Differential Interpretation in the Survey of Professional Forecasters

Sebastiano Manzan

Department of Economics & Finance, Baruch College, CUNY

55 Lexington Avenue, New York, NY 10010

phone: +1-646-312-3408, email: `sebastiano.manzan@baruch.cuny.edu`

September 2010

Abstract

In this paper we estimate a simple Bayesian learning model to expectations data from the Survey of Professional Forecasters. We reformulate the model in terms of forecast revisions, which allows to abstract from differences in priors and to focus the analysis on the relationship between revisions and signal. The model depends on two parameters, the forecaster's belief about the signal bias and its weight. The empirical analysis shows that there is significant heterogeneity in the parameters among forecasters, in particular at longer forecast horizons. The cross-sectional distribution of the estimated bias parameter has a median close to zero, while its dispersion decreases with the horizon. A similar result is obtained for the weight parameter, with the median across forecasters increasing toward one and the dispersion decreasing when approaching the target date. Furthermore, the results indicate that a possible explanation for the persistence in dispersion, even at short horizons, is the heterogeneity of interpretation of new information, in the sense that agents with optimistic (relative to the other forecasters) priors are also likely to believe that the signal underestimates the future realization of the variable, and the opposite for forecasters with pessimistic views. We also find that the latter type of forecasters are more likely to assign a low weight to the signal, while optimistic forecasters incorporate the new information faster.

JEL code: C14, E17, E37
Keywords: differential interpretation, survey expectations, Bayesian learning, disagreement

1 Introduction

Policy-makers regularly look at public's expectations in order to learn how they impact current economic decisions and thus the future outlook of the economy. Although a proxy such as the *consensus* forecast is a useful indication of the expected direction of an economic variable, it obscures the significant disagreement that occurs among agents. This disagreement can have real consequences when the effectiveness of a policy relies on coordinating the expectations of the public. Recently, this topic has received a lot of attention in macroeconomics. Kandel and Zilberfarb (1999) investigate differential interpretation in inflation forecasts for Israel and find that there is significant dispersion in beliefs which they attribute to disagreement regarding the incoming information. Mankiw *et al.* (2003) examine disagreement in inflation expectations of professional forecasters and consumers; they show that it increases with the level of inflation and when agents experience large changes in the inflation rate. Lahiri and Sheng (2008) explore dispersion in GDP growth forecasts and find that dispersion reduces closer to the forecast target date as agents incorporate real-time information in revising their prior expectations. Patton and Timmermann (2008) find that the most relevant factor driving forecast disagreement is dispersion in prior expectations, rather than heterogeneity in the information signal. Capistran and Timmermann (2009a) argue that the evidence of disagreement is consistent with the assumption that forecasters use asymmetric loss functions with heterogeneous degrees of asymmetries. This literature extends on the earlier debate on the rationality of consensus and individual forecasts (see Keane and Runkle, 1990, and Bonham and Cohen, 2001, among others). Pesaran and Weale (2006) provide a comprehensive survey of the extensive literature on the empirical analysis of survey expectations and Hommes (2006) reviews the role of heterogeneous beliefs in economic models.

Several theories have been proposed to explain why disagreement among professional forecasters exists[1]. A standard argument is that agents observe a private signal, and thus have heterogeneous expectations since these are formed based on different information sets. However, this argument seems to be less relevant when considering agents forming expectations

---

[1]The discussion will be mostly focused on models of the behavior of professional forecasters, such as economists employed at financial institutions, large companies or research institutes. Other models try to explain the dispersion of (inflation) expectations among households. Carroll (2003) and Mankiw and Reis (2002) assume stickiness in the information processing of households, in the sense that they update their expectation infrequently. In this sense, new information spreads slowly across the public and creates dispersion of expectations because, at any point in time, only a fraction of households incorporate the most recent information, while the remaining households hold on to their out-of-date expectations. Branch (2004) finds supporting evidence for a model in which consumers switch between different predictors of inflation based on their recent forecast accuracy.

about macroeconomic variables, such as inflation or GDP growth. Other explanations for the existence of disagreement rely on the strategic behavior of forecasters. These models typically assume that forecasters observe the public signal and share the same expectation about the future value of the variable conditional on the signal. However, they might have an incentive to report a distorted forecast in order to generate publicity or for reputational considerations (Ehrbeck and Waldmann, 1996, Laster *et al.*, 1999, and Ottaviani and Sorensen, 2006). An alternative argument for the presence of heterogeneous expectations is provided by the Bayesian learning model of Kandel and Pearson (1995) in which agents update their prior expectation based on a public signal, but they have biased beliefs about its informational content for the future realization of the variable being forecast. In this model, differential interpretation of the signal arises when agents have different values for the bias, with some forecasters, for example, expecting the signal to over-estimate the future realization while others expect the opposite. Despite starting from the same prior expectation, agents might disagree on the interpretation of the signal and thus hold different posterior beliefs. Kandel and Pearson derive testable implications from the model in terms of revisions that are inconsistent with homogeneous interpretation of the signal. One type of inconsistency occurs when, comparing two forecasters, their posterior expectations are farther away compared to their priors. In this case, despite both forecasters observing the same (public) signal, the more optimistic among the two forecasters would revise the prior upward while the other forecaster would revise it downward resulting in their posteriors diverging compared to their priors. Another case inconsistent with homogeneous interpretation of the signal occurs when the forecaster with more pessimistic priors becomes, after observing the signal, the more optimistic one in terms of posteriors.

Kandel and Pearson empirically measure the deviation from homogeneity using the fraction of (pairwise) inconsistent revisions and find strong evidence to support differential interpretation. The aim of this paper is to extend their analysis to investigate the dispersion of macroeconomic forecasts from the Survey of Professional Forecasters (SPF). We propose an alternative empirical strategy to evaluate differential interpretation which consists in directly estimating the model on expectations allowing for forecaster-specific parameters. The Bayesian learning model implies a simple relationship between forecast revision and the unexpected (with respect to the prior) component of the signal and depends on two parameters: the interpretation bias and the weight of the signal which represents the belief of the forecaster about its precision (in comparison to the precision of the prior). The estimation of the learning parameters for each forecaster thus allows us to statistically evaluate the

hypothesis that the bias coefficient is equal across forecasters. This approach offers several advantages compared to the approach followed in Kandel and Pearson. First, it provides a sharper test of the differential interpretation hypothesis because it is able to distinguish cases of heterogeneity that are not detected by the pairwise comparison of revisions. An example is helpful in explaining when this situation might occur. Take two forecasters that, upon observing the signal, revise their priors in a way that the distance in their posterior forecasts is smaller than the distance between their priors. While the forecasts are converging and this situation would be classified as homogeneous using the approach of Kandel and Pearson, it can still be the case that forecasters have different interpretation bias. For example, the forecaster with more optimistic priors might believe that the signal underestimates the future realization of the variable, while the other forecaster might believe the signal is too optimistic. Since we estimate the bias of each forecaster, we are able to distinguish heterogeneous interpretation also in a situation like the discussed example, as opposed to the method of Kandel and Pearson. An additional advantage of our approach is that we are also able to estimate the weight that a forecaster assigns to the signal in revising the prior expectation. Differences in the weight represent a further source of disagreement driving dispersion in forecasts that was not considered by Kandel and Pearson (1995). This situation can be illustrated as follows. Assume two forecasters have the same prior, observe the same public signal, and they also have the same interpretation bias but assign different weights to the signal; then their revision, and thus their posterior forecast, will be different entirely because of different beliefs about the precision of the signal compared to the prior expectation. In this sense, we are able to test differences in both the mean and weight parameters used by forecasters to interpret information.

The estimation of the Bayesian learning model for expectations for six macroeconomic variables provides evidence of heterogeneous interpretation both in terms of the bias and weight of the signal. The dispersion in the interpretation bias of the signal is largest at the longest forecast horizon of the SPF (three quarters) and decreases toward zero approaching the target date. A statistical test of homogeneity of the parameter across forecasters rejects the hypothesis for most variables considered, namely at the longest forecast horizons. This provides further evidence that forecasters have differential interpretations about the information contained in the signal for the future realization of the variable. A similar result is found for the weight assigned to the signal, with significant dispersion across forecasters at the longer horizons while it is reduced as the horizon shortens. At the three-quarter horizon, a typical range of values for the contribution of real-time information to the forecast revi-

sion is between 30 and 70%, while it increases to between 70 and 90% for the forecast made at the target quarter (before the data is officially released in the next quarter). Also for the weight parameter, the hypothesis of homogeneity across forecasters is rejected for most variables and horizons, providing further evidence in favor of the differential interpretation hypothesis.

Furthermore, we characterize forecasters by relating their prior expectations to the interpretation parameters. We find that forecasters with optimistic prior views are likely to have optimistic interpretation of the news, in the sense that they expect the signal to underestimate the future realization. Hence, they revise their prior expectation in the direction of the signal, although the bias in its interpretation prevents them from fully incorporating such information. A similar argument can be made for forecasters that have pessimistic priors who are characterized by negative bias, meaning that they believe the signal overestimates the future realization of the variable. If we compare the revisions of these two types of forecasters following the approach of Kandel and Pearson, they would be classified as having homogeneous interpretation, despite the fact that they interpret differently the information contained in the signal in terms of its mean. The effect of the biased interpretation of the signal is a delay in the convergence of the forecasts as the horizon shortens compared to a situation with homogeneous beliefs. In addition, we find evidence that forecasters with optimistic views of the signal are more likely to assign higher weight to the incoming news when revising their prior forecasts, compared to forecasters with negative views of the signal. In this sense, forecasters with pessimistic views for the variable are likely to hold to their prior expectations and incorporate in their posterior forecast only a relatively small fraction of the signal. Another issue that we consider is whether the interpretation parameters affect the predictive accuracy of forecasters. We find that larger bias (either positive or negative) and lower weight assigned to the signal affect negatively the forecaster's performance in terms of accuracy of their predictions for most variables, with the only exception of CPI inflation where the interpretation parameters are irrelevant to explain performance.

The paper is organized as follows. Section (2) introduces the Bayesian learning model of Kandel and Pearson (1995) and Kandel and Zilberfarb (1999) and the estimation and testing strategy, Section (3) presents the survey data used in the paper and the estimation results, and Section (4) discusses various implications for forecaster's characteristics and predictive accuracy. Finally, Section (5) concludes.

5

2 Model

2.1 Bayesian Learning Model

Assume a forecaster i at time $t-h$ (h indicates the forecast horizon) has a prior distribution for the realization of a variable at target date t that is normally distributed with mean $Y^i_{t-h-1,t}$ and precision $\psi_{i,h+1}$. The prior is interpreted as the belief of the forecaster regarding the distribution of the variable at the target date t based on the information available up to time $t-h-1$. At the beginning of time $t-h$, the forecaster observes a public signal given by $L_{t-h} = Y_t + \epsilon_{t-h,t}$, and believes that the distribution of the signal error is normally distributed with mean $\mu_{i,h}$ and precision $\phi_{i,h}$². The term $\mu_{i,h}$ acts as bias in the interpretation of the signal. A positive value of $\mu_{i,h}$ indicates that the forecaster believes the signal, L_{t-h}, overestimates the variable at the target date, Y_t. On the other hand, the forecaster expects the variable (at time t) to be higher compared to the signal if $\mu_{i,h}$ is negative. The confidence that the forecaster attributes to the signal is measured by the precision $\phi_{i,h}$. We assume that both the precisions of the prior and the signal might vary across horizons, since at short horizons forecasters might regard the signal to be more precise compared to their prior, and vice versa at long forecasting horizons. After the signal at $t-h$ is observed, the forecaster updates the prior expectation and the posterior expectation for the realization of the variable at time t, denoted by $Y^i_{t-h,t}$, is given by

$$Y^i_{t-h,t} = (1 - \rho_{i,h})Y^i_{t-h-1,t} + \rho_{i,h}(L_{t-h} - \mu_{i,h}) \tag{1}$$

where the weight $\rho_{i,h} = \phi_{i,h}/(\psi_{i,h+1} + \phi_{i,h})$ is within the unit interval, and the variance of the posterior is $(\psi_{i,h+1} + \phi_{i,h})^{-1}$. Equation (1) shows that the posterior expectation can be expressed as a weighted average of the prior expectation, $Y^i_{t-h-1,t}$, and the biased signal $L_{t-h} - \mu_{i,h}$, that is, the signal adjusted for the interpretation bias of forecaster i. A positive value of the bias indicates that the forecaster believes the signal is too optimistic and he/she corrects the signal downward to form the posterior expectation. The weight $\rho_{i,h}$ is determined by the ratio of the precision of the signal to the precision of the posterior distribution (given by the sum of the prior and signal precisions). The weight is in the lower end of the unit interval if a forecaster believes that the signal has low precision relatively to the prior expectation. In this case, the posterior expectation is mostly determined by the prior. On the other hand, if the forecaster believes that the signal carries relevant

---

²This can occur both in a behavioral model where forecasters use different models to interpret the signal (Kandel and Pearson, 1995, and Kandel and Zilberfarb, 1999) or in a strategic model in which forecasters aim to promote their reputation as well-informed or in order to generate publicity (Ehrbeck and Waldmann, 1996, Laster *et al.*, 1999, and Ottaviani and Sorensen, 2006).

information (relatively to the prior) $\rho_{i,h}$ lies in the upper end of the unit interval and the biased signal receives a higher weight compared to the prior expectation.

Equation (1) can be expressed as follows:

$$Y^i_{t-h,t} - Y^i_{t-h-1,t} = \rho_{i,h}(L_{t-h} - Y^i_{t-h-1,t} - \mu_{i,h}) \qquad (2)$$

where $Y^i_{t-h,t} - Y^i_{t-h-1,t}$ represents the revision of the expectation at time $t - h$ (about the realization at time t), and $L_{t-h} - Y^i_{t-h-1,t}$ corresponds to the news or unexpected component of the signal with respect to the prior mean. The news is adjusted for the bias in the forecaster's interpretation of the signal. In case a forecaster believes the signal has high precision, then the weight $\rho_{i,h}$ is large (within the unit interval) and the revision reflects, to a large extent, the biased news in the signal. On the other hand, if the forecaster believes that the signal has low precision (compared to the prior), then the weight is at the lower end of the unit interval and the revision is small compared to the biased news. In this case, the forecaster is anchored to its prior expectation and the forecast is relatively stable as the horizon shortens. Equation (2) can also be interpreted as an adaptive expectation mechanism as in Nerlove (1958), although in that literature the bias term $\mu_{i,h}$ is typically restricted equal to zero.

Overall, this simple model of learning with heterogeneous interpretation of public information provides a plausible representation of expectation formation and serves as a useful tool in understanding the dynamics of forecast disagreement and its persistence even at short horizons. Compared to the analysis in Kandel and Pearson (1995), that detects heterogeneity resulting from inconsistent revisions, the aim of this paper is, more generally, to detect also differences in the interpretation bias that might not lead to inconsistent revisions, in addition to gain insights on the role played by divergence in the interpretation weight as an explanation for forecasts dispersion.

2.2   Model Estimation

The empirical validation of the Bayesian learning model requires to specify the signal that is used by forecasters to form and revise their expectations. This is a difficult task since forecasters are likely to consider a large set of economic and financial variables that might play a role in predicting the evolution of the variable of interest. The approach in the recent literature (see Kandel and Zilberfarb, 1999, Lahiri and Sheng, 2008, and Patton and Timmermann, 2008) contemplates the assumption that the signal follows a latent process. Instead, we relax this assumption and conjecture that a component of the signal is indeed

observable by the econometrician and that it can be plausibly measured by the latest realization of the variable being forecast. This seems to be a reasonable assumption since most macroeconomic variables are (to some degree) persistent, so that forecasters might consider the current realization of the variable as an essential predictor of its future value. Moreover, the unobserved part of the signal captures the effect of news about other macroeconomic variables that the forecaster considers relevant in revising the prior forecast. We thus interpret the signal L_{t-h} in Equation (2) as the time $t-h$ announcement of the variable Y_{t-h}, while the effect of the unobserved part of the signal is captured by an error term as follows:

$$R_{t-h,t}^i = \alpha_{i,h} + \rho_{i,h} N_{t-h,t}^i + \epsilon_{t-h,t}^i \qquad (3)$$

where: $R_{t-h,t}^i$ represents the forecast revision at time $t-h$ about the realization of the variable at time t (that is, $Y_{t-h,t}^i - Y_{t-h-1,t}^i$), and $N_{t-h,t}^i$ indicates the news for forecaster i at time $t-h$ given by the difference between the latest announcement of the variable (the observable component of the signal) and the forecaster's prior expectation ($Y_{t-h} - Y_{t-h-1,t}^i$). The intercept $\alpha_{i,h}$ is equal to $-\rho_{i,h}\mu_{i,h}$, while $\rho_{i,h}$ represents the weight on the observable signal and carries a similar interpretation to the previous Section[3]. Equation (3) includes an error term $\epsilon_{t-h,t}^i$ that accounts for the effect of unobservable macroeconomic factors (denoted by F_{t-h}) on the revision, in addition to allowing for errors in the reporting of forecasts (denoted by $u_{t-h,t}^i$). We thus assume that the error term is given by $\epsilon_{t-h,t}^i = \gamma_{i,h}' F_{t-h} + u_{t-h,t}^i$, where $\gamma_{i,h}$ represents a vector of forecaster-specific parameters and $u_{t-h,t}^i$ is assumed to be normally distributed with mean zero, variance $\sigma_{i,h}^2$[4], and independent across forecasters, horizons, and time[5].

As mentioned earlier, the model assumes that the interpretation parameters vary across individuals and across horizons. The motivation for allowing them to vary across horizons

---

[3]The identification of $\mu_{i,h}$ requires $\rho_{i,h}$ to be larger than zero. This seems to be a reasonable assumption since a weight equal to zero implies that the forecaster believes the signal is totally uninformative (its precision is zero) and the posterior expectation is equal to the prior. We regard this as a possible issue at long forecast horizons when the informational content of the latest variable release might be minimal and forecasters are more likely to hold on to their prior beliefs. However, the first revision in the SPF occurs three quarters before the target date and it seems reasonable to assume that forecasters respond to the observed signal in revising their expectation.

[4]We assume that the variance of the reporting error varies across individuals and horizons. The justification for allowing the variance to change across horizons is that forecasters might be more careful in reporting their forecast as the target date is closer (in particular for the forecast made at time $t+1$ for target date t when the first release of the variable is publicly known).

[5]The assumption of no serial correlation does not contradict the earlier evidence of serial correlation in revisions based on efficiency tests for fixed-event forecasts as in Nordhaus (1987), Clements (1995), Clements (1997), and Isiklar et al. (2006) among others. While efficiency tests are typically based on evaluating the significance of past information (e.g., past revisions), in this paper we explain current revisions based on current information, such as the latest release of the variable of interest. We also evaluated empirically the assumption and concluded that there is no evidence of serial correlation (similarly to Lahiri and Sheng, 2008).

is due to the fixed-event nature of the SPF that allows to observe revisions for several horizons of the same target date. Hence, it is possible that, as the target approaches, forecasters give more weight to incoming information and that the bias may be reduced. In addition, the model allows the parameter to be forecaster-specific (in addition to horizon-specific), and this is justified by the aim of testing whether there is evidence to support the hypothesis that forecasters have homogeneous interpretation parameters.

To estimate the model we use the Common Correlated Effect (CCE) estimator proposed by Pesaran (2006). This estimator is suitable in this context since it accounts for the correlation of the errors across forecasters and across horizons due to unobservable factors, in addition to allowing the forecaster-specific news variable to be correlated with such unobserved factors (as it is likely to be the case). The approach consists of estimating Equation (3) by OLS with the individual-specific regressor augmented by auxiliary variables, such as the cross-sectional average of the dependent and independent variables (at time $t-h$). This allows to filter out the common effects of the unobserved factors and to obtain consistent estimates of the regression parameters. We can thus express the model in terms of a system of Seemingly Unrelated Regressions for each horizon h $(= -1, 0, 1, 2,$ and $3)$, with n indicating the number of forecasters as follows:

$$R_h = N_h \theta_h + U_h \tag{4}$$

where

$$R_h = \begin{bmatrix} R_{1,h} \\ R_{2,h} \\ \vdots \\ R_{n,h} \end{bmatrix}, \quad N_h = \begin{bmatrix} N_{1,h} & 0 & \cdots & 0 \\ 0 & N_{2,h} & \cdots & 0 \\ \vdots & & \ddots & \\ 0 & \cdots & 0 & N_{n,h} \end{bmatrix}, \quad \text{and } \theta_h = \begin{bmatrix} \theta_{1,h} \\ \vdots \\ \theta_{n,h} \end{bmatrix}$$

and the elements of R_h represent the revisions at horizon h of the n forecasters given by $R_{i,h} = (\cdots, R^i_{t-h,t}, \cdots)'$ (for $i = 1, \cdots, n$). Instead, the rows of $N_{i,h}$ are given by $(1, N^i_{t-h|t}, \overline{R}_{t-h|t}, \overline{N}_{t-h|t})$, where $\overline{R}_{t-h|t}$ and $\overline{N}_{t-h|t}$ represent the cross-sectional averages of the dependent and independent variables. The number of rows in R_h and N_h is equal to $T_h = \sum_{i=1}^n T_{i,h}$, where $T_{i,h}$ indicates the number of revisions available for forecaster i at horizon h[6]. The dimension of R_h is T_hx1 while for N_h is T_hx4n. The rows of $R_{i,h}$ (and $N_{i,h}$) might span different time periods across forecasters since they might not have reported forecasts in all available quarters. The 4n vector of parameters θ_h is equal to $(\theta_{1,h}, \cdots, \theta_{n,h})'$, where $\theta_{i,h} = (\alpha_{i,h}, \rho_{i,h}, \gamma^1_{i,h}, \gamma^2_{i,h})'$ (for $i = 1, \cdots, n$), and $\gamma^1_{i,h}$ and $\gamma^2_{i,h}$ denote the parameters

---

[6]Due to the entry and exit of forecasters and lack of reporting in some quarters, not all forecasters have a complete record of revisions. See Capistran and Timmermann (2009b) for a detail discussion of the issue.

of the cross-sectional averages.

In addition to estimating the forecaster-specific interpretation parameters, the aim of this paper is to test for the homogeneity of the parameters across forecasters for a given forecast horizon. More specifically, we are interested in testing the following hypotheses:

1. *Homogeneity of the bias parameter $\mu_{i,h}$*: since $\alpha_{i,h} = -\rho_{i,h}\mu_{i,h}$, the hypothesis can be stated as $H_0 : \alpha_{1,h}/\rho_{1,h} = \cdots = \alpha_{n,h}/\rho_{n,h}$

2. *Homogeneity of the weight parameter $\rho_{i,h}$*: in this case, the hypothesis to test can be simply stated as $H_0 : \rho_{1,h} = \cdots = \rho_{n,h}$

We test these hypotheses using a Wald-type test. The hypothesis of homogeneity of the interpretation bias can be stated in vector form as

$$
f(\theta_h) = \begin{bmatrix} \frac{\alpha_{1,h}}{\rho_{1,h}} - \frac{\alpha_{2,h}}{\rho_{2,h}} \\ \frac{\alpha_{2,h}}{\rho_{2,h}} - \frac{\alpha_{3,h}}{\rho_{3,h}} \\ \vdots \\ \frac{\alpha_{n-1,h}}{\rho_{n-1,h}} - \frac{\alpha_{n,h}}{\rho_{n,h}} \end{bmatrix}
$$

and the Jacobian matrix $R(\theta_h) = \partial f(\theta_h)/\partial \theta'_h$ is given by

$$
R(\theta_h) = \begin{bmatrix} 1/\rho_{1,h} & -\alpha_{1,h}/\rho_{1,h}^2 & 0 & 0 & -1/\rho_{2,h} & \alpha_{2,h}/\rho_{2,h}^2 & 0 & 0 & & \cdots & \cdots & 0 \\ 0 & 0 & 0 & 0 & 1/\rho_{2,h} & -\alpha_{2,h}/\rho_{2,h}^2 & 0 & 0 & -1/\rho_{3,h} & \alpha_{3,h}/\rho_{3,h}^2 & \cdots & 0 \\ \vdots & & & & & & \ddots & & & & & \\ 0 & \cdots & & & 1/\rho_{n-1,h} & -\alpha_{n-1,h}/\rho_{n-1,h}^2 & 0 & 0 & -1/\rho_{n,h} & \alpha_{n,h}/\rho_{n,h}^2 & 0 & 0 \end{bmatrix}
$$

The hypothesis can then be tested using the statistic $f(\hat{\theta}_h)' \left[R(\hat{\theta}_h)(N'_h \Sigma_h N_h)^{-1} R(\hat{\theta}_h)' \right]^{-1} f(\hat{\theta}_h)$ (where Σ_h denotes the variance-covariance matrix with diagonal elements $\sigma_{i,h}^2$) which is distributed as χ_{n-1}^2. For the hypothesis of homogeneity of the interpretation weight, the test is performed similarly with $f(\hat{\theta}_h) = R\hat{\theta}_h$ and the $(n-1)$x$4n$ restriction matrix R given by

$$
R = \begin{bmatrix} 0 & 1 & 0 & 0 & 0 & -1 & 0 & \cdots & \cdots & & & 0 \\ 0 & & \cdots & & 0 & 1 & 0 & 0 & 0 & -1 & \cdots & 0 \\ \vdots & & & & & & \ddots & & & & & \\ 0 & \cdots & \cdots & \cdots & & 1 & 0 & 0 & 0 & -1 & 0 & 0 \end{bmatrix}
$$

3 Application

3.1 Data

The Survey of Professional Forecasters (SPF) is the source of expectation data for this study. It was conducted by the National Bureau of Economic Research and the American Statistical Association (NBER-ASA) starting in the forth quarter of 1968 and since 1990 it is managed by the Federal Reserve Bank of Philadelphia (see Croushore, 1993, for a detailed discussion of the Survey). The SPF collects and publishes forecasts for many macroeconomic variables from a large panel of professional forecasters[7] in business, research institutes, and academia. Forecasters provide a number of 6 forecasts every quarter, ranging from the realization of the variable last quarter up to four quarters ahead, and they are formed after the realization of the variable for the previous quarter is released. In the case of the forecast for the previous quarter, it is actually formed after forecasters observe the first release of the variable for the target date and, to a large extent, the forecasts agree with the first release[8]. The forecast horizon h in the SPF takes values between -1 to 4, with $h = -1$ indicating the forecast made in quarter $t + 1$ for target quarter t. The revision is obtained from two consecutive forecasts for the same target date t. The first revision occurs at time $t - 3$ and the last revision happens in quarter $t + 1$ (when the first release of the variable is publicly known). The horizon h for revisions ranges thus between -1 and 3. In calculating the news variable, we use the real-time macroeconomic dataset collected by the Federal Reserve Bank of Philadelphia, see Croushore and Stark (2001) for a detailed discussion. The news variable in the current quarter is thus given by the first release of the variable (about last quarter) minus last quarter forecast about quarter t. Another relevant issue in the analysis of the SPF is that some forecasters might enter or exit the panel, while others might not provide forecasts in all quarters of their participation to the Survey. To avoid including forecasters with a limited number of forecasts, we consider only those forecasters that provided at least 180 revisions at all five horizons. This requirement reduces the sample of forecasters to between 18 and 22 forecasters with at least 36 forecasts per horizon[9].

---

[7]Contrary to other similar expectation surveys, the SPF does not reveal the identity of the forecaster, although it allows to track the forecaster over time. Since the beginning of 2000, the SPF provides a classification of forecasters by industry (either "financial service provider" or "non-financial service provider").

[8]There are various reasons why this might not be the case. First, forecasts are typically rounded to one decimal place for rates and no decimals for levels and indices. However, the data released might use a different rounding rule, for example to the second or higher decimal place. Second, forecasters might expect the first release to be revised in later quarters, as it happens regularly for NIPA variables and for the Industrial Production Index. In this case, the question is whether the forecasts represent predictions for the first announcement or the final value of the variable. However, although the forecasts might diverge from the first release data, the difference is typically quite modest and in an interval of $\pm$ 0.1% from the announced value.

[9]We also considered a lower threshold of 120 revisions and found qualitatively similar results.

In this paper we consider a subset of the macroeconomic variables available in the SPF that are the most often followed indicators of economic activity and price level. The variables are: real GDP (RGDP), Industrial Production Index (INDPROD), civilian Unemployment Rate (UNEM), CPI Inflation (CPI), GDP deflator (PGDP), and the T-bill rate (TBILL). Forecasts for UNEM and INDPROD are available since the start of the SPF in 1969, while the remaining variables were added to the SPF in 1982. In the empirical part, we estimate the Bayesian learning model on the revisions on the common to all variable period that starts in the first quarter of 1982 and ends in the last quarter of 2007. We decided to drop from the analysis the second quarter of 1990 because the Federal Reserve Bank of Philadelphia assumed then the management of the Survey and the forms were sent out late with a limited number of responses.

For some variables (such as RGDP, INDPROD, and PGDP) forecasters were asked to predict the level of the variables, which are commonly considered as non-stationary. We thus define the revision and news variable in percentage changes as follows

$$R^i_{t-h,t} = 100 \frac{Y^i_{t-h,t} - Y^i_{t-h-1,t}}{Y^i_{t-h-1,t}} \qquad \text{and} \qquad N^i_{t-h,t} = 100 \frac{Y_{t-h} - Y^i_{t-h-1,t}}{Y^i_{t-h-1,t}}$$

Instead, UNEM, TBILL, and CPI inflation can be considered stationary and we use the definition of revision and news introduced in Section (2). In addition, the base-year for NIPA variables and Industrial Production Index changed several times in the period considered. We thus re-base forecasts and released data to a common 2000 dollars base.

3.2 Empirical results

Figure (1) and (2) show the evolution of $\hat{\mu}_{i,h}$ (obtained as $-\hat{\alpha}_{i,h}/\hat{\rho}_{i,h}$) and $\hat{\rho}_{i,h}$ for each forecaster as the horizon decreases from 3 to -1 (gray lines) together with a box-plot of the cross-sectional distribution of the estimated parameters. Some features of the evolution of the cross-sectional distribution of the estimates are common across the six variables considered. The median of the estimated interpretation bias ($\hat{\mu}_{i,h}$) is close to zero at all horizons for most variables, with the only exception of RGDP and PGDP for which it is slightly negative at the longer horizons. In addition, its cross-sectional dispersion tends to decrease as the horizon shortens (for $h = -1$ forecasters know the first release of the variables and the dispersion is negligible). The dispersion of the distribution of the $\hat{\mu}_{i,h}$ for $h \geq 0$ provides evidence that forecasters have differential interpretation of the signal, with approximately half of the forecasters having positive bias and the other half negative. The magnitude of the bias can also be substantial (for some forecasters as large as $\pm 1.4\%$ at longer horizons).

The distribution of $\hat{\rho}_{i,h}$ across the different variables also shares some common characteristics, such as the median (of the cross-sectional distribution) converging toward one as the revision date approaches the target date, and the dispersion among forecasters decreasing with the horizon. These results suggest that at two and three quarter horizons forecasters assign a weight to the latest variable release between 40 and 80%, while prior expectations still account for approximately 20 to 60%. However, approaching the target date the weight on the news increases and, for $h = -1$, the role of the prior expectation is irrelevant[10]. In addition, the considerable dispersion of the estimated weights across forecasters indicates that differences in the interpretation weight might play a significant role in driving the heterogeneity of forecasts, and thus revisions.

Despite these general features, the distribution of the interpretation parameters has characteristics specific to some of the macroeconomic variables considered. Figure (1c) shows the results for the unemployment rate (UNEM). Two elements seem to characterize UNEM compared to the other variables. First, the dispersion of $\hat{\mu}_{i,h}$ appears to be relatively small, ranging from $\pm 0.2\%$ at the three quarters horizon, and even smaller closer to the target date. This is probably due to the higher persistence of the unemployment rate in comparison to the other variables considered. Second, forecasters seem to disagree more (compared to other variables) on the weight given to the news: at the three quarters horizon the estimated $\rho_{i,h}$ ranges between 0.3 and 0.75 and between 0.6 and 1 for the quarter before the preliminary estimate is released ($h = 0$). This suggests that forecasters interpret differently the data release, mostly in terms of the precision they assign to the signal rather than in terms of a systematic (horizon-specific) bias. Another variable with specific features is CPI. Considering first the interpretation bias $\hat{\mu}_{i,h}$, Figure (2a) suggests that the cross-sectional dispersion decreases only marginally close to the target date, instead of rapidly decreasing as for the other variables. For $h = 0$, most forecasters have an estimated bias within an interval that ranges between -0.4% and 0.4%, which can be considered relevant given that the preliminary estimate will be released the following quarter. Furthermore, the cross-sectional distribution of the interpretation weights shows some unique features: the median weight for $h = 3$ is 0.35, with some forecasters having a weight as low as 0.1 and others close to 0.4. This indicates that, three quarters before the target date, the (median) forecaster assigns 65% weight to the prior expectation and 35% to the CPI data recently released, while it increases to 50% the quarter before the target date. In addition, although the median

---

[10]These findings confirm those in Lahiri and Sheng (2008) and Patton and Timmermann (2008) that, at longer horizons, forecasters give higher weight to their prior expectation relative to the real-time signal, while the latter becomes more important closer to the target date.

weight assigned to the signal increases for $h = 0$ to 80%, the dispersion among forecasters is substantial with some $\rho_{i,h}$ as low as 0.4. On the other hand, when considering the GDP deflator (PGDP) as an alternative measure of the price level, we find that the estimated weights behave similarly to most other variables, i.e., the bias converges to zero and the weights toward one. When considering the TBILL, it can be noticed that for revisions made at $h = 0$ and 1 the interpretation weight is larger than one for few forecasters, while they converge to one for $h = -1$[11].

Statistical testing of the hypothesis of homogeneous interpretation is provided in Table (1) which reports the test statistic and p-value of the null hypothesis of homogeneity of the interpretation bias and weight (at all forecast horizons). Overall, the hypothesis of homogeneity of the interpretation bias and weight is rejected very strongly for $h \geq 1$ for most variables and both participation levels, thus supporting the assumption that forecasters interpret public information differently. At the shorter horizons ($h = 0, -1$) the evidence of heterogeneity is weaker, as in the case of INDPROD, where the null of homogeneity for both parameters is not rejected. For revisions at horizon $h = -1$ the only rejections of homogeneity in the bias $\mu_{i,h}$ occur for RGDP and CPI[12]. Furthermore, the evidence of heterogeneity in the weight indicates that it represents a relevant source driving the dispersion of revisions (and forecasts), not only at long-horizons but it also persists for revisions made at the target date ($h = 0$). While heterogeneity in the bias among forecasters is limited close to the target date, the dispersion in weights is still relevant and constitutes a main force behind the persistence of beliefs dispersion among forecasters.

4   Discussion

The results in the previous Section point to the significance of differential interpretation of the real-time signal, both in terms of the specific bias of a forecaster as well as the weight given to the signal. An interesting issue is characterizing the forecasters in terms of their interpretation parameters and how these parameters affect their forecast precision.

---

[11]The fact that some forecasters have interpretation weight larger than one is not consistent with the Bayesian learning model discussed earlier, and it can be due to either estimation error or interpreted as over-reaction to news in revising their expectations close to the target date.

[12]Since the revision at $h = -1$ occurs after the forecaster observes the release of the variable for the target quarter, the news ($N_{t+1|t}^i$) should equal the revision ($R_{t+1|t}^i$) and the bias should be homogeneous and equal to zero. However, looking at Figure (1a) the estimated bias for RGDP is limited to a few forecasters and smaller than $\pm$ 0.1% in magnitude. For CPI (shown in Figure 2a) the heterogeneity found at $h = -1$ seems to be due to a small portion of forecasters that have a systematic bias of the order of $\pm$ 0.1 to 0.2%. A possible explanation is that for some variables, in particular CPI, a few forecasters might have misinterpreted the variable definition which was being forecast.

4.1 Interpreting the interpretation parameters

Figure (3) and (4) show scatter plots of the average (over time) deviation of the news of a forecaster from the cross-sectional median[13] against the estimated bias. They show a positive relationship between the variables for UNEM, CPI, and TBILL while the dependence seems to hold only at the shortest horizons for RGDP and INDPROD.

We can illustrate graphically the implications of this finding. In Figure (7) we present a situation in which two forecasters have prior expectations denoted by $Y^1_{t-h-1,t}$ for forecaster 1, and $Y^2_{t-h-1,t}$ for forecaster 2. We assume that forecaster 2 is more optimistic about the future realization of the variable in quarter t compared to the other forecaster ($Y^1_{t-h-1,t} < Y^2_{t-h-1,t}$). At the beginning of quarter $t - h$ the announcement for the previous quarter is released. In the graph we assume that the announcement falls between the prior expectation of the forecasters, such that forecaster 1 receives a positive news while forecaster 2 has a negative news. The finding of a positive relation between average news and bias suggests that, in this example, we would expect forecaster 1 to have positive interpretation bias (denoted by $\mu_{1,h}$) and forecaster 2 a negative one (similarly denoted as $\mu_{2,h}$). Positive (negative) interpretation bias means that the forecaster believes the announcement is an overly optimistic (pessimistic) signal for the realization of the variable in quarter t. The posterior expectation of forecaster 1 thus lies between the prior ($Y^1_{t-h-1,t}$) and the biased signal ($L_{t-h} - \mu_{1,h} < L_{t-h}$), while for forecaster 2 the posterior will lie between $L_{t-h} - \mu_{2,h}$ ($> L_{t-h}$) and $Y^2_{t-h-1,t}$ (the exact location of the posteriors will depend on the value of the weight given to the signal). It is thus the case that both forecasters revise their priors in the direction of the announcement, although the heterogeneity in their interpretation bias delays the process of convergence of expectations compared to a situation in which forecasters interpret homogeneously the signal. In the discussed example, homogeneity in the interpretation bias (i.e. $\mu_{1,h} = \mu_{2,h} = \mu_h$) implies that the posterior expectations will lie in the interval between the priors and the common biased signal $L_{t-h} - \mu_h$ so that it is more likely that posterior expectations converge more rapidly compared to the heterogeneous bias situation (and depending on the weights assigned to the announcement).

Figure (5) and (6) show scatter plots of the estimated interpretation weight $\hat{\rho}_{i,h}$ against the estimated bias $\hat{\mu}_{i,h}$. At longer horizons of two and three quarters, there seems to be negative dependence between the two parameters, namely for RGDP, INDPROD, PGDP

---

[13]The average news deviation from the median is intended to characterize the prior sentiment of a forecaster with respect to the median forecaster in the sense of being, on average, more optimistic or pessimistic than the median.

and, to a lesser extent, CPI. In other words, it appears that forecasters that have positive bias are more likely to give a low weight (in the unit interval) to the signal, and vice versa for those with negative bias. The group of forecasters with positive bias and low weight (and likely to have pessimistic priors as we discussed above) can be described as doubtful of the information content of the signal: the low weight means that they believe the precision of the signal to be low compared to their prior, while the positive bias indicates that they expect the signal to overestimate the future realization. This skepticism means that these forecasters adjust downward the signal (since they believe it is biased) and assign a low weight to this information, preferring to hold on to their prior expectation. On the other hand, the forecasters with negative bias and (relatively) high weight tend to adjust upward the signal (since they believe it underestimates the true value of the variable) and assign high precision to the incoming information relatively to their priors.

4.2   Heterogeneity and forecast error

Until now the focus was on the learning process used by forecasters to incorporate new information as it becomes available. An interesting question to evaluate is to what extent the parameters of the likelihood function of the individuals affect the accuracy of their forecasts. We use the estimated parameters as explanatory variables in the following regression:

$$rmse_{i,h} = \gamma_0 + \gamma_1|\hat{\mu}_{i,h}| + \gamma_2\hat{\rho}_{i,h} + \xi_{i,h} \tag{5}$$

where $rmse_{i,h}$ represents the Root Mean Square Error of forecaster i at horizon h, $\hat{\mu}_{i,h}$ and $\hat{\rho}_{i,h}$ are the estimated interpretation bias and weight, and γ_k ($k = 0, 1, 2$) are parameters to be estimated. We include the bias parameter in absolute value since it is likely that the magnitude of the bias determines the accuracy of forecasters' expectations. In addition, we also consider the case of fixed forecaster effects in the regression, that is,

$$rmse_{i,h} = \gamma_{0,i} + \gamma_1|\hat{\mu}_{i,h}| + \gamma_2\hat{\rho}_{i,h} + \xi_{i,h} \tag{6}$$

In this case, we account for the forecaster's characteristics that are constant across horizons, such as the optimism/pessimism in the prior expectation relative to the other forecasters in the sample.

The results are reported in Table (2). For most variables, we find that γ_1 and γ_2 are strongly significant (see column F for the joint test of their significance), except in the case of CPI. For this variable, forecast accuracy is unrelated to the interpretation parameters. For all other variables we find that larger (absolute) bias is associated with worse forecast accuracy

while faster incorporation of recent news (high $\rho_{i,h}$) improves accuracy. Overall, the results indicate that the more accurate forecasters are characterized by a bias close to zero and higher (in the unit interval) weight given to the incoming information.

5 Conclusion

The Bayesian learning model adopted in this paper predicts that dispersion in forecasts is driven by two forces, heterogeneity in prior expectations and differential interpretation of news across forecasters. In a fixed-event setting, differences in priors and thus dispersion in forecasts tend to vanish as the horizon shortens when forecasters interpret homogeneously current events. On the other hand, heterogeneous interpretation of the incoming information might delay the convergence or even exacerbate the dispersion of forecasts. The model assumes that the forecast revision, following a public signal, depends upon two parameters: the forecaster's interpretation of the signal as a predictor of the future realization of the variable, and a weight that captures the confidence that they assign to the signal.

Estimation of these parameters shows considerable diversity in interpretation across forecasters at the longer horizons while it attenuates close to the target date. In terms of interpretation of the mean of the signal, the results suggest that some forecasters expect the variable at the target date to be higher compared to the signal, while others have a pessimistic interpretation of the signal. We also notice that there is, for some of the macroeconomic variables considered, a positive relationship between interpretation of the mean signal and the prior sentiment. An individual with prior forecast above the median is more likely to expect the signal to provide a conservative indication of the variable at the target date, while a forecaster with prior below the median believes the opposite. Thus, the effect of the interpretation bias implies a delay in the convergence of forecasts, in comparison to a situation in which agents have homogeneous interpretations of the signal. This mechanism can therefore be associated with the observed stickiness of forecasts, in the sense that the persistence in dispersion might be attributed to agents holding on to their (heterogeneous) beliefs when interpreting the new information available.

Furthermore, the results for the weight parameter also suggest that forecasters are heterogeneous, in the sense of attributing different weights to the signal when revising their prior forecast. As the horizon shortens, the cross-sectional distribution of the weights tends toward one and its dispersion attenuates, although for CPI inflation we find that it widens close to the target date. This is probably related to the finding of "excess dispersion" in

inflation forecasts at short horizons in Patton and Timmermann (2008). Based on our analysis, the (relatively) high dispersion in forecasts could be explained by heterogeneity in the interpretation of the relevance of the news, with some forecasters assigning a 40% weight while others assign close to 100%, meaning that they fully incorporate the signal in their posterior forecast. Another interesting result is that forecasters that interpret the signal pessimistically (that is, they expect the future realization to be lower than today's signal) are reluctant to incorporate a large fraction of the signal in their forecast, and rather hold on to their prior expectation. On the other hand, those agents with an optimistic interpretation of the signal are more likely to assign a large weight to the signal. Therefore, there seems to be an asymmetry in the forecast stickiness aforementioned, in the sense that pessimistic forecasters are slower in accounting for new information compared to those with rosier interpretations of the signal.

Overall, these findings provide evidence to suggest that differential interpretation of new information is an important element in explaining dispersion in expectations amongst agents, in addition to the dispersion due to diversity in prior beliefs. Although the main focus of the paper is to investigate heterogeneity in the way forecasters interpret the incoming information and revise their expectations, it can be argued that the Bayesian learning model provides a simplistic representation of the way forecasters form their expectations. In particular, our assumption that identifies the signal with the latest announcement of the variable neglects to consider several other macroeconomic indicators that forecasters might regard as relevant, especially at the longest horizons. For example, when revising their four-quarter ahead Treasury-bill rate expectation forecasters are likely to consider also the recent announcements for real GDP growth and inflation measures. In this paper we decided to focus on the news about the variables being forecast and treated every other news as unobserved factors common to all forecasters. While we still achieve consistent estimates of the effect of news, we are not able to bring to light the possibly more complicated way in which forecasters use available information (beyond news about the predicted variable) to form and revise their expectations. Furthermore, we allow the individual-specific interpretation parameters to vary across forecasters and horizons, but we restrict them to be constant over time. On the other hand, there could be time-variation in the forecasters' interpretation of the news in the sense that they weight the current news differently depending, e.g., on the state of the business cycle. These considerations suggest possible extensions toward a more elaborate model of expectation formation in economics that we intend to consider in future research.

References

Bonham, C.S. and Cohen, R.H. (2001). To aggregate, pool, or neither: Testing the rational-expectations hypothesis using survey data. *Journal of Business & Economic Statistics*, **19,** 278–291.

Branch, W.A. (2004). The theory of rationally heterogeneous expectations: Evidence from survey data on inflation expectations. *Economic Journal*, **114,** 592–621.

Capistran, C. and Timmermann, A. (2009a). Disagreement and biases in inflation expectations. *Journal of Money, Credit, and Banking*, **41,** 365–396.

Capistran, C. and Timmermann, A. (2009b). Forecast combination with entry and exit of experts. *Journal of Business and Economic Statistics*, **27,** 428–440.

Carroll, C.D. (2003). Macroeconomic expectations of households and professional forecasters. *Quarterly Journal of Economics*, **118,** 269–298.

Clements, M.P. (1995). Rationality and the role of judgment in macroeconomic forecasting. *Economic Journal*, **105,** 410–420.

Clements, M.P. (1997). Evaluating the rationality of fixed-event forecasts. *Journal of Forecasting*, **16,** 225–239.

Croushore, D. (1993). Introducing: The Survey of Professional Forecasters. *Federal Reserve Bank of Philadelphia Business Review*, **11/12,** 3–13.

Croushore, D. and Stark, T. (2001). A real-time data set for macroeconomists. *Journal of Econometrics*, **105,** 111–130.

Ehrbeck, T. and Waldmann, R. (1996). Why are professional forecasters biased? Agency versus behavioral explanations. *Quarterly Journal of Economics*, **111,** 21–40.

Hommes, C.H. (2006). Heterogeneous agents models in economics and finance. In *Handbook of Computational Economics, Vol 2* (eds L. Tesfatsion and K.L. Judd). Elsevier.

Isiklar, G., Lahiri, K. and Loungani, P. (2006). How quickly do forecasters incorporate news? Evidence from cross-country surveys. *Journal of Applied Econometrics*, **21,** 703–725.

Kandel, E. and Pearson, N.D. (1995). Differential interpretation of public signals and trade in speculative markets. *Journal of Political Economy*, **103,** 831–872.

Kandel, E. and Zilberfarb, B-Z. (1999). Differential interpretation of information in inflation forecasts. *Review of Economics and Statistics*, **81,** 217–226.

Keane, M.P. and Runkle, D.E. (1990). Testing the rationality of price forecasts: new evidence from panel data. *American Economic Review*, **80,** 714–735.

Lahiri, K. and Sheng, X. (2008). Evolution of forecast disagreement in a Bayesian learning model. *Journal of Econometrics*, **144,** 325–340.

Laster, D., Bennett, P. and Geoum, I.S. (1999). Rational bias in macroeconomic forecasts. *Quarterly Journal of Economics*, **114,** 293–318.

Mankiw, N.G. and Reis, R. (2002). Sticky information versus sticky prices: A proposal to replace the New Keynesian Phillips curve. *Quarterly Journal of Economics*, **117,** 1295–1328.

Mankiw, N.G., Reis, R. and Wolfers, J. (2003). Disagreement about inflation expectations. *NBER Macroeconomic Annual*, **18,** 209–248.

Nerlove, M. (1958). Adaptive expectations and cobweb phenomena. *Quarterly Journal of Economics*, **72,** 227–240.

Nordhaus, W.D. (1987). Forecasting efficiency: concepts and applications. *Review of Economics and Statistics*, **69,** 667–674.

Ottaviani, M. and Sorensen, P.N. (2006). The strategy of professional forecasters. *Journal of Financial Economics*, **81,** 441–466.

Patton, A.J. and Timmermann, A. (2008). Why do forecasters disagree? Lessons from the term structure of cross-sectional dispersion. *working paper*.

Pesaran, M.H. (2006). Estimation and inference in large heterogeneous panels with a multifactor error structure. *Econometrica*, **74,** 967–1012.

Pesaran, M.H. and Weale, M.R. (2006). Survey expectations. In *Handbook of Economic Forecasting* (eds G. Elliott, C.W.J. Granger and A. Timmermann). Elsevier, North-Holland.
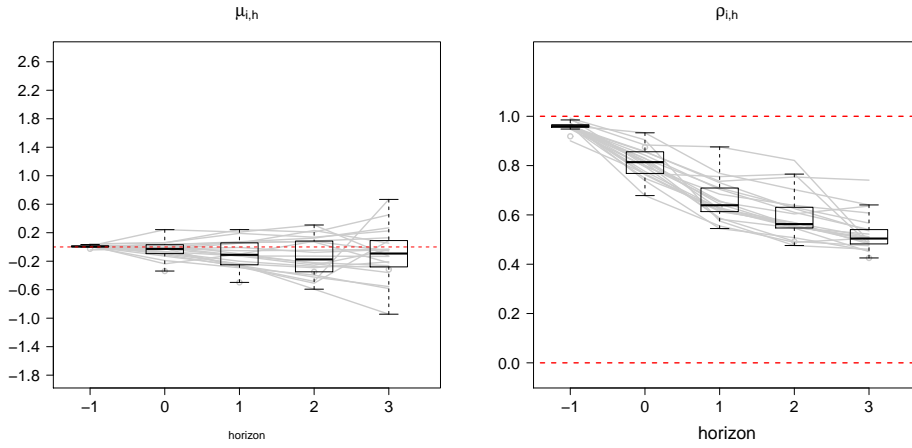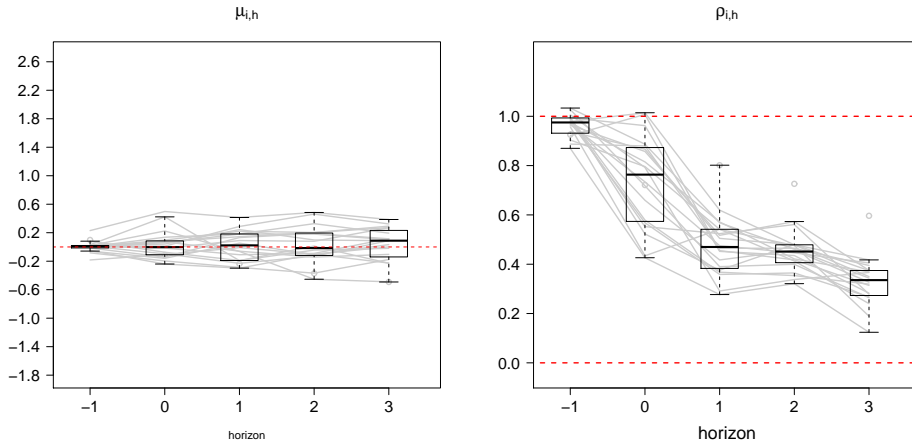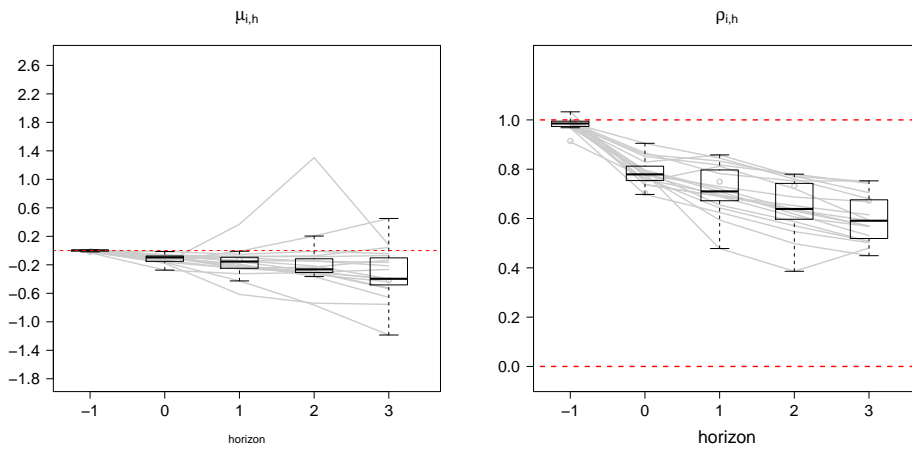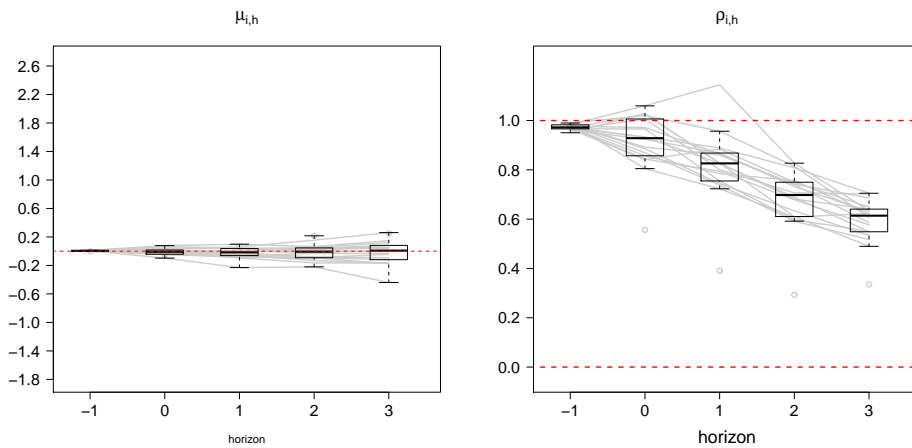
(a) RGDP
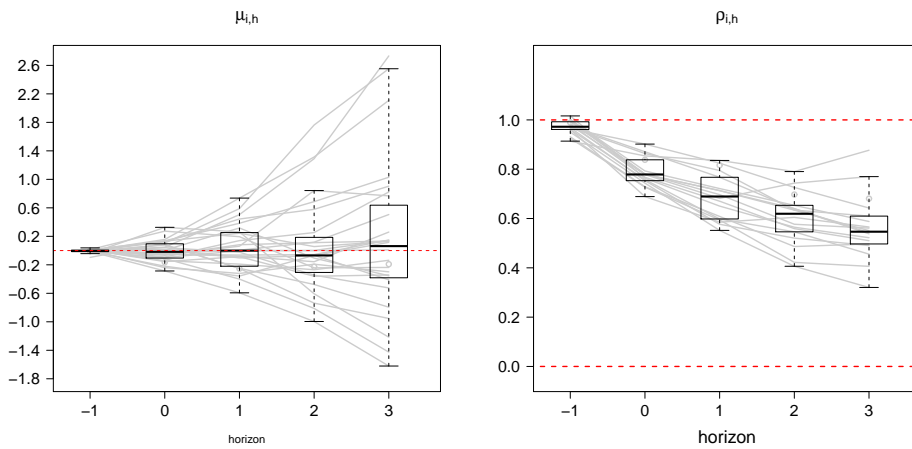
(b) INDPROD
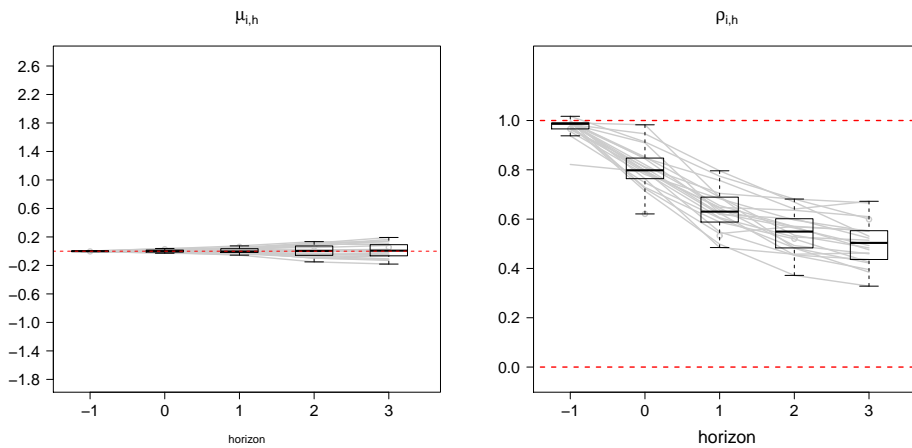
(c) UNEM

Figure 1: The plots show the evolution of the (estimated) interpretation parameters (bias on the left column and weight on the right) as the horizon varies from 3 (quarters ahead) to -1 (the quarter following the target date when the first release of the macroeconomic variable is known). Each gray line represents the evolution of $\hat{\mu}_{i,h}$ and $\hat{\rho}_{i,h}$ for a forecaster as h converges to -1. The forecasters included are those that have reported at least 180 revisions in the period 1982-2007. The box plot shows the median and the interquartile range, while the whiskers represent a two standard deviation interval around the mean.

21

(a) CPI

(b) PGDP

(c) TBILL

Figure 2: The plots show the evolution of the (estimated) interpretation parameters (bias on the left column and weight on the right) as the horizon varies from 3 (quarters ahead) to -1 (the quarter following the target date when the first release of the macroeconomic variable is known). Each gray line represents the evolution of $\hat{\mu}_{i,h}$ and $\hat{\rho}_{i,h}$ for a forecaster as h converges to -1. The forecasters included are those that have reported at least 180 revisions in the period 1982-2007. The box plot shows the median and the interquartile range, while the whiskers represent a two standard deviation interval around the mean.

| Variable	N	Test	h = -1	h = 0	h = 1	h = 2	h = 3
RGDP	22	hom. $\mu_{i,h}$	37.13	83.26	50.07	37.81	39.63
			[0.016]	[0.000]	[0.000]	[0.014]	[0.008]
		hom. $\rho_{i,h}$	18.48	34.72	44.08	39.98	32.51
			[0.618]	[0.030]	[0.002]	[0.007]	[0.052]
INDPROD	18	hom. $\mu_{i,h}$	10.00	14.74	24.52	26.83	57.42
			[0.904]	[0.614]	[0.106]	[0.061]	[0.000]
		hom. $\rho_{i,h}$	17.02	12.59	30.78	41.30	64.12
			[0.453]	[0.763]	[0.021]	[0.001]	[0.000]
UNEM	22	hom. $\mu_{i,h}$	20.35	27.44	29.55	55.39	57.89
			[0.499]	[0.157]	[0.102]	[0.000]	[0.000]
		hom. $\rho_{i,h}$	46.53	28.99	31.17	33.11	36.33
			[0.001]	[0.114]	[0.071]	[0.045]	[0.020]
CPI	20	hom. $\mu_{i,h}$	29.56	21.99	41.88	51.21	25.24
			[0.058]	[0.285]	[0.002]	[0.000]	[0.153]
		hom. $\rho_{i,h}$	45.98	62.58	48.97	30.21	28.89
			[0.000]	[0.000]	[0.000]	[0.049]	[0.068]
PGDP	19	hom. $\mu_{i,h}$	17.31	18.16	29.99	17.34	19.76
			[0.502]	[0.445]	[0.038]	[0.500]	[0.347]
		hom. $\rho_{i,h}$	21.56	26.27	43.41	43.92	34.36
			[0.252]	[0.094]	[0.001]	[0.001]	[0.011]
TBILL	18	hom. $\mu_{i,h}$	16.85	70.71	35.82	30.93	43.13
			[0.465]	[0.000]	[0.005]	[0.020]	[0.000]
		hom. $\rho_{i,h}$	26.99	181.64	172.18	86.66	39.70
			[0.058]	[0.000]	[0.000]	[0.000]	[0.001]

Table 1: Results of the test for homogeneity of $\mu_{i,h}$ and $\rho_{i,h}$. Column N reports the number of forecasters with at least 180 revisions across all forecast horizons, and the last five columns report the test statistic (and p-value in parenthesis) for horizon $h = -1, 0, 1, 2, 3$.
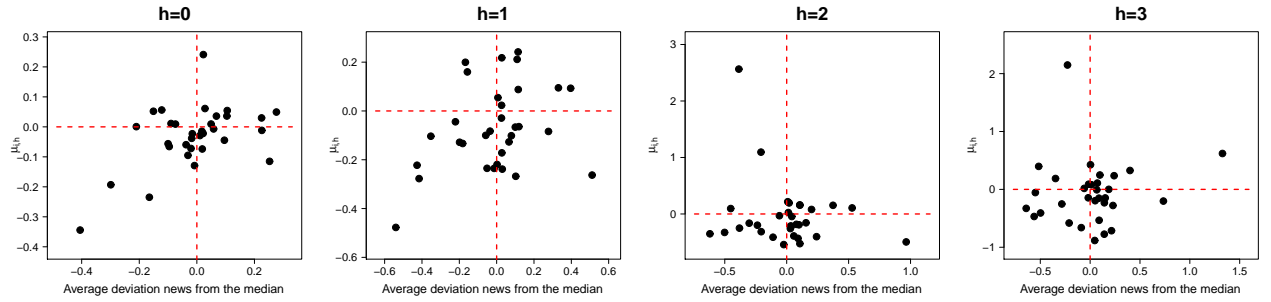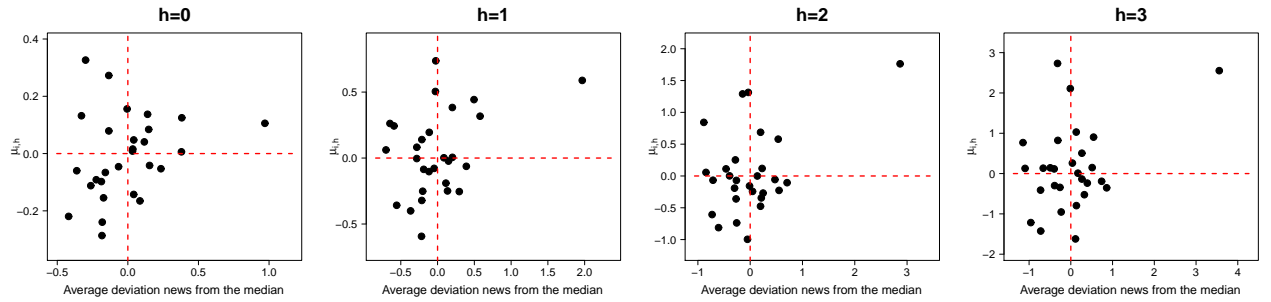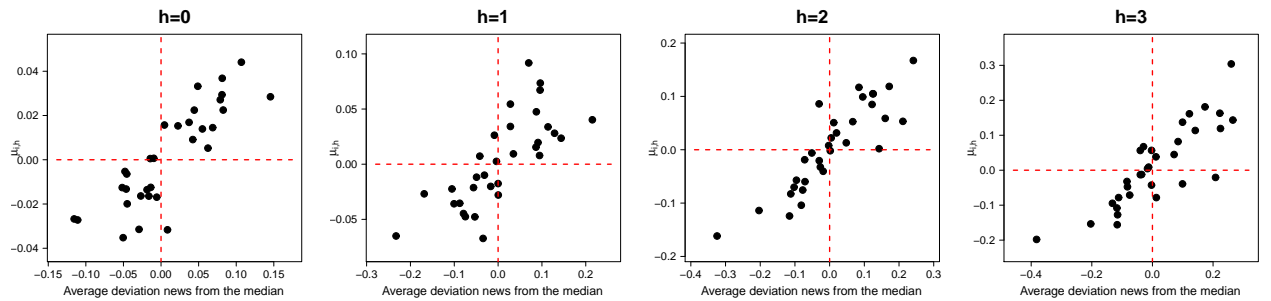
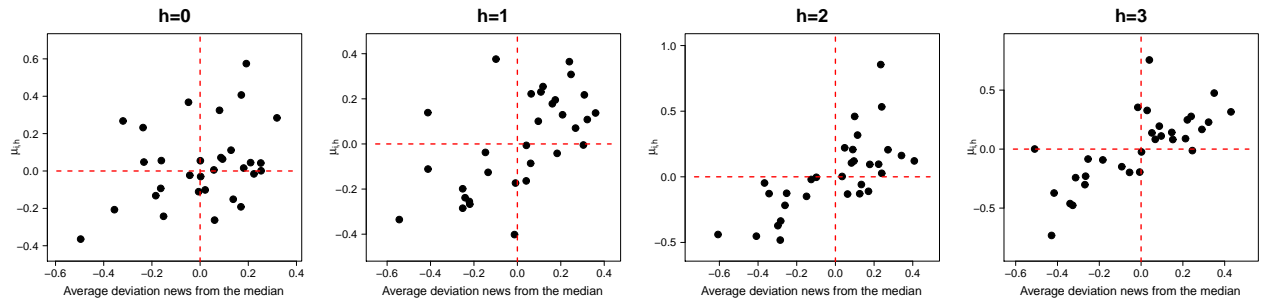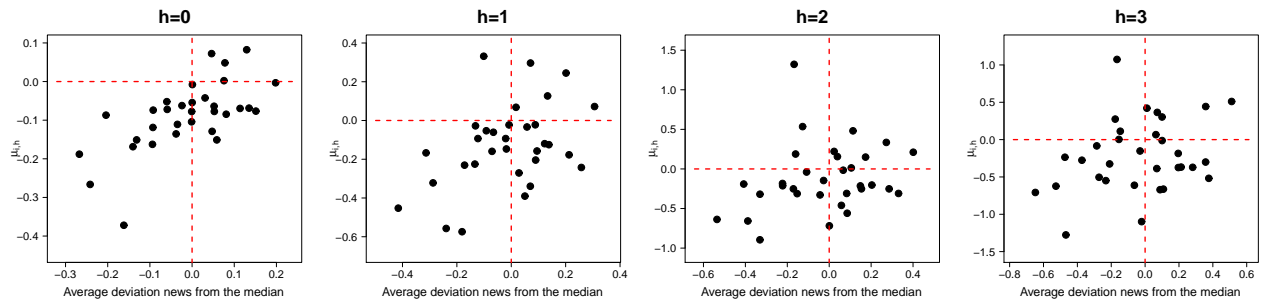Figure 3: For each variable, the Figure shows scatter plots of the estimated bias, $\hat{\mu}_{i,h}$, against the average (over time) deviation of the forecaster's prior from the median. A positive value for the deviation indicates that the forecaster has, on average, higher prior expectation compared to the median. In all plots, the vertical dashed line indicates zero - the median prior - while the horizontal dashed line shows zero for the bias.
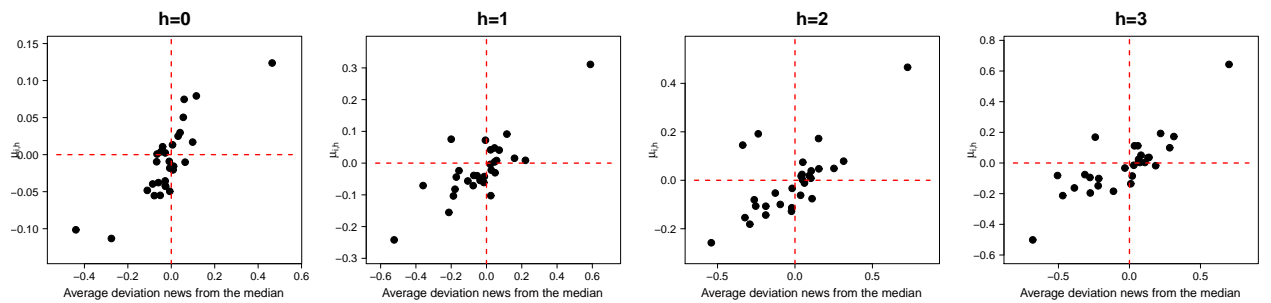
(a) CPI

(b) PGDP

(c) TBILL

Figure 4: For each variable, the Figure shows scatter plots of the estimated bias, $\hat{\mu}_{i,h}$, against the average (over time) deviation of the forecaster's prior from the median. A positive value for the deviation indicates that the forecaster has, on average, higher prior expectation compared to the median. In all plots, the vertical dashed line indicates zero - the median prior - while the horizontal dashed line shows zero for the bias.
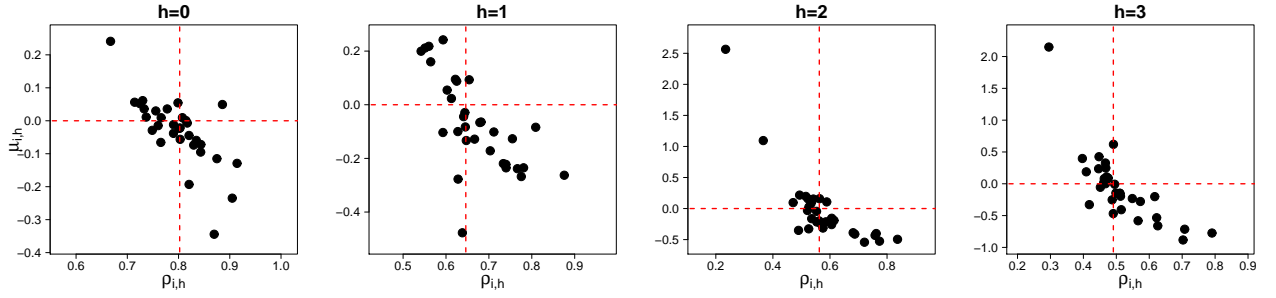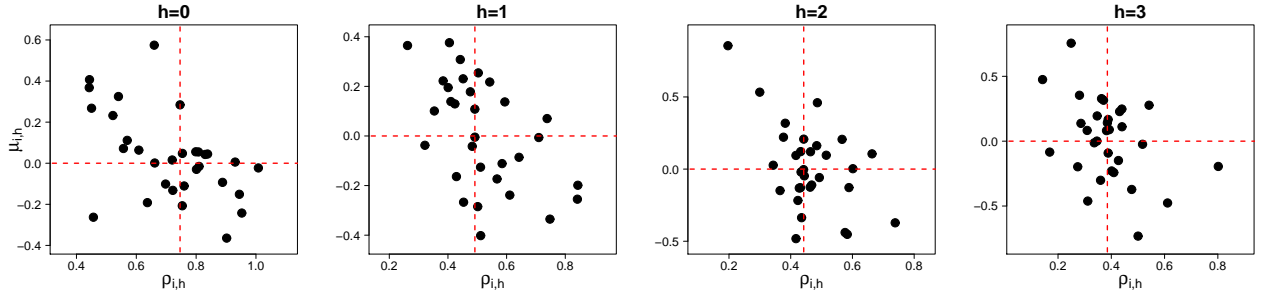
(a) RGDP
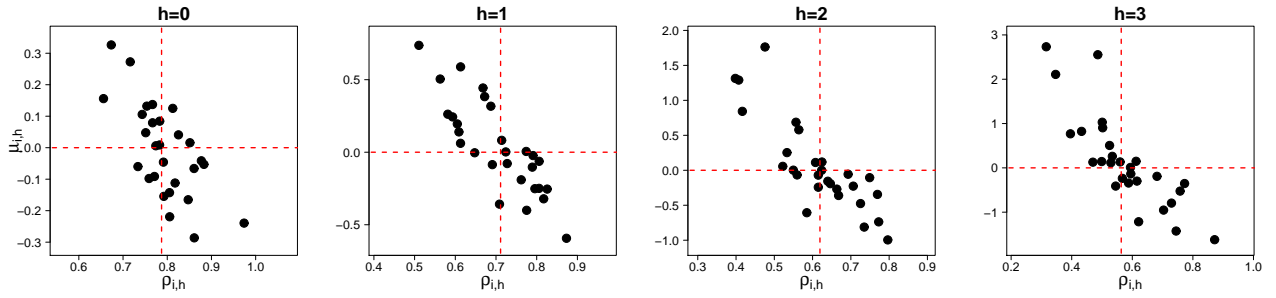
(b) INDPROD
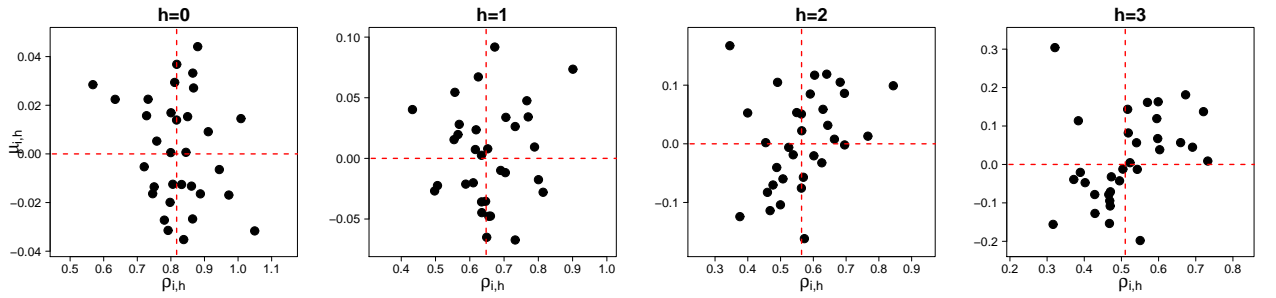
(c) UNEM

Figure 5: The Figure shows scatter plots of the weight, $\hat{\rho}_{i,h}$, against the bias, $\hat{\mu}_{i,h}$, for each forecaster included in the sample with the participation threshold equal to 180 revisions. The vertical dashed line represents the median $\rho_{i,h}$ and the horizontal dashed line is drawn at zero.
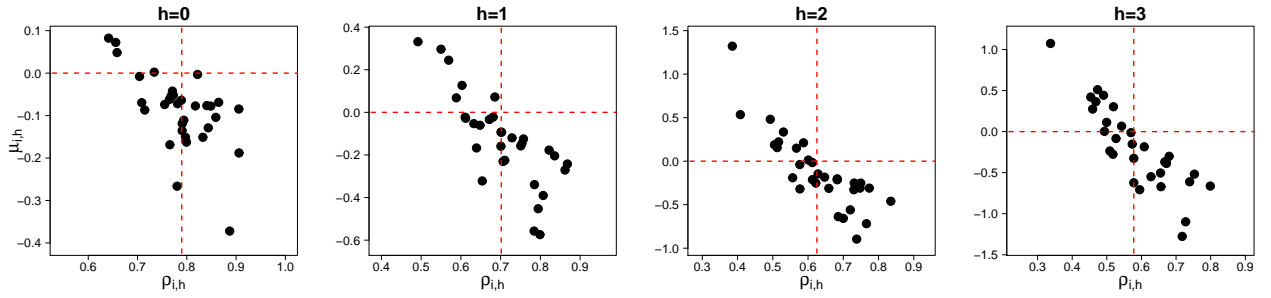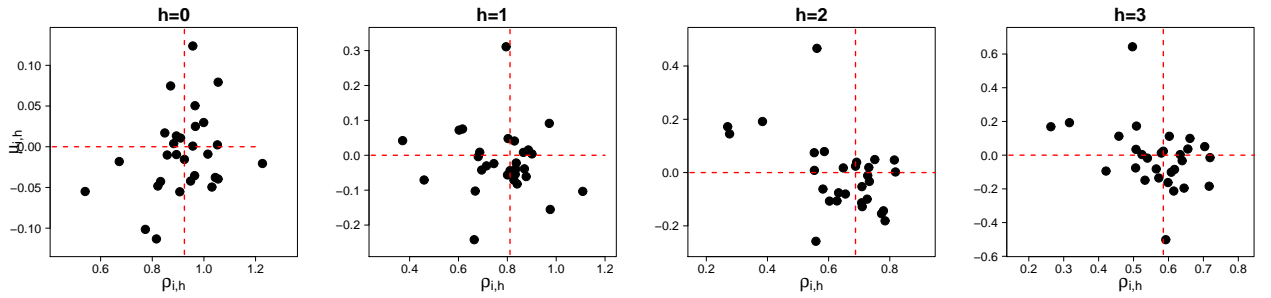
26

(a) CPI

(b) PGDP

(c) TBILL

Figure 6: The Figure shows scatter plots of the weight, $\hat{\rho}_{i,h}$, against the bias, $-\hat{\mu}_{i,h}$, for each forecaster included in the sample with the participation threshold equal to 180 revisions. The vertical dashed line represents the median $\rho_{i,h}$ and the horizontal dashed line is drawn at zero.

	$|\alpha_h^i|$	$\rho_{i,h}$	FE	F
RGDP	0.208	-3.006	no	219.61
	[0.04]	[0]		[0]
	0.104	-3.475	yes	15.8
	[0.25]	[0]		[0]
INDPROD	1.416	-4.039	no	172.79
	[0]	[0]		[0]
	1.167	-4.765	yes	7.26
	[0]	[0]		[0.001]
UNEM	1.386	-0.723	no	176.84
	[0]	[0]		[0]
	1.457	-0.791	yes	16.26
	[0]	[0]		[0]
CPI	0.392	0.179	no	3.77
	[0.008]	[0.082]		[0.025]
	0.127	0.034	yes	0.01
	[0.205]	[0.593]		[0.99]
PGDP	0.341	-1.548	no	156.85
	[0]	[0]		[0]
	0.232	-1.972	yes	14.08
	[0.007]	[0]		[0]
TBILL	1.475	-1.462	no	140.33
	[0]	[0]		[0]
	0.581	-1.856	yes	6.82
	[0.024]	[0]		[0.002]

Table 2: Estimation results for the model $rmse_{i,h} = \gamma_{0,i} + \gamma_1|\hat{\mu}_{i,h}| + \gamma_2\hat{\rho}_{i,h} + \xi_{i,h}$. We estimate the model by pooling across forecasters (in which case $\gamma_{0,i} = \gamma_0$) and with fixed forecaster effects. The last column reports the F-test for the null hypothesis that $\gamma_1 = \gamma_2 = 0$. The results are based on the estimated $\mu_{i,h}$ and $\rho_{i,h}$ with a participation threshold of 120 revisions.
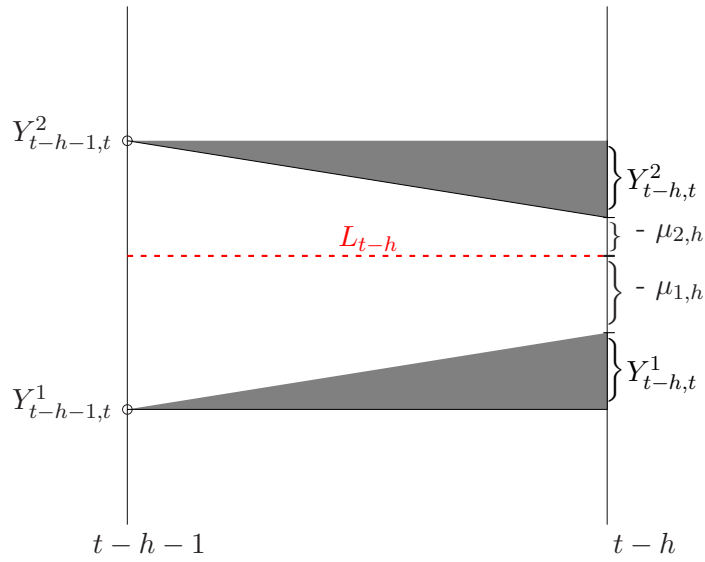
Figure 7: Illustration of the Bayesian learning model where forecaster 1 has prior expectation at time $t-h-1$ lower than forecaster 2 $(Y^1_{t-h-1,1} < Y^2_{t-h-1,1})$, the signal is denoted by L_{t-h} and the forecasters have different bias $(\mu_{2,h} < 0 < \mu_{1,h})$. The posterior expectation of forecaster 1 and 2 are denoted by $Y^1_{t-h,t}$ and $Y^2_{t-h,t}$, respectively.