

Adaptive Learning and Survey Data

Agnieszka Markiewicz
Erasmus University Rotterdam

Andreas Pick
Erasmus University Rotterdam
and De Nederlandsche Bank

February 28, 2013

Abstract

This paper investigates the ability of the adaptive learning approach to replicate the expectations of professional forecasters. For a range of macroeconomic and financial variables, we compare constant and decreasing gain learning models to simple, yet powerful benchmark models. We find that both, constant and decreasing gain models, provide a good fit for the expectations of professional forecasters for a range of variables. These results suggest that, instead of relying only on the the most recent observation, agents use more complex models to form their expectations even for financial variables where random walk forecasts are often difficult to beat.

Keywords: expectations, survey of professional forecasters, adaptive learning, bounded rationality

JEL Codes: E37, E44, G14, G15

1 Introduction

Expectations are a key ingredient of many economic and financial models. They reflect behavior of agents and influence economic outcomes. The macroeconomic and finance literature usually endows agents with rational expectations. Although unlikely in reality, they present some advantages over bounded rationality. Under rational expectations, agents' subjective probability distribution coincides with the true distribution of the model economy. This model consistency makes rational expectations unique to the model. In contrast, there is an infinite number of non-rational ways to form expectations.

One widely cited alternative to rational expectations has been suggested by Bray (1982), Bray and Savin (1986), Marcet and Sargent (1989) and Sargent (1993). This approach assumes that agents, modeled by economists, have at most the same knowledge as these economists themselves and hence they behave as econometricians. This approach, called adaptive learning, presents an advantage of imposing modeling discipline. Similar to rational expectations, adaptive learning generates forecasts which are optimal given agents' information at the time.¹

While an extensive finance and macroeconomic literature has been using adaptive learning to explain observed dynamics of financial and macroeconomic variables, to our best knowledge, no study systematically tests for empirical validity of this approach.²

This paper attempts to fill this gap. Specifically, we investigate in how far agents' expectations can be approximated by adaptive learning. For this purpose, we use a set of financial and macroeconomic survey data and fit adaptive learning laws of motions. The formal test of this approach is conducted by assessing the out-of-sample forecasting performance of such estimated models.

¹By optimal we mean here that the forecasts are orthogonal to forecast errors.

²There is an extensive literature that uses adaptive learning to understand numerous financial and macroeconomic unexplained stylized facts. Here, we mention only a few of studies. Adam et al. (2009) and Evans and Branch (2010), among others, show that adaptive learning helps in replicating a variety of asset pricing puzzles. Milani (2011) and Preston and Eusepi (2013) argue that learning dynamics can explain business cycle fluctuations.

We assume that economic agents use simple time series models to make forecasts and they update their parameters using recursive least squares. Both, empirical and experimental evidence motivates the use of simple forecasting algorithms that require very little computational or informational resources. In forecasting applications difficulties in observing the driving variable, together with parameter instability could lead to large and persistent forecast errors. Numerous empirical studies have demonstrated that fundamentals-based models can frequently underperform forecasts derived from simple univariate time series models, such as a random walk, or AR and MA models.³ A further reason for agents to use simple forecasting models is the cost related to gathering and processing information. Experimental evidence collected for instance by Adam (2007) indicates that agents use very simple forecasting models, thereby conditioning forecasts only on one explanatory variable, even if information on other relevant variables is available to them.

Accordingly, in this paper, we consider simple time-series models. These are represented by autoregressive specifications updated by recursive least squares with constant and decreasing gains, adaptive expectations and the random walk model. Additionally, for each variable we use macroeconomic and financial regressors. These specifications represent more structural approach.

We first estimate parameters in an initial training sample and assess their forecasting accuracy out-of-sample. We also conduct forecasting accuracy tests against the simple, yet powerful benchmark of the random walk. We find that estimated gain parameters are of the range 0.01 to 0.15 across variables and specifications, which suggests that agents use a relatively long history of observations to make their forecasts. The comparison between constant gain and decreasing gain specifications leaves no clear winner. Constant gain models perform better for inflation, the unemployment rate and the yen-dollar exchange rate but decreasing gain specifications are

³For instance Orphanides and Van Norden (2005) and Stock and Watson (2007) show that univariate times-series models provide better forecasts for inflation. Meese and Rogoff (1983) and Cheung, Chinn, and Pascual (2005) demonstrate forecasting superiority of random walk over fundamentals-based models for exchange rates.

better representations of expectations for interest rates, GDP growth and the pound-dollar exchange rate. Out-of-sample Diebold-Mariano tests show that for most variables at least one model performs significantly better than a naïve random walk forecast. This suggests that, instead of relying only on the last available observation, agents use more complex models to form their expectations.

The remainder of the paper is organized as follows. We start with a review of the related literature in Section 2. The models of expectations and their estimation are described in Section 3. Section 4 presents the results of in-sample estimation and out-of-sample forecasts. Finally, Section 5 concludes.

2 Related literature

A large body of literature in macroeconomics and finance employed survey data to examine expectation formation by economic agents. This literature can be broadly divided in two strands: studies testing rational expectations hypothesis and studies proposing alternative modeling approaches to match survey expectations.⁴ The first branch consists of a group of studies seeking to verify hypothesis of rationality for inflation survey expectations, including Bonham and Dacy (1991), Bonham and Cohen (2001), Croushore (1997), and Evans and Gulamani (1984). This body of research largely documents the failure of the rational expectations hypothesis for inflation expectations.⁵

Similar results have been found using survey expectations of foreign exchange market traders by Frankel and Froot (1987b, 1990a) and Ito (1990). Froot (1989), Friedman (1990), and Jeong and Maddala (1996) also reject hypothesis of rationality for interest rate forecasts. More recently, Bacchetta, Mertens and van Winkoop (2009) investigate the link between the predictability of excess returns and expectational errors in a set of financial markets, using survey data. They find predictability of excess returns and expectational errors in foreign exchange, stock and bond markets.

⁴Clearly some of the studies do both: they test for rationality and propose alternative modeling approaches if the latter is rejected.

⁵An important exception is work by Keane and Runkle (1990), which fails to reject the rational expectations hypothesis.

As survey data reject rational forecasts hypothesis for a large number of markets, new ways of modeling agents' beliefs have been proposed. Allen and Taylor (1990), Ito (1990), and Frankel and Froot (1987a, 1990a, and 1990b) argue that the dynamics of survey expectations of foreign exchange traders display behavioral rather than rational features. Specifically, they found forecasts heterogeneity across individuals and over time. Ito (1990) argues that, in line with behavioral bias of 'wishful thinking', exporters tend to anticipate a currency depreciation while importers anticipate an appreciation. Frankel and Froot (1987a, 1990a, and 1990b), Taylor and Allen (1992) show that at short horizons traders tend to use extrapolative chartist rules, whilst at longer horizons they tend to use mean reverting rules based on fundamentals. The time-varying heterogeneity of expectations has been modeled by Brock and Hommes (1997) within a framework of an adaptively rational equilibrium dynamics (ARED). Branch (2004) used the ARED to test whether inflation survey data exhibited rationally heterogeneous expectations. Similarly Jongen, Verschoor, Wolff, and Zwinkels (2012) explained dispersion in exchange rate survey data using ARED framework.

The average dynamics displayed by economic agents' beliefs have been often modeled by adaptive learning mechanism, which presumes that agents behave as econometricians in the sense that they estimate model parameters using econometric techniques (for example, Evans and Honkapohja, 2001). Numerous applications and extensions of a standard adaptive learning have been suggested by recent literature. Initially, recursive least squares parameter learning, as proposed by Bray (1982) and Marcet and Sargent (1989) presumed small departures from rationality by assuming that economic agents knew the model economy and updated only its parameters. Additional extensions of this approach include constant gain learning occasionally leading to escape dynamics (Cho, Williams and Sargent, 2002).

This paper is closely related to the work by Branch and Evans (2006). Within a VAR framework and using survey data, they estimate various laws of motion for expectation formation. They find that the constant gain learning rule fits best in sample and performs well in the out-of-sample exercise. While Branch and Evans (2006) studied the survey data for GDP

growth and inflation, our aim is to cover a large set of variables and markets to provide a systematic overview of agents' beliefs. Furthermore, we use real time data, which represent the information set available to the professional forecasters.

3 Modeling expectations

Rational forecasts are built on strong assumptions about the representative agent's information set. We relax these informational assumptions by postulating that, instead of using a structural model based forecasts, agents use simple time-series models to make predictions. This type of forecasting rule embodies very low computational cost as it is based on the past values of observed variables only. Agents are assumed to update parameters of their forecasting rules using recursive least squares.

3.1 Recursive least squares

The forecasting rule is governed by adaptive learning where agents behave as econometricians and for a linear model $y_t = \boldsymbol{\theta}_t' \mathbf{x}_t + \varepsilon_t$ they update model parameters using recursive least squares,

$$\begin{aligned}\boldsymbol{\theta}_t &= \boldsymbol{\theta}_{t-1} + \gamma_t \mathbf{R}_t^{-1} \mathbf{x}_t (y_t - \boldsymbol{\theta}_{t-1}' \mathbf{x}_t), \quad t = 1, 2, \dots, T, T + 1, \\ \mathbf{R}_t &= \mathbf{R}_{t-1} + \gamma_t (\mathbf{x}_t \mathbf{x}_t' - \mathbf{R}_{t-1})\end{aligned}\quad (1)$$

where $\boldsymbol{\theta}_t$ is a $k \times 1$ vector of parameters, \mathbf{x}_t is the $k \times 1$ vector of regressors, and γ_t is the gain parameter.

The gain parameter γ_t can take two forms. First, when $\gamma_t = 1/t$, the estimator converges to the OLS estimator. The finite-sample difference between (1) and the OLS estimator is determined by the initialization of \mathbf{R}_0 , which we set to $\mathbf{0}$ to obtain OLS forecasts from expanding windows. Under parameter stability OLS forecasts are optimal in the mean square forecast error (MSFE) sense. Since the weight associated with an observation decreases with the amount of available data, this learning scheme is sometimes referred to as "decreasing gain" learning. Second, when the parameter γ_t is set to a constant, $\gamma_t = c$ with $c \in (0, 1]$, observations receive decaying

weights that decrease with the distance from the most recent observation. This updating rule is known as a “constant gain learning” or “perpetual learning” and the estimator is referred to as “constant gain least squares” (CGLS). CGLS can be shown to deliver good forecasts under parameter instability. In a simplified model (adaptive expectations, discussed below), it can be shown to be the optimal forecasting strategy in the MSFE sense when the model’s parameter evolves as a random walk (Pesaran, Pick and Pranovich, 2013).

The vector or parameters $\boldsymbol{\theta}_t^{(i)}$ is used to make forecast for the next period:

$$\hat{y}_{t+1}^{(i)} = \boldsymbol{\theta}_t^{(i)'} \mathbf{x}_{t+1}$$

where the superscript $i \in \{\text{ls}, \text{cg}\}$ denotes the particular forecasting scheme, and $\hat{y}_{t+1}^{(\text{ls})}$ is the forecast using a decreasing gain and $\hat{y}_{t+1}^{(\text{cg})}$ that using a constant gain.

Below we use a range of regressor sets. First, we experiment with autoregressive specifications, in particular, AR(1) to AR(4). Additionally, we include one or more variables that may improve forecasts with details provided below. The final regressor set we consider contains only the intercept. When using only an intercept, that is, $x_t = 1, \forall t$, the expectation of the dependent variable is the intercept, $\hat{y}_t = \theta_t$. Setting $R_0 = 1$, it is easy to see that CGLS reduces to what is called adaptive expectations or exponential smoothing:

$$\hat{y}_{t+1}^{(a)} = \gamma y_t + (1 - \gamma) \hat{y}_t^{(a)}. \quad (2)$$

where superscript (a) denotes the forecast from adaptive expectations. As already mentioned, this specification has optimality properties under certain data generating processes. Furthermore, in practice, it has been shown to have good forecasting properties for a wide range of variables (Hyndman, Koehler, Ord and Snyder, 2008, Pesaran and Pick, 2008, Pesaran, Pick and Pranovich, 2013).

Under decreasing gain learning and setting $y_0 = 0$, adaptive expectations yield the simple mean forecast

$$\hat{y}_t = t^{-1} \sum_{i=1}^t y_i$$

Finally, when setting the downweighting parameter, γ , to unity, adaptive expectations (2) yield the random walk forecast

$$\hat{y}_{t+1}^{(\text{rw})} = y_t. \quad (3)$$

When the random walk is differenced, the optimal forecast is

$$\Delta \hat{y}_{t+1}^{(\text{rw})} = 0. \quad (4)$$

Whether (3) or (4) delivers the better forecast will depend on the nature of the variable under consideration.

3.2 Estimation of the model parameters

The specifications of the constant gain least square forecasting rule and the adaptive expectations forecasting rule require estimates of the gain parameter, γ , and initialization of the time varying parameters. We seek to identify the law of motion that agents use to make their forecasts, which are represented by survey data. Specifically, we assume that agents employ the observed data to make their forecasts. Accordingly, the parameter of the forecasting rule, γ , is estimated by minimizing the mean square error between the forecasts resulting from the gain parameters and the survey forecasts:

$$\hat{\gamma}^{(i)} = \arg \min_{\delta \in (0,1]} \left[\frac{1}{T_{\text{pre}} - T_{\text{init}}} \sum_{t=T_{\text{init}}+1}^{T_{\text{pre}}} \left(\hat{y}_t^{(i)} - y_t^{(\text{spf})} \right)^2 \right], \quad (5)$$

where superscript i denotes the different forecast specifications, and $y_t^{(\text{spf})}$ is the forecasts provided by professional forecasters. We use the forecasts for observations after T_{init} and up to T_{pre} to determine the gain parameters, where we set $T_{\text{init}} = 20$ and $T_{\text{pre}} = 60$.

The first T_{init} observations are used to estimate the initial values for the time varying parameters, $\boldsymbol{\theta}_0$ of CGLS, where we use reverse downweighting:

$$\boldsymbol{\theta}_0 = \left(\sum_{t=1}^{T_{\text{init}}} (1 - \kappa)^{t-1} \mathbf{x}_t \mathbf{x}_t' \right) \sum_{t=1}^{T_{\text{init}}} (1 - \kappa)^{t-1} \mathbf{x}_t y_t.$$

We use this reverse downweighting in order to reduce the dependence of the starting value on the later observations in the initialization sample, and we

use $\kappa = 0.05$. We set $\mathbf{R}_0 = \mathbf{I}_k$ for the CGLS forecasts. For the decreasing gain forecasts we estimate parameters by expanding window OLS, which implies $\mathbf{R}_0 = \mathbf{0}$. Finally, we start with $\hat{y}_a^{(a)} = y_1$ for the adaptive expectations forecast.

We use the observations after T_{init} until the end of the sample to construct out-of-sample forecasts, which we will use to evaluate the different learning models against the median survey of professional forecasters. In order to evaluate the forecasts, we use the mean square forecast error between the median forecasts in the survey of professional forecasters and the forecasts we obtain from each model

$$\text{MSFE} = \frac{1}{T - T_{\text{init}}} \sum_{t=T_{\text{init}}}^T \left(e_t^{(i)} \right)^2 \quad (6)$$

where

$$e_t^{(i)} = \left(\hat{y}_t^{(i)} - y_t^{(\text{spf})} \right)$$

Additionally, we test for predictive accuracy using the Diebold and Mariano (1995) test statistic for the loss differential

$$l(i) = \left(e_t^{(\text{rw})} \right)^2 - \left(e_t^{(i)} \right)^2$$

where $e_t^{(\text{rw})}$ is the forecast error of the random walk specification.

3.3 Forecasting models

The forecasts from recursive least squares use a number of different models. The first ones are pure autoregressive models, where we have $\mathbf{x}_t = (1, y_{t-1})$ for AR(1) model and up to $\mathbf{x}_t = (1, y_{t-1}, y_{t-2}, y_{t-3}, y_{t-4})$ for the AR(4) model. We also use a range of structural variables in the forecasts. When forecasting the inflation rate we use the unemployment rate as a predictor, following the literature on the Phillips curve (for example, Atkeson and Ohanian, 2001). We add the ten year government bond yield to the model for the real GDP growth following the literature on the predictive power of asset prices for the business cycle (Stock and Watson, 2003). In the regressions for the interest rates we incorporate the unemployment rate and the inflation rate in the spirit of Taylor rule forecasts. For the corporate

Table 1: Summary of the data

Variable	Explanatory variables	Freq.	Sample	1st fore	no.fore
Inflation	Unemployment	Q	1981.1–2012.3	1996.4	64
3 month T-bill yield	Unemployment, inflation	Q	1981.1–2012.3	1996.4	64
10 year gov't bond yield	Unemployment, inflation	Q	1992.1–2012.3	2007.2	22
Corporate bond yield	Unemployment, inflation, corporate profits	Q	1981.1–2012.3	1996.4	64
GDP growth	10 year gov't bond rate	Q	1970.2–2012.3	1985.2	110
Unemployment	GDP growth rate	Q	1969.1–2012.3	1984.1	115
EUR/USD	Inflation differential, interest differential	M	1996.4–2009.11	2001.4	104
JPY/USD	Inflation differential, interest differential	M	1996.4–2009.11	2001.4	104
GBP/USD	Inflation differential, interest differential	bi-M	1996.5–2009.11	2006.5	22

The first column gives the variables and the second column reports the variables that we use as structural variables in the forecasting exercise. The frequencies are in column headed 'Freq.' are quarterly (Q), monthly (M), and bi-monthly (bi-M). The column headed 'Sample' reports the sample used for estimation and forecasting after taking account of the presample necessary for differencing and lags. The column with heading '1st fore' shows the first period for which an out-of-sample forecast is made. Finally, the column with heading 'no.fore' gives the number of out-of-sample forecasts per variable.

bond yield we additionally include aggregate corporate profits growth as a regressor. Real GDP growth is added to the regressions for the unemployment growth rate as an indicator of the business cycle. In the exchange rate specifications we incorporate inflation differentials and one month interbank interest differentials. These variables are derived from exchange rate parities: uncovered interest rate parity and purchasing power parity. All the regressors are reported in Table 1.

4 Data

We use a range of macroeconomic and financial data on monthly and quarterly frequencies. The variables are reported in the first column of Table 1. The quarterly data are the U.S. CPI inflation rate, the three month U.S. T-bill rate, the ten year U.S. government bond rate, the Moody's AAA corporate bond yield, the U.S. real GDP growth rate, and the U.S. unemployment rate. The survey of professional forecasters (SPF) data for these

series were obtained from the Philadelphia Fed’s Real Time Data Research Center and we use the median response in our analysis. For actual observations on U.S. real GDP and the U.S. unemployment rate we use real time data also obtained from the Philadelphia Fed’s Real Time Data Research Center. The remaining quarterly data are from the St. Louis Fed FRED database.

We transform the interest rates, real GDP and the unemployment rate into annual growth rates

$$y_{jt} = \ln(Y_{jt}/Y_{j,t-4})$$

where a lower case variable denotes the growth rate and an upper case variable a level variable, j indicates the variable, and t time. We also transform survey data into expected growth rates

$$y_{jt}^{(\text{spf})} = \ln(Y_{jt}^{(\text{spf})}/Y_{j,t-4}).$$

For applications with the real time data, to calculate growth rates, we use the vintages available to the forecasters at the time of the forecast. After allowing for pre-samples to calculate growth rates and for four lags for the autoregressive models, we are left with the samples’ lengths reported in the fourth column of Table 1.

The variables used as structural regressors are reported in the second column of Table 1. An additional variable is the growth rate of U.S. corporate profits, which we use as an explanatory variable for the corporate bond yields. This series is from the St. Louis Fed FRED data base.

In addition we analyze three exchange rate series: the euro-U.S. dollar exchange rate (EUR/USD), the Japanese yen-U.S. dollar exchange rate (JPY/USD), and the British pound-U.S. dollar exchange rate (GBP/USD). We use survey data collected by Consensus Economics of London, and their data set also provides the actual exchange rates. We transform exchange rate series into monthly growth rates

$$y_{jt} = \ln(Y_{jt}/Y_{j,t-1})$$

and

$$y_{jt}^{(\text{spf})} = \ln(Y_{jt}^{(\text{spf})}/Y_{j,t-1})$$

where we use the same notation as for the quarterly data. The survey data for the U.S. dollar-British pound exchange rates are only available on a bi-monthly basis.

The structural variables that we use to explain the exchange rate are the differentials in CPI inflation and the differentials in the one month interbank interest rates. For the interest rates we use the LIBOR rate, the dollar LIBOR rate, the yen LIBOR rate, and the EURIBOR rate. These series are provided by the St. Louis Fed's FRED database.

Figures 1–2 plots the series and the median survey expectations. The solid lines correspond to median expectations and the dotted lines represent the realized data. While the survey expectations do a good job predicting the actual series, it is notable that, for all series, the actual data are substantially more volatile. This is particularly so for the inflation series and for the three exchange rate series.

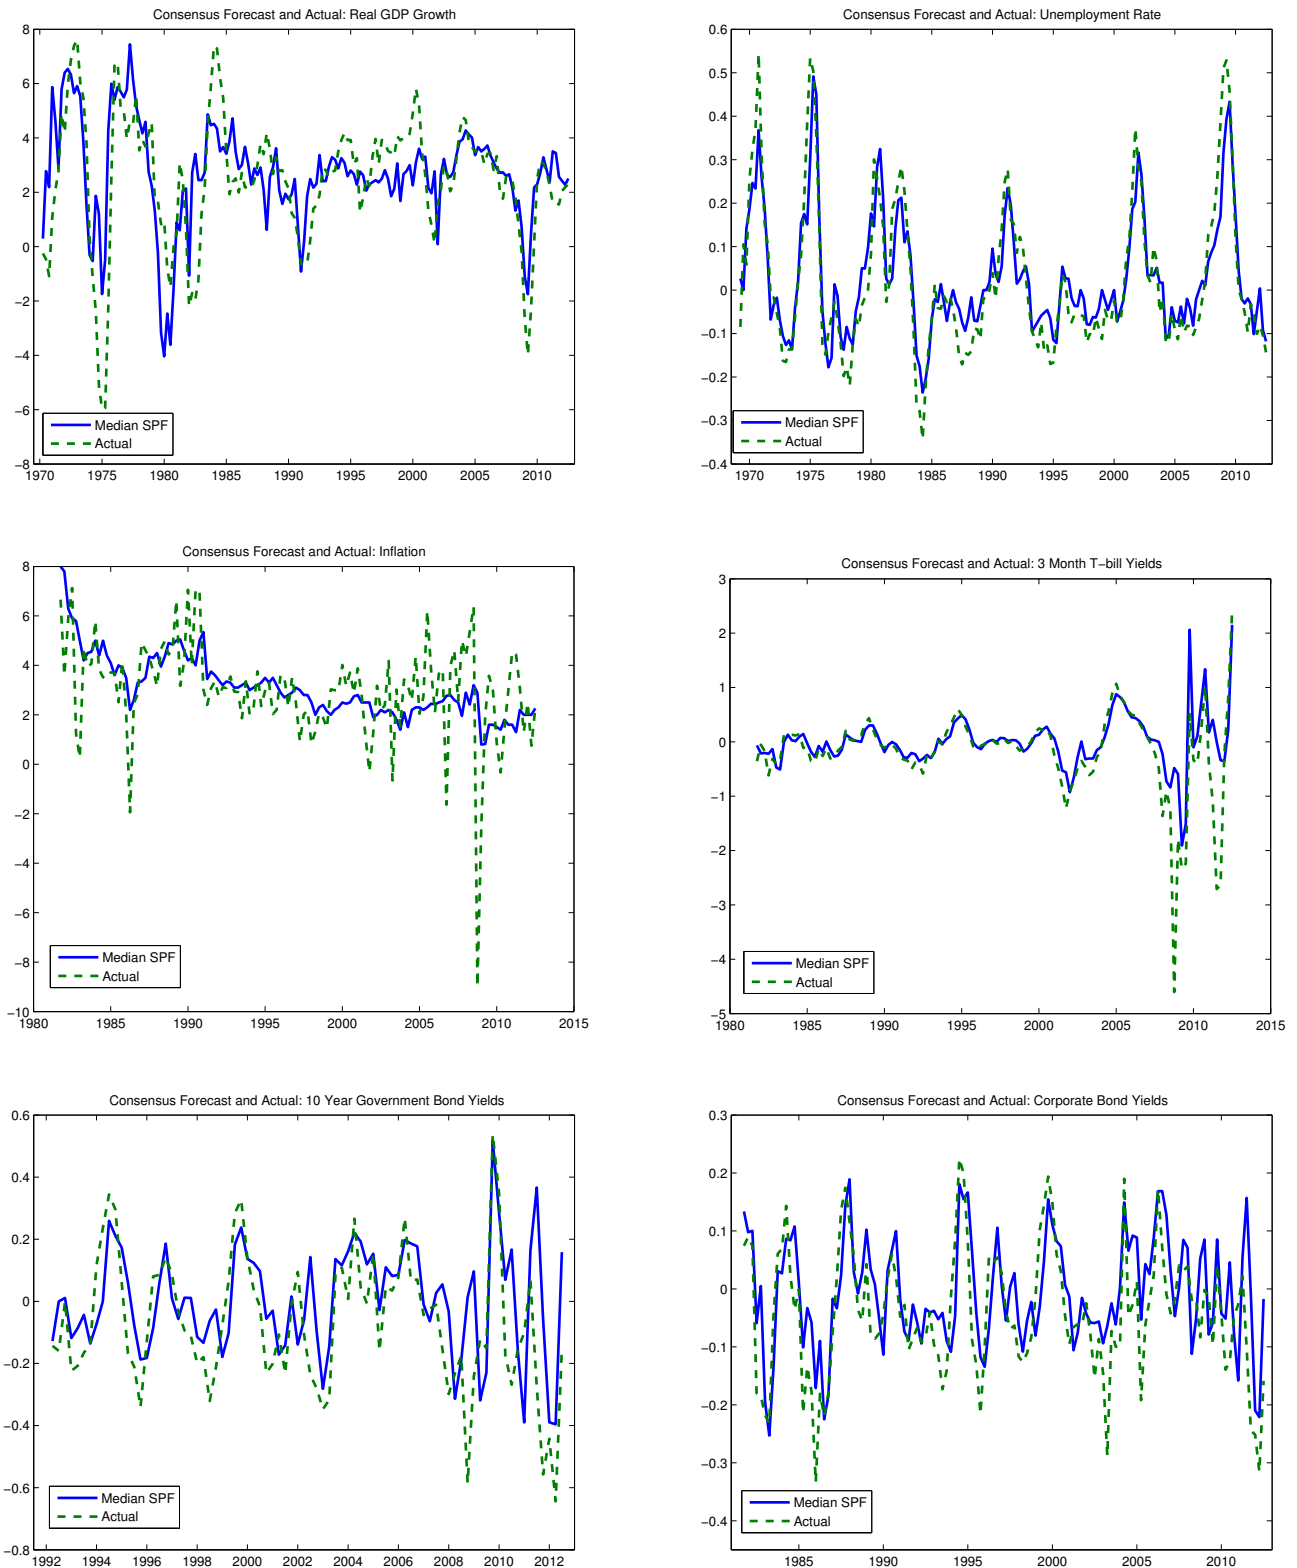
5 Results

Table 2 reports the estimates of the constant gain parameters. The first column reports the value of gain parameter of the adaptive expectations model. The next three columns report the values of gain parameters for autoregressive specifications, and the following three columns display values of the gain for models including both autoregressive components and structural variables. Finally, the last column reports the gain parameter of the model containing only structural predictors.

The estimates for the adaptive learning parameter vary considerably between the series. The parameter is relatively small for inflation, GDP growth and the exchange rates, which implies that the expectations change slowly. In contrast, γ 's for the interest rates and the unemployment rate are relatively high.



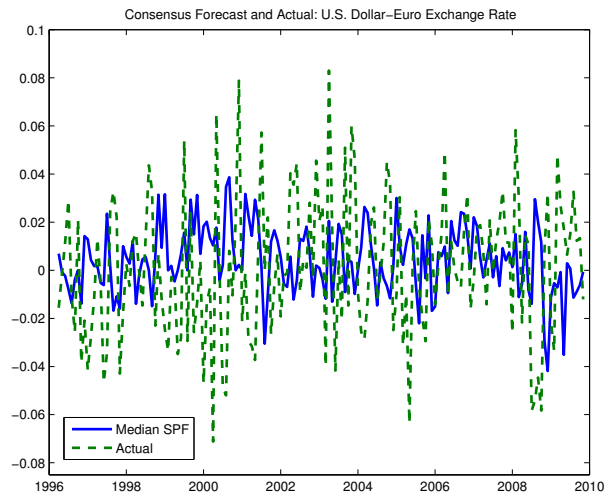
The values of estimated constant gain parameters of the autoregressive and structural models range between 0.001 for GBP/USD exchange rate under AR(1) specification and 0.151 for the structural model of the unemployment growth rate. The small values of the gain are similar to the ones

Figure 1: Median SPF forecasts and actual outcomes



The figures display the median one quarter ahead forecasts from the survey of professional forecasters as the solid line and the actual outcome of each variable as the dashed line. The variables are: real GDP growth as the (1,1) element of the matrix, the growth rate of the unemployment rate as the (1,2) element, the inflation rate (2,1), the growth rate of the three month t-bill rate (2,2), the growth rate of the ten year government bond yields (3,1), and the average AAA corporate bond yields (3,2).

Figure 2: Median SPF forecasts and actual outcomes



The figures display the median one month ahead forecasts from a survey among forecasters collected by Consensus Economics of London as the solid line and the actual outcome of each variable as the dashed line. The variables are top to bottom: the euro-dollar exchange rate, the yen-dollar exchange rate, and the pound-dollar exchange rate.

Table 2: Estimates of the gain parameter, γ

Variable	AE	AR(1)	AR(2)	AR(4)	AR(1)SM	AR(2)SM	AR(4)SM	SM
Inflation	0.295	0.052	0.057	0.019	0.057	0.017	0.035	0.103
3 month yields	0.887	0.083	0.109	0.111	0.109	0.081	0.055	0.130
10 year bond yield	0.512	0.001	0.001	0.009	0.001	0.012	0.047	0.094
Corp bond yields	0.886	0.084	0.003	0.004	0.003	0.072	0.001	0.004
GDP growth	0.052	0.010	0.030	0.045	0.030	0.008	0.039	0.001
Unemployment	0.900	0.010	0.129	0.120	0.129	0.010	0.009	0.151
EUR/USD	0.001	0.127	0.045	0.0410	0.045	0.127	0.117	0.060
JPY/USD	0.070	0.026	0.106	0.108	0.106	0.025	0.094	0.107
GBP/USD	0.040	0.001	0.011	0.085	0.011	0.117	0.126	0.019

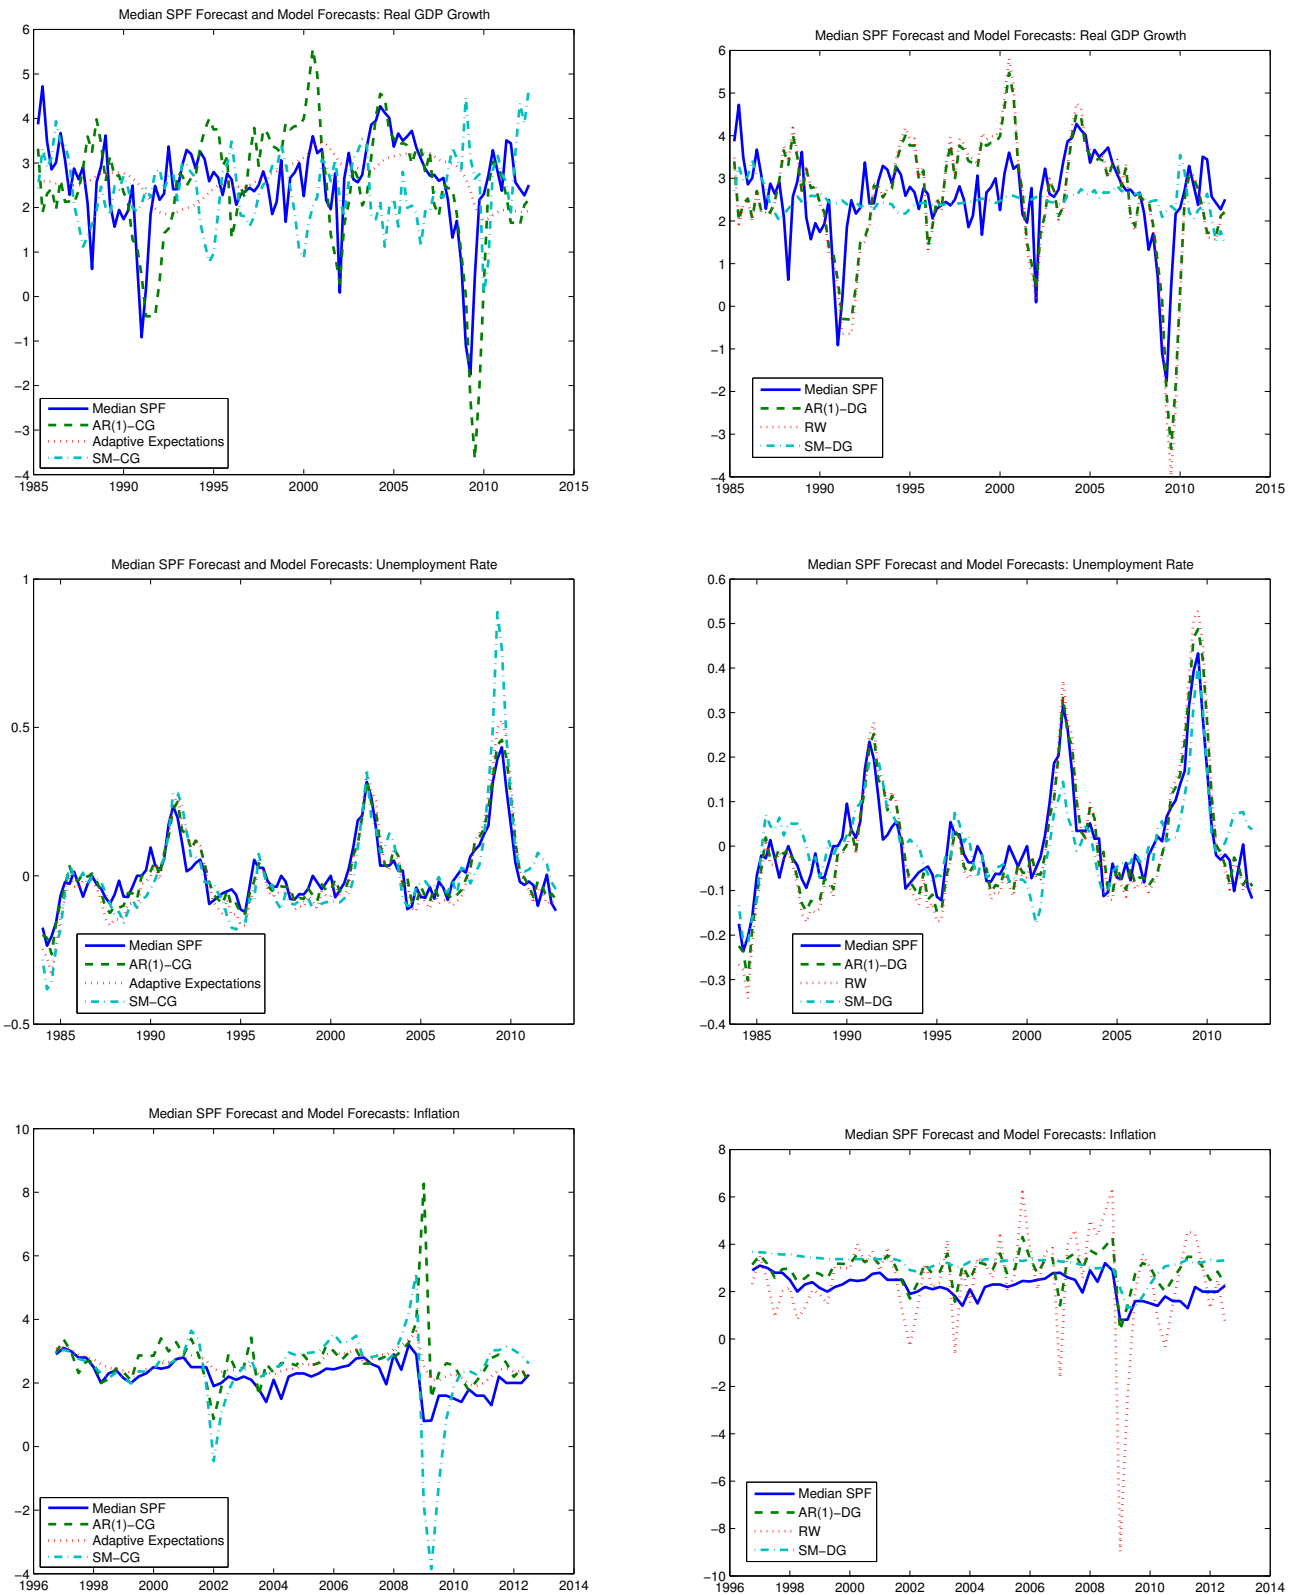
Table reports estimates of γ s obtained using the training sample. AE stands for adaptive expectations, AR(1)–AR(4) are autoregressive specifications, and SM for structural model. Predictors included in structural models for each variable are reported in Table 1.

previously found in the literature by Branch and Evans (2006a) and Milani (2007) for instance. They imply a long memory of agents as they put relatively low weight on more recent observations. The average estimated gain across all the variables and specifications equals 0.060 and suggests that agents use a long history of observations to make their forecasts.

Table 3 reports the MSFEs from the out-of-sample forecasts. The estimated gain parameters reported in Table 2 are used to generate these forecasts. The top panel of the table compares the performance of constant gain learning specifications to adaptive expectations and the random walk forecast, $\hat{y}_{T+1} = y_T$. The lower panel of Table 3 shows the MSFE results for the decreasing gain learning specifications in addition to the differenced random walk forecast, $\hat{y}_{T+1} = 0$. In order to make the comparison easier, the lowest MSFEs for each variable and gain specification are reported in bold font. The corresponding forecasts are plotted in Figures 3–5: green, dashed lines represent the forecasts generated by the AR(1) specifications, red, dotted lines those by constant gain adaptive expectations in the plots on the left side and the random walk model in the plots on the right side, light blue dashed lines those by structural models, and the actual median forecast is the solid, blue line. The left panels plot forecasts for constant gain updating while the right panels show the decreasing gain learning.

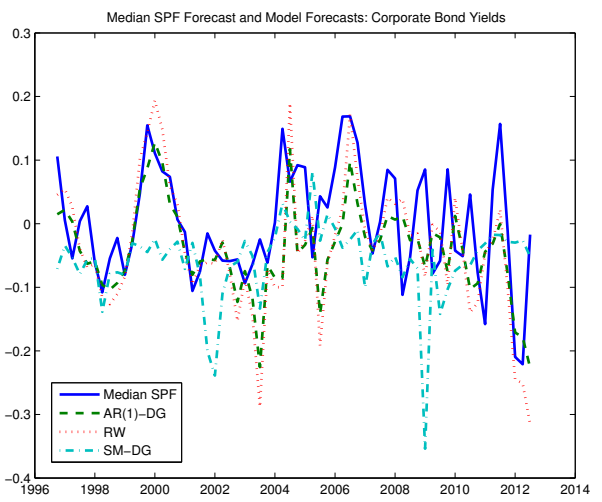





The top panel indicates that for five out of nine variables adaptive ex-

Figure 3: Median SPF forecasts and model forecasts



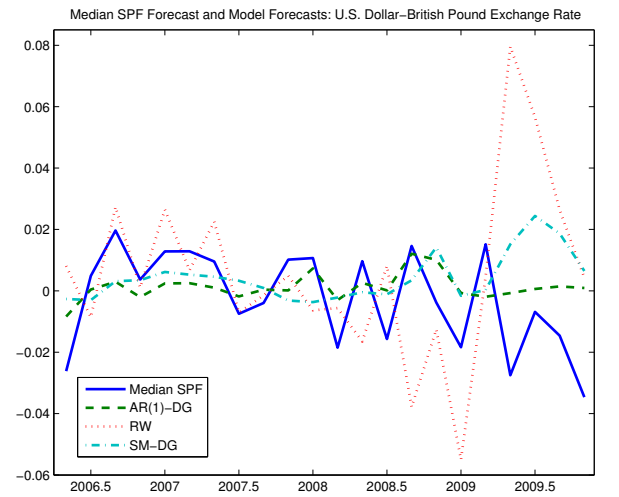





The figures display the median one quarter ahead forecasts from the survey of professional forecasters as the solid line together with the forecasts of selected model: dashed lines represent the AR(1) forecasts with constant gain learning on the left and decreasing gain learning on the right, dotted lines on the left graphs correspond to adaptive expectations and the level random walk on the right, dash-dotted lines give the structural models with constant gain learning on the left and decreasing gain learning on the right. The variables per row top to bottom are: real GDP growth, unemployment growth, and the inflation rate.

Figure 4: Median SPF forecasts and model forecasts



The variables per row top to bottom are: the yields of three month U.S. t-bills, U.S. ten year government bonds, and AAA corporate bonds. For the other details see the footnote of Figure 3.

Figure 5: Median SPF forecasts and model forecasts



The variables per row top to bottom are: the euro-dollar exchange rate, the yen-dollar exchange rate, and the British pound-dollar exchange rate. For the other details see the footnote of Figure 3.

Table 3: MSFE for decreasing and constant gain learning

Constant Gain Least Squares									
	RW	AE	AR(1)	AR(1)SM	AR(2)	AR(2)SM	AR(4)	AR(4)SM	SM
Inflation	3.989	0.474	1.219	0.999	1.160	1.028	1.703	1.281	1.154
3 month yields	0.878	0.815	1.285	0.942	1.655	1.112	2.052	2.398	0.884
10 year yield	0.100	0.075	0.072	0.071	0.084	0.091	0.054	0.068	0.083
Corp bond yields	1.015	0.931	0.878	0.952	0.853	0.923	0.692	0.856	1.619
GDP growth	1.143	1.134	1.382	1.515	1.321	1.282	1.486	1.371	1.974
Unemployment	0.387	0.385	0.203	0.672	0.194	0.600	0.297	0.598	0.828
EUR/USD	0.141	0.022	0.030	0.047	0.030	0.051	0.059	0.080	0.039
JPY/USD	0.220	0.039	0.024	0.058	0.024	0.068	0.053	0.076	0.067
UKP/USD	0.121	0.027	0.032	0.042	0.047	0.043	0.063	0.084	0.040
Decreasing Gain Least Squares									
	Δ RW	Mean	AR(1)	AR(1)SM	AR(2)	AR(2)SM	AR(4)	AR(4)SM	SM
Inflation	5.180	1.264	0.872	0.875	0.890	0.858	1.221	1.273	1.065
3 month yields	0.399	0.424	0.786	0.696	0.925	0.795	0.858	0.768	0.408
10 year yield	0.062	0.061	0.068	0.069	0.074	0.077	0.046	0.047	0.072
Corp bond yields	0.755	0.904	0.781	1.013	0.879	1.052	0.659	0.781	1.226
GDP growth	7.571	1.079	1.296	1.179	1.265	1.110	1.323	1.161	0.987
Unemployment	1.275	1.384	0.255	0.215	0.283	0.236	0.241	0.214	0.460
EUR/USD	0.020	0.022	0.024	0.034	0.024	0.034	0.030	0.043	0.034
JPY/USD	0.037	0.037	0.035	0.044	0.037	0.046	0.042	0.051	0.047
UKP/USD	0.025	0.026	0.021	0.036	0.024	0.038	0.027	0.045	0.039

The table reports the MSFE for constant and decreasing gain learning models, where the MSFE is defined in (6). AE stands for adaptive expectations, AR(1)-AR(4) are autoregressive specifications. The column ‘Mean’ reports the mean forecast, $\hat{y}_{T+1} = T^{-1} \sum_{t=1}^T y_t$. SM stands for structural model. Predictors included in structural models are reported in Table 1. The smallest MSFE for each variable is reported in bold font.

expectations exhibit lower MSFEs than the constant gain specifications in the other columns. Figures 3–5 show that adaptive learning forecasts are substantially less volatile than the other forecasts. For this reason they do particularly well after unstable periods, such as the recent financial crisis, where the other models forecast a longer or more extreme deviations from the mean than the median professional forecaster. Constant gain adaptive expectations generate particularly good forecasts in case of the unemployment rate and ten year government bond yields, where the survey forecasts are very well tracked by their model predictions.

When comparing adaptive expectations in the plots in the left panels with its special case, the random walk, in the plots in the right panels (both

are in dotted, red lines), one can see that γ different from one provides much better forecasts. However, the differenced random walk forecasts, whose MSFE results are in the first column of the lower panel, provides the most precise forecasts for the three month T-bill yields and the euro-dollar exchange rate.

Among the constant gain models the AR specifications, generally, do better than those that are augmented with structural predictors. All three AR models provide the best forecast for at least one series and provide competitive forecasts for all variables. Models with structural predictors, in contrast, never deliver the lowest MSFE. The plots in the left panels of Figures 3–5 show that structural models usually predict much larger movements than are observed in the survey data.

When a decreasing gain is employed, fluctuations predicted by structural models are much less pronounced, as suggested by the light blue, dashed lines plotted in the right panels of Figures 3–5. This is mirrored in the lower panel of Table 3, which shows findings with decreasing gain learning. To contrast with the constant gain forecasts, there is no single dominant model in the class of decreasing gain forecasts. Also, models with structural predictors provide the most precise forecasts for four out of the nine variables.

Finally, when we compare the top and bottom panels of Table 3, we find that in five out of nine cases the decreasing gain forecasts perform better than constant gain forecasts. However, in most cases the forecast performance is comparable. Notable differences are in the cases of the inflation rate, where a constant gain provides better forecasts, and the three month t-bill yield, where decreasing gain forecasts are substantially more precise. The plots in the first rows of Figures 3 and 4 show that the differences are largely due to the extreme forecasts during the recent crisis. In both cases, forecasts that are not overly influenced by the most recent observations perform relatively well.

Table 4 reports Diebold-Mariano test statistics for forecast accuracy against the random walk forecast, $\hat{y}_{T+1} = y_T$, where significant differences are in bold font. The positive, significant entries of Table 4 indicate that the alternative specification forecasts significantly better than the random

Table 4: Test for predictive accuracy against random walk forecast, $\hat{y}_{T+1} = y_T$

Constant Gain Least Squares								
	AE	AR(1)	AR(1)SM	AR(2)	AR(2)SM	AR(4)	AR(4)SM	SM
Inflation	2.306	3.732	2.337	1.896	1.991	1.533	1.712	1.965
3 month yields	1.289	-0.826	-0.245	-1.316	-0.930	-0.751	-0.980	1.289
10 year yield	1.785	2.087	1.379	1.242	0.449	2.221	1.703	0.618
Corp bond yields	3.005	1.545	0.086	1.775	0.580	2.172	0.485	-1.649
GDP growth	1.167	4.499	0.414	1.870	1.866	0.532	1.427	-0.825
Unemployment	0.271	6.398	-2.557	6.144	-1.887	1.277	-1.932	-1.966
EUR/USD	5.453	7.067	6.558	6.876	6.326	4.339	3.677	6.549
JPY/USD	6.933	6.851	6.019	6.830	5.611	6.099	5.196	6.227
UKP/USD	1.828	1.641	1.441	1.512	1.439	1.082	0.578	1.779
Decreasing Gain Least Squares								
	Mean	AR(1)	AR(1)SM	AR(2)	AR(2)SM	AR(4)	AR(4)SM	SM
Inflation	1.860	2.046	2.091	2.030	2.083	1.794	1.723	1.921
3 month yields	1.260	0.816	2.295	-0.252	1.074	0.063	0.412	1.623
10 year yield	1.077	2.032	1.558	1.661	1.177	2.286	2.278	0.767
Corp bond yields	0.443	2.386	-0.025	1.466	-0.252	2.388	1.300	-0.933
GDP growth	1.333	4.067	3.828	1.972	2.969	1.478	2.634	1.648
Unemployment	-4.106	7.058	5.735	3.024	3.832	3.867	4.371	-0.897
EUR/USD	7.207	7.232	6.512	7.188	6.453	6.779	5.772	6.479
JPY/USD	6.927	6.620	6.898	6.609	6.782	6.457	6.660	6.963
UKP/USD	1.856	1.800	1.809	1.666	1.725	1.473	1.768	1.856

The table reports the Diebold-Mariano (1995) test statistic for predictive accuracy against the forecast $\hat{y}_{T+1} = y_T$. AE stands for adaptive expectations, AR(1)-AR(4) are autoregressive specifications. ‘Mean’ is the mean forecast. SM stands for structural model. Numbers in bold font indicate significance at the 5% level.

walk model. For the constant gain forecast, 31 out of 64 comparisons are significant and 29 of them positive, implying that many models produce significantly more accurate forecasts than the random walk. For the decreasing gain forecasts, 36 out of 64 prediction differences are significant and 35 of them positive.

The Diebold-Mariano test statistics for forecast accuracy against the difference random walk forecast, $\hat{y}_{T+1} = 0$, are in Table 5. Again, significant differences are in bold font. The constant gain forecasts produce 25 significantly more accurate forecasts and 11 significantly less accurate forecasts than the differenced random walk forecast, and the decreasing gain 24 significantly more accurate and 13 significantly less accurate forecasts.

Table 5: Test for predictive accuracy against random walk forecast, $\hat{y}_{T+1} = 0$

Constant Gain Least Squares								
	AE	AR(1)	AR(1)SM	AR(2)	AR(2)SM	AR(4)	AR(4)SM	SM
Inflation	16.301	4.040	8.346	9.381	10.323	4.399	6.928	7.368
3 month yields	-1.269	-1.177	-1.455	-1.549	-1.689	-1.093	-1.305	-1.471
10 year yields	-0.443	-0.391	-0.355	-0.777	-0.909	0.329	-0.197	-0.856
Corp bond yields	-0.738	-0.560	-1.071	-0.471	-0.746	0.326	-0.678	-2.086
GDP growth	12.434	12.371	10.577	13.736	12.445	12.216	12.703	9.731
Unemployment	3.457	4.027	3.140	4.057	3.216	3.556	3.210	2.327
EUR/USD	-0.655	-2.829	-4.678	-2.985	-4.614	-3.586	-5.313	-3.719
JPY/USD	-1.130	4.991	-1.865	4.793	-2.153	-1.487	-2.311	-4.187
GBP/USD	-1.103	-0.568	-0.909	-1.914	-1.446	-1.713	-1.499	-1.771
Decreasing Gain Least Squares								
	Mean	AR(1)	AR(1)SM	AR(2)	AR(2)SM	AR(4)	AR(4)SM	SM
Inflation	10.083	13.197	12.929	12.997	13.059	9.032	7.619	12.168
3 month yields	-1.209	-1.044	-0.930	-1.119	-1.046	-1.282	-1.292	-0.058
10 year yields	0.110	-0.262	-0.288	-0.482	-0.588	0.757	0.757	-0.589
Corp bond yields	-1.801	-0.143	-1.170	-0.649	-1.363	0.548	-0.137	-1.945
GDP growth	13.059	12.844	12.743	14.159	14.144	14.104	14.224	13.538
Unemployment	-1.947	3.874	4.027	3.764	3.931	3.919	4.011	3.173
EUR/USD	-2.233	-4.110	-3.750	-3.971	-3.602	-4.589	-5.148	-3.673
JPY/USD	0.463	2.820	-2.343	0.456	-2.615	-2.073	-3.622	-3.097
GBP/USD	-0.969	1.556	-1.300	0.250	-1.642	-0.480	-1.899	-1.846

The table reports the Diebold-Mariano (1995) test statistic for predictive accuracy against the forecast $\hat{y}_{T+1} = 0$. For the remaining details see the footnote of Table 4.

For inflation all models are more accurate than both random walk specifications and most of them significantly so. In contrast, no model can significantly beat the differenced random walk when forecasting any of the interest rates. While the differenced random walk specification has the lowest MSFE for the three month interest rate and the AR(4) the lowest MSFE for the ten year rate and the corporate bond rate, although the difference is not statistically significant.

For the macroeconomic variables, GDP growth and unemployment growth, many models beat both specifications of the random walk. It appears that survey expectations are well represented by more complex models than the random walk.

The literature on exchange rates suggests that random walks are difficult to beat, and we would imagine that survey expectations should reflect that.⁶

⁶See Cheung, Chinn, and Pascual (2005) for the literature overview and recent findings

The results for the euro-dollar exchange rate confirm this, where the differenced random walk has the lowest MSFE and is significantly better than all models except adaptive expectations. However, this is not the case for the other two exchange rates. In both cases AR(1) models have the lowest MSFE: constant gains in the case of the yen-dollar and decreasing gains in the case of the pound-dollar exchange rate. Furthermore, in the case of the yen the AR(1) model is significantly better than the random walk, which suggests that professional forecasters are not convinced about the random walk nature of this currency.

6 Conclusion

In this paper, we evaluate in how far adaptive learning models can replicate the expectations of professional forecasters. We used constant gain and decreasing gain specifications of models with differing complexity. When we fit constant gain models to the survey data, we find that relatively small gains (values smaller than 0.15) provide the best fit, which implies that professional forecasters use rather long data samples.

The comparison between constant gain and decreasing gain specifications leaves no clear winner. Constant gain models do better for inflation, the unemployment rate and the yen-dollar exchange rate but decreasing gain specifications are better representations of professional forecasters' expectations for interest rates, GDP growth and the pound-dollar exchange rate. Compared to other constant gain models, the highly parsimonious adaptive expectations rule does particularly well, in that it is the closest representation of the SPF expectations in five cases and a reasonably close representation in the remaining four.

An additional interesting finding is that professional forecasters agree with the academic economists that exchange rates are best represented by random walks only in the case of the euro-dollar exchange rate but not in the cases of the yen-dollar and the pound-dollar exchange rates.

on exchange rate predictability.

References

- [1] Adam, K. (2007) ‘Experimental Evidence on the Persistence of Output and Inflation.’ *Economic Journal* 117, 603-635.
- [2] Adam, K., A. Marcet, and J. P. Nicolini (2009) ‘Stock Market Volatility and Learning.’ mimeo, University of Manheim.
- [3] Allen, H. and M. Taylor (1990) ‘Charts, Noise in Fundamentals in the London Foreign Exchange Market.’ *Economic Journal* 100, 49–59.
- [4] Atkeson, A. and L. E. Ohanian (2001) ‘Are Phillips Curves Useful for Forecasting Inflation?’ *Quarterly Review, Federal Reserve Bank of Minneapolis* Winter, 2–11.
- [5] Bacchetta, P., E. Mertens, and E. van Wincoop (2009) ‘Predictability in Financial Markets: What Do Survey Expectations Tell Us?’ *Journal of International Money and Finance* 28, 406–426.
- [6] Bonham, C., and R. Cohen (2001) ‘To Aggregate, Pool, or Neither: Testing the Rational Expectations hypothesis Using Survey Data.’ *Journal of Business and Economic Statistics* 19, 278–291.
- [7] Bonham, C., and D. Dacy (1991), ‘In Search of a Strictly Rational Forecast.’ *Review of Economics and Statistics* 73, 245–253.
- [8] Branch, W. (2004) ‘The Theory of Rationally Heterogeneous Expectations: Evidence from Survey Data on Inflation Expectations.’ *Economic Journal* 114, 592–621.
- [9] Branch, W., and G. W. Evans (2006a) ‘A Simple Recursive Forecasting Model.’ *Economics Letters* 91, 158–66.
- [10] Branch, W., and G. W. Evans, (2006b) ‘Intrinsic Heterogeneity in Expectation Formation.’ *Journal of Economic Theory* 127, 264–95.
- [11] Bray, M. (1982) ‘Learning, Estimation, and the Stability of Rational Expectations.’ *Journal of Economic Theory* 26, 318–39.

- [12] Bray, M., and N. Savin (1986) ‘Rational Expectations Equilibria, Learning and Model Selection.’ *Econometrica* 54, 1129–1160.
- [13] Brock, W.A., and C. H. Hommes (1997) ‘A Rational Route to Randomness.’ *Econometrica* 65, 1059–1160.
- [14] Cheung, Y., M. D. Chinn, and A. G. Pascual (2005) ‘Empirical Exchange Rate Models of the Nineties: Are Any Fit to Survive?’ *Journal of International Money and Finance* 24, 1150–1175.
- [15] Cho, I.-K., N. Williams, and T. Sargent (2002) ‘Escaping Nash Inflation.’ *Review of Economic Studies* 69, 1–40.
- [16] Croushore, D. (1997), ‘The Livingston Survey: Still Useful After All These Years.’ *Federal Reserve Bank of Philadelphia Business Review* March/April, 15–26.
- [17] Diebold, F. X., and R. S. Mariano, (1995) ‘Comparing Predictive Accuracy.’ *Journal of Business and Economic Statistics* 13, 253–263.
- [18] Evans, G. W., and W. Branch (2010) ‘Asset Return Dynamics and Learning.’ *Review of Financial Studies* 23, 1651–1680.
- [19] Evans, G.W. and R. Gulamani (1984) ‘Tests for rationality of the Carlson-Parkin inflation expectations data.’ *Oxford Bulletin of Economics and Statistics* 46, 1–20.
- [20] Evans, G. W., and S. Honkapohja (2001) *Learning and Expectations in Macroeconomics*, Princeton, NJ: Princeton University Press.
- [21] Frankel, J., and K. Froot (1987a) ‘Short-term and Long-term Expectations of the Yen/Dollar Exchange Rate: Evidence from Survey Data.’ *Journal of the Japanese and International Economies* 1, 249–274.
- [22] Frankel, J., and K. Froot (1987b) ‘Using Survey Data to Test Standard Propositions Regarding Exchange Rate Expectations.’ *American Economic Review* 77, 133–153.
- [23] Frankel, J., and K. Froot (1990a) ‘Chartists, Fundamentalists and the Demand for Dollars.’ in A.S. Courakis and M. P. Taylor, ed., *Private*

Behaviour and Government Policy in Interdependent Economies, Oxford University Press, Oxford, pp.73–126.

- [24] Frankel, J. and K. A. Froot (1990b) ‘The Rationality of the Foreign Exchange Rate: Chartists, Fundamentalists and Trading in the Foreign Exchange Market.’ *American Economic Review, Papers and Proceedings* 80, 181–185.
- [25] Friedman, B. (1980) ‘Survey Evidence on the “Rationality” of Interest Rate Expectations.’ *Journal of Monetary Economics* 6, 453–465.
- [26] Froot, K. A. (1989) ‘New Hope for the Expectations Hypothesis of the Term Structure of Interest Rates.’ *Journal of Finance* 44, 283–305.
- [27] Hommes, C., and G. Sorger (1998) ‘Consistent Expectations Equilibria.’ *Macroeconomic Dynamics* 2, 287–321.
- [28] Hyndman, R. J., A. Koehler, J. K. Ord, and R. D. Snyder (2008) *Forecasting with Exponential Smoothing: The State Space Approach*. Berlin: Springer Verlag.
- [29] Ito, T., (1990) ‘Foreign Exchange Expectations: Mirco-Survey Data.’ *American Economic Review* 80, 434–449.
- [30] Jeong, J., and G. S. Maddala (1996) ‘Testing the Rationality of Survey Data Using the Weighted Double-bootstrapped Method of Moments.’ *Review of Economics and Statistics* 78, 296–302.
- [31] Jongen, R., W. Verschoor, C. Wolff, and R. Zwinkels (2012) ‘Explaining Dispersion in the Foreign Exchange Market: A Heterogeneous Agent Approach.’ *Journal of Economic Dynamics and Control* 36, 719–735.
- [32] Keane, M. P., and D. E. Runkle (1990) ‘Testing the Rationality of Price Forecasts: New Evidence from Panel Data.’ *American Economic Review* 80, 714–735.
- [33] Marcet, A., and J. P. Nicolini (2003) ‘Recurrent Hyperinflations and Learning.’ *American Economic Review* 93, 1476–1498.

- [34] Marcet, A., and T. Sargent (1989) ‘Convergence of Least Squares Learning Mechanisms in Self-Referential Linear Stochastic Models.’ *Journal of Economic Theory* 48, 337–68.
- [35] Meese, R., and K. Rogoff (1983) ‘Empirical Exchange Rate Models of the Seventies: Do They Fit Out of Sample?’ *Journal of International Economics* 14, 3–24.
- [36] Milani, F. (2007) ‘Expectations, Learning and Macroeconomic Persistence.’ *Journal of Monetary Economics* 54, 2065–2082.
- [37] Milani, F. (2011) ‘Expectation Shocks and Learning as Drivers of the Business Cycle.’ *Economic Journal* 121, 379–401.
- [38] Orphanides, A., and S. van Norden (2005) ‘The Reliability of Inflation Forecasts Based on Output Gap Estimates in Real Time.’ *Journal of Money, Credit and Banking* 37, 583–601.
- [39] Pesaran, M. H., and A. Pick (2011) ‘Forecast Combination across Estimation Windows.’ *Journal of Business and Economic Statistics* 29, 307–318.
- [40] Peseran, M. H., A. Pick, and M. Pranovich (2013) ‘Optimal Forecasts in the Presence of Structural Breaks.’ *Journal of Econometrics* (forthcoming).
- [41] Preston, B., and S. Eusepi (2013) ‘Expectations, Learning and Business Cycle Fluctuations.’ *American Economic Review* (forthcoming).
- [42] Sargent, T. J. (1993) *Bounded Rationality in Macroeconomics*. New York: Oxford University Press.
- [43] Stock, J. H., and M. W. Watson (2007) ‘Why Has Inflation Become Harder to Forecast?’ *Journal of Money, Credit, and Banking* 39, 3–34.
- [44] Stock, J. H., and M. W. Watson (2003) ‘Forecasting Output and Inflation: The Role of Asset Prices.’ *Journal of Economic Literature* 41, 788–829.