of our series displayed in Figure B.1. A comparison between the first parameters throughout our sample and the ones reported by Weber (2010), reveals that our parameters exceed her estimates in the case of Germany and the Netherlands, while they are lower for Italy and Spain. Specifically, her gains are 0.0051 for France, 0.001 for Germany, 0.0024 for Italy, 0.001 for the Netherlands, and 0.053 for Spain, while our results are 0.001, 0.0142, 0.001, 0.0071, and 0.0003 for the same order of countries. The difference in outcome between the two studies may be an indicator that the level of the gain is quite sensitive to the initialization method and the underlying learning rule.

¹⁷We calculate the confidence bands using a sandwich estimate of the variance matrix for extremum estimators as described in Cameron and Trivedi (2005).


¹⁸The optimal LS rules are LS-CE for Belgium, LS-PF for Italy, and LS-FI for Spain.




(a) Belgium




(b) France




(c) Germany



(d) Italy



(e) Netherlands



(f) Spain

Figure 6.1: Time-varying parameter estimation with constant information sets

The figure displays the time-varying gain parameters estimated individually for each country over a rolling window of 60 subsequent periods based on information set 1 (red solid line) or information set 3 (brown cross line) together with the 95% confidence bands (dashed line). The time-constant gain parameters used for calculations in Section 5 are displayed by the horizontal line for comparison. Parameters are shown at the date of the forecast.


For Spain and Italy, we furthermore observe a relatively large variation in the gain parameter when LS learning is considered. This finding additionally emphasizes the importance of the underlying sample period for the gain estimation. It also supports the assumption that households were aware of the institutional changes through the creation of the common monetary union as already identified for Germany. Moreover, a similarity between the learning behavior of Spanish and German households becomes apparent. Consumers have downweighted observations more rapidly at the beginning of the EMU as well as during the currency changeover which indicates that those events have in fact influenced the learning behavior. In addition, we do not find any support for the assumption that consumers have used higher gains during the recent financial turmoil.

6.2 Time variation in the gain parameter and the information set


Next, we investigate to what extent the optimal gain parameters depend on the underlying information set and if the information set considered by the consumer changes in times that are characterized by pronounced variation in the economic environment. We conduct the gain selection in a similar manner as in Section 6.1, this time however we keep only the learning algorithm constant. That is, we select the combination of gain parameter and information set that most closely replicates the household's inflation forecast at time t . Here again, we consider a rolling window of 60 periods and keep more than 12 years to investigate potential changes in γ^{HF} . Results of this exercise are shown individually for each country in Figure 6.2 involving both, LS and SG learning. We display the optimal gain which has been calculated based on information set 1 ($1, \pi_{t-1}$) by a solid line, based on information set 2 ($1, \pi_{t-1}, y_{t-1}$) by circles, based on information set 3 ($1, \pi_{t-1}, i_{t-1}$) by a cross line, and finally, based on information set 4 ($1, \pi_{t-1}, y_{t-1}, i_{t-1}$) by triangles. Vertical lines indicate those periods when the optimal information set has changed.

For France, we can observe with one exception that the information set consisting of past inflation remains optimal over the entire sample period. While we detect several changes in the optimal information sets employed by German and Italian consumers, those changes do not seem to affect the level of the gain parameter nor the uncertainty associated with its estimation. However, variations in the information sets are remarkable for Belgium, Spain, and the Netherlands and they also have a considerable impact on the level of the gain parameters. Moreover, we do not find any support for the assumption that measures for economic activity have been relevant for the inflation forecast during the latest financial turmoil. Overall, we only obtain little evidence that economic activity has had an effect on the formation of inflation expectations. Agents seem to favor rather simple models based on past inflation and short-term interest rates.


To allow for a meaningful comparison, we also conduct this exercise considering only the optimal LS based learning rules. Results are displayed in Figure B.2 in the Appendix. Again, we observe more variation in the gain parameters over time compared to the case where we also consider SG learning. This however does not seem to be due to the variation in the information sets but rather to the shorter estimation periods as the results are quantitatively not very different from the ones presented in Figure B.1 where we held the information set constant. Thus, considering different time periods for the gain estimation seems to affect the level of the gain parameters in some countries, while allowing for different information sets does not appear to be too much of a concern.




(a) Belgium




(b) France




(c) Germany



(d) Italy



(e) Netherlands



(f) Spain

Figure 6.2: Time-varying parameter estimation with time-varying information sets

The figure displays the time-varying gain parameters estimated individually for each country over a rolling window of 60 subsequent periods based on the time-varying optimal information set. Gains are calculated either based on information set 1 (red solid line), or information set 2 (purple circle line), or information set 3 (brown cross line), or information set 4 (orange triangle line) and are displayed together with the 95% confidence bands (dashed line). Vertical lines indicate a change in the optimal information set. Parameters are shown at the date of the forecast.

7 Concluding remarks

In this paper, we empirically assess the ability of several constant gain learning rules to replicate household inflation expectations in six key countries of the euro area, that is, Belgium, France, Germany, Italy, the Netherlands, and Spain. The question of how agents form their expectations is particularly relevant for price stability oriented central banks such as the ECB. Obtaining a more precise understanding of this process allows policymakers to better design policies geared towards medium- to long-term price stability.

For all countries we present evidence that learning rules are a good description of the expectation formation process. The learning algorithms we consider are modifications of those that are traditionally used in the learning literature. They are either more backward-looking or, alternatively, allow for different interpretations of the inflation term through the introduction of more flexibility in the updating process. We identify a pronounced heterogeneity in the optimal learning mechanisms on the country level throughout the euro area. In particular, we find evidence that households in Germany, Italy and the Netherlands incorporate some type of forward-looking information in their expectations, either in the form of professional forecasts or some other variable which is best approximated by the future inflation rate. This finding is in line with results presented by Doepke et al. (2008) and supports the epidemiological model brought forward by Carroll (2003). The prominent role of future information therefore highlights the necessity of ensuring that expectations of professional forecasters are well anchored as they feed in the expectation formation process of private consumers. Households in other countries such as France rely more on energy price information. This highlights the importance for central banks to also analyze and closely monitor the development of price components beyond the ones contained in core inflation. High energy price inflation can result in high inflation expectations and thereby affect core inflation.

A comparison of the optimal gain parameters on the national level does not reveal a systematic difference between northern and southern European countries. Thus, our findings do not support the evidence brought forward by Weber (2010) that consumers which have experienced high levels of inflation use higher constant gains than households which have not experienced pronounced increases in the price level. We also do not observe a systematic relationship between the magnitude of the gain parameters resulting from a forecasting exercise of the actual inflation rate and the optimal gains for tracking the household survey data. However, our findings do highlight the sensitivity of the gain parameters to the underlying estimation period, method, and learning rule which complicates a meaningful comparison with other studies.

In addition, we conduct a systematic comparison between LS and SG based learning algorithms with respect to their ability to replicate consumer forecasts and their ability to forecast future inflation. This is particularly relevant for modeling inflation expectations by means of adaptive learning which has become a popular way particularly in large macroeconomic models. Our evidence suggests that the LS version, which is widely used in applied work, significantly outperforms its SG counterpart. The finding applies to both, tracking the survey measures as well as forecasting changes in the price level.

From those findings, several extensions emerge for future research. First of all, as more and more quantitative survey data becomes available for households in different European countries (such as the data employed in Easaw et al., 2013), it would be of great interest to fit the discussed algorithms to those expectation series. This micro data would also allow to further explore the question to which extend demographic differences account for heterogeneity in the learning mechanisms as

proposed for example by Malmendier and Nagel (2015) or Pfajfar and Santoro (2008). Second, as we find evidence that households in some countries possess information on future inflation beyond the one coming from professional forecasters, it would be interesting to further explore the nature of this information. Finally, allowing for switches in the underlying algorithms over time, would enable the policymaker to gain more insights in the relationship between consumer learning and changes in the economic environment.

References

- Batchelor, R. A. and Orr, A. B. (1988). Inflation expectations revisited. *Economica*, 55(219):317–331.
- Berardi, M. and Galimberti, J. K. (2012). 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, The University of Manchester.
- Berardi, M. and Galimberti, J. K. (2013). A note on exact correspondences between adaptive learning algorithms and the Kalman filter. *Economics Letters*, 118(1):139–142.
- Berardi, M. and Galimberti, J. K. (2014). A note on the representative adaptive learning algorithm. *Economics Letters*, 124(1):104–107.
- Berk, J. M. (1999). Measuring inflation expectations: A survey data approach. *Applied Economics*, 31(11):1467–1480.
- Bernanke, B. S. (2007). Inflation expectations and inflation forecasting. Speech presented at the Monetary Economics Workshop of the National Bureau of Economic Research Summer Institute, Cambridge, MA <http://www.federalreserve.gov/newsevents/speech/bernanke20070710a.htm>, Accessed: 2015-02-21.
- Bewley, R. and Fiebig, D. G. (2002). On the herding instinct of interest rate forecasters. *Empirical Economics*, 27(3):403–425.
- Branch, W. A. and Evans, G. W. (2006). A simple recursive forecasting model. *Economics Letters*, 91(2):158–166.
- Bruine de Bruin, W., Potter, S., Rich, R., Topa, G., and Van der Klaauw, W. (2010). Improving survey measures of household inflation expectations. *Current Issues in Economics and Finance*, 16(7):1–7.
- Cameron, A. C. and Trivedi, P. K. (2005). *Microeconometrics: Methods and applications*. Cambridge University Press.
- Carceles-Poveda, E. and Giannitsarou, C. (2007). Adaptive learning in practice. *Journal of Economic Dynamics and Control*, 31(8):2659–2697.
- Carlson, J. A. and Parkin, J. M. (1975). Inflation expectations. *Economica*, 42(166):123–38.
- Carroll, C. D. (2003). Macroeconomic expectations of households and professional forecasters. *The Quarterly Journal of Economics*, 118(1):269–298.
- Clark, T. E. and Davig, T. (2008). An empirical assessment of the relationships among inflation and short- and long-term expectations. Research Working Paper 08-05, Federal Reserve Bank of Kansas City.
- Cukierman, A. and Wachtel, P. (1979). Differential inflationary expectations and the variability of the rate of inflation: Theory and evidence. *American Economic Review*, 69(4):595–609.
- Dias, F., Duarte, C., and Rua, A. (2010). Inflation expectations in the euro area: Are consumers rational? *Review of World Economics*, 146(3):591–607.

- Diebold, F. X. and Mariano, R. S. (2002). Comparing predictive accuracy. *Journal of Business & Economic Statistics*, 20(1):134–144.
- Doepke, J., Doern, J., Fritsche, U., and Slacalek, J. (2008). The dynamics of European inflation expectations. *The BE Journal of Macroeconomics*, 8(1):1–23.
- Easaw, J., Golinelli, R., and Malgarini, M. (2013). What determines households inflation expectations? Theory and evidence from a household survey. *European Economic Review*, 61:1–13.
- European Commission (2007). *The Joint Harmonised EU Programme of Business and Consumer Surveys - User guide*. http://ec.europa.eu/economy_finance/db_indicators/surveys/method_guides/index_en.htm, Accessed: 2013-12-20.
- Evans, G. W. and Honkapohja, S. (2001). *Learning and expectations in macroeconomics*. Princeton University Press.
- Evans, G. W., Honkapohja, S., and Williams, N. (2010). Generalized stochastic gradient learning. *International Economic Review*, 51(1):237–262.
- Forsells, M. and Kenny, G. (2002). The rationality of consumers’ inflation expectations: Survey-based evidence for the euro area. ECB Working Paper 163, European Central Bank.
- Galimberti, J. K. (2013). *Adaptive learning for applied macroeconomics*. PhD thesis. <https://www.escholar.manchester.ac.uk/uk-ac-man-scw:213190>, Accessed: 2015-02-21.
- Giannone, D., Lenza, M., Momferatou, D., and Onorante, L. (2014). Short-term inflation projections: A Bayesian vector autoregressive approach. *International Journal of Forecasting*, 30(3):635–644.
- Harvey, D., Leybourne, S., and Newbold, P. (1997). Testing the equality of prediction mean squared errors. *International Journal of Forecasting*, 13(2):281–291.
- Heinzel, S. and Wollmershäuser, T. (2005). Quantifying inflation expectations with the Carlson-Parkin method: A survey-based determination of the just noticeable difference. *Journal of Business Cycle Measurement and Analysis*, 2(3):321–352.
- Lynch, M. and Trehan, B. (2013). Household inflation expectations and the price of oil: It’s a déjà vu all over again. *FRBSF Economic Letter*, 16:1–4.
- Malmendier, U. and Nagel, S. (2015). Learning from inflation experiences. *Quarterly Journal of Economics*, forthcoming.
- Mankiw, N. G., Reis, R., and Wolfers, J. (2004). Disagreement about inflation expectations. In Gertler, M. and Rogoff, K., editors, *NBER Macroeconomics Annual 2003*. The MIT Press.
- Markiewicz, A. and Pick, A. (2014). Adaptive learning and survey data. *Journal of Economic Behavior & Organization*, 107:685–707.
- Milani, F. (2007). 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. *Economic Journal*, 121(552):379–401.

- Nardo, M. (2003). The quantification of qualitative survey data: As critical assessment. *Journal of Economic Surveys*, 17(5):645–668.
- Orphanides, A. and Williams, J. (2004). Imperfect knowledge, inflation expectations, and monetary policy. In Bernanke, B. S. and Woodford, M., editors, *The inflation-targeting debate*. University of Chicago Press.
- Orphanides, A. and Williams, J. C. (2005). 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 Santoro, E. (2008). Asymmetries in inflation expectation formation across demographic groups. Cambridge Working Papers in Economics 824, University of Cambridge.
- Pfajfar, D. and Santoro, E. (2010). Heterogeneity, learning and information stickiness in inflation expectations. *Journal of Economic Behaviour & Organization*, 75(3):426–444.
- Pons-Novell, J. (2003). Strategic bias, herding behaviour and economic forecasts. *Journal of Forecasting*, 22(1):67–77.
- Siklos, P. L. (2013). Sources of disagreement in inflation forecasts: An international empirical investigation. *Journal of International Economics*, 90(1):218–231.
- Weber, A. (2010). Heterogeneous expectations, learning and European inflation dynamics. In Cobham, D., Eitrheim, Ø., Gerlach, S., and Qvigstad, J. F., editors, *Twenty years of inflation targeting: Lessons learned and future prospects*.

Appendix A Tables

Table A.1: Optimal gains and MSEs for Belgium

The table reports optimal gain parameters together with the associated MSEs in parentheses for SG and LS learning in Belgium. The first four columns assume that learning is based on current forecasts errors (CE) or, alternatively, on previous forecast errors (PE), and errors are updated to current inflation. The remaining columns provide results for updating with respect to professional forecasts (PF), future inflation (FI), and energy prices (EN) whereas learning takes place based on previous forecast errors. The combination of gain parameter and forecasting algorithm that yields the lowest MSE is highlighted. * indicates the algorithm that is more accurate than its counterpart (either LS or SG) according to the modified Diebold-Mariano two-sided test at the 20% significance level.

	SG-CE	LS-CE	SG-PE	LS-PE	SG-PF	LS-PF	SG-FI	LS-FI	SG-EN	LS-EN
Household forecast										
Set 1	0.0031 (0.3408)	0.0098 (0.3568)	0.0012 (0.4130)	0.0021* (0.3840)	0.0443 (0.6237)	0.0023 (0.3778)	0.0017 (0.3474)	0.0036 (0.3434)	0.0003 (0.7410)	0.0004 (0.6692)
Set 2	0.0021 (0.5846)	0.0081 (0.4873)	0.0005 (1.4181)	0.0020* (0.5909)	0.0012 (0.4765)	0.0020 (0.4523)	0.0010 (0.6011)	0.0034* (0.4454)	0.0001 (2.0802)	0.0001 (2.0735)
Set 3	0.0030 (0.2828)	0.0089 (0.2839)	0.0012 (0.4290)	0.0021 (0.3880)	0.0013 (0.3661)	0.0025 (0.3664)	0.0017 (0.3638)	0.0032 (0.3768)	0.0003 (0.7689)	0.0004 (0.6976)
Set 4	0.0020 (0.5733)	0.0082 (0.3976)	0.0005 (1.4061)	0.0020* (0.5726)	0.0012 (0.4632)	0.0020 (0.4507)	0.0010 (0.6280)	0.0032* (0.4560)	0.0001 (2.0600)	0.0001 (2.0527)
	SG-CE	LS-CE	SG-PE	LS-PE	SG-PF	LS-PF	SG-FI	LS-FI	SG-EN	LS-EN
CPI inflation										
Set 1	0.0029 (2.1211)	0.0075 (1.9514)	0.0011 (3.0129)	0.0156* (2.2355)	0.0543 (2.2105)	0.0300* (1.9325)	0.0630 (0.6526)	0.4252 (0.2801)	0.0002 (3.7727)	0.0003 (3.6849)
Set 2	0.0018 (2.9203)	0.0075* (1.8843)	0.0003 (3.7527)	0.0214 (2.1027)	0.0118 (2.3121)	0.0320* (1.9704)	0.0085 (1.2151)	0.2501 (0.3894)	0.0001 (3.9575)	0.0001 (3.9218)
Set 3	0.0026 (2.2712)	0.0063 (2.1578)	0.0011 (3.0804)	0.0146* (2.1219)	0.0372 (2.2425)	0.0256* (1.9739)	0.0740 (0.4101)	0.3480 (0.2844)	0.0002 (3.8537)	0.0003 (3.7346)
Set 4	0.0017 (3.0072)	0.0063 (2.0781)	0.0003 (3.8047)	0.0144 (2.3308)	0.0117 (2.3079)	0.0269* (2.0189)	0.0097 (0.9041)	0.3629 (0.4006)	0.0001 (3.9654)	0.0001 (3.9296)

Table A.2: Optimal gains and MSEs for France

The table reports optimal gain parameters together with the associated MSEs in parentheses for SG and LS learning in France. For further explanations see notes in Table A.1.

	SG-CE	LS-CE	SG-PE	LS-PE	SG-PF	LS-PF	SG-FI	LS-FI	SG-EN	LS-EN
Household forecast										
Set 1	0.0038 (0.4094)	0.0127 (0.4539)	0.0018 (0.4136)	0.0035 (0.4086)	0.0018 (0.3978)	0.0035 (0.4130)	0.0026 (0.3839)	0.0063 (0.3762)	0.0008 (0.3671)	0.0012* (0.3211)
Set 2	0.0027 (0.8137)	0.0086 (0.5313)	0.0021 (0.7333)	0.0033 (0.5426)	0.0019 (0.4449)	0.0036 (0.4314)	0.0027 (0.4261)	0.0069 (0.4111)	0.0002 (2.1405)	0.0006 (1.8724)
Set 3	0.0037 (0.4079)	0.0098 (0.4747)	0.0017 (0.4186)	0.0033 (0.4166)	0.0019* (0.4031)	0.0036 (0.4250)	0.0026 (0.3681)	0.0065 (0.3581)	0.0007 (0.4887)	0.0010 (0.5213)
Set 4	0.0026 (0.7663)	0.0085 (0.5345)	0.0021 (0.7464)	0.0030 (0.5795)	0.0019 (0.4515)	0.0037 (0.4426)	0.0027 (0.4172)	0.0072 (0.3916)	0.0002 (2.1916)	0.0005 (1.9573)
	SG-CE	LS-CE	SG-PE	LS-PE	SG-PF	LS-PF	SG-FI	LS-FI	SG-EN	LS-EN
CPI inflation										
Set 1	0.0040 (0.7002)	0.0094 (0.6791)	0.0016 (1.0645)	0.0144* (0.9272)	0.1513 (0.8951)	0.2795* (0.8755)	0.1501 (0.1646)	0.3916* (0.0911)	0.0005 (1.4917)	0.0008 (1.4043)
Set 2	0.0026 (1.2024)	0.0080 (0.7463)	0.0023 (1.4100)	0.0121* (1.1156)	0.0018 (1.0089)	0.1660 (0.8547)	0.0106 (0.7007)	0.3027 (0.1184)	0.0001 (2.4682)	0.0002 (2.4262)
Set 3	0.0037 (0.6990)	0.0084 (0.7073)	0.0015 (1.1000)	0.0028* (1.0491)	0.1046 (0.8938)	0.1373 (0.8876)	0.1090 (0.1881)	0.3098* (0.0907)	0.0005 (1.5991)	0.0007 (1.5743)
Set 4	0.0025 (1.1402)	0.0079 (0.7477)	0.0023 (1.4461)	0.0029 (1.2231)	0.0018 (0.9894)	0.1511 (0.8509)	0.0112 (0.6670)	0.2884* (0.1035)	0.0001 (2.5090)	0.0002 (2.4809)

Table A.3: Optimal gains and MSEs for Germany

The table reports optimal gain parameters together with the associated MSEs in parentheses for SG and LS learning in Germany. For further explanations see notes in Table A.1.

	SG-CE	LS-CE	SG-PE	LS-PE	SG-PF	LS-PF	SG-FI	LS-FI	SG-EN	LS-EN
Set 1	0.0170 (0.4229)	0.0134* (0.3935)	0.0067 (0.4166)	0.0104 (0.4183)	0.0823 (0.3928)	0.0305 (0.3665)	0.0030 (0.4839)	0.0071 (0.4433)	0.0020 (0.9014)	0.0066* (0.7936)
Set 2	0.0015 (1.0528)	0.0072 (0.4997)	0.0018 (0.8752)	0.0027 (0.5726)	0.0036 (0.4768)	0.1441* (0.4033)	0.0020 (0.6899)	0.0101 (0.6300)	0.0012 (1.0413)	0.0056* (0.7630)
Set 3	0.0027 (0.4072)	0.0089 (0.3857)	0.0056 (0.4845)	0.0019 (0.5417)	0.0792 (0.4038)	0.0232 (0.3865)	0.0018 (0.5587)	0.0038 (0.5244)	0.0022 (0.8631)	0.0057 (0.8796)
Set 4	0.0015 (1.0814)	0.0067 (0.5229)	0.0015 (0.7299)	0.0025 (0.6045)	0.0029 (0.4861)	0.0300* (0.4091)	0.0017 (0.7123)	0.0116 (0.6837)	0.0012 (1.0422)	0.0034 (0.8800)
	SG-CE	LS-CE	SG-PE	LS-PE	SG-PF	LS-PF	SG-FI	LS-FI	SG-EN	LS-EN
CPI inflation										
Set 1	0.0024 (0.5595)	0.0073 (0.5574)	0.0013 (0.9096)	0.0176* (0.6946)	0.0796 (0.6751)	0.1847 (0.6910)	0.0759 (0.1730)	0.4880 (0.1348)	0.0018 (1.2481)	0.0047 (1.1918)
Set 2	0.0015 (1.0790)	0.0065 (0.5660)	0.0004 (1.7021)	0.0017* (0.9418)	0.0013 (0.7648)	0.0461 (0.6196)	0.0049 (0.6101)	0.2846* (0.1791)	0.0009 (1.5383)	0.0031* (1.1922)
Set 3	0.0023 (0.5490)	0.0069 (0.5654)	0.0009 (1.0104)	0.0190 (0.9221)	0.0780 (0.6606)	0.0584 (0.6168)	0.0733 (0.1988)	0.2814 (0.1421)	0.0018 (1.2725)	0.0035 (1.3385)
Set 4	0.0016 (0.8655)	0.0063 (0.6197)	0.0017 (1.5929)	0.0016* (0.9911)	0.0013 (0.7573)	0.0531 (0.6095)	0.0046 (0.7858)	0.2189* (0.1631)	0.0008 (1.5733)	0.0019 (1.3761)

Table A.4: Optimal gains and MSEs for Italy

The table reports optimal gain parameters together with the associated MSEs in parentheses for SG and LS learning in Italy. For further explanations see notes in Table A.1.

	SG-CE	LS-CE	SG-PE	LS-PE	SG-PF	LS-PF	SG-FI	LS-FI	SG-EN	LS-EN
Household forecast										
Set 1	0.0011 (0.4483)	0.0045 (0.4514)	0.0004 (0.3173)	0.0007 (0.3137)	0.0005 (0.3042)	0.0010 (0.3042)	0.0063 (0.2891)	0.0013 (0.3018)	0.0003 (0.4147)	0.0004* (0.3611)
Set 2	0.0010 (0.6881)	0.0042 (0.5559)	0.0008 (0.7446)	0.0007 (0.6976)	0.0005 (0.2930)	0.0010 (0.2984)	0.0007 (0.3476)	0.0014 (0.3980)	0.0002 (1.2566)	0.0004* (1.0200)
Set 3	0.0011 (0.4471)	0.0046 (0.4394)	0.0004 (0.3146)	0.0007 (0.3154)	0.0005 (0.3180)	0.0010 (0.3192)	0.0006 (0.3191)	0.0012 (0.3221)	0.0003 (0.4201)	0.0004* (0.3839)
Set 4	0.0010 (0.6552)	0.0041 (0.5509)	0.0008 (0.7376)	0.0007 (0.7079)	0.0005 (0.3051)	0.0010 (0.3103)	0.0006 (0.3702)	0.0014 (0.4160)	0.0002 (1.2924)	0.0003* (1.0636)
	SG-CE	LS-CE	SG-PE	LS-PE	SG-PF	LS-PF	SG-FI	LS-FI	SG-EN	LS-EN
CPI inflation										
Set 1	0.0013 (0.7932)	0.0057 (0.7255)	0.0005 (1.2482)	0.0009* (1.2175)	0.0525 (1.0366)	0.3313 (1.0153)	0.0665 (0.1277)	0.5006 (0.0592)	0.0003 (1.4285)	0.0005* (1.3217)
Set 2	0.0010 (1.2654)	0.0051 (0.8455)	0.0004 (2.1679)	0.0155 (1.3359)	0.0007 (1.1615)	0.1356 (1.0200)	0.0062 (0.6733)	0.3722* (0.0603)	0.0002 (1.8261)	0.0005 (1.4896)
Set 3	0.0013 (0.7927)	0.0057* (0.7187)	0.0005 (1.2615)	0.0009 (1.2288)	0.0502 (1.0032)	0.1816 (1.0623)	0.0452 (0.3618)	0.3190* (0.0634)	0.0003 (1.4464)	0.0005 (1.3718)
Set 4	0.0010 (1.2110)	0.0051 (0.8282)	0.0004 (2.1618)	0.0010* (1.5869)	0.0006 (1.1380)	0.1031 (1.0659)	0.0069 (0.6164)	0.3049* (0.0686)	0.0002 (1.8849)	0.0005 (1.6058)

Table A.5: Optimal gains and MSEs for the Netherlands

The table reports optimal gain parameters together with the associated MSEs in parentheses for SG and LS learning for the Netherlands. For further explanations see notes in Table A.1.


	SG-CE	LS-CE	SG-PE	LS-PE	SG-PF	LS-PF	SG-FI	LS-FI	SG-EN	LS-EN
Household forecast										
Set 1	0.0087 (0.3906)	0.0197* (0.3201)	0.0018 (0.3587)	0.0034* (0.3145)	0.0654 (0.2629)	0.0049 (0.3011)	0.0025 (0.3251)	0.0060 (0.3001)	0.0005 (0.8876)	0.0008* (0.7449)
Set 2	0.0036 (0.4785)	0.0205* (0.2943)	0.0017 (0.5468)	0.0050 (0.3435)	0.0106 (0.2776)	0.0062 (0.2727)	0.0025 (0.4596)	0.0068* (0.2933)	0.0002 (2.2703)	0.0004 (1.9756)
Set 3	0.0093 (0.3047)	0.0137 (0.3342)	0.0018 (0.2986)	0.0032* (0.2585)	0.0607 (0.2586)	0.0368 (0.2434)	0.0025 (0.2846)	0.0062* (0.2406)	0.0005 (0.8580)	0.0007 (0.8129)
Set 4	0.0035 (0.4465)	0.0152 (0.3186)	0.0018 (0.5352)	0.0038 (0.3823)	0.0096 (0.2742)	0.0337* (0.2549)	0.0024 (0.4507)	0.0064* (0.2924)	0.0002 (2.3683)	0.0003 (2.1629)
	SG-CE	LS-CE	SG-PE	LS-PE	SG-PF	LS-PF	SG-FI	LS-FI	SG-EN	LS-EN
CPI inflation										
Set 1	0.0073 (0.7948)	0.0161 (0.8036)	0.0022 (0.9786)	0.0116 (0.8354)	0.0968 (0.7118)	0.2843 (0.7016)	0.1031 (0.1588)	0.3605 (0.1286)	0.0005 (1.6747)	0.0009 (1.4779)
Set 2	0.0036 (0.9486)	0.0167* (0.7685)	0.0016 (1.1903)	0.0075* (0.7294)	0.0032 (0.6921)	0.0114 (0.6512)	0.0155 (0.3620)	0.3101* (0.1577)	0.0002 (3.0164)	0.0006 (2.4475)
Set 3	0.0145 (0.6517)	0.0120 (0.8873)	0.0022 (0.9279)	0.0045 (0.8500)	0.0846 (0.6397)	0.0771 (0.6548)	0.0960 (0.1597)	0.2628 (0.1395)	0.0005 (1.6741)	0.0007 (1.6229)
Set 4	0.0037 (0.9063)	0.0129 (0.8437)	0.0016 (1.1816)	0.0045 (0.7969)	0.0038 (0.6800)	0.0571 (0.6451)	0.0141 (0.3688)	0.2310* (0.1644)	0.0002 (3.1404)	0.0004 (2.7580)

Table A.6: Optimal gains and MSEs for Spain


The table reports optimal gain parameters together with the associated MSEs in parentheses for SG and LS learning for Spain. For further explanations see notes in Table A.1.

	SG-CE	LS-CE	SG-PE	LS-PE	SG-PF	LS-PF	SG-FI	LS-FI	SG-EN	LS-EN
Household forecast										
Set 1	0.0013 (1.1286)	0.0081 (1.1973)	0.0018 (0.9498)	0.0010 (1.1043)	0.0006 (1.1112)	0.0018 (1.1362)	0.0009 (0.9916)	0.0028 (1.0739)	0.0004 (1.2719)	0.0009 (1.2014)
Set 2	0.0010 (2.5831)	0.0059 (1.4329)	0.0007 (1.6746)	0.0012 (1.4175)	0.0007 (1.2406)	0.0021* (1.1526)	0.0007 (1.6370)	0.0025 (1.4027)	0.0002 (3.1005)	0.0008 (1.7499)
Set 3	0.0012 (1.1518)	0.0072 (1.2413)	0.0005* (1.0306)	0.0010 (1.0908)	0.0007 (1.1167)	0.0020 (1.1441)	0.0021* (0.9649)	0.0030 (1.0747)	0.0004 (1.2343)	0.0009 (1.1448)
Set 4	0.0010 (2.1293)	0.0061 (1.3757)	0.0006 (1.7991)	0.0010 (1.4880)	0.0007 (1.2440)	0.0022 (1.1359)	0.0007 (1.6259)	0.0027 (1.3393)	0.0002 (3.0794)	0.0008 (1.7334)
	SG-CE	LS-CE	SG-PE	LS-PE	SG-PF	LS-PF	SG-FI	LS-FI	SG-EN	LS-EN
CPI inflation										
Set 1	0.0013 (2.1940)	0.0065 (1.9421)	0.0005 (2.7557)	0.0155 (2.4294)	0.0520 (2.0093)	0.0560 (1.9669)	0.0540 (0.3294)	0.4757* (0.1722)	0.0004 (3.0762)	0.0009* (2.7723)
Set 2	0.0011 (3.1617)	0.0051 (2.3429)	0.0004 (3.2644)	0.0009 (2.7926)	0.0054 (2.3351)	0.0609* (1.9080)	0.0047 (1.6633)	0.2662* (0.2331)	0.0002 (4.3391)	0.0006* (3.5528)
Set 3	0.0013 (2.1857)	0.0059 (2.0042)	0.0004 (2.7770)	0.0008 (2.6873)	0.0479 (2.0447)	0.0482* (2.0022)	0.0443 (0.5879)	0.4083 (0.1840)	0.0004 (2.9739)	0.0009 (2.7076)
Set 4	0.0010 (2.9854)	0.0056 (2.0874)	0.0004 (3.3274)	0.0007 (2.8502)	0.0006 (2.3102)	0.0510 (1.9536)	0.0047 (1.5651)	0.2091* (0.2377)	0.0002 (4.2416)	0.0006* (3.4979)


Appendix B Figures




(a) Belgium




(b) France




(c) Germany



(d) Italy




(e) Netherlands




(f) Spain

Figure B.1: Time-varying parameter estimation with constant information sets for optimal LS based algorithms


The figure displays the time-varying gain parameters estimated individually for each country over a rolling window of 60 subsequent periods for the optimal LS based algorithm. Gains are computed either based on information set 1 (red solid line), information set 2 (purple square line), or information set 3 (brown cross line) and are displayed together with the 95% confidence bands (dashed line). The time-constant gain parameters used for calculations in Section 5 are indicated by the horizontal line for comparison. Parameters are shown at the date of the forecast.




(a) Belgium




(b) France




(c) Germany



(d) Italy



(e) Netherlands



(f) Spain

Figure B.2: Time-varying parameter estimation with time-varying information sets for optimal LS based algorithms

The figure displays the time-varying gain parameters estimated individually for each country over a rolling window of 60 subsequent periods based on the time-varying optimal LS based information set. Gains are calculated either based on information set 1 (red solid line), or information set 2 (purple circle line), or information set 3 (brown cross line), or information set 4 (orange triangle line) and are displayed together with the 95% confidence bands (dashed line). Vertical lines indicate a change in the optimal information set. Parameters are shown at the date of the forecast.

Previous DNB Working Papers in 2015

- No. 454 **Mauro Mastrogiacomo and Rob Alessie**, Where are the retirement savings of self-employed? An analysis of 'unconventional' retirement accounts
- No. 455 **Clemens Bonner and Paul Hilbers**, Global liquidity regulation - Why did it take so long?
- No. 456 **Leo de Haan, Jan Willem van den End and Philip Vermeulen**, Lenders on the storm of wholesale funding shocks: Saved by the central bank?
- No. 457 **Wilko Bolt and David Humphrey**, Assessing bank competition for consumer loans
- No. 458 **Robert Beyer and Michael Stemmer**, From progress to nightmare - European regional unemployment over time
- No. 459 **Stijn Claessens and Neeltje van Horen**, The impact of the global financial crisis on banking globalization
- No. 460 **Francisco Blasques, Falk Bräuning and Iman van Lelyveld**, A dynamic network model of the unsecured interbank lending market
- No. 461 **Carin van der Crujsen, Lola Hernandez and Nicole Jonker**, In love with the debit card but still married to cash
- No. 462 **Koen van der Veer**, Loss shocks and the quantity and price of private export credit insurance: Evidence from a global insurer
- No. 463 **Cenkhan Sahin and Jakob de Haan**, Market reactions to the ECB's Comprehensive Assessment
- No. 464 **Melanie de Waal, Floor Rink and Janka Stoker**, How internal and external supervisors influence employees' self-serving decisions
- No. 465 **Andrea Colciago and Lorenza Rossi**, Firms entry, oligopolistic competition and labor market dynamics
- No. 466 **Peter Heemeijer and Ronald Heijmans**, Central bank intervention in large value payment systems: An experimental approach
- No. 467 **Natalya Martynova**, Effect of bank capital requirements on economic growth: a survey
- No. 468 **Melle Bijlsma and Robert Vermeulen**, Insurance companies' trading behaviour during the European sovereign debt crisis: Flight home or flight to quality?
- No. 469 **Robert Vermeulen, Marco Hoeberichts, Bořek Vašíček, Diana Žigraiová, Kateřina Šmídková and Jakob de Haan**, Financial stress indices and financial crises
- No. 470 **Nicole Jonker, Mirjam Plooi and Johan Verburg**, Does a public campaign influence debit card usage? Evidence from the Netherlands
- No. 471 **Carin van der Crujsen en Mirjam Plooi**, Changing payment patterns at point-of-sale: their drivers

