Inattentive professional forecasters*

Philippe Andrade
Banque de France & CREM

Hervé Le Bihan
Banque de France

November 2009
(Preliminary draft)

Abstract

Using the ECB Survey of Professional Forecasters we find evidence that those forecasters draw systematically biased predictions and disagree even if they forecast the same variable. Recent theoretical advances in the macroeconomics of imperfect information relate these bias and disagreement to theories of inattention. We provide a micro data estimation of the extent of inattention among professional forecasters. We show that, on our sample, about 20% of professional forecasters are inattentive to new information released each quarter. However, a formal test reveals that this observed inattention cannot generate the extent of systematic forecasting errors and disagreement among forecasters characterizing the data. There is more stickiness in expectations than the one the mere inattention is able to generate.

Keywords: Expectations, information, disagreement, inattention, business cycle

*We thank Anil Kashyap, Noburo Kiyotaki, Bartosz Maćkowiak, Ernesto Pastén, Giorgio Primiceri, Jonathan Willis, Alexander Wolman, Michael Woodford and seminar participants at the ECB and Banque de France for useful comments. We are also grateful to Sylvie Tarrieu for superb research assistance as to Claudia Marchini and Ieva Rubene for their help with the SPF data. This paper does not reflect the views of the Banque de France. e-mails: philippe.andrade@banque-france.fr, herve.lebihan@banque-france.fr

1 Introduction

There has been recently a renewal of interest in the macroeconomic literature for models in which the formation of expectations can act as a transmission mechanism in economic fluctuations. In particular Woodford (2001), Mankiw and Reis (2002), Carroll (2003) and Sims (2003) describe how the persistent effects of unpredictable macroeconomic shocks on the economy may be related to the inattention of agents to new information. These approaches have the potential to account for two recurrent characteristics of individual expectations as measured by survey data, namely that forecasting errors are predictable and that forecasts differ across forecasters. The present paper uses such survey of individual forecasts to assess whether models of inattention accurately describe the behavior of forecasters and therefore may provide a better understanding of the business cycle.

More precisely, we exploit the ECB's survey of professional forecasters, a quarterly panel, to characterize how expectations are formed. Professional forecasters may not be representative of all agents' forecasts in the economy, since professionals obviously allocate substantial time, human, collecting and computing resources to the task of forecasting. However professional forecasts have been shown to spread to decision makers and therefore influence expectations and decisions of firms and households (see Carroll, 2003). Furthermore we may expect professional forecasters to be the agents in best position to update their information set. As a result, any evidence of information stickiness among professionals would be decisive evidence in favor of theories of information stickiness.

A first contribution of the paper is to highlight four important regulari-

ties from these data. To start with, the forecasts of experts exhibit persistent errors and systematic bias. Second, even though they forecast for the same variables for identical horizons, they report different forecasts: experts disagree. These results confirm several previous studies: bias and heterogeneity of forecasts have been identified for a long time (see among others Zarnowitz, 1985, Keane and Runkle, 1990 and the survey of Pesaran and Wheale, 2006)¹ while time variation in disagreement across forecasters has been highlighted only more recently (see among others, Mankiw *et al.*, 2003, Branch, 2007 or Coibion and Gorodnichenko, 2008b).

Our third empirical regularity is that professional forecasters are inattentive. The originality of our approach is to rely on one specificity of the European SPF, which is to provide sequences of individual forecasts for the same event (variable and date). We can therefore construct a direct micro-data estimate of the degree of inattention, *i.e.* the frequency of non-updating its forecast despite new information has been released. We find that, each quarter, on average 20% professional forecasters do not update their 1 year or 2 year forecasts despite the environment has evolved. That inattention parameter is central to a lot of imperfect information models and the value of that parameter has been recently calibrated or estimated based on aggregate time series by a series of contributions (Mankiw *et al.* 2003, Kiley 2007, Klenow and Willis 2007, Döpke *et al.* 2008, Coibion 2008, Coibion and Gorodnichenko 2008a). These studies obtain values which usually range from 90 to 70%. The fact that we rely on a panel of professionals, which may

¹See also Andrade and Gregoir (2008) who build and implement an efficient testing strategy to assess whether forecasts are biased when forecasters have access to both individual and also common information.

update more frequently their forecasts than other agents, explains part of the discrepancy. However, a lot of these recent studies (for instance Mankiw *et al.*, 2003, Döpke *et al.* 2008, Coibion and Gorodnichenko 2008a) also incorporate professional forecasters expectations. The discrepancy is thus also explained by the methodology. Studies that use macroeconomic time series typically relying on auxiliary assumptions on the economy, and are possibly subject to aggregation biases. We provide a direct, arguably more reliable, micro-data based estimate of inattention. Note also that, even if the degree of attention we estimate is very high, it is still striking that professional forecasters are somewhat inattentive.

Exploiting the panel dimension of the data furthermore allows us to calculate both the time and individual average frequencies of non-updating its forecast, *i.e.* the time and individual specific degree of inattention. This leads to our fourth and last empirical regularity, which is that the degree of inattention varies over time and among individuals. Interestingly, the cross-section distribution of individual inattention parameters show that the preceding results are not driven by a small group of professional forecasters with very unfrequent forecast revision.

Our second contribution in this paper is to assess the empirical content of inattention models. More precisely, we check whether the observed forecasting errors, disagreement and inattention models jointly behave as predicted of the sticky information models of Mankiw and Reis (2002) and Reis (2006) and the imperfect information models of Woodford (2001) and Sims (2003). Because imperfect information models cannot rationalize why the frequency of forecast updating is lower from one, we focus on a quantitative assessment

of the sticky information model *à la* Mankiw-Reis. We carrying out both a VAR-based exercise along the lines of Mankiw *et al.* (2003) and a minimum distance estimation and testing exercise. We find that the degree of inattention that we compute from the SPF data is too low to generate the smoothness observed in the average SPF forecasts. Likewise, the inattention degree that would be able to generate the degree of forecast stickiness observed in the SPF would lead to much more disagreement than what is observed among those forecasters. Our conclusion is therefore that more than the inattention than the one put forth by Mankiw and Reis is needed in order to describe both the low disagreement among professional forecasters and the high persistence in their forecasting errors.

The rest of the paper is organized as follows. In section 2 we describe the European SPF data and characterize how errors, disagreement and inattention behave. In section 3 we put the sticky information model of Mankiw and Reis (2002) into a test and provide robustness check for rounding effects.

2 European Professional forecast: some facts

2.1 Data: the ECB Survey of Professional Forecasters

The ECB's survey of professional forecasters is conducted every quarter since 1999. A detailed presentation and discussion of the data can be found in Bowles *et al.* (2007). The survey covers around 90 institutions involved in forecasting and operating in the Euro zone. Each institution is asked report, among other things, forecasts for the (year-on-year) inflation rate, real GDP growth rate and unemployment rate for a forecasting horizon of one year

and two years.² Respondents provide two types of forecasts. The first one is a ‘rolling forecast’, with an horizon of one and two years ahead of the last available observations. The second type of forecast is a ‘calendar horizon’ forecast: in each quarter, forecasters are surveyed about their forecast for current year and next year. This question is specific to the European SPF and allows to calculate a direct estimate of the inattention parameter that is central in several models of inattention. At the time of writing this article, the last available survey round is 2009Q2, so that we have 42 time periods available. These data are matched with the corresponding realization of the forecasted variable along with the last value known at the date of the survey.

We introduce some useful notations. In the case of the rolling forecast, the horizon is set one or two years posterior to the last available observation of each series. We will let $f_{it,t+k}^x$ denotes individual i forecast for variable x at date t for a k quarters ahead horizon. Variable x is either π (the year-on-year inflation rate), Δy (the year on year GDP growth rate) or u (the unemployment rate). Importantly there is an observation lag τ_x : the survey collects the forecasts $k - step$ ahead of the last observation available at date t , thus the forecasts for the horizon $t - \tau_x + k$. The observation lag varies across variables: inflation is observed with a one month lag, unemployment with a two month lag and GDP growth with a two quarter lag. The survey thus actually collects:

$$f_{it,t-\tau_x+k}^x$$

In practice the SPF reports data for $k = 4$ or 8 . In our baseline analysis, we focus on the one year (4 quarters) horizon, so, except when there is ambiguity,

²Longer horizon forecasts are also provided but will not be emphasized here.

to simplify we drop the horizon index and note $f_{it}^x = f_{it,t-\tau_x+4}^x$.

For the calendar horizon case, the forecast horizon is either for the current or the next calendar year. The horizon is not adjusted for information lag. Assuming time starts at $t = 1$, year h ends with quarter $4h$. Thus the current year forecast is:

$$f_{it,\text{current year}}^x = f_{it,4\text{int}[t/4]+4}^x$$

where $\text{int}[\cdot]$ is the integer part function.

2.2 Fact 1: Professionals are biased

The consensus forecast, measured either by the time mean or median across individuals' forecasts, broadly follows the evolution of realizations. These results are illustrated in Figure (1) which shows that both average and median forecasts track the trend of the forecasted variable.

Yet, there has been some systematic average bias over the sample period. Inflation has been underestimated. Unemployment has been overrated. Only the growth rate forecasts do not exhibit any marked systematic bias. This is shown in Table (1) which provides some basic descriptive statistics. It can also be noted that these differences in forecasting performances do not seem to be related to differences in the average number of respondents, which are broadly the same for all three variables.

Figure (1) shows that average forecasts are smoother than realizations. Consequently, forecast errors are more volatile than forecasts themselves. It is noticeable that forecasters did a better job at forecasting GDP growth rate in terms of level but also in terms of variance. Whereas the standard deviation of the GDP growth rate is 45% higher than that for the unemployment rate,

the RSME of the forecast errors is only 25% greater for RGDP than for the unemployment rate. The same picture emerges when one compares GDP with inflation.

Furthermore, column (1) of Table (5) shows the result of a regression of the average forecast error on the last error known at the date where the forecast is made (that is h quarters ago, with $h = 4$). For the three variables (inflation, unemployment and output), the coefficients significantly differ from zero meaning that errors are predictable on the basis of the information set available at the date of forecast. These results confirm earlier references on the topic which actually spurred the informational rigidities approach to expectation modeling.

Relationship with models of information rigidity The fact that forecasters regularly do biased forecasts has been already largely documented in the literature. This may result however from a particular sample and a particular history. Empirical evidence of a systematic forecasting bias can be a result of small sample and very persistent errors. As Pesaran and Wheale (2006) recall, when large sample sizes are considered, the evidence in favor of systematic bias tends to disappear. In short, periods where forecasters have a tendency to overestimate a variable are followed by periods where they have a tendency to underestimate that same variable. This kind of phenomena can be generated by models in which individuals are rational but have an imperfect information on the current state of the economy.

In particular the sticky information model proposed by Mankiw and Reis (2002) and the imperfect information model developed by Woodford (2001)

may account for persistent forecast biases. In the following sections we exploit the panel dimension of the data set to document two other characteristics of individual forecasts that these approaches imply: disagreement and inattention. While the first has been already stressed in recent papers (in particular by Mankiw *et al.*, 2003), to our knowledge, using individual data to document the second has not been yet investigated in the literature.

2.3 Fact 2: Professionals disagree

The database brings to light a second clear pattern: forecasts differ among agents. The striking result here is the population who disagree are experts who presumably have access to the same information set and who forecast the same set of variables (therefore should have the same objective). Figures (3) to (4) plot the sequence of distributions of forecasts for all survey vintages in the sample and illustrates that disagreement is present at any date. To save space, we only present the distribution of individual forecasts for GDP growth but qualitatively similar results hold for either unemployment or inflation forecasts.

Furthermore, there is a systematic pattern in disagreement since the average forecast bias differs across experts. This can be seen in Figure (2) which reports the cross section distribution of individual average forecast in-sample bias for inflation. On average, professionals underestimated the inflation rate but some of them were on average more optimistic and some others were on average more pessimistic. The dispersion between individuals is substantial, the individual average bias ranging broadly from .1% to 1%. The histogram enlightens that disagreement does not come from some outlier professionals.

Lastly, it is striking that disagreement across forecasters is not constant over time. This regularity appears in Figure (6) which shows the time series of disagreement measured by the inter-individuals forecasts standard deviation, σ_t^x , namely

$$\sigma_t^x = \left[\frac{1}{n_t} \sum_{i=1}^{n_t} (f_{it}^x - f_t^x)^2 \right]^{1/2},$$

where f_{it}^x denotes individual i forecast for variable x at date t , n_t the number of respondents at that date and f_t^x the average forecast.

Relationship with models of information rigidity The results that forecasters disagree and that the extent of disagreement varies over time confirm recent evidence provided by Mankiw *et al.* (2003), Branch (2007) or Coibion and Gorodnichenko (2008b). Mankiw *et al.* (2003) claim this stylized fact is in accordance with a sticky-information model of expectations as developed in Mankiw and Reis (2002). In particular, they stress that the extent of disagreement evolves with the magnitude of the shocks hitting the economy. In this model each forecaster has a fixed probability, λ , to be able to update its information set at each date. Increases in the degree of disagreement therefore derives from the occurrence of large shocks. In the case of a large shock hitting the economy, individuals that do update their information set produce forecasts which are quantitatively very different from individuals who do not update. These differences are less pronounced in the wake of a small shock. By contrast, and as stressed by Coibion and Gorodnichenko (2008b), in the imperfect information models of Woodford (2001) or Sims (2003), the extent of disagreement does not change with the size of the shocks hitting the economy. In these setups, disagreement is initiated by

the realizations of the idiosyncratic noise that prevent individuals from perfectly observing the true state of nature. By definition, these idiosyncratic noises do not bring any information on the state of nature and therefore disagreement does not react to shocks.

We can test more formally for these implications by regressing time t disagreement, σ_t^x , on several measures of shocks hitting the economy: the last forecasting error absolute value, $|e_{t-1}|$, the current and the last variation of the forecasted variable, $|\Delta x_t|$, $|\Delta x_{t-1}|$. Columns (1) to (3) in Table (2) reports the results. The coefficients are all positive and most of them are significant. This evidence corresponds to the prediction of the sticky information setup and contrasts with the imperfect information one. This applies especially to inflation forecasts that were investigated in Coibion and Gorodnichenko (2008b). These authors find that the dispersion of forecasts does not react to structural shocks. But they also recognize that the results are sometimes mixed and seems to rely on the fact that these structural shocks only account for small fractions of the total inflation variance (section 5.2 and appendix C).³ By contrast, here we analyze the unconditional reaction to a mixture of all events that can shock the economy.

Another implication of the sticky-information model is that, after a shock has hit the economy, forecasts distribution would spread out before narrowing to a new average level. When economic conditions change, it takes time for everybody to acknowledge it. Indeed, this phenomenon can be seen from

³Among the 6 structural shocks investigated, monetary policy, technology, fiscal policy, oil price, information *à la* Beaudry and Portier (2006) and confidence *à la* Barsky and Sims (2008), none account for more 35% of total inflation variance, even in the long-run. Moreover, depending on the horizon, the share of total variance that these shocks cannot altogether account for ranges from 67% to 47%.

the sequence of forecasts' histograms given in Figures (3) and (4). Consider for instance the episode of the cyclical downturn of 2001. The distribution of growth forecast, which was very concentrated around 3% spreads out toward zero leading to a larger dispersion. Then the distribution narrows again during year 2002. Even clearer is the impact of the current recession (year 2008-2009): along with the recession, the support of the distribution of forecast widens substantially, and several local modes appear.

Column (4) in Table (2) reports the results of a formal test for the link between disagreement and economic news. We regress the disagreement observed at the current date, σ_t^x , over the absolute change in average forecast, $|f_t^x - f_{t-1}^x|$. According to the sticky information approach, the information implying a change in the level of forecast takes time to spread out among agents. Therefore, one expects a positive relationship between those two variables. This is actually the case, the coefficient estimate being significant in 2 over 3 cases.

All in all, these features suggest that the rational inattention sticky information model may be a good candidate to understand and describe how expectations are formed. The next two subsections go further in assessing the capacity of sticky information model to account for the SPF patterns by documenting inattention among professionals.

2.4 Fact 3: Professionals are inattentive

Taking advantage of the panel structure of the SPF dataset, we are able to build different measures of information stickiness or inattention. Basically, these measures give the average frequency with which an individual does

not revises its forecasts. Of course it can be derived from the probability to revise, which we call the degree of attentiveness.

The indicator most closely related to the model is the quarterly probability of revising the $k - step$ ahead forecast at date t . Using the notation above, $f_{it,t-\tau_x+k}^x$ the individual i forecast for variable x at time t for horizon $t + k$, this probability is

$$\lambda_t^q = P(f_{it,t-\tau_x+k}^x \neq f_{it-1,t-\tau_x+k}^x).$$

The SPF database does not provide a forecast for all horizons at all dates. We are however able to compute a direct counterpart to λ_t^q in two particular cases.

Our first indicator uses the calendar horizon forecast. At all dates, each forecaster is surveyed about the current calendar year and the next calendar year. In other terms assuming time starts at $t = 1$, year h ends with quarter $4h$, and forecast horizons are such that $(t + k)/4$ is equal to some integer h . For calendar year h forecast, the probability of revision is thus

$$\lambda_{t,h}^{cal} = P(f_{it,4h}^x \neq f_{it-1,4h}^x).$$

In practice the forecasters are surveyed for their expectations of current calendar year and next calendar year. Overall, for each calendar year h (horizon $4h$), we have a sequence of 8 forecasts (starting from quarter 1 of year $(h - 1)$). Formally those are forecasts for all dates t such that $int[t/4] + 1 = h$ or $int[t/4] + 2 = h$, where $int[.]$ is the integer part function.

The second measure of attentiveness we can derive from the survey exploits the fact that the rolling horizon forecast is provided both for a 4 quarters horizon and for a 8 quarters horizon. Consider the 8 quarter horizon

forecast released at date t by forecaster i , $f_{it,t+8}^x$. This can be compared to the 4 quarter forecast horizon released 4 quarter later $f_{it+4,t+8}^x$. Indeed both forecasts refer to the same object (since all forecasts refer to year-on-year increases). We thus can evaluate the probability of a revision between the two forecasts, i.e.

$$\lambda_t = P(f_{it,t-\tau_x+4}^x \neq f_{it-4,t-\tau_x+4}^x).$$

To construct an empirical estimate of λ_t , we consider the indicator function d_{it} , such that $d_{it} = 1$ if $f_{it,t-\tau_x+4}^x \neq f_{it-4,t-\tau_x+4}^x$ and $d_{it} = 0$ otherwise. The degree of attentiveness in period t can be measured as:

$$\lambda_t = \frac{1}{n_t} \sum_{i=1}^{n_t} (d_{it}) = \frac{1}{n_t} \sum_{i=1}^{n_t} I(f_{it,t-\tau_x+4}^x \neq f_{it-4,t-\tau_x+4}^x)$$

A similar formula is used for estimating $\lambda_{t,h}^{cal}$. Note that indicator λ_t compares two forecasts that are 4 quarters away from each other so to recover a quarterly adjustment rate λ_t^q we use the formula (valid under the Mankiw et al. assumption of constant attention): $(1 - \lambda_t) = (1 - \lambda_t^q)^4$.

Parameter $1 - \lambda_t$ and $1 - \lambda_t^{cal}$ are direct micro-level indicators of information stickiness in a Mankiw Reis (2002) model, i.e. the parameter estimated by Mankiw et al (2003). Note that while disagreement σ_t^x can be assessed using data sets of repeated cross-sections, measures of attentiveness can be built only owing to the panel data structure of the SPF data set (at least two forecast sets must be observed for each individual). To our knowledge, previous assessments (e.g. Kiley, 2007) have all relied on macroeconomic estimates, implying some auxiliary assumptions about the economy, such as the rational expectation hypothesis, or assumption on price-setting behavior by firms.

Interpretation of $1 - \lambda_t$ and $1 - \lambda_t^{cal}$ as indicators of inattention deserves further discussion. It could in theory be the case that a forecaster chooses not to revise his forecast ($d_{it} = 0$) in spite of having updated his forecast. However, given the vast information set available to professional forecasters we deem unlikely that updating the information set leads to an exactly unchanged forecast after one quarter. Rather, those cases more plausibly correspond to cases when the forecaster chooses not to run a new forecast exercise, i.e. not to make use of the new set of information, for instance due to the the cost of processing information. It is relevant to characterize this situation as inattention.

Finally, we also consider another "naive" indicator of inattention based on the probability of revision between two adjacent (rolling horizon) forecasts. Given our data set, the frequency of forecast revision in period t can be measured as:

$$\theta_t = \frac{1}{n_t} \sum_{i=1}^{n_t} I(f_{it,t-\tau_x+k}^x \neq f_{it-1,t-\tau_x+k-1}^x)$$

where $k = 4$ or $k = 8$ in our data set. This can be computed at all dates, and we are able to use a much larger set of observation than for computing λ_t and λ_t^{cal} . This is a "naive" measurement of inattention because even in the Mankiw et al. model, θ_t should be equal to one at every date, reflecting the fact that forecasts are related to a different horizon. Indeed, when comparing $f_{it-1,t+k-1}^x$ and $f_{it,t+k}^x$ the forecast horizon is moved one period ahead. As a result, if at date $t - 1$ a non constant future trajectory of variable x was expected, f_{it-1}^x should differ from f_{it}^x even if the information set was not updated. Consequently, our measure θ_t is an upper bound to the degree of attention. In other words, any evidence θ_t lower than 1 is found in the data,

would suggest that there is a form of inattention in the data.⁴

Figure (7) plots the evolution of attention, θ_t and λ_t , over time for the forecasts of the three variables, and shows that there is a significant degree of fluctuations in these parameter.⁵ The proportion of forecasters changing their forecasts ranging from 75% to 90% with an average of around 80%. As Table (4) indicates, the average probability of forecast revision λ is respectively 81% for inflation, 83% for unemployment and 88% for GDP. Converted into quarterly figure these would imply quarterly frequencies of 34%, 36% and 41%. These figures suggest a degree of inattention that is substantial, although lower than the estimate provided by the existing estimate with aggregate data. Mankiw et al. (2003) found $\hat{\lambda} = 0.10$ with monthly data (corresponding to $\hat{\lambda} = 0.27$ when converted into quarterly data). Our estimate of the naive indicator θ_t also fluctuates over time around an average of around 80%, also suggesting some form of (deep) inattention.

Figure (8) presents the estimate of indicator of inattention λ_t^{cal} . Each line in the figure reports a sequence of forecast pertaining to the same calendar year. There are thus 9 lines corresponding to calendar horizons from 2001 to 2009. At each point in time λ_t^{cal} is the proportion of forecasters that revise their forecast for a given target calendar year. Note that forecast partially overlap since each vintage of the survey respondent are asked their forecast for both current and next calendar year. Two main patterns stand out. First the degree of attention is high, although attention is not complete, since the

⁴We did not choose to associate missing data with an instance of non-revision. Doing the reverse would of course have increased the degree of inattention.

⁵It can also be noted that there is a seasonality in the response rate (we do not report the graph to save space). This may as well be interpreted as an indication of inattention, if failure to respond reflect failure to update forecast at given dates.

proportion of forecast revision varies over time between 60% and 100%. The average λ_t^{cal} is 79%. Second all lines are upward sloping. Two factors can explain this pattern: first, mean reversion implies that long run forecasts are close to the unconditional average of the process. So that news that lead to revising short run forecast may leave the forecast at a long horizon unchanged. Second, it may be the case that forecasters put more attention on revising their forecast for closest forecast horizons. Experiments in the next section suggest both factors are present.

We observe that there is a quantitative conflict between our two alternative indicators of attention λ_t^{cal} (close to 0.80) and λ_t^q (close to 0.40). Our assessment is that λ_t^{cal} is presumably the most reliable indicator since it compares forecast set that are formed in adjacent quarters, so we have a larger number of observations and are surer that we are comparing forecasts delivered by the same forecaster. Also the 8 quarter ahead rolling forecast is possibly a less conventional forecast horizon forecast than the calendar year. In any case, this conflict is evidence against the basic sticky information model, which would predict both measures to be equal.

Finally the SPF data set also allows to study the degree of information stickiness across individuals. We can compute a measure of attention for each individual forecaster as the percentage of quarters in which his forecasts are revised:

$$\lambda_i = \frac{1}{T_i} \sum_{t=1}^{T_i} (d_{it})$$

with T_i the number of observations available for individual i . We observe some heterogeneity in attention. For instance figure (9) plots the distribution of θ_i across agents for the inflation forecast. Though there is a mass of

forecasters with values of θ_i above 80%, the dispersion is substantial. The less attentive among forecasters revise their forecasts every 3 quarters on average and the most attentive adjust every quarter.

Relationship with models of information rigidity Observing that forecasters are inattentive in the sense that they do not update their forecasts at each period is a result that again is implied by Mankiw and Reis (2002). On the contrary, the presence of some degree of inattention in the data is difficult to reconcile with the imperfect information model of Woodford (2001) and Sims (2003). In these approaches individuals receive an imperfect signal on the true state of the economy, but at each period they try to infer this imperfectly observed state incorporating the news they received. Forecasts are thus revised at each period. Observing a degree of attention that is lower than one is therefore at odds with this approach. Moreover, and as Coibion and Gorodnichenko (2008b) underline, the model of imperfect information does not feature time varying disagreement either. It is always equal to the standard deviation of the idiosyncratic signal. However it can still be the case that the two forms of inattention coexists.

A difference though between our results and the sticky information model of Mankiw and Reis (2002) is that λ_t is not constant as it is postulated in that model. Observing fluctuations in aggregate attention and heterogeneity of the attention among forecasters is not a direct evidence against the sticky information model of Mankiw and Reis. Indeed both could be reconciled with the more general model of inattention proposed by Reis (2006) which, under some restrictions, reduces to a Mankiw-Reis model where inattention can be

described by a single parameter. In the following section we analyze further what drives the variation of inattention. We also provide some insights on whether these variations satisfy the restrictions stated in Reis (2006) under which inattention can be aggregated across time and individuals.

2.5 Fact 4: Inattention varies

In Reis' model, each forecaster faces a specific cost to process new information and decides to revise or not his expectations and therefore his optimal plan depending on the cost of processing information and the cost of not reviewing expectations and plans. The author characterizes the optimal length of inattention (i.e. the predetermined duration of non-revision) which can be written in a general form

$$h_i^*(s_t) = h \left[\frac{K_i(s_t)}{l_i(s_t)} \right],$$

where s_t denotes the state of the economy at date t , K_i is individual i 's cost of processing new information and updating the forecast and l_i is the loss associated with non-updating for a realization of the state variable s_t . Then individual i 's probability to revise at date is $\lambda_i = 1/E_s[h_i^*(s)]$, where $E_s[\cdot]$ is the expectation integrating on the values of the state of the economy. Likewise, the time t probability to revise is $\lambda_t = 1/E_i[h_i^*(s_t)]$, with $E_i[\cdot]$ the expectation with respect to individuals.

An implication of the rational inattention model of Reis (2006) is that the level of attention rises with the forecasted variable variance. Table (3) provides loose evidence that this is the case. The probability of forecast updating is larger when the variance of the forecasted variable is large. However, whereas RGDP growth rate has a standard deviation which is almost twice

the inflation rate one, the degree of attention is only 10% higher for RGDP than for inflation. This small difference is very efficient in term of prediction bias, the RGDP bias being nearly 10 times smaller than the inflation one.

The rational inattention model of Reis (2006) could be also invoked to account for the variations in λ_t . In whirling times, agents would revise more compare to quiet episodes. We provide a first test of this assumption in column (7) of Table (5), where we regress the measure of attention λ_t to the last forecast absolute error observed, $|e_t|$. This link is significantly positive for inflation and unemployment forecasts. However it is non-significant for RGDP. This underlines that forecasters always pay more attention to GDP growth forecasts than to inflation and unemployment ones.

One may wonder if professional forecasters decide to allocate limited computational capacities, by shifting attention toward the most fluctuating variable at the time they forecast. However, Table (4) gives evidence that the level of attention given to one variable is positively correlated to the level of attention given to the others. Forecasters do not re-balance their attention between variables. On average forecasters increase their attention to all variables. Therefore at first glance the determinants of inattention seem to be corresponds to the inattention model of Reis (2006). However, as we will see now, the effects of inattention on forecasting performances are difficult to reconcile with this set-up.

Column (2) of Table (5) shows that the effect of disagreement on forecast errors is ambiguous: either positive, negative or insignificant. Yet this is not completely surprising since more disagreement means both more adjustment to news but also bigger shocks and therefore bigger potential errors for those

who did not incorporate the new piece of information. More puzzling are the results appearing in columns (3) and (5). First, there is no positive effect of attention, λ_t on incoming forecast errors $|e_{t+h}|$. More attention does not lead to less forecast errors. And second, the correlation between current attention, λ_t , and disagreement, σ_t is positive. More aggregate attention to information is associated with more disagreement. Here again two conflicting effects are at stake. As shown above, and as shown in column (4), more disagreement can come from more uncertainty or larger shock as captured by the positive effect of $|e_t|$ on σ_t . Yet, in a sticky information approach, more attention therefore more forecast updates, should also lead to less differences among forecasters, the fraction of forecasters relying on the most updated information increasing. It is thus striking that the first effect always dominates. Table (6) provides results for the same kind of analysis than above but using cross-section rather than time-series variations. Columns (2) and (3) show that individuals with higher attention rate, λ_i , exhibit less forecast precision, σ_i and higher systematic bias, e_i . All in all, it appears that more attention is not associated to better forecast performances.

Reis (2006) also shows that under certain conditions on the density of the optimal inattention length (which is conditional on the date when the individual last adjusted), $f_i(t)$ say, these individual optimal decision rules can be aggregated in the form of the sticky information model of Mankiw and Reis (2002), namely that $E(\lambda_t) = E(\lambda_i) = \lambda$. Among the required conditions is the fact that the $f_i(t)$'s should be independent across time and individuals. A simple test of that requirement can be implemented by regressing the current degree of attention on its lagged value λ_{t-1} . Column (6) of Table

(5) provides the results and shows that there is significant time-dependence in the attention for unemployment and GDP growth. The condition of aggregation are not satisfied for 2 over 3 cases. The conclusion is therefore that the condition in which individual rational inattention behavior can be described by a single aggregate parameter does not hold in general. However the inattention model of Reis (2006) may be adapted to model the formation of expectation at the individual level. In the following section, we investigate whether this is the case by analyzing to what determinants inattention reacts and how attention is related to forecast errors and disagreement.

3 Errors, disagreement and inattention: a test of the sticky information model

The previous section has shown, although with some caveats, that the SPF data features the three basics ingredients of a sticky information model *à la* Mankiw-Reis: biased forecasts, disagreement and inattention. In this section we assess whether the inattention observed in the data is able to generate the errors and disagreement we observe, in a way that is consistent with the sticky information model. We perform an informal assessment, and then turn to a formal testing procedure.

3.1 A VAR-based assessment

In this section, we implement an exercise close to what Mankiw *et al.* (2003) did to assess the importance of sticky information in order to explain the disagreement observed among US forecasters. We first estimate inflation, output growth and unemployment rate forecasts using a weak linear repre-

sensation of these data, *i.e.* a Vector Auto-Regressive model

$$A(L)X_t = c + \epsilon_t,$$

where $X_t = (\pi_t \ y_t \ u_t)'$, L is the lag operator, $A(L) \equiv \sum_{j=0}^p A_j L^j$ has all its roots outside the unit-circle and where $A_0 = I$, c is a vector of constants and $E\{\epsilon_t | X_{t-1}, \dots, X_{t-p}\} = 0$. Estimating $A(L)$ by OLS gives us sequences of (quadratic distance optimal) forecasts for the forecasting horizon h

$$\hat{f}_{t,t+h} = \hat{E}\{X_{t+h} | X_t, \dots, X_{t-p+1}\} = \hat{c} - \sum_{j=h+1}^p \hat{A}_j X_{t+h-j}$$

We now ask what will be the average of individual forecasts and the disagreement between forecasters under sticky information that is if each individual in the population that has a probability λ to update its information set at each date. The average of individuals' forecasts will follow

$$\text{AVG}(f_{t,t+h}, \lambda) = \sum_{j=0}^{\infty} (1-\lambda)^j \lambda \hat{f}_{t-j,t+h},$$

whereas the disagreement across individuals will be

$$\text{DIS}(f_{t,t+h}, \lambda) = \left\{ \sum_{j=0}^{\infty} \left[(1-\lambda)^j \lambda \hat{f}_{t-j,t+h} - \text{avg}(f_{t,t+h}) \right]^2 \right\}^{1/2}$$

Figure (10) plots the evolution of the HICP inflation rate realizations, the forecasts generated by a VAR model, the SPF consensus (averaging across individuals) and the sticky-information average with $\lambda = .10$ representative (though even lower) of previous aggregate estimates. This degree of attention is much lower than the one obtained from the micro-data information in previous section. However, a rather high degree of information stickiness is

needed if one wants to reproduce the observed flatness of the SPF consensus. This is a first dimension in which the sticky-information model departs from the data. Figure (11) plots alternatively the disagreement between forecasters that is observed, for the HICP rate, in the SPF and the disagreement that will be generated by a sticky information model with a attention parameter $\lambda = .1$. This provides evidence of a second dimension in which the sticky-information model departs from the data. If the attention is set to a low level, so as to account for persistent forecasting errors, this leads to a counterfactual high level of disagreement. We conclude from this exercise that much less attention (*i.e.* more inattention) than observed in the SPF data is needed to account for the persistence of the SPF average forecast errors. It remains to be understood why forecasters agree on sluggish revisions of their forecasts. Equivalently one needs more than just inattention of the Mankiw & Reis kind to account for the stickiness of forecasts.

3.2 A formal test

We complete the previous informal empirical assessment by a more formal test relying on a Minimum Distance Estimation (MDE) procedure. More precisely, we first estimate the attention parameter λ by minimizing the following distance criteria:

$$[\hat{\mu} - \mu(\lambda)]' \hat{\Omega}^{-1} [\hat{\mu} - \mu(\lambda)]$$

where $\hat{\mu}$ is a vector of S data-moments, such as average disagreement or average forecast error, which are computed from the SPF panel data set. Also $\mu(\lambda)$ are the equivalents of these data moments generated by the functions $\text{AVG}(f_{t,t+h}, \lambda)$ and $\text{DIS}(f_{t,t+h}, \lambda)$ defined above, and $\hat{\Omega}$ is a consistent

estimator of Ω , the variance associated with the asymptotic convergence $\sqrt{T}(\hat{\mu} - \mu) \rightarrow \mathcal{N}(0, \Omega)$ when $T \rightarrow \infty$.

Let $\hat{\lambda}$ be the minimum distance estimator of λ . Tests on the true parameter value λ can be implemented using the property that, when $T \rightarrow \infty$ (see for example Ch14 in Wooldridge, 2002)

$$\sqrt{T}(\hat{\lambda} - \lambda) \rightarrow \mathcal{N}[0, (H'\Omega^{-1}H)^{-1}]$$

with $H \equiv H(\lambda) = \nabla_{\lambda}\mu(\lambda)$ the Jacobian of $\mu(\lambda)$ with respect to λ evaluated at $\hat{\lambda}$. An estimator of $H'\Omega^{-1}H$ is given by $\hat{H}'\hat{\Omega}^{-1}\hat{H}$ where $\hat{H} \equiv H(\hat{\lambda})$. MDE also allows to test for over-identifying restrictions, i.e. the fact that one can accurately describe the set of S moments μ with the P parameters to be estimated (in our case a single parameter, λ) relying on (see for example Ch14 in Wooldridge, 2002)

$$[\hat{\mu} - \mu(\lambda)]'\hat{\Omega}^{-1}[\hat{\mu} - \mu(\lambda)] \rightarrow \chi^2_{S-P}$$

Table (7) presents the results for two cases. In the first case, which results are reported in Column (1), we chose for μ a set of two moments: the average disagreements observed for the inflation and the real GDP growth rates. We get an $\hat{\lambda} = .675$. This figure is substantial above a similar estimate provided by Mankiw et al (2003): on US data they find 0.25 to be the best fitting value for λ (converting their monthly estimate into quarterly ones). This exercise however suggests there is less attention than suggested by our preferred estimate based on micro data (recall λ_t^{cal} is found to be close to 0.80). We then test for the fact that $\hat{\lambda}$ is not significantly different from both the average attention observe in the SPF for the inflation and the real GDP growth rates. This null hypothesis is strongly rejected. We also test

for over-identifying constraints in that setup. The test rejects the null that the distance between the estimated moments and the observed one is non-significantly different from zero. We did the same for a set of four moments composed of the average disagreements and also the average errors observed for the inflation and the real GDP growth rates. Column (2) in Table (7) gives the results. The attention estimate falls sharply (*i.e.* inattention increases a lot) since we have $\hat{\lambda} = .154$. The persistence of the error tend to lower the estimation of the attention degree. Not surprisingly, the same two tests as above lead to rejecting the null. The sticky information model is not supported by the data in that the inattention is not able to generate both disagreement and forecast errors that are in line with the observations.

3.3 A robustness experiment: the influence of rounding

Our results are presumably influenced by rounding. Indeed virtually all forecasts (97% of forecast figures) in our data set are reported with one digit only. Failure to revise a forecast may merely reflect rounding by forecasters. Rounding in our context can be interpreted in two different ways. First one may consider rounding as a form of inattention. Indeed it is not formally requested by the ECB survey that figures are provided with a single digit, so providing rounded figures may reflect cost of processing the information. Under an alternative interpretation, there is widespread acknowledgement that higher order digits are not economically meaningful, or the forecasters expect that a rounded figure is requested by the survey. Under the latter interpretation, part of the (already low) degree of inattention is possibly due

to rounding. Thus, our estimates of the degree of attention, which fall in the range 40% to 80% are biased downwards.

To get a quantitative assessment of the influence of rounding on our results, we perform the following experiment. We use the estimated VAR model of Section 3.1 above as a simulation and forecasting device, drawing shocks from a multivariate distribution with a covariance matrix equal to that of the estimated innovations. We start by simulating a large time series of artificial observations of length S . We then generate a large number of recursive sets of forecasts from each point of the simulated data set. We carry $S = 2000$ forecast simulations, each of an horizon of H periods ahead. Consistently with the way figures are reported in the SPF survey, we aggregate forecasts for quarterly inflation and GDP growth into forecast for annual year-on-year growth rates. We then round each forecast to the first digit. Using the data set of rounded forecasts, we compute the probability that two adjacent forecast, corresponding to same forecast horizon (i.e. the same target date) are different. We are able to compute such a quarterly frequency of forecast revision for all horizons $h = 1$ to $h = 8$. Crucially, in our simulation exercise there is no inattention at all from the part of the forecaster: actual forecast are changed every period in line with the VAR model. As a result, the probability of not updating the forecast is here an estimate of the bias to our measurement of attention that is due to rounding. The magnitude of this bias will obviously depend on several parameters of the exercise: the horizon considered (we expect more bias at longer horizons due to mean reversion), the size of the innovation variance (larger shock will imply lower bias since forecast revision will be too large to be wiped out by

rounding) and the persistence of the process (low persistence of shocks will imply fast mean reversion thus less forecast revision).

The results we obtain are as follows. The estimated probability of updating a forecast is 92% for the horizon $h = 8$ quarters and rises to 96% for the horizon $h = 1$. These figures are lower than one, which suggest there is actually a rounding bias. The probability also increases when the target date get closer, rationalizing the pattern of Figure 8. However these figures are at all horizons markedly above than the estimates of attention we recover from actual SPF micro data. Thus, independently of rounding effects, there is a degree of actual inattention in professional forecasts. Finally, it is important to note that this exercise makes the gap between the data and the sticky information model discussed above even wider. Indeed should we attempt to correct for the rounding bias in the exercise of section 3.2 above, this would yield a larger degree of attention than seen in the micro data, and thus lead to an even stronger rejection of the sticky information model.

4 Conclusion

We have analyzed an original data set, the European Survey of Professional Forecasters, in the perspective of testing sticky and imperfect information models. In particular we have provided micro data estimates of the degree of inattention based on individual observations in that data set, and found that professionals are, albeit mildly, inattentive. A formal test however rejects the basic sticky information model. Indeed, the observed degree of inattention is too low to generate using that model the extent of systematic forecasting errors and disagreement among forecasters present in the data. There is thus

more stickiness in experts expectations than the one the mere inattention is able to generate.

At least two avenues for future research are worthwhile considering. First, the methodology implemented in this paper may be used to investigate other panel data sets of expectations such as consumer and firms expectations. Second, results in the present paper suggest that the sticky information model should be altered in order to replicate the main features of the data namely the coexistence of disagreement and a significant degree of attention, together with smooth forecasts.

References

- [1] Andrade, P. and Gregoir, S. (2008). Assessing the rational expectation hypothesis when agents have common and private information. Mimeo.
- [2] Barsky, R. B. and Sims, E. R. (2008). Information, animal spirits and the meaning of innovation in consumer confidence. Mimeo.
- [3] Beaudry, O. and Portier, F. (2006) News, stock prices and economic fluctuations *American Economic Review*, 96:1293–1307.
- [4] Bowles C., Friz, R., Genre, V., Geoff, K., Meyler, A. and Rautanen, T. (2007). The ECB survey of professional forecasters (SPF). A review after eight years’ experience. ECB Occasional Paper Series, n. 59.
- [5] Branch, W. A. (2007). Sticky information and model uncertainty in survey data of inflation expectations. *Journal of Economic Dynamic and Control*, 31:269–98.
- [6] Carroll, C. (2003). Macroeconomic expectations of households and professional forecasters. *Quarterly Journal of Economics*, 118:269–298.
- [7] Coibion, O (2008). Testing the sticky information Phillips’ curve forth. *Review of Economic and Statistics*.
- [8] Coibion, O. and Gorodnichenko, Y. (2008a). Strategic interaction among heterogenous price-setters in an estimated DSGE model NBER Working Paper (14323).
- [9] Coibion, O. and Gorodnichenko, Y. (2008b). What can survey forecasts tell us about informational rigidities? NBER Working Paper (14586).

- [10] Jörg D., J. Dovern, U. Fritsche and J. Slacalek. (2008). Sticky Information Phillips Curves: European Evidence. *Journal of Money, Credit and Banking*, 40(7): 1513-1520.
- [11] Keane, M. P. and Runkle, D. E. (1990). Testing the rationality of price forecasts: New evidence from panel data. *American Economic Review*, 80:714–35.
- [12] Kiley, M. T. (2007). A quantitative comparison of sticky-price and sticky-information models of price setting. *Journal of Money Credit and Banking*, 39:101–25.
- [13] Klenow, P. J. and Willis, J. L. (2007). Sticky information and sticky prices. *Journal of Monetary Economics*, 54:79–99.
- [14] 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*, 117:1295–328.
- [15] Mankiw, N. G., Reis, R. and Wolfers, J. (2003). Disagreement about inflation expectations. *NBER Macroeconomic Annuals*, 209–247.
- [16] Pesaran, M. H. and Weale, M. (2006). Survey Expectations. *Handbook of Economic Forecasting*. C. Elliot, C.W.J. Granger and A. Timmermann (eds). North-Holland Press.
- [17] Reis, R. (2006). Inattentive producers. *Review of Economic Studies*, 73:793–821.

- [18] Sims, C. (2003). Implications of rational inattention. *Journal of Monetary Economics*, 50:665–690.
- [19] Woodford, M. (2001). Imperfect common knowledge and the effect of monetary policy. *Knowledge, Information and Expectations in Modern Macroeconomics: In Honor of Edmund S. Phelps*. Aghion, P., Frydman, R., Stiglitz, J. and Woodford, M. eds. Princeton University Press.
- [20] Zarnowitz, V. (1985). Rational expectations and macroeconomic forecasts. *Journal of Business and Economic Statistics*, 3:293–311.

Table 1: Descriptive Statistics

<i>ECB's Survey of Professional Forecasters^a</i>			
Sample period: 1999Q1–2008Q4			
	HICP	UNEM	RGDP
REALIZATIONS			
Mean	2.32	8.23	2.18
Std Dev	.63	.60	1.07
FORECASTS			
(Forecasting horizon: 1 year)			
# of individuals	92	90	92
Mean	1.79	8.39	2.01
Std Dev	.22	.86	.63
Resp. rate	81	83	87
FORECAST ERRORS			
Mean	.56	-.27	.06
RMSE	.73	.69	.95
MAE	.57	.47	.78

^aAll variables (except number of respondents) are in %; HICP = Euro area HICP inflation rate (yoy); UNEM = Euro area unemployment rate; RGDP = Euro area real GDP growth rate (yoy)

Table 2: Understanding time variations in disagreement

Time-series bivariate regressions					
Dependent variable	σ_t				
	(1)	(2)	(3)	(4)	(5)
HICP^a					
$ \Delta x_t $	.053** (.021)	-	-	-	-
$ \Delta x_{t-1} $	-	.065* (.039)	-	-	-
$ e_{t-1} $	-	-	.080*** (.019)	-	-
$ \Delta f_t $	-	-	-	.034*** (.011)	-
λ_t	-	-	-	-	.291** (.117)
UNEM					
$ \Delta x_t $	.033 (.044)	-	-	-	-
$ \Delta x_{t-1} $	-	.094 (.114)	-	-	-
$ e_{t-1} $	-	-	.048** (.020)	-	-
$ \Delta f_t $	-	-	-	.066 (.065)	-
λ_t	-	-	-	-	.176* (.094)
RGDP					
$ \Delta x_t $	.072** (.036)	-	-	-	-
$ \Delta x_{t-1} $	-	.054 (.059)	-	-	-
$ e_{t-1} $	-	-	.012 (.029)	-	-
$ \Delta f_t $	-	-	-	.257*** (.063)	-
λ_t	-	-	-	-	.288* (.162)

^aAll variables are in %; HICP = Euro area HICP inflation rate (yoy); UNEM = Euro area unemployment rate; RGDP = Euro area real GDP growth rate (yoy); σ = forecasters disagreement, λ = fraction of people changing their forecast; x = forecasted variable, f = average forecast; e = forecast errors; all regressions include a constant term; numbers in brackets are std error of estimates; ***, **, * indicates significance at the 1, 5 and 10% levels respectively

Table 3: Errors, disagreement & inattention

Comparing unconditional moments

ECB's SPF^a

Sample period: 1999Q1–2008Q4

	HICP	UNEM	RGDP
Mean(e)	.56	-.27	.06
Mean(λ)	81	83	88
Mean(σ)	.26	.26	.33
Std Dev(x)	.62	.60	1.07

^aAll variables are in %; HICP = Euro area HICP inflation rate (yoy); UNEM = Euro area unemployment rate; RGDP = Euro area real GDP growth rate (yoy); e forecast error; λ = frequency of changing forecast; σ disagreement; x = forecasted variable

Table 4: Errors, disagreement & inattention

Correlations across variables

<i>ECB's SPF^a</i>			
Sample period: 1999Q1–2008Q4			
	HICP	UNEM	RGDP
REALIZATIONS, x			
HICP	1		
UNEM	-.495	1	
RGDP	-.033	-.300	1
ERRORS, e			
HICP	1		
UNEM	.128	1	
RGDP	-.078	-.610	1
DISAGREEMENT, σ			
HICP	1		
UNEM	.668	1	
RGDP	.475	.611	1
INATTENTION, λ			
HICP	1		
UNEM	.426	1	
RGDP	.210	.243	1

^aAll variables are in %; HICP = Euro area HICP inflation rate (yoy); UNEM = Euro area unemployment rate; RGDP = Euro area real GDP growth rate (yoy); e forecast error; λ = frequency of changing forecast; σ disagreement; x = forecasted variable

Table 5: Errors, disagreement & inattention

Time-series bivariate regressions

Dependent variable	(1) $ e_{t+1} $	(2) $ e_{t+1} $	(3) $ e_{t+1} $	(4) σ_t	(5) σ_t	(6) λ_t	(7) λ_t
HICP^a							
$ e_{t-h+1} , e_t $	.650*** (.151)	-	-	.081*** (.016)	-	-	.053** (.026)
σ_t	-	-3.67** (1.41)	-	-	-	-	-
λ_t	-	-	-.335 (1.09)	-	.291** (.117)	-	-
λ_{t-1}	-	-	-	-	-	-.085 (.174)	-
UNEM							
$ e_{t-h+1} , e_t $	.842*** (.055)	-	-	.055*** (.019)	-	-	.059* (.036)
σ_t	-	3.09** (1.25)	-	-	-	-	-
λ_t	-	-	1.89** (.74)	-	.176* (.094)	-	-
λ_{t-1}	-	-	-	-	-	.316** (.164)	-
RGDP							
$ e_{t-h+1} , e_t $	.595*** (.139)	-	-	.025 (.028)	-	-	-.013 (.027)
σ_t	-	.040 (1.03)	-	-	-	-	-
λ_t	-	-	1.71 (1.09)	-	.288* (.162)	-	-
λ_{t-1}	-	-	-	-	-	.417** (.158)	-

^aAll variables are in %; HICP = Euro area HICP inflation rate (yoy); UNEM = Euro area unemployment rate; RGDP = Euro area real GDP growth rate (yoy); σ = forecasters disagreement, λ = fraction of people changing their forecast; x = forecasted variable, f = average forecast; e = forecast errors; all regressions include a constant term; numbers in brackets are std error of estimates; ***, **, * indicates significance at the 1, 5 and 10% levels respectively

Table 6: Errors, disagreement & inattention

Cross-section bivariate regressions
(Individuals weighted by their response rate)

Dependent variable	(1) $ e_i $	(2) $ e_i $	(3) σ_i
HICP^a			
σ_i	1.09*** (.102)	-	-
λ_i	-	.521*** (.035)	.384*** (.021)
UNEM			
σ_i	.283*** (.078)	-	-
λ_i	-	.127*** (.042)	.484*** (.025)
RGDP			
σ_i	-.024 (.041)	-	-
λ_i	-	.002 (.031)	.699*** (.027)

^aAll variables are in %; HICP = Euro area HICP inflation rate (yoy); UNEM = Euro area unemployment rate; RGDP = Euro area real GDP growth rate (yoy); σ = forecasters disagreement, λ = fraction of people changing their forecast; e = average forecast errors; all regressions include a constant term; all observations are weighted by the correspondent response rate, numbers in brackets are std error of estimates; ***, **, * indicates significance at the 1, 5 and 10% levels respectively

Table 7: Minimum Distance Estimation and Tests

<i>Attention rate, λ^a</i>			
	(1)	(2)	(3)
Estimate $\hat{\lambda}$	.675	.154	
$H_0 : \hat{\lambda} = \lambda^{\text{SPF}}$	.000	.000	
$H_0 : \text{Over-identification}$	.000	.000	

^aNumber in lines reporting tests are the significance level. See Section 3.2 for more details

Figure 1: Consensus forecast & realizations

Figure 2: Distribution of individuals' average inflation forecasts errors

Figure 3: Distribution of RGDP forecasts (evolution over time)

Figure 4: Distribution of inflation forecasts over time

Figure 5: Distribution of inflation forecasts over time

Figure 6: Evolution of disagreement

Disagreement, 1-year forecasts

Figure 7: Evolution of inattention

Inattention (λ), 1-year forecasts

Figure 8: Evolution of inattention

Figure 9: Distribution of inattention across individuals

HICP, 1-year: distribution of λ_i
(Individual prob. of revising)

Figure 10: Comparing Forecasts: 1-year, HICP

Inflation rate forecasts
Realized, VAR, Sticky & SPF

Figure 11: Implied disagreement: 1-year, HICP

Disagreement on inflation forecasts
Sticky information model & SPF