
DOCUMENT
DE TRAVAIL
N° 307

**INATTENTIVE PROFESSIONAL
FORECASTERS**

Philippe Andrade and Hervé Le Bihan

December 2010



**INATTENTIVE PROFESSIONAL
FORECASTERS**

Philippe Andrade and Hervé Le Bihan

December 2010

Les Documents de travail reflètent les idées personnelles de leurs auteurs et n'expriment pas nécessairement la position de la Banque de France. Ce document est disponible sur le site internet de la Banque de France « www.banque-france.fr ».

Working Papers reflect the opinions of the authors and do not necessarily express the views of the Banque de France. This document is available on the Banque de France Website “www.banque-france.fr”.

Inattentive professional forecasters*

Philippe Andrade
Banque de France & CREM

Hervé Le Bihan
Banque de France

*We thank our discussants Bartosz Maćkowiak and Ernesto Pastén as well as Carlos Carvalho, Olivier Coibion, Christian Hellwig, Anil Kashyap, Noburo Kiyotaki, Juan Pablo Nicolini, Giorgio Primiceri, Sergio Rebelo, Jonathan Willis, Alexander Wolman, Michael Woodford, Tao Zha and seminar participants at the Banque de France, ECB, New-York Fed, Philadelphia Fed, San-Francisco Fed, University Paris 1 and at the conferences ESEM 2009, AFSE 2010, CESifo on “Macroeconomics and Survey Data” and SED 2010 for useful comments. All remaining errors are ours. We are also grateful to Sylvie Tarrieu for superb research assistance as well as to Claudia Marchini and Ieva Rubene for their help with the SPF data. This paper does not reflect necessarily the views of the Banque de France. e-mails: philippe.andrade@banque-france.fr, herve.lebihan@banque-france.fr

Abstract

We use the ECB Survey of Professional Forecasters to characterize the dynamics of expectations at the micro level. We find that forecasters (i) have predictable forecast errors; (ii) disagree; (iii) fail to systematically update their forecasts in the wake of new information; (iv) disagree even when updating; and (v) differ in their frequency of updating and forecast performances. We argue that these micro data facts are qualitatively in line with recent models in which expectations are formed by inattentive agents. However building and estimating an expectation model that features two types of inattention, namely sticky information à la Mankiw-Reis and noisy information à la Sims, we cannot quantitatively generate the error and disagreement that are observed in the SPF data. The rejection is mainly due to the fact that professionals relatively agree on very sluggish forecasts.

Keywords: Expectations, imperfect information, inattention, forecast errors, disagreement, business cycle

JEL Classification: D84, E3, E37

Résumé

Nous utilisons l'enquête « Survey of Professional Forecasters » conduite par la BCE, un panel trimestriel, afin de caractériser la formation des anticipations des prévisionnistes professionnels. Nous mettons en évidence plusieurs faits : (i) le caractère prévisible des erreurs de prévision; (ii) un degré significatif de désaccord entre prévisionnistes; (iii) l'absence de révision systématique : chaque trimestre, environ un quart des prévisions ne sont pas révisées; (iv) un niveau de désaccord significatif même parmi les experts qui révisent leurs prévisions (v) une hétérogénéité entre les experts du panel dans la fréquence de révision des anticipations et dans leur performance prédictive. Ces faits sont qualitativement cohérents avec les modèles théoriques reposant sur l'hypothèse d'"inattention" comme ceux de Mankiw et Reis, ou de Sims.

Dans un second temps, nous élaborons et estimons un modèle empirique qui englobe ces deux types de spécification de l'inattention. Nous obtenons toutefois que ce modèle est quantitativement rejeté : il ne permet pas de rendre compte simultanément du niveau de désaccord et du niveau de persistance de l'erreur de prévision présents dans les données. Compte tenu du degré de désaccord observé, le modèle prédit des prévisions beaucoup moins inertes que celle observées : à l'origine de ce rejet est donc un certain consensus des experts sur des prévisions inertes.

Mots-clefs: Anticipations, prévisions d'experts, information imparfaite, inattention, erreur de prévision, désaccord, cycle macroéconomique.

Codes **JEL:** D84, E3, E37

1 Introduction

Models in which imperfect information and the formation of expectations act as a transmission mechanism of economic fluctuations—in the spirit of Friedman (1968), Phelps (1968) and Lucas (1972)—have recently regained interest in the macroeconomic literature.¹ Imperfect information has, in particular, been related to the inattention of agents to new information, a behavior that can be rationalized by costly access to information and limited processing capacities.² One appeal of these models is to provide an alternative channel to sticky prices to explain the persistent effects of transitory shocks, and in particular monetary shocks, on the economy. Moreover, this approach can parsimoniously account for patterns of individual expectations observed in survey data that are at odds with the standard perfect information rational expectation framework, namely that forecast errors are predictable and forecasts differ across forecasters.³

In this paper, we exploit the panel dimension of such a survey of forecasts, namely the ECB Survey of Professional Forecasters (SPF), to produce new micro facts characterizing the formation of expectations. We then elaborate on those characteristics to assess whether models of inattention accurately describe the behavior of forecasters and thus may contribute to a better understanding of business cycle fluctuations. To be consistent with the recent literature we consider two types of inattention models. On the one hand, *sticky information* models developed by Mankiw & Reis (2002) and Reis (2006a,b) in which agents update their information set infrequently but get perfect information once they do. On the other hand *noisy information* models proposed by Woodford (2002), Sims (2003) and Maćkowiak & Wiederholt (2009) in which agents continuously update their information but have an imperfect access to it at each period.

The ECB-SPF is a quarterly panel starting in 1999 surveying around 90 forecasting units in either public or private institutions. Professional forecasters may not be representative of less sophisticated agents, since professionals obviously allocate substantially more time, human, collecting and computing resources to the task of forecasting macroeconomic variables. However professionals'

¹See, among others, Woodford (2002), Hellwig & Veldkamp (2008), Angeletos & La'O (2009) and Lorenzoni (2009). Imperfect information is also crucial in the recent welfare analysis of information. See, among others, Morris & Shin (2002), Angeletos & Pavan (2007) or Adamor & Weill (2009). Veldkamp (2009) and Mankiw & Reis (2010) provide surveys.

²See Mankiw & Reis (2002), Sims (2003), Reis (2006a,b) and Maćkowiak & Wiederholt (2009).

³See for example Mankiw, Reis & Wolfers (2003).

opinion has been shown to spread to other types of agents and therefore influence expectations and decisions of firms and households (Carroll, 2003). Furthermore we may expect professional forecasters to be the agents in the best position to pay attention to the relevant macroeconomic information. As a result, the extent of attention to news among professional forecasters can be seen as an upper bound for other agents' attention to aggregate conditions.

We highlight five main categories of micro facts from these SPF data that we argue are consistent with professional forecasters behaving as if they were inattentive. First, the forecasts of experts exhibit predictable errors and systematic bias. Second, experts disagree as they report different predictions for the same variable at the same forecast horizon. Moreover, the disagreement between forecasters evolves over time. Third, agents do not systematically update their forecasts even when new information is released. Fourth, the forecasters who update also disagree on their forecasts. Five, the frequency of updating a forecast, the average forecast error and the revision of forecast vary across individuals.

To our knowledge, the present study is the first to document infrequent updating in survey forecasts. The originality of our approach is to exploit the fact that the European SPF provides sequences of individual forecasts for the same event (variable and date). We can therefore construct a direct micro-data estimate of the frequency of updating a forecast. The results show that, on average, each quarter, only 75% of professional forecasters do update their 1-year or 2-year forecasts while the macroeconomic environment evolves. The frequency of updating has a structural interpretation: it corresponds to the degree of attention, a key parameter in a sticky-information type of model. Furthermore, identifying forecasters that update their forecasts, we also uncover that they too disagree. Thus the lack of information updating is not the sole responsible for disagreement among experts. Even when updating their forecasts, they may not have access to the same information. This result is in line with the predictions of a noisy-information model. Lastly, the individual dimension of the data allows us to analyze the cross section distribution of the degree of attention. We find a minimum at 50% and a maximum at 100%. Moreover, the shape of the distribution reveals that the implied average inattention is not driven by a specific group of professional forecasters.

We then turn to a formal empirical assessment of inattention models. More precisely, we argue that the previous results *qualitatively* support a model featuring two types of inattention namely

sticky-information *à la* Mankiw-Reis and noisy-information *à la* Sims. We therefore develop such an expectation model and then use it along with the SPF data to carry out a Minimum Distance Estimation (MDE). We find that this inattention model fails to *quantitatively* reproduce the observed persistence of the average forecasting errors together with the relatively small disagreement between forecasters. Moreover, the smoothness observed in the average SPF forecasts would require a much lower attention degree than our micro data estimates. Such a low attention would in turn lead to much more disagreement than observed in the SPF data. Therefore, elements others than the mere inattention included in our expectation model are needed to reconcile both the relatively low disagreement among professionals and the relatively high persistence of the aggregate forecasting error.

Our paper relates to the vast literature, mostly relying on US data, that studies the behavior of survey forecasts and compares it with the implications of theoretical expectation models. Numerous studies (see Pesaran & Weale, 2006 for a recent survey) found systematic aggregate forecast errors and disagreement in these data, at odds with the perfect information rational expectation framework. We confirm these results for a recent sample period and for European SPF data. We also complement these results by providing new empirical micro evidence on individual expectations.

Our work is also related to Mankiw *et al.* (2003), Branch (2007), Coibion & Gorodnichenko (2008) or Patton & Timmermann (2009) who rely on the characteristics of survey expectations to assess inattention and, more generally, imperfect information theories. Mankiw *et al.* (2003) and Branch (2007) focus on the cross-section distribution of forecasts to calibrate the sticky-information attention parameter mentioned above. By comparison, we underline the importance of investigating the consistency of this parameter values with both the cross-section dispersion of forecasts and the aggregate forecast errors. Furthermore, we improve on their approach by considering a model which can explain the disagreement among forecasters who update their information set. Lastly, rather than calibrate it, we estimate the attention parameter using either micro-data estimates or a MDE procedure. Coibion & Gorodnichenko (2008) look at the conditional response to various structural shocks of the aggregate error and disagreement implied by surveys to disentangle the sticky-information and the noisy-information models of inattention. They find mixed support in favor of the two, as we do. The distinctive feature of our analysis is that we estimate a model

featuring simultaneously the two types of inattention. Patton & Timmermann (2009) rely on the evolution of forecasts over different forecast horizons to stress that differences in the interpretation of information, rather than different information sets are the main culprit for forecasters' disagreement. They leave the source of these different interpretations unexplained. The model we consider is an alternative approach to generate disagreement, that does not rely on "deep" heterogeneity among forecasters.

Our paper is moreover closely linked to several recent contributions that rely on aggregate time series to estimate the attention degree in a sticky-information model of inflation dynamics (Kiley 2007, Döpke *et al.* 2008, Coibion 2010). When significant, the results imply a frequency of updating a forecast ranging from 10% to 30%, well below the figure of 75% we obtain. That we rely on a panel of professionals could explain part of the discrepancy, since these agents may update more frequently their forecasts than other types. However many of these recent studies (for instance Döpke *et al.* 2008, Coibion, 2010) also use professional forecasters' expectations to perform their estimation. The discrepancy is thus also related to the methodology. Studies that use macroeconomic time series typically rely on auxiliary assumptions on the economy and are subject to aggregation biases. By contrast, we provide a direct, arguably more reliable, micro-data estimate of this parameter that is key in sticky information models. Even if we find an attention degree that is much higher than these previous works, it still remains remarkable that professionals do not systematically update their forecasts.

Finally, our work is related to Klenow & Willis (2007) and Maćkowiak *et al.* (2009). They rely on structural models of price setting under inattention and show that several of the distinctive implications of these models are supported by the dynamics of individual prices (Klenow & Willis, 2007) or sectoral price index data (Maćkowiak *et al.*, 2009). By comparison, looking at the characteristics of expectations, we find more direct evidence in favor of inattention theories. We also stress some quantitative limits of the existing inattention models.

The rest of the paper is organized as follows. In section 2, we describe the European SPF data. We turn to the facts in section 3. In section 4, we develop a model of expectations that incorporate both sticky and noisy information and put it to a test. We give some concluding remarks in section 5.

2 Data: the ECB Survey of Professional Forecasters

The ECB's survey of professional forecasters has been conducted every quarter since 1999. The survey covers around 90 institutions involved in forecasting and operating in the euro area. Each institution is asked to report, among other things, forecasts for the (year-on-year) euro area inflation rate, real GDP growth rate and unemployment rate for forecasting horizons of one year and two years.⁴ Respondents provide two types of forecasts. The first one is a 'rolling forecast', with a fixed horizon of one or two years ahead of the last available observation. The second type of forecast is a 'calendar horizon' forecast: in each quarter, forecasters are surveyed about their forecast for a fixed event, namely the current and next years. In this case, the forecast horizon shrinks as time goes by. At the time of the writing of this article, the last available survey round is 2010Q2, so that we have 46 time periods available. These data are matched with the corresponding realization of the forecasted variable.

The ECB SPF has been rarely used for research purposes so far. It was indeed launched with the euro and therefore was bound to cover a too short period of time for several years. However, the ECB SPF has some specific advantages compared to some other survey expectation data. As will be detailed below (Section 3.4.1), combining these specificities is crucial to observe the individual expectation characteristics that relate to models of inattention, in particular the ones associated with forecast revisions. To start with, the data base is a panel which implies that one can track the response of a particular individual institution over time. Moreover, the responses are quantitative rather than qualitative. By contrast, many of the survey that cover firms or households are typically repeated cross-sections, and report qualitative data.⁵ Furthermore, the number of actual respondents is relatively high (around 60) compared with other surveys of professional forecasters. This number is for example twice as large as the number provided in the more widely used US SPF, and leads to arguably more precise estimates of statistics based on individual expectations (the forecast average, the forecast error variance, the disagreement between forecasters, the individual probability to update a forecast). Finally, the ECB SPF provides both

⁴Longer horizon forecasts are also provided but will not be emphasized here. See Bowles *et al.* (2007) for a thorough presentation and discussion of the survey.

⁵Professionals' forecasts relate to macroeconomic conditions whereas individual agents may be concerned with forecasting their future idiosyncratic conditions which we cannot assess here. However, in most models imperfect macroeconomic information will also have an impact on optimal individual decisions.

‘calendar’ and ‘rolling’ forecasts along with several forecast horizons, namely 4 and 8 quarters, for both of them. By contrast, the US SPF, provides only 4 quarter ‘rolling’ horizon forecasts.

We now introduce some useful notations. In the case of the rolling horizon forecasts, we denote $f_{it,t+h}^x$ individual’s i forecast for the variable x at date t , h quarters ahead. The variable x is either π (the year-on-year inflation rate), Δy (the year on year GDP growth rate) or u (the unemployment rate). The forecast horizon h is set to 4 or 8 quarters ahead of the last observation of variable x available at the date of the response to the survey. Importantly there is an observation lag, τ_x , between the date of the response to the survey t and the date of the last publicly released figure of each macro variable, $t - \tau_x + h$. 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:

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

In practice the SPF reports data for $h = 4$ or $h = 8$. For notation simplicity, in the following, we drop the reference to the observation lag and simply refer to $f_{it,t+h}^x$.

For the calendar horizon case, the forecast horizon is either the current or the next calendar year. The horizon is not adjusted for information lag. Letting T be the last quarter of a calendar year in the sample, calendar horizon forecast can be written as

$$f_{it,T}^x = f_{it,t+(T-t)}^x$$

The second term in the equality makes clear that, for each year in the sample, the forecast horizon $T - t$ decreases with t . For notation simplicity again, we use the notation $f_{it,T}^x$.

Let x_{t+h} be the realization of the forecasted variable at date t , we denote $e_{it,t+h}^x$ individual’s i forecast error at date $t + h$, namely $e_{it,t+h}^x = x_{t+h} - f_{it,t+h}^x$.⁶ The average forecast error is defined by $e_{t,t+h}^x = \frac{1}{n_t} \sum_i^{n_t} e_{it,t+h}^x$, with n_t the number of respondents to the survey at date t . It can alternatively be defined with reference to the average, or consensus forecast, $f_{t,t+h}^x = \frac{1}{n_t} \sum_i^{n_t} f_{it,t+h}^x$ simply through the equality $e_{t,t+h}^x = x_{t+h} - f_{t,t+h}^x$.

⁶For rolling forecasts, the date of the realization takes into account the observation lag τ_x mentioned above in order to be consistent with the date for which the forecast is made.

3 Some facts about individuals' expectations

3.1 Fact 1: Forecasters are biased

The professional forecasters make, on average, predictable errors. This emerges from Figure (1) which shows, in the left panel, the time series of the so-called consensus forecast for the 1-year horizon⁷ together with the realizations of the predicted variable, and, in the right panel, the time-series of the corresponding average forecast error. It is striking to see that periods of under/overestimation of the target variable realizations are very persistent, and last for more than 1 year, *i.e.* over time periods that are longer than the forecast horizon.⁸

Table (1) further asserts these points by providing descriptive statistics on realizations, average forecasts and errors and a test on the predictability of forecast errors. Inflation has been underestimated by an average of .56% (on an annual rate). Real GDP growth has been overrated by an average of .27%, mainly due to the recent crisis. Finally unemployment exhibits a small underestimation of about .04% over the period. These systematic biases explain why the root-mean-squared-error of the forecasts is larger than the variance of the forecasted series.

Systematic biases go along with very persistent forecast errors. Their first order auto-correlations range from .735 for inflation to .860 for unemployment.⁹ The bottom panel of Table (1) also shows the result of a regression of the average forecast error on the last error known at the date when the forecast was made (that is h quarters before, with $h = 4$). For the three variables (inflation, unemployment and output), the coefficients significantly differ from zero. Thus, errors are predictable on the basis of the information set available at the date of forecast. These results confirm other references on the topic.¹⁰

Relationship with models of information rigidity Predictable forecast errors are a prediction of both *sticky information* and *noisy information* inattention models. Forecasters are rational

⁷The consensus is defined as the average of individual forecasts for each date in the sample. The median forecast across forecasters is very close to this average.

⁸This is strikingly the case for inflation with a systematic average underestimation up to 2006.

⁹These 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.

¹⁰See, for example, Andrade & Gregoir (2010) and Coibion & Gorodnichenko (2010) for recent references.

but exhibit biased average forecasts because they have an imperfect information on the current state of the economy.

More precisely, in a *sticky information* model, agents update their information set infrequently with a probability of updating given and constant across dates and individuals. Therefore, each date, only a fraction of the population has access to the last vintage of macroeconomic news. The average forecast is therefore partly made of individual forecasts that are predictable with respect to the new information set.

In a *noisy information* model agents update their information but they know that the news they get is imperfect and therefore only partly pass it onto their forecast. The average forecast thus incorporates only part of this new information, which makes the forecast error predictable with respect to the (perfect) information available.¹¹

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.

3.2 Fact 2: Forecasters disagree

A second pattern in the data is that forecasts differ among forecasters even though they forecast the same object. This is illustrated in Figure (2) which plots the sequence of cross-section distributions of 1-year forecasts for all survey vintages in the sample.¹² The distribution never degenerates to a single peak, which illustrates that disagreement is present at any date. The result is particularly striking given that forecasters in our sample are experts who presumably have access to the same information set.

An empirical assessment of the extent of disagreement among forecasters is usually done by calculating the cross-section standard deviation of forecasts at each date, σ_t^x , namely, using the notations introduced in the previous section

$$\sigma_{t,h}^x = \left[\frac{1}{n_t} \sum_{i=1}^{n_t} (f_{it,t+h}^x - f_{t,t+h}^x)^2 \right]^{1/2} .$$

¹¹Here it is implicitly assumed that the econometrician can have access ex-post to perfect macroeconomic information.

¹²To save space, we only present the distribution of individual forecasts for real GDP growth but qualitatively similar results hold for either unemployment or inflation forecasts.

Table (3) reports the time average disagreement, σ_h^x , for the 1-year horizon rolling forecasts. This average disagreement is equal to .26 for inflation, .26 for unemployment and .33 for real GDP (see Table 3). Expressing these values as a fraction of the underlying variable standard deviation over time we obtain a normalized indicator of disagreement of 42%, 43% and 31% for, respectively, inflation, unemployment and real GDP.

Figure (3) shows the time series of disagreement, $\sigma_{t,h}^x$, for the 1-year horizon rolling forecasts of the three macroeconomic variables of interest. Disagreement across forecasters is not constant over time and the three time-series of disagreement have a strong positive correlation.¹³

It is worth investigating whether this disagreement is completely random—and in particular represents measurement errors only—or whether it stems from information imperfections, for instance a slow diffusion of macroeconomic news among forecasters. In that second case, the cross-section distribution of forecasts would spread out after the shocks and then narrow to a new average level. This phenomenon seems consistent with a crude event study applied to the sequence of forecasts' histograms of Figure (2). Consider for instance the episode of the cyclical downturn of 2001. The distribution of growth forecasts, which was very concentrated around 3%, spreads out toward zero leading to a larger dispersion. Then the distribution narrows again over 2002 and 2003 but around a lower average of about 2%. Even clearer is the impact of the current recession (years 2008-2009): along with the recession, the support of the distribution of forecasts widens substantially to the left, and several local modes appear.

We go further in investigating the link between disagreement and economic news using formal statistical tests. More precisely, we regress disagreement at time t , $\sigma_{t,h}^x$, on several rough measures of shocks hitting the economy: the last absolute variation in the forecasted variable, $|\Delta x_{t-1}|$, the squared last forecast error, $(e_{t-1,t+h-1}^x)^2$, and the current absolute change of the forecast, $|\Delta f_{t,t+h}^x| = |f_{t,t+h}^x - f_{t-1,t+h-1}^x|$. Results are presented in Table (2) and show that the coefficients are all positive and most of them are significant: disagreement is an increasing function of the amplitude of the shocks hitting the economy.

This last evidence somehow contrasts with Coibion & Gorodnichenko (2008) who find that the

¹³Relying on rolling-horizon forecasts to evaluate forecasters' disagreement is important to avoid the seasonal patterns emerging from the resolution of uncertainty as time gets closer to the forecasted event one gets when relying on calendar-horizon forecasts instead.

dispersion of inflation forecasts does not react to structural shocks. They however also recognize that the results are sometimes mixed and seem to come from the fact that these structural shocks only account for small fractions of inflation total time variance.¹⁴ By contrast, here we analyze the unconditional reaction to a mixture of all events that can shock the economy and find that the information implying a change in the level of forecast takes time to spread out among agents.

Relationship with models of information rigidity The results that forecasters disagree and that the extent of disagreement varies over time can be rationalized by both *sticky-information* or *noisy-information* models of expectations. In both cases, disagreement is a consequence of imperfect information which implies that agents do not have the same information set.

In a *sticky-information* model, when a large shock hits the economy, individuals who update their information set produce forecasts which are quantitatively very different from individuals who do not. These differences are less pronounced in the wake of a small shock. Consequently, the extent of disagreement evolves with the magnitude of the shocks hitting the economy.

In a *noisy-information* setup, forecasters think differently because they randomly get different perception of reality. In simple versions of this approach, and by contrast with simple *sticky-information* models, disagreement does not change with the size of the shocks hitting the economy.¹⁵ In these setups, disagreement is initiated by the realizations of the idiosyncratic noise that prevent individuals from perfectly observing the true state of nature and that are assumed to have a variance that is constant over time and across individuals. The disagreement is thus not affected by the size of the macroeconomic shocks hitting the economy. It is however possible to refine this simple version of a noisy information model such as to generate time-varying disagreement also in this setup. This would be the case if, for instance, one introduced conditional time variance of the noisy signal correlated with the size of the true shock, or heterogeneity in the noise among the population of forecasters so that the ones that monitor the shock the best disagree more with the ones that have very imprecise signals after a big shock than after a small one.

¹⁴Among the several structural shocks they investigate, 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 altogether cannot account for ranges from 67% to 47% (see Section 5.2 and Appendix C in their paper).

¹⁵This distinction is stressed by Coibion and Gorodnichenko (2008).

All in all, predictable and biased forecasts errors and disagreement among forecasters suggest that both *sticky information* and the *noisy information* models may be good candidates to understand and describe how expectations are formed.

The next three subsections go further by documenting other patterns in the SPF data, that are related to the unfrequent updating of forecasts and, as we argue below, can be related to the inattention of professional forecasters.

3.3 Fact 3: Forecasters do not systematically revise their forecasts

To our knowledge, it is an original contribution of the present paper to document the attention of professional using individual data. While average forecast errors can be constructed using data on the average or consensus forecast, and time-specific disagreement can be assessed using data sets of repeated cross-sections, our measure of the degree of attention 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). In this section, we first describe this measure, then turn to the results we obtain on our sample, and finally discuss the robustness of our indicator.

3.3.1 Measuring the degree of attention

The panel dimension of the SPF data set allows us to build different measures of the average frequency with which an individual revises its forecasts over time. This frequency of updating a forecast provides a measure of the extent with which agents pay attention to new macroeconomic information by incorporating it in their forecasts. It corresponds to a structural parameter in *sticky information* models, called the degree of attention, *i.e.* the frequency with which agents update their information set and therefore, in these models, their forecasts.

More precisely, we consider the quarterly probability of revising the h -step ahead forecast between date $t - 1$ and t . Letting $f_{it,t+h}^x$ be individual's i forecast h quarters ahead for variable x at time t , the probability we aim at estimating is

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

Assuming that the probability is homogeneous across agents, $\lambda_{it}^x(h) = \lambda_t^x(h)$ for all i , we can deliver

empirical counterparts to this attention degree using the SPF data.

Our first indicator relies on calendar horizon forecasts. At all dates, each forecaster is surveyed about a given calendar year. Recalling that T is the index of the last quarter of a given calendar year in the sample, the survey brings information on the probability of revision on a quarterly basis given by

$$\lambda_t^x(T-t) = P(f_{it,T}^x \neq f_{it-1,T}^x).$$

In practice the ECB-SPF surveys forecasters about their expectations for the current calendar and the next calendar years, as well as, in the third and fourth quarter vintage of the survey, about their expectations for two years ahead. Therefore, for each calendar year, say Y , ending in quarter T , we have a sequence of 10 forecasts: two sets of forecasts made at the third and the fourth quarters of year $Y-2$, and 8 sets from the first quarter of year $Y-1$ onwards. Thus we can build a sequence of 9 forecasts revisions, for the same event T , and forecast horizons $h = T-t = 1, \dots, 9$. The degree of attentiveness for the calendar year ending in T and the horizon h can be estimated using the empirical frequency

$$\widehat{\lambda}_{t,cal}^x(h) = \frac{1}{n_t} \sum_{i=1}^{n_t} I(f_{it,T}^x \neq f_{it-1,T}^x),$$

with $h = T-t = 1, \dots, 9$, n_t the number of respondents to the survey at date t and $I(f_{it,T}^x \neq f_{it-1,T}^x)$ an indicator function equal to 1 if $f_{it,T}^x \neq f_{it-1,T}^x$ and 0 otherwise.

A second measure of the probability to update a forecast can be derived from rolling horizon forecasts, exploiting the fact that they are provided for both a 4-quarter and a 8-quarter horizons. Consider the 8-quarter horizon forecast released at date $t-4$ by forecaster i , $f_{it-4,t+4}^x$. This can be compared to the 4-quarter horizon forecast released 4 quarters later $f_{it,t+4}^x$ so as to define the probability of a updating on a yearly basis

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

An empirical estimate is obtained using

$$\widehat{\lambda}_{t,rol}^x = \frac{1}{n_t} \sum_{i=1}^{n_t} I(f_{it,t+4}^x \neq f_{it-4,t+4}^x),$$

with $I(f_{it,t+4}^x \neq f_{it-4,t+4}^x)$ an indicator function equal to 1 if $f_{it,t+4}^x \neq f_{it-4,t+4}^x$ and 0 otherwise. By contrast with the previous measure, the horizon is bound to 4 quarters due to data limitations. We therefore skip the horizon index in the notation, referring to $\widehat{\lambda}_{t,rol}^x = \widehat{\lambda}_{t,rol}^x(4)$.

This probability to update compares two forecasts that are 4 quarters away from each other. It can be converted to a quarterly adjustment rate, $\widehat{\lambda}_{t,rol,q}^x$ if one assumes that the probability of *not* updating is constant over the 4 quarters. In that case, the probability of not updating over the whole year is $(1 - \lambda_t^x(4)) = (1 - \lambda_t^x(1))^4$ so that a quarterly attention rate estimate is $\widehat{\lambda}_{t,rol,q}^x = [1 - (1 - \widehat{\lambda}_{t,rol}^x)^{\frac{1}{4}}]$

Furthermore, assuming $\lambda_t^x(h)$ is constant over time and across horizons, we can also recover micro-data based estimates of the average attention degrees $\widehat{\lambda}_{cal}^x$ and $\widehat{\lambda}_{rol,q}^x$ by simply taking the time average of the empirical frequencies defined above. These estimates can be compared with the macro-data estimates of Kiley (2007), Coibion (2010), Döpke *et al.* (2008) or calibration in Mankiw *et al.* (2003). In addition to assuming a constant attention degree, these previous works also needed auxiliary macroeconomic assumptions, such as the price-setting behavior of firms, to achieve their estimation. By contrast we provide a more direct estimate based on micro data.

Interpretation of $\lambda_t^x(h)$ as a measure of attention deserves further discussion. It could in theory be the case that a forecaster chooses not to revise his forecast in spite of having updated his information set. However, given the vast information set available to professional forecasters we deem unlikely that updating the information set leads to an exactly unchanged optimal forecast after one quarter. Rather, those cases more plausibly correspond to cases where either the forecaster chooses not to run a new forecast exercise, i.e. not to pay the cost of processing the new information by running a full statistical exercise, or not to pay the cost of communicating the new information within or outside the institution. In both situations, the forecaster chooses to leave his public forecast unchanged because of information cost. It is therefore relevant to characterize this lack of reaction to news as (optimal) inattention.¹⁶

3.3.2 Results

The first result is that the estimated attention degree is lower than 100%: forecasters do not update systematically their forecasts on a quarterly basis. This can be seen from Figure (4) which

¹⁶Alvarez, Lippi & Paciello (2010) describe a model where firms pay an information cost to calculate the optimal reset price (*i.e.* implement a price review) and then decide to pay or not the usual menu cost of changing the price. Here the forecaster could be seen as paying the price to calculate the optimal forecast but not the one to change the optimal forecast because of the cost of communicating the reasons of this change.

presents the estimate of attention $\widehat{\lambda}_{t,cal}^x(h)$ for the HICP variable. Figures for real GDP growth and unemployment are not reported to save space but have a similar pattern. Each line in the figure reports a sequence of forecast pertaining to the same calendar year T_Y . There are thus 11 lines corresponding to calendar horizons from $Y = 2001$ to 2011. At each point in time, t , $\widehat{\lambda}_{t,cal}^\pi(h)$ is the proportion of forecasters revising their forecast for a given target calendar year T_Y ending with quarter T . The associated forecasting horizon is thus $h = T - t$. Sequences of forecast revisions partially overlap since, in each vintage of the survey, respondents are asked their forecast for both current and next calendar year.

Although attention is not complete, with a probability of forecast revision varying between 60% and 100% depending on the date, t , it is much higher than the values provided by previous empirical studies. Table (3) shows that the average $\widehat{\lambda}_{cal}^x$ across horizons and dates is 71% for HICP, 74% for the unemployment rate and 80% for GDP growth. Averaging across variables the typical degree of attention is thus $\widehat{\lambda}_{cal} \simeq 75\%$. By comparison, Mankiw *et al.* (2003) calibrated a value of $\widehat{\lambda} = 10\%$ for monthly data, *i.e.* a corresponding $\widehat{\lambda} = 27\%$ when converted to quarterly data, to reproduce the disagreement in US-SPF inflation rate forecasts.

A second result that stands out from our estimation is that the average probability to revise a forecast increases when the forecast horizon decreases: all lines in Figure (4) 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.

A third pattern is that the average attention level varies over time. This can be seen from the average level of each curves in Figure (4) but is even clearer when one looks at Figure (5) which plots the evolution of the alternative attention indicator $\widehat{\lambda}_{t,rol}^x$, over time for the forecasts of the three variables x and shows that there is a significant degree of fluctuations in these parameter.

A fourth and last result stemming from the $\widehat{\lambda}_{t,rol}^x$ estimate is that forecasters do not update systematically their forecasts even on a yearly basis. This confirm our first result that attention is not equal to 100% when looking at $\widehat{\lambda}_{t,cal}^x(h)$ estimates. It is somewhat even more striking since

here some forecasters choose not updating their one-year ahead forecast even after having learned one year of macroeconomic news. Table (3) reports the time average attention degree, $\widehat{\lambda}_{rol}^x$ for the three macroeconomic variables x in the sample. It is equal to respectively 88% for inflation, 95% for unemployment and 94% for GDP. Converting these average probability of forecast revision to quarterly figures gives frequencies, $\widehat{\lambda}_{rol,q}^x$, of 41%, 53% and 51%. Averaging across variables x one gets $\widehat{\lambda}_{rol,q} \simeq 48\%$. The difference between this result and the average $\widehat{\lambda}_{cal} \simeq 75\%$ is further evidence that the frequency of updating is not constant over forecast horizons. Our assessment is that $\widehat{\lambda}_{cal}$ is the most reliable indicator. Indeed, in practice, 8-quarter rolling horizon forecasts are usually not a conventional exercise implemented by professionals. Moreover, since calendar forecasts are available for adjacent quarters, it is more likely that they compare forecasts delivered by the same forecaster.

Relationship with models of information rigidity That forecasters do not update their forecasts at each period is a prediction of *sticky-information* models of expectations. In this approach, agents updating their information systematically also update their forecasts so that not revising a forecast is equivalent to not revising information.¹⁷ On the contrary, the presence of non-updating in the data is not a direct consequence of *noisy-information* models. 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.

Our results also show that the attention degree varies over time. By contrast, the sticky information model of Mankiw & Reis (2002) postulates that λ is constant. Observing fluctuations in aggregate attention is not a direct evidence against the sticky information model of Mankiw-Reis. Indeed it could be reconciled with the more general model of inattention proposed by Reis (2006a) which, under some restrictions, aggregates to a Mankiw-Reis model in which inattention can be described by a single parameter. Note however that also we find a degree of attention that also varies with the forecasting horizon as shown in Figure (6) which violates these conditions of aggregation.

¹⁷A more stringent prediction of these models is that the full stream of future forecasts is unchanged when a forecaster does not update his information set. However, this model does not predict that a forecaster may only revise part of this stream, so that revising a forecast for a particular horizon implies revising for the whole stream.

3.3.3 The influence of rounding

Almost all forecasts (97% of forecast figures) in our data set are reported with one digit only. Failure to revise a forecast may thus merely reflect the fact that these institutions report only rounded figures.

Rounding in our context can be interpreted in two different ways. First, one can consider rounding as a form of inattention. Indeed it is not formally requested by the ECB SPF questionnaire that figures are provided with a single digit, so providing rounded figures may reflect the cost of processing and communicating the information discussed above. 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. Rounding would then lead to underestimate the degree of attention.

To obtain a quantitative assessment of the influence of rounding on our estimation results, we perform the following experiment. We use an estimated VAR model as a simulation and forecasting device, drawing shocks from a multivariate distribution with a covariance matrix equal to that of the estimated innovations.¹⁸ We start by simulating a large sample of artificial observations of length S using this VAR model ($S = 2000$). We then generate a number S of recursive sets of forecasts, up to a 12 quarter horizon, for each date in the simulated data set. 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 forecasts, corresponding to the 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 $h = 9$. Crucially, in our simulation exercise, inattention is absent from the forecasting model: actual forecasts are updated 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

¹⁸See Section 4.1. for a description of the estimated VAR model.

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 for inflation are plotted in Figure (6). The estimated probability of updating a forecast is 83% for the horizon $h = 9$ quarters and rises to 91% 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 gets closer, rationalizing the pattern of Figure (4). However these figures are at all horizons markedly above the estimates of attention we recover from actual SPF micro data, as seen in Figure (6) for the calendar forecast revision case. Thus, independently of rounding effects, there is a degree of supplementary inattention in professional forecasts.

3.4 Fact 4: Forecasters who revise disagree

The panel data set allows us to observe at each date the number of forecasters updating their forecast. We further exploit this specificity of the database in order to document the behavior of the forecasters when they update their forecasts.

We computed at each date in the sample the disagreement among forecasters that do revise their forecast. To get a sense of its importance, we compare it with the disagreement among all forecasters. Figure (7) gives a scatter plot of the level of disagreement in the whole population of forecasters on the x axis against disagreement among the set of forecasters that have revised their forecast at the same date on the y axis. Each observation in the plot corresponds thus to one date and one variable. They are all close to the 45 degree line. The correlation between the two disagreements is thus strongly positive, as high as 97%. At each period, the overall level of disagreement is very similar to the level of disagreement among revisers.

We also investigate whether forecasters update forecasts for several macroeconomic variables at the same time or balance their updating across dates and variables. More precisely we compute (over all observations) and for each variable (inflation, GDP growth, unemployment) the probability that the forecast of this variable is revised, conditional on the fact that the forecast of another variable (inflation, GDP growth, unemployment) has been itself revised. Figures are reported in the bottom panel of Table (3). The notion of forecast revision considered here is the rolling horizon forecast (not converted to quarterly figures). Figures for the conditional probability of revision range from

0.877 (probability of revising inflation forecast given that unemployment forecast was changed) to 0.951 (probability of revising unemployment forecast given that inflation forecast was changed). Recalling that the unconditional probability of forecast revision range from 0.88 (for inflation) to 0.95 (for unemployment), the probability of revising a forecast is thus of the same order whenever forecasts of other variables have been revised.

Relationship with models of information rigidity A prediction of *sticky-information* models is that forecasters who revise their information set will draw the same optimal forecast and thus should not disagree. Therefore this approach cannot explain the large degree of disagreement among forecasters updating their forecast we find. By contrast, *noisy-information* models can generate disagreement between forecasters who revise since every of them has a specific information due to the heterogenous signals on the true state they receive.

In both approaches, the inattention can be rationalized by resorting to the assumptions of limited computational capacity or costly information. In this logic of rational inattention, forecasters could face a trade-off in the degree of attention devoted to the alternative variables. The strong positive correlation at the individual level between the forecast revisions of the different variables shows this is not the case. Once a forecaster update his information set, he tends to update his forecasts of every macroeconomic variable.

3.5 Fact 5: Forecasters are heterogeneous

The SPF data set finally allows to investigate the cross-section properties of individual forecasts. We find that forecasters are heterogenous in terms of their average forecast error, their average attention degree and their average forecast revision.

We first look at individual average forecast errors, namely

$$e_{i,h}^x = \frac{1}{T_i} \sum_{t=1}^{T_i} e_{it,t+h}^x,$$

with T_i the number of forecast observations available for individual i . The top panel in Figure (8) reports the cross section distribution of these individual average bias for 1-year ahead rolling inflation forecasts. On average, professionals underestimated the inflation rate but some of them

were more optimistic and some others were more pessimistic. There is thus a systematic pattern in disagreement since the average forecast bias differs across experts, with substantial dispersion between individuals. The individual average bias ranges broadly from -1% to 1%.¹⁹ The histogram also enlightens that disagreement does not come from some outlier professionals.

We also compute a measure of the attention degree for each individual forecaster. To do so, we assume an attention degree $\lambda_{it}^x(h) = P(f_{it,t+h}^x \neq f_{it-1,t+h}^x)$ that is homogenous across dates t , $\lambda_{it}^x(h) = \lambda_i^x(h)$. We can then estimate this individual specific attention degree as the percentage of quarters in which his forecasts for a given horizon h are being revised,

$$\widehat{\lambda}_i^x(h) = \frac{1}{T_i} \sum_{t=1}^{T_i} I(f_{it,t+h}^x \neq f_{it-1,t+h}^x).$$

The middle panel of Figure (8) plots the cross-section distribution of $\lambda_i^x(h)$ for the 1-year rolling inflation forecasts. The conclusion is that the degree of attention varies across forecasters. Though there is a mass of forecasters with values of λ_i above 80%, the dispersion is substantial. The less attentive among forecasters revise on average their forecasts every 2 quarters while the most attentive ones adjust every quarter.²⁰

Moreover, as a measure of the extent of forecast revision, we also calculate the average absolute of forecast revision between two subsequent quarters at the individual level

$$|\Delta f_{i,h}^x| = \frac{1}{T_i} \sum_{t=1}^{T_i} (f_{it,t+h}^x - f_{it-1,t+h}^x).$$

The cross-section distribution of the average revision for the 1-year rolling inflation forecasts is depicted in the bottom panel of Figure (8), underlining again substantial heterogeneity around a mode of .28%.

Finally, we investigate whether differences in the individual attention degree generate differences in both the individual average forecast errors and revisions. We regress the absolute value of the individual average forecast error, $|e_{i,h}^x|$, and of individual average forecast revision, $|(\Delta f_{i,h}^x)|$, on the individual attention rate, $\lambda_i^x(h)$. Table (4) provides the results for the 1-year rolling forecasts of

¹⁹This systematic disagreement pattern can be linked to disagreement stemming from heterogeneity in priors or in models as Patton & Timmermann (2010) put forth.

²⁰The fact that some forecasters systematically adjust their forecasts also shows that rounding is not the mere responsible for observing a frequency of forecast revision lower than one. Would it be the case, no forecaster would systematically update.

inflation, unemployment and real GDP. Strikingly, attention is positively correlated with forecast errors and forecast revisions.

Relationship with models of information rigidity The cross-section heterogeneity observed in average individual forecast errors can be generated both by *sticky-information* and *noisy-information* models, postulating that agents have either different frequencies of updating their forecasts or different precisions in their noisy signals. This heterogeneity could be rationalized by resorting to different information capacities constraints. However neither model generates systematic forecast bias so that the unconditional individual average errors should be zero. The observed non-zero average forecast errors can only be associated with non-zero in-sample bias due to the particular history of news over the time sample studied.

Again, only *sticky-information* models generate a lack of systematic forecast revision. Postulating heterogeneity in the frequency of updating has the potential to rationalize the heterogeneity in forecast errors that is also observed. Allowing for heterogeneity in the frequency of updating the information set may give rise to an aggregation issue. The frequency of updating the information that is consistent with the pattern of the average forecast across the whole population of heterogeneous forecasters would be lower than the average of the individual probability of updating (see Carvalho, 2006).

Finally, both *sticky-information* and *noisy-information* models can generate heterogeneity in the cross-section of the average size of forecast revisions. In the *sticky-information* setup, this revision will decrease with the degree of attention. When forecasters update more frequently their information set, they incorporate more frequently the news in their forecast and therefore are subject to a lower probability of a large revision in the future. This is at odds with the results of Table (4). By contrast, in the *noisy-information* framework and for low attention level, *i.e.* very noisy individual signals, the forecast revision can become an increasing function of inattention. Very noisy signals produce almost no updating of forecasts because forecasters know about their imprecision. However, this property of the model does not help explain why the frequency of revision increases the size of the forecast error and the forecast revisions. These results could hint to a situation where the causal link is reversed: attention increases when forecasters had bad previous forecast

performances implying large forecast errors and forecast revisions.

4 Models of inattention: a quantitative assessment

The previous section has shown that the SPF data features the two basics ingredients that are in line with either a sticky information model of expectations *à la* Mankiw-Reis or a noisy information model *à la* Sims: biased forecasts and disagreement among forecasters. Moreover, forecasts are not always updated, a prediction of Mankiw-Reis' model and agents disagree on their forecasts even when they update their information/prediction, a prediction of Sims' model. These *qualitative* results suggest a model of expectations in which both type of inattention coexist: agents infrequently update their information set and when do they get a noisy perception of the true information.²¹

In this section we develop such a model of individual expectation and then assess whether it is able to *quantitatively* match the average forecast errors, disagreement and probability to update a forecast that we observe in the data using a formal testing procedure.²²

4.1 A sticky and noisy information model

4.1.1 DGP and information structure

We assume that the economy can be described by the following reduced form VAR(p) model

$$A(L)X_t = \epsilon_t, \quad t \in \mathbb{Z},$$

where X_t is a set of n macroeconomic variables, centered on their average, L is the lag operator, $A(L) \equiv \sum_{k=0}^p A_k L^k$ has all its roots outside the unit-circle and $A_0 = I$ and where $E\{\epsilon_t | X_{t-1}, \dots, X_{t-p}\} = 0$. The VAR(p) model can be rewritten in the usual more compact first-order VAR companion form

$$Z_t = FZ_{t-1} + \eta_t,$$

²¹This approach shares some features with the price-setting model of Woodford (2009). In this setup, firms have to choose on implementing a price review knowing that they will get a noisy information on the true state of the economy. This leads to optimal non-systematic price review and, when it happens, to reset prices that are determined conditional on the noisy information.

²²Given that we do not have a clear-cut interpretation of what lies behind the heterogeneity of forecasters we consider a model that disregards this deep heterogeneity.

with $Z_t = (X_t' X_{t-1}' \cdots X_{t-p+1}')$ and $\eta_t = (\epsilon_t' 0 \cdots 0)'$.

Let i be an individual in the population of forecasters, $i = 1, \dots, n$. At each date, every forecaster may update his information set or not. Along the lines of Mankiw & Reis (2002), we model this updating as a Poisson random variable, $\mathcal{P}(\lambda)$. Therefore λ is the probability to update and pay attention to the news. It is also called the *degree of attention*. Assuming that the population of forecasters is large, by the law of large numbers, each date t a fraction λ of the total population updates its forecast. Consequently, at each point in time, the whole population of forecasters is split into groups, generations say, within which each forecaster refers to the same vintage of information. We let j denote the generation of forecasters that last updated their information set j periods before the current one, *i.e.* in $t - j$. Thus j is also an index of the vintage of information they use. The fraction of generation j in the whole population is given by $\lambda(1 - \lambda)^j$. In Mankiw & Reis (2002) model, every agent updating his information set gets a perfect signal on the state of the economy Z_t . The optimal forecast at horizon h is given by the expectation conditional on this perfect information, $E(Z_{t+h}|Z_t) = F^h Z_t$. It is therefore identical for every forecaster who updates. Likewise, the forecasters who last updated in $t - j$ receive a perfect signal on Z_{t-j} and their optimal forecast for $t + h$ is the conditional expectation with respect to that information vintage, $E(Z_{t+h}|Z_{t-j}) = F^{h+j} Z_{t-j}$. So, the Mankiw & Reis (2002) information structure implies no disagreement within a generation j of forecasters. As the previous section made clear, this is at odds with the facts.

We therefore extend Mankiw and Reis' model to include imperfect perception of the information when updating. More precisely, we assume that, when an agent i updates his information, he gets a noisy perception of the true state, Z_t , namely a signal Y_{it} that follows

$$Y_{it} = H'Z_t + v_{it}, \quad v_{it} \sim \text{iid}(0, \Sigma_v),$$

where H is a matrix that selects the state variables that are observed with a noise.²³ Remark that the average over forecasters gives the right signal: $E_i(Y_{it}) = H'Z_t$. A simple way to rationalize this noise is to rely on the fact that forecasters use real-time data which they know are prone to measurement errors.

²³A typical case is $H = I$, in which case forecasters have noisy perceptions of all the state variables. An alternative would be to consider that forecasters have perfect access to past realizations of X_t .

4.1.2 Optimal individual forecasts

Let $f_{it,t+h}$ denotes agent's i optimal forecast for the vector X at date $t+h$, with respect to date t information (and $f_{it,t+h}^x$ the corresponding expectation for the specific component x in X). Every individual updating his information set faces the problem of inferring the true state of the economy, Z_t , from his imperfect signal Y_{it} . Let $Z_{it|t}$ denotes the date t state of the economy perceived by agent i at date t : $Z_{it|t} = E(Z_t|Y_{it}, Y_{it-1}, \dots)$. Its optimal vector of forecast is given by

$$f_{it,t+h} = E(Z_{t+h}|Z_{it|t}) = F^h Z_{it|t}.$$

A solution to this signal extraction problem is given by the Kalman filter namely,

$$Z_{it|t} = Z_{it|t-1} + G_{it} (Y_{it} - H' Z_{it|t-1}),$$

with $Z_{it|t-1} = E(Z_t|Y_{it-1}, Y_{it-2}, \dots) = F Z_{it-1|t-1}$, and G_{it} the gain of the Kalman filter, defined by

$$G_{it} = P_{it|t-1} H (H' P_{it|t-1} H + \Sigma_v)^{-1},$$

where $P_{it|t-1}$ denotes the variance of the perceived forecast error,

$$P_{it|t-1} = E [(Z_t - Z_{it|t-1})(Z_t - Z_{it|t-1})'].$$

Noticeably, the gain is common across agents, $G_{it} = G_t$, as soon as one postulates the usual initial conditions of the recursion, namely that $Z_{i1|0} = E(Z_1)$ and $P_{i1|0} = E\{[Z_1 - E(Z_1)][Z_1 - E(Z_1)]'\}$, for every agent in the population. Indeed, the conditional variance of the forecast error in a Kalman filter follows

$$P_{it+1|t} = F [P_{it|t-1} - P_{it|t-1} H (H' P_{it|t-1} H + R)^{-1} H' P_{it|t-1}] F' + \Sigma_v,$$

and the recursion is therefore identical across agents when they start from the same *a priori* $P_{i1|0}$.²⁴

Individuals who cannot update their information set at date t stick to their old one, like in a Mankiw-Reis type model, but with the difference that past information vintages are noisy. Let $f_{it-j,t+h}$ be the optimal expectation of an agent i using information vintage j , his optimal forecast for date $t+h$ is given by $f_{it-j,t+h} = E(Z_{t+h}|Z_{it-j|t-j})$ with $Z_{it-j|t-j} = E(Z_t|Y_{it-j}, Y_{it-j-1}, \dots)$ the

²⁴That would not be the case if, in particular, one considered a "deep" disagreement involving differences in the perception of the parameters value underlying the true state DGP, for instance $E(Z_{i1|0}) \neq E(Z_1)$.

state Z_{t-j} perceived by an agent i that sticks to the information vintage j . The optimal forecast of individuals that updated j periods ago is therefore given by

$$f_{it-j,t+h} = \mathbf{E}(Z_{t+h}|Z_{it-j|t-j}) = F^{h+j}Z_{it-j|t-j}.$$

4.1.3 Average forecast and error

Let $\mathbf{E}_{ij}(\cdot)$ be the expectation over i and j . The average forecast over the whole population of forecasters, also called the *consensus*, is given by

$$\mathbf{E}_{ij}(f_{it-j,t+h}) = \mathbf{E}_j[\mathbf{E}_i(f_{it-j,t+h}|j)].$$

That forecast average can be split into the average over individuals using the same information vintage j , and the average over the different generations of forecasters or information vintages.

We consider, without loss of generality,²⁵ individuals using the information vintage $j = 0$ (*i.e.* the generation of forecasters able to update information at the current date). Combining the Kalman filter expressions above, we can rewrite the optimal vector of forecasts of individuals in this generation as

$$f_{it,t+h} = F^h Z_{it|t-1} + F^h G_t (Y_{it} - H' Z_{it|t-1}). \quad (1)$$

Remarking that $\mathbf{E}_{ij}(Y_{it}|j = 0) = H' Z_t$, this leads to the following forecast average within this generation

$$\mathbf{E}_{ij}(f_{it-j,t+h}|j = 0) = F^{h+1} \mathbf{E}_i(Z_{it-1|t-1}) + F^h G_t H' [Z_t - F \mathbf{E}_i(Z_{it-1|t-1})].$$

Useful insights can be derived from the comparison with the perfect information case. In that case every individual would have an optimal forecast given by

$$f_{t,t+h}^* = F^{h+1} Z_{t-1} + F^h (Z_t - F Z_{t-1}) = F^{h+1} Z_{t-1} + F^h \eta_t. \quad (2)$$

Comparing equations (1) and (2), the difference between the noisy information and the perfect information cases stem from two channels. First, the perceived state one period backward is not the true one: $Z_{t-1} \neq \mathbf{E}_i(Z_{it-1|t-1})$. This also implies that the perceived innovation is not the

²⁵Forecasts that were generated at date $t-j$ for the $t+h$ horizon date can be rewritten as forecast at date τ for the $\tau+l$ horizon date, with $\tau = t-j$ and $l = h+j$.

true one; $Z_t - FE_i(Z_{it-1}|Z_{t-1}) \neq \eta_t$. Second, because it is acknowledged to be noisy, the perceived innovation is not completely incorporated into the forecast: $G_t H' \neq I$. Consequently, the forecast error associated with the average within a generation is predictable with respect to the information available at date t . Indeed, let $\mathcal{E}_{t,t+h}^j = Z_{t+h} - E_i(f_{it-j,t+h})$ be that forecast error, we have

$$\begin{aligned}\mathcal{E}_{t,t+h}^{j=0} &= F^h Z_t + \sum_{k=0}^h F^{h-k} \eta_{t+k} - \left[F^h E_i(Z_{it|t-1}) + F^h G_t H' (Z_t - E_i(Z_{it|t-1})) \right] \\ &= F^h (I - G_t H') [Z_t - E_i(Z_{it|t-1})] + \sum_{k=0}^h F^{h-k} \eta_{t+k},\end{aligned}$$

which implies $E(\mathcal{E}_{t,t+h}^{j=0}|Z_t) \neq 0$ as long as $I \neq G_t H'$.

The optimal forecast for an individual updating his information set at time t is given by $f_{it,t+h}$ and by $f_{it-j,t+h}$ if the information set was last updated j periods ago. The average of individuals' across generations of forecasters then follows

$$E_j [E_i(f_{it-j,t+h}|j)] = \sum_{j=0}^{\infty} (1-\lambda)^j \lambda E_i(f_{it-j,t+h}).$$

Let $\mathcal{E}_{t,t+h} = E_j(\mathcal{E}_{t,t+h}^j)$, the average forecast error, *i.e.* the error associated with the consensus forecast, we have

$$\begin{aligned}\mathcal{E}_{t,t+h} &= Z_{t+h} - E_j [E_i(f_{it-j,t+h}|j)] \\ &= F^{h-1} (I - G_t H') \sum_{j=0}^{\infty} (1-\lambda)^j \lambda [Z_t - E_i(Z_{it|t-j})] + \sum_{k=0}^{h-1} F^{h-k} \eta_{t+k+1}.\end{aligned}$$

Taking the expectation with respect to the true state, Z_t leads to

$$E(\mathcal{E}_{t,t+h}|Z_t) = F^{h-1} (I - G_t H') \sum_{j=0}^{\infty} (1-\lambda)^j \lambda [Z_t - E_i(Z_{it|t-j})], \quad (3)$$

which shows that the average forecast error is predictable, *i.e.* $E(\mathcal{E}_{t,t+h}|Z_t) \neq 0$, as long as $I \neq G_t H'$ and/or $\lambda < 1$.

4.1.4 Disagreement between forecasters

The extent of differences in opinion can be assessed by the cross-section variance of point forecasts over individuals, i , and information vintages, j , that we denote $V_{ij}(f_{it-j,t+h})$. Using the standard

variance decomposition formula leads to

$$V_{ij}(f_{it-j,t+h}) = E_j \{V_i(f_{it-j,t+h}|j)\} + V_j \{E_i(f_{it-j,t+h}|j)\}. \quad (4)$$

This expressing underlies that the disagreement across individuals stems from two sources.

The first source is the noise in individuals' signal leading to differences in perception *within* a given generation of forecasters j , *i.e.* $V_i(f_{it-j,t+h}|j) \neq 0$. Within a generation of forecasters disagreement is only generated by the variance of the individual noise, Σ_v . Indeed it holds that

$$V_i(f_{it-j,t+h}|j) = F^{h+j} \{V_i(Z_{it-j|t-j-1}|j) + G_t [\Sigma_v + H'V_i(Z_{it-j|t-j-1}|j) H] G_t'\} (F^{h+j})'. \quad (5)$$

This cross section variance evolves with the forecast horizon, shrinking progressively to zero with h . The first term in total disagreement averages the vintage- j specific 'within' components of disagreement described by equation(5). Each generation j is weighted by its relative share, $\lambda(1-\lambda)^j$, in the total population, so that

$$E_j \{V_i(f_{it-j,t+h}|j)\} = \sum_{j=0}^{\infty} (1-\lambda)^j \lambda V_i(f_{it-j,t+h}|j). \quad (6)$$

It is important to remark that this component of disagreement is not equal to zero, even when $\lambda = 1$.

The second source comes from the differences in average opinion due to the different information vintages used by the forecasters, that is differences *between* generations of forecasters. Indeed, within generation heterogeneity averages out in $E_i(f_{it-j,t+h}|j)$. So the only cross-section dispersion remaining in the second term of equation (4) is due to differences in information vintage, j . More precisely, we have

$$V_j \{E_i(f_{it-j,t+h}|j)\} = \sum_{j=0}^{\infty} \{(1-\lambda)^j \lambda E_i f_{it-j,t+h} - E_j [E_i(f_{it-j,t+h}|j)]\}^2. \quad (7)$$

Unlike the first component of disagreement of equation (6), this second component is equal to zero when $\lambda = 1$.

4.1.5 Model properties

The above hybrid sticky/noisy information model can match several of the empirical regularities documented in the previous section, namely predictable forecast errors and disagreement among

forecasters updating their information. We now discuss the properties of some moments characterizing the forecast error and disagreement and their relation to the inattention parameters, λ and Σ_v .²⁶

Properties of the average forecast errors The average forecast does not fully incorporate the news released at date t , η_t . This leads to persistent forecast errors, as can be seen from equation (3). A decrease in the attention degree, λ , increases both the persistence and the variance of forecast error. An increase in a term on the first diagonal of the noise variance matrix, Σ_v , increases the persistence of the forecast error of the corresponding forecasted variable. It also increases the variance of the forecast errors. All in all, less attention generates more persistence and variance of the forecast errors.

Properties of the disagreement As equations (6) and (7) above show, the model generates disagreement even under full information updating, $\lambda = 1$. In line with Mankiw-Reis model, disagreement increases when λ decreases. Another important feature of the model is to generate time-varying disagreement even when there is no time conditional heteroscedasticity in the measurement error, v_{it} so that Σ_v is constant across dates. This comes from the differences across generations of forecasters. The degree of disagreement depends on the difference between the new vintage of information and the previous one. If the last innovation is large compared to the average, the difference of opinion between the individuals revising and the others, therefore the disagreement, will be larger than the average one. This time variance of disagreement increases when λ decreases.

An increase in any diagonal element of Σ_v has two different effects on the disagreement within a generation of forecasters. On the one hand it increases the amount of noise, thus raises the differences of opinions within the subgroup of forecasters that refer to the same vintage of information. On the other hand, because individuals know that the signal is very imprecise, they incorporate less the news to their forecast. In the extreme case, when the signal is completely uninformative, the optimal forecast is the unconditional mean of the process for all forecasters, implying zero disagreement.

²⁶The discussion is based on simulations not reported in the paper to save space but available upon request.

Lastly, the model structure also implies that the time variance of disagreement decays with the variance of the noise. The less informative the news, the less they are incorporated in the optimal forecast, therefore the less disagreement there is between generations of forecasters. In the extreme case when precision of the signal approaches zero, news are uninformative so are not reflected in the forecasts, and the time variance of the disagreement shrinks to zero.

4.2 Estimation procedure

We perform an estimation and a test of the previous model relying on a Minimum Distance Estimation (MDE) procedure.²⁷ The building blocks in this estimation procedure are first, $\hat{\mu}$, a vector of K data-moments, such as average disagreement or average forecast error, which are computed from the SPF panel data set, and, second, a corresponding set of model-generated moments $\mu(\lambda, \Sigma_v)$ which are a function of the parameters λ and Σ_v . Estimates of the parameters λ and Σ_v are produced by minimizing the following distance criteria:

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

where $\hat{\Omega}$ is a consistent estimator of Ω , the asymptotic variance of $\hat{\mu}$ defined by $\sqrt{T}(\hat{\mu} - \mu) \rightarrow \mathcal{N}(0, \Omega)$ when $T \rightarrow \infty$. In addition, letting $\theta = (\lambda, \text{vec}(\Sigma_v))$ denote the parameters of interest, and $\hat{\theta}$ the minimum distance estimator of θ , standard errors on the parameter estimate $\hat{\theta}$ can be implemented using the property that, when $T \rightarrow \infty$

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

with $D \equiv D(\theta) = \nabla_{\theta}\mu(\theta)$ the Jacobian of $\mu(\theta)$ with respect to θ evaluated at $\hat{\theta}$. An estimator of $D'\Omega^{-1}D$ is given by $\hat{D}'\hat{\Omega}^{-1}\hat{D}$ where $\hat{D} \equiv D(\hat{\theta})$.

The MDE procedure also allows to test for over-identifying restrictions, i.e. the null hypothesis that the set of K moments μ can accurately be described with the P parameters to be estimated i.e. the elements of λ and Σ_v . The test statistic is given by the objective function $[\hat{\mu} - \mu(\lambda, \Sigma_v)]' \hat{\Omega}^{-1} [\hat{\mu} - \mu(\lambda, \Sigma_v)]$. Here, rather than relying on the asymptotic Chi-square distribution, we will use Monte Carlo simulations of the model to approximate the exact small distribution of the test statistics.

²⁷See for example Chapter 14 in Wooldridge (2002) for a general description of this method.

The moments we select to define the objective function and estimate the model are guided by the structural model of the above section. We consider moments related to forecast errors and disagreement, since they are functions of attention rate and measurement error variance, thus entailing econometric identification. More precisely we select two moments related to forecast errors: the mean square of forecast errors (MSE), $E[e_{t,t+h}^x]^2$ and the first-order autocorrelation of forecast errors, $\rho_e^x(1)$. We also consider two moments that are associated with disagreement: first the average level of disagreement, $E(\sigma_{t,h})$ and second the time variance of disagreement, $V(\sigma_{t,h})$. Note that both the sticky information and the noisy information components of the model raise the level of disagreement. On the contrary, only the sticky information model can predict a high level of across time variance in disagreement. Thus this latter moment is particularly crucial in order to discriminate among the two models. We carry out estimation using five alternative sets of these moments, allowing to perform some robustness checks. The first set is made of all four moments mentioned above for both output and inflation, giving a total of eight target moments. The second and third sets include these four moments respectively for inflation only and for output only. The last two sets of moments are made of the mere disagreement moments, for output and inflation respectively. These last two cases are thus just-identified.

The simulation of the model generated moments is carried out as follows. In a preliminary step, we estimate an auxiliary model, a Vector Auto Regressive (VAR) model for inflation, output growth and unemployment. The VAR model include 4 lags and all variables are centered prior to estimation. We use quarterly euro area data over the estimation period 1987:Q1 to 2009:Q2.²⁸ The choice of the sample period reflects a trade-off between conflicting objectives: first, using a sufficiently long time series and second, considering an homogenous period in terms of monetary policy and the inflation regime. Then taking the VAR parameters as given we perform simulation of moments in the following way. For a set of structural parameters (λ, Σ_v) , we use the Kalman filter and the hybrid model of the previous section to generate average forecasts and forecasts errors given by equation (3). At each date, we also use the formulae (4) to (7) to compute the level of disagreement for all variables. We can thus deduce the standard deviation in disagreement over time. These moments are compared to actual moments in order to compute a distance. Estimation is performed using a numerical routine to minimize the objective function over the space of parameters. We restrict the

²⁸We use the ECB AWM database.

covariance matrix of the measurement error to $\Sigma_v = \sigma_v^2 I_3$. Thus we assume that the variance of the measurement error is common for inflation, unemployment and output growth. Experiments relaxing the variance of the measurement errors for each of the variable led to broadly similar results.

4.3 Results

Table (5) gives the estimation results. Five different sets of estimates are presented, each corresponding to a particular set of target moments. The first case, reported in Column (1), uses the full set of eight moments: for each variable, the considered moments are the average squared forecast errors, $E((e_{t,t+h}^x)^2)$, the degree of persistence, $\rho_e^x(1)$, the average disagreement $E(\sigma_{t,h}^x)$ and the time variance in disagreement $V(\sigma_{t,h}^x)$. The estimated value for σ_v the measurement noise, is 0.225 and that for $\hat{\lambda}$, the attention parameter, is 0.054. This latter figure is lower than existing estimates of the attention parameter obtained from time series evidence. For instance Mankiw *et al.* (2003), on US data find the best fitting value for λ to be 0.27 (where we convert their monthly estimate into a quarterly one). More strikingly, this figure is well below the micro estimate of around 0.75 obtained from the micro data in Section 3 above.²⁹ Thus, along this first dimension, the model is at variance with the SPF micro facts. We also test for over-identifying constraints using the test outlined above. The test rejects the null that the distance between the estimated moments and the observed ones is zero, with a p-value of 0.04.

Column (2) presents results obtained when fitting only the inflation moments as a target for estimation. The point estimates turn out to be very similar to the previous case, and the model is rejected by the formal test at the 10% level, with a p-value of 0.07. Column (3) presents estimation results obtained with only GDP moments. Here, the parameter estimates are even more extreme with $\hat{\lambda} = 0.023$ and $\hat{\sigma}_v = 0.563$. A larger degree of noise is indeed needed to fit the degree of disagreement in output forecast, which is higher. In that exercise, which restricts to GDP moments, the model is not rejected by the data (the p-value of the test is 0.96). However, the estimated value for λ is well below our micro-data estimate. Lastly, columns (4) and (5) present models

²⁹Due to the heterogeneity in λ_i we document in Section 3.5., and following the argument in Carvalho (2006) in the case of the frequency of price adjustment, our estimation of the average attention rate may be prone to an aggregation bias. However, since the lower bound in the λ_i is 50%, that aggregation bias would not be enough to reconcile the MDE results with the micro data evidence.

that are just-identified. They provide similar estimates of the attention degree as in the baseline case, and larger estimates of the noise variance than the baseline case ($\hat{\sigma}_v = 0.631$ and $\hat{\sigma}_v = 0.441$ respectively). Overall, these estimates and tests suggest that our flexible noisy information/sticky information model fails to quantitatively fit the data.

To illustrate what lies behind this formal rejection, Table (6) reports the set of target moments in the SPF data and the same set of moments as generated by the model for different values of the parameters.³⁰ Column (2) reports the moments generated by the model for the baseline estimated values of parameters. For inflation, and even more for output, the level and the time variance of disagreement predicted by the model is higher than that observed in the data (for instance average disagreement about output is 0.362 in the data and 0.484 in the model). Conversely, for both variables, the mean squared forecast error generated by the model is lower than that observed in the data. The over-identifying restriction test indicates that these differences are significant, so even if flexible the model cannot jointly fit the two sets of moments. In essence, to fit the observed large degree of forecast error variance and persistence the model needs a very low degree of attention λ . For this level of attention the degree of disagreement generated by the model turns out to be much larger than that observed in the data.

The additional experiments reported in columns (3) and (4) of Table (6) further illustrate this point. First, in column (3) we keep σ_v at the estimated value $\hat{\sigma}_v$, but set the attention degree λ to a much higher value, here the value observed in the micro data $\lambda = 0.75$. We observe that the disagreement for both output and inflation is lowered as compared with column (2). For instance, disagreement in GDP moves from $\sigma_t^{GDP} = 0.484$ to $\sigma_t^{GDP} = 0.443$, therefore closer to the data value of 0.362. However, at the same time, the mean squared forecast error generated by this degree of attention degree is much too low to be in line with the data: for inflation, setting $\lambda = 0.75$ rather than $\lambda = \hat{\lambda}$ moves the model-generated mean squared forecast error from 0.529 to 0.145 while the data value is 0.819. The variance of the noise is helpful in our model to generate high mean squared forecast error. However there is a limit to this margin in terms of fitting the data, as indicated by our last experiment reported in Column (4). In this experiment we increase the standard deviation of the noise to the value of $\sigma_v = 1$. As expected the MSE of forecast increases: for inflation it rises

³⁰Values of the across-time variance of disagreement have been multiplied by 100 so that moments presented in that table are of comparable magnitude.

to 0.599. We might have expected that increasing σ_v would also increase disagreement because each forecaster typically gets a signal farther from the true value of the state variable. However, another mechanism, discussed in Section 4.1.5., counteracts this effect: with a higher variance of the noise, forecasters tend to shrink more their estimate of the state variable and put less emphasis on the signal they perceive. It turns out that here the latter effect dominates. For instance, in the case of output, disagreement decreases from $\sigma_{t,h}^{GDP} = 0.484$ in the estimates to $\sigma_{t,h}^{GDP} = 0.406$. It is still noticeable that such a level of noise variance is not optimal in terms of minimizing the distance with observed moments. As indicated by the rows pertaining to the moment $V(\sigma_{t,h})$, the gap with the data that emerges which such a level of noise variance is that it flattens too much the time pattern of disagreement. The model-generated values for time variance of disagreement are too low in this experiment, while they were too high for more moderate values, in experiments of Columns (2) and (3).

All in all, the sticky/noisy information model is rejected by our quantitative test because the optimal (minimal distance) estimates of attention and noise parameters tend to generate too much disagreement and not enough variance and persistence in forecast errors as compared with the data. It remains to be understood why forecasters agree on sluggish revisions of their forecasts.

5 Conclusion

In this paper, we analyze an original data set, the European Survey of Professional Forecasters, to characterize the formation of individual expectations with a particular emphasis on providing micro facts relating to the sticky and noisy imperfect information models recently introduced in the macroeconomic literature. In particular, we provide an estimate for the degree of attention based on individual observations, and find that professionals are, albeit mildly, inattentive. We also document that, to a large extent, the disagreement among forecasters stays the same no matter whether they revise their forecasts or not.

A formal test rejects a model of expectation featuring both sticky and noisy information. Indeed, that model is not able to account for the strong persistence of the forecast errors together with the relatively low level of disagreement between the forecasters observed in the data. There is 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 more elaborate versions of inattention models could be investigated. One could for example study how much of the discrepancy between the data and our hybrid expectation model could be filled in by considering a degree of attention that varies across individuals and over forecasting horizons, two features of the SPF data we highlighted. One could also consider noisy signals that are common to every forecasters, as a way to account for the low relatively low disagreement in the data. Second, and beyond the mere framework of inattention theories, another avenue is to investigate whether alternative forms of deviations from the perfect information rational expectation setup, for instance model uncertainty or strategic interactions between forecasters, provide a better match of the empirical patterns documented in this paper.

References

- [1] Amador, M. and Weill, P. (2009). Learning from private and public observations of others' actions. Mimeo, Stanford University & UCLA.
- [2] Alvarez, F., Lippi, F. and Paciello, L. (2010). Optimal price setting with observation and menu costs. NBER Working Paper (15852).
- [3] Andrade, P. and Gregoir, S. (2010). Can heterogenous information explain the gap between survey expectations and the rational expectation hypothesis? Test and evidence from exchange rate forecasts. Mimeo.
- [4] Angeletos, G. M. and Pavan, A. (2007). Efficient use of information and social value of information. *Econometrica*, 75: 1103–42.
- [5] Angeletos, G. M. and La'O, J. (2009). Noisy business cycles. *NBER Macroeconomics Annual*, 24.
- [6] 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.
- [7] Branch, W. A. (2007). Sticky information and model uncertainty in survey data of inflation expectations. *Journal of Economic Dynamic and Control*, 31:269–98.
- [8] Carroll, C. (2003). Macroeconomic expectations of households and professional forecasters. *Quarterly Journal of Economics*, 118:269–298.
- [9] Coibion, O (2010). Testing the sticky information Phillips' curve, *Review of Economic and Statistics*, 92: 87–101.
- [10] Carvalho, C. (2006). Heterogeneity in price stickiness and the real effect of monetary shocks. *The BE Journal of Macroeconomics (Frontiers)*, 2.
- [11] Coibion, O. and Gorodnichenko, Y. (2008). What can survey forecasts tell us about informational rigidities? NBER Working Paper (14586).

- [12] Coibion, O. and Gorodnichenko, Y. (2010). Information rigidity and the expectations formation process: A simple framework and new facts. Mimeo UC Berkeley.
- [13] Döpke J., Doovern J., Fritsche U. and Slacalek J. (2008). Sticky Information Phillips Curves: European Evidence *Journal of Money, Credit and Banking*, 40(7): 1513-1520.
- [14] Friedman, M. (1968). The role of monetary policy. *American Economic Review*, 58:1–17.
- [15] Hellwig, C. & Veldkamp, L. (2009). Knowing what others know: Coordination motives in information acquisition. *Review of Economic Studies*, 76:223–51.
- [16] 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.
- [17] Klenow, P. J. and Willis, J. L. (2007). Sticky information and sticky prices. *Journal of Monetary Economics*, 54:79–99.
- [18] Lorenzoni, G. (2009). A theory of demand shocks. *American Economic Review*, 99:2050–84.
- [19] Lucas, R. E. Jr. (1972). Expectations and the neutrality of money. *Journal of Economic Theory*, 4:103–124.
- [20] Maćkowiak, B. and Wiederholt, M. (2009). Optimal sticky prices under rational inattention. *American Economic Review*, 99:769–803.
- [21] 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.
- [22] Mankiw, N. G., Reis, R. and Wolfers, J. (2003). Disagreement about inflation expectations. *NBER Macroeconomic Annuals*, 209–247.
- [23] Mankiw, N. G. and Reis, R. (2010). Imperfect information and aggregate supply. forth. *Handbook of Monetary Economics*.
- [24] Morris, S. and Shin, H. S. (2002). Social value of public information. *American Economic Review*, 92:1521–32.

- [25] Patton A. and Timmermann, A. (2010) Why do forecasters disagree? Lessons from the term structure of cross-sectional dispersion. *Journal of Monetary Economics*, 57:803–820.
- [26] 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.
- [27] Phelps, E. S. (1968). Money-wage dynamics and labor market equilibrium. *Journal of Political Economy*, 76:678–711.
- [28] Reis, R. (2006a). Inattentive producers. *Review of Economic Studies*, 73:793–821.
- [29] Reis, R. (2006b). Inattentive consumers. *Journal of Monetary Economics*, 53:1761–1800.
- [30] Sims, C. (2003). Implications of rational inattention. *Journal of Monetary Economics*, 50:665–690.
- [31] Veldkamp, L. (2009). Information choice in macroeconomic and finance. Mimeo, NYU.
- [32] Woodford, M. (2002). 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.
- [33] Woodford, M. (2009). Information-constrained state-dependent pricing. *Journal of Monetary Economics*, 56:S100–24.
- [34] Wooldridge, J. (2002). *The econometrics of individual and panel data*. MIT press.

Table 1: Realizations, average forecasts & errors

Descriptive statistics & tests			
Sample period: 1999Q1–2010Q2			
	HICP	UNEM	RGDP
REALIZATIONS, x			
Mean	2.32	8.23	2.18
Std Dev	.63	.60	1.07
AVERAGE FORECASTS, f (Forecasting horizon: 1 year)			
# of individuals	92	90	92
Mean	1.73	8.65	1.74
Std Dev	.27	1.06	1.05
Resp. rate	.57	.53	.56
AVERAGE FORECAST ERRORS, e			
Mean	.31	.04	-.36
RMSE	.89	.75	1.62
MAE	.66	.53	1.11
$\rho_e(1)$.735	.860	.841
AVERAGE FORECAST EFFICIENCY TEST			
$\rho_e(h)$	-.906	.468	.218
p-value	.052	.056	.253

Notes: All 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); $\rho_e(h)$ = h-order autocorrelation of the error; p-values are computed using a robust Newey-West covariance matrix.

Table 2: Time variations in disagreement

Time-series bivariate regressions			
Sample period: 1999Q1–2010Q2			
	HICP	UNEM	RGDP
DEPENDENT VARIABLE: σ_t			
$ \Delta x_{t-1} $.076 (.141)	.571 (.838)	2.179** (.161)
e_{t-1}^2	.018* (.012)	.041* (.027)	.009*** (.002)
$ \Delta f_t $.006* (.004)	.039 (.057)	.047** (.003)

Notes: All variables are in %; HICP = euro area HICP inflation rate (yoy); UNEM = euro area unemployment rate; RGDP = euro area real GDP growth rate (yoy); σ_t = disagreement across forecasters; Δx_t = change in forecasted variable; Δf_t change in the average forecast; e_t = average forecast error; all regressions include a constant term; numbers in brackets are standard errors of estimates using a robust Newey-West procedure; ***, **, * indicates significance at the 1, 5 and 10% levels respectively.

Table 3: Errors, disagreement & inattention

Sample period: 1999Q1–2010Q2			
	HICP	UNEM	RGDP
UNCONDITIONAL MOMENTS			
Mean(e)	.31	.04	-.36
Mean(λ^{cal})	.71	.74	.80
Mean(λ^{rol})	.88	.95	.94
Mean($\lambda^{rol,q}$)	.41	.53	.51
Mean(σ)	.26	.26	.33
Std Dev(x)	.62	.60	1.07
CONDITIONAL PROBA OF UPDATING			
$P(\Delta f^y \Delta f^x)$			
HICP	1	.878	.877
UNEM	.951	1	.944
RGDP	.939	.938	1

Notes: All 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; Δf = revision of forecast; $P(\Delta f^y | \Delta f^x)$ = probability to revise y forecast (in line) conditional on revising x (in column).

Table 4: Individual errors, revision & inattention

Cross-section bivariate regressions			
	HICP	UNEM	RGDP
DEPENDENT VARIABLE: $ e_i $			
λ_i	.581** (.245)	1.92*** (.672)	1.01 (1.06)
DEPENDENT VARIABLE: $ \Delta f_i $			
λ_i	.721*** (.154)	3.17** (1.34)	.989* (.577)

Notes: All variables are in %; HICP = euro area HICP inflation rate (yoy); UNEM = euro area unemployment rate; RGDP = euro area real GDP growth rate (yoy); λ_i = individual probability to update its forecast; e_i = average forecast errors; Δf_i = average forecast revision; 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 5: Minimum Distance Estimation and Tests

Selected moments					
	(1)	(2)	(3)	(4)	(5)
$\hat{\lambda}$.054 (.028)	.052 (.080)	.023 (.046)	.064 (.176)	.023 (.048)
$\hat{\sigma}_v$.255 (.019)	.221 (.040)	.563 (.099)	.631 (.244)	.441 (.139)
TEST: Over-identifying restrictions					
J-stat	135.2	29.5	.22	-	-
p-value	.04	.07	.96	-	-

Notes: Table reports MDE estimates (with standard errors in brackets) & test statistics. e_t denotes date t average error, $\rho_e(1)$ its first-order autocorrelation and σ_t the date t disagreement across forecasters. Matched moments are **(1)**: $E(e_t^2)$, $\rho_e(1)$, $E(\sigma_t)$, $V(\sigma_t)$ for π and Δy ; **(2)**: $E(e_t^2)$, $\rho_e(1)$, $E(\sigma_t)$, $V(\sigma_t)$ for π ; **(3)**: $E(e_t^2)$, $\rho_e(1)$, $E(\sigma_t)$, $V(\sigma_t)$ for Δy ; **(4)**: $E(\sigma_t)$, $V(\sigma_t)$ for π ; **(5)**: $E(\sigma_t)$, $V(\sigma_t)$ for Δy .

Table 6: Model generated moments

Moments under various configurations				
	(1)	(2)	(3)	(4)
Data	$\lambda = \hat{\lambda}$	$\lambda = .75$	$\lambda = \hat{\lambda}$	$\lambda = \hat{\lambda}$
	$\sigma = \hat{\sigma}$	$\sigma = \hat{\sigma}$	$\sigma = \hat{\sigma}$	$\sigma = 1$
INFLATION, π				
$E(e_t^2)$.819	.529	.145	.599
$\rho_e(1)$.781	.806	.537	.822
$E(\sigma_t)$.272	.278	.228	.235
$V(\sigma_t)$.624	.880	.458	.421
REAL GDP, Δy				
$E(e_t^2)$	2.722	2.194	.413	2.539
$\rho_e(1)$.873	.783	.431	.786
$E(\sigma_t)$.362	.484	.443	.406
$V(\sigma_t)$	2.685	6.238	3.653	1.481

Notes: Table reports various moments of inflation, π , and real GDP growth rate, Δy , obtained with MDE estimates or alternate parameters values. e_t denotes date t average error, $\rho_e(1)$ its first-order autocorrelation and σ_t the date t disagreement across forecasters.

Figure 1: Consensus forecast & realizations

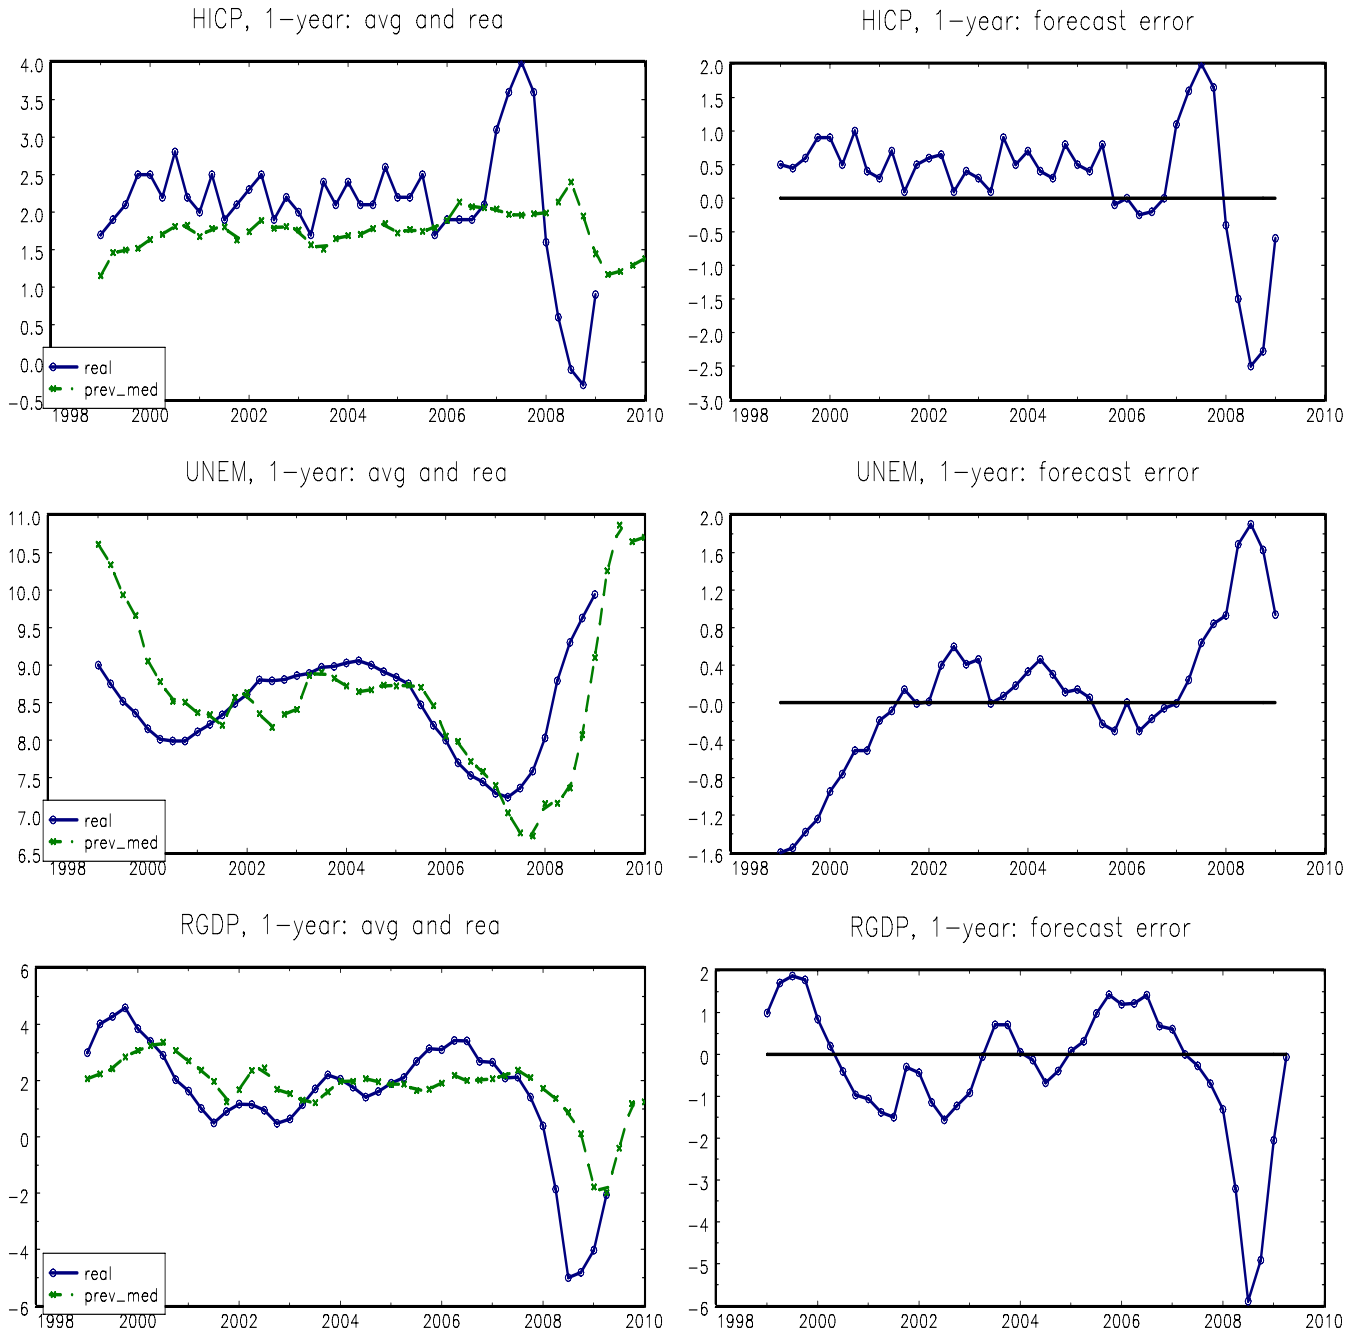


Figure 2: Cross-section distribution of RGDP 1-year forecasts (1999:Q1–2003:Q4)




Figure 2: Cross-section distribution of RGDP 1-year forecasts (2004:Q1–2008:Q4)




Figure 2: Cross-section distribution of RGDP 1-year forecasts (2009:Q1–2010:Q2)

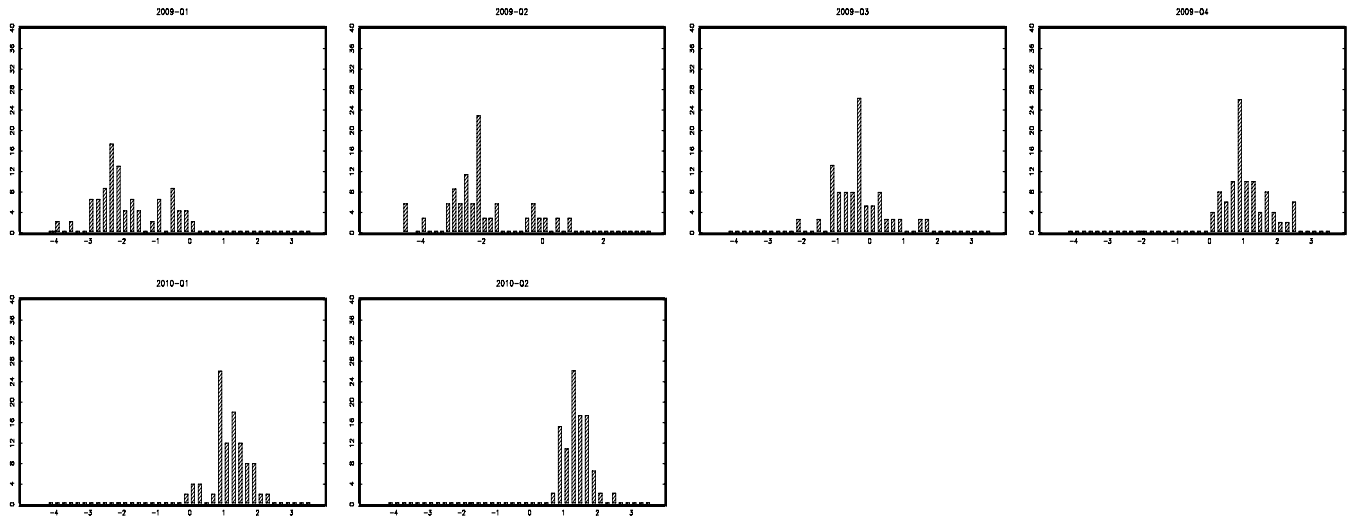


Figure 3: Evolution of disagreement

Disagreement, 1-year forecasts




Figure 4: Evolution of attention

Attention degree (calendar forecasts)
HICP




Figure 5: Evolution of attention

Attention degree
rolling horizon forecasts




Figure 6: The effect of rounding: 1-year, HICP

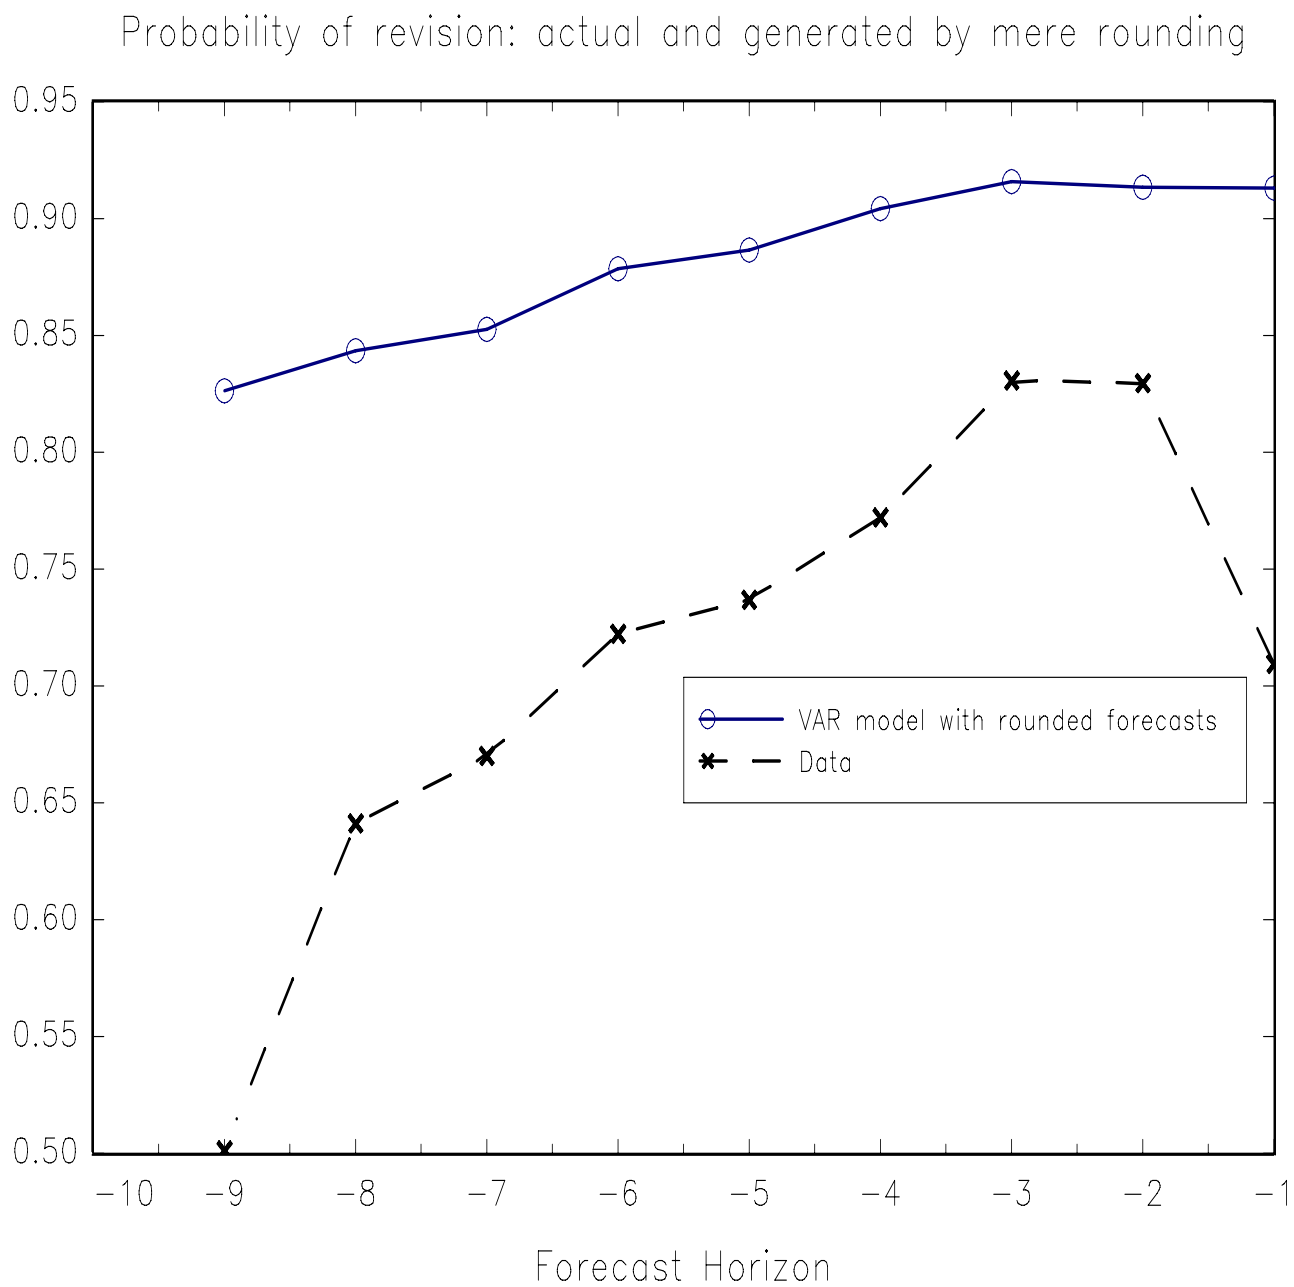


Figure 7: Disagreement among revisers

Disagreement and disagreement among revisers





Figure 8: Cross-section distribution of errors, attention rates & revisions; HICP, 1-year



Documents de Travail

290. C. Bordes and L. Clerc, "The ECB art of central banking and the separation principle," August 2010
291. R. Jimborean and J-S. Mésonnier, "Banks' financial conditions and the transmission of monetary policy: a FAVAR approach," September 2010
292. G. Dufrenot and L. Paul, "Fiscal development in the euro area beyond the crisis: some lessons drawn from fiscal reaction functions," October 2010
293. R. Cooper, H. Kempf and D. Peled, "Insulation impossible: monetary policy and regional fiscal spillovers in a federation," October 2010
294. C. Célérier, "Compensation in the financial sector: are all bankers superstars?," October 2010
295. O. de Bandt and S. Malik, "Is there evidence of shift-contagion in international housing markets?," October 2010
296. F. Ferroni, "Did Tax Policies mitigate US Business Cycles?," October 2010
297. E. Challe and X. Ragot, "Fiscal policy in a tractable liquidity-constrained economy," October 2010
298. P. Cahuc and E. Challe, "Produce or speculate? Asset bubbles, occupational choice and efficiency," October 2010
299. H. Kempf and G. Rota Graziosi, "Endogenizing leadership in tax competition: a timing game perspective," October 2010
300. X. Ragot, "The Case for a Financial Approach to Money Demand," October 2010
301. E. Challe, F. Le Grand and X. Ragot, "Incomplete markets, liquidation risk, and the term structure of interest rates," October 2010
302. F. Le Grand and X. Ragot, "Prices and volumes of options: A simple theory of risk sharing when markets are incomplete," October 2010
303. D. Coulibaly and H. Kempf, "Does Inflation Targeting decrease Exchange Rate Pass-through in Emerging Countries?," November 2010
304. J. Matheron, « Défisicalisation des heures supplémentaires : une perspective d'équilibre général », Décembre 2010
305. G. Horny and P. Sevestre, "Wage and price joint dynamics at the firm level: an empirical analysis," December 2010
306. J. Coffinet and S. Lin, "Stress testing banks' profitability: the case of French banks," December 2010
307. P. Andrade and H. Le Bihan, "Inattentive professional forecasters," December 2010

Pour accéder à la liste complète des Documents de Travail publiés par la Banque de France veuillez consulter le site :
http://www.banque-france.fr/fr/publications/documents_de_travail/documents_de_travail_10.htm

For a complete list of Working Papers published by the Banque de France, please visit the website:
http://www.banque-france.fr/fr/publications/documents_de_travail/documents_de_travail_10.htm

Pour tous commentaires ou demandes sur les Documents de Travail, contacter la bibliothèque de la Direction Générale des Études et des Relations Internationales à l'adresse suivante :

For any comment or enquiries on the Working Papers, contact the library of the Directorate General Economics and International Relations at the following address :

BANQUE DE FRANCE
49- 1404 Labolog
75049 Paris Cedex 01
tél : 0033 (0)1 42 97 77 24 ou 01 42 92 62 65 ou 48 90 ou 69 81
email : marie-christine.petit-djemad@banque-france.fr
jeannine.agoutin@banque-france.fr
michael.brassart@banque-france.fr
veronique.jan-antuoro@banque-france.fr
[nathalie.bataille-salle@banque-france.f](mailto:nathalie.bataille-salle@banque-france.fr)