Forecasting ECB Monetary Policy: Accuracy Is a Matter of Geography*

This version: December 2008

Helge Berger

Free University Berlin
Department of Economics
Boltzmannstr. 20
14195 Berlin, Germany
& CESifo
Helge.Berger@ fu-berlin.de

Michael Ehrmann

European Central Bank
Kaiserstr. 29
60311 Frankfurt am Main
Germany
Michael.Ehrmann@ecb.int

Marcel Fratzscher

European Central Bank
Kaiserstr. 29
60311 Frankfurt am Main
Germany
Marcel.Fratzscher@ecb.int

Abstract

Monetary policy in the euro area is conducted within a multi-country, multi-cultural, and multi-lingual context. How does this heterogeneity affect the ability of economic agents to understand and to anticipate monetary policy by the ECB? Using a database of surveys of professional ECB policy forecasters in 24 countries, we find remarkable differences in forecast accuracy, and show that these have important repercussions on market behaviour. Explaining the differences in forecast accuracy, we provide evidence that they are partly related to geography and clustering around informational hubs, as well as to country-specific economic conditions. In large part this heterogeneity can be traced to differences in forecasting models.

JEL No.: E52, E58, G14
Keywords: monetary policy; ECB; forecast; geography; heterogeneity; Taylor rule; survey data.

* We would like to thank Reuters for providing data on their ECB monetary policy polls and Ann-Kristin Koch for able research assistance. Berger thanks the ECB's Research Department and the IMF's European Department for their hospitality. Two anonymous referees, Martin Cihak, Jörg Decressin, Hamid Faruquee, Elmer Sterken, Bob Traa, and seminar participants at the ECB, the IMF and the ifo Institute provided helpful comments. This paper presents the authors' personal opinions and does not necessarily reflect the views of the European Central Bank or the IMF.

1. Introduction

Monetary policy in the euro area is conducted within a multi-country, multi-cultural, and multi-lingual context. With the formation of the European Economic and Monetary Union (EMU), countries with markedly different histories of inflation, monetary policy strategies, and central banking traditions, delegated the conduct of monetary policy to a single entity--the European Central Bank (ECB). Since this transition, questions that have come to the fore include whether economic agents in different member countries have been able to adjust and understand equally well the European perspective of monetary policy? To what extent are expectations about monetary policy still influenced by different national and cultural backgrounds? Moreover, to what extent may differences in the ability to forecast monetary policy decisions reflect more permanent information asymmetries related to geographic proximity to the ECB, and country-specific factors?

While random heterogeneity in forecasts will characterize many uncertain environments, the presence of systematic heterogeneity is important from a policy perspective. Expectations are a crucial transmission channel for monetary policy, and systematic heterogeneity linked, for instance, to geographical factors can imply that monetary policy in the euro area exerts differential effects in the various EMU countries. Furthermore, systematic differences in expectations among financial market participants imply that their trading behaviour differs, thus leading to heterogeneous investment profitability. Furthermore, a systematically different perception of the ECB's monetary policy could prompt country-specific political controversy about the ECB's role.

This paper analyzes the ability of economic agents in EMU and non-EMU countries to forecast monetary policy decisions by the ECB. In particular, we investigate to what extent expectations are related to geographic location of forecasters and country-specific characteristics. We develop a novel database of monetary policy expectations by 120 financial institutions in 24 countries between 1999 and 2005. The data stem from surveys conducted by Reuters, and provide information on the expected ECB policy rate for the upcoming Governing Council meeting, the probability distribution around forecasters' point estimates, as well as their expectations about the timing of future monetary policy moves. The survey responses accurately represent forecasters' expectations for two reasons. First, they are generally in the public domain, which implies that they must be in line with the recommendations given by the institutions to their clients. Second, as most institutions participate regularly, clients have the possibility to evaluate the respective forecasting performance of the various institutions.

We find that differences in forecast accuracy are substantial, as the top 10 percent of all institutions have a forecast error that is on average 8 basis points smaller than the error made by the worst 10 percent of performers. These differences are significant in economic terms, both from a financial market perspective and from a policy point of view, reaching a level of about one-third of the typical ECB policy rate change of 25 basis points during the sample period. Furthermore, they have repercussions on actual trading behaviour, as larger heterogeneity in market expectations causes an increase in volatility in asset prices following the announcement of monetary policy decisions.

What explains this large heterogeneity in anticipating ECB monetary policy decisions? A first result of our empirical analysis is that *geography* matters for forecast accuracy. There is a surprising amount of cross-country variance in expectations about ECB policy rates. But the pattern of forecast accuracy exceeds the concept of nationality. Frankfurt, which hosts the ECB headquarters and the German Bundesbank (one of the euro area's National Central Banks and, arguably, one of the ECB's early role models), also is Germany's financial centre. Being close to this informational hub tends to improve forecast accuracy of forecasters. The advantage of being located in an informational hub is corroborated by equivalent results for forecasters based in London/UK or working for institutions with a subsidiary in the City.

We find furthermore that national *macroeconomic conditions* tend to influence forecast accuracy, as deviations of national inflation from the euro area average affect the quality of forecasts. Forecasters in countries with higher (lower) than euro area inflation tend to produce more hawkish (dovish) forecasts, suggesting that the national inflation environment biases forecasters' views of ECB monetary policy.

We also show that the observed heterogeneity is *systematic* rather than based on differences in the judgmental component among forecasters. To extract the systematic component from observed predictions, we estimate bank-specific reaction functions. These Taylor-type rules, which tend to be different across institutions, capture a significant part of the underlying structure of the published forecast.

To our knowledge, the focus on explaining the heterogeneity in monetary policy expectations is novel. It relates to earlier literature on the differences in the transmission of the ECB's monetary policy. In the run-up to EMU, several papers asked whether a change in policy rates would affect national economies in a heterogeneous fashion, possibly due to differences in expectations (Dornbusch et al. 1998; Cecchetti 2001; Mihov 2001). However, results are contradictory across studies (Mojon and Peersman 2003). Evidence using data obtained under EMU is scarce, for the most part because long time samples are needed to estimate the full transmission path from interest rates to inflation. Accordingly, the few studies available analyze only elements of the transmission process (e.g., Angeloni and Ehrmann 2003). No study has yet been conducted on the homogeneity of interest rate expectations in the euro area, and this is where the present paper attempts to contribute.

In addition, our work is broadly related to the literature on trade in goods and in financial assets, as well as the literature on home bias in the allocation of financial portfolios. For instance, there is substantial empirical evidence that information asymmetries and information frictions are fundamental in explaining trade in goods and financial assets as well as financial investment decisions (e.g. Froot and Stein 1991; Gordon and Bovenberg 1996, Portes, Rey, and Oh 2001; Hau 2001; Dvorak 2005). Such information asymmetries can take various forms and can be related to language, cultural ties, common legal origins, and institutions, among others. As the literature on home bias and capital flows emphasizes, information is also a key factor inhibiting "optimal" investment decisions based on portfolio theory. Moreover, location decisions by financial firms point to the importance of information-based agglomeration effects. Even though centrifugal factors in the sense of Krugman (1998) exist—for instance, the need to be close to locally dispersed customers—and advances in communication technology continue to lower transaction costs, centripetal forces seem to matter more in the financial services (Tschoegl 2000, Clark 2002, Cook et al. 2007).

In particular, geographic proximity and common socio-cultural attitudes remain key when it comes to the realization of information spillovers and economies of scale in information processing (Thrift 1994, Grote 2004). Faulconbridge (2003, p. 237) counts "face-to-face contact facilitated by social proximity" among the arguments why financial companies agglomerate in international financial centres. More generally, Strauss-Kahn and Vives (2005) show that firms tend to locate their headquarters preferably in close vicinity to other headquarters in the same sector of activity. Another related strand of the literature emphasizes that investors are more profitable when investing in firms that are located in geographic proximity. For instance, Coval and Moskowitz (1999, 2001) show that mutual fund managers earn significantly more on investments in firms with headquarters located geographically near to the mutual fund's offices. The present paper suggests that information asymmetries and agglomeration effects along these lines might also influence the quality of forecasting of ECB monetary policy.

Finally, our results are linked to the literature on the importance of location for the quality of financial markets forecasts—but there are also important differences. Bae et al. (2009) provide international evidence for a significant advantage of local analysts even when controlling for the quality of information provided by firms, and Malloy (2005) shows that within the US equity analysts perform better when investing in firms located in close geographic proximity. But these findings may not necessarily extend to the domain of monetary policy forecasting in the case of a large cross-national currency union. Financial analysts are said to profit from geographic proximity because of similarities in culture, language, closeness to local market conditions in which firms operate, and potentially better links to decision makers (e.g. Malloy 2005). In contrast to a typical firm, however, the ECB does not have a predominantly national character. Instead, it is a very international institution focusing on the euro area in its entirety. In addition, national central banks, located in each member country, take part in all monetary policy decisions, which should further reduce the relative advantage of ECB watchers located in the proximity of Frankfurt. In that sense, our result that geography matters also in the context of monetary policy is indeed surprising.

Our findings have important policy implications. The results indicate that the ECB operates in an environment where economic agents have yet to converge on a common expectation-formation process when it comes to monetary policy. This heterogeneity is systematically related to differences in forecasting models, significant in size, and closely related to geographic and country-specific factors. Given the importance of expectations for the transmission of monetary policy, continued heterogeneity along these lines might prove problematic. Furthermore, a systematically different perception of the ECB's monetary policy could prompt country-specific political controversy. One implication of these findings is that there is room for policies that foster the convergence to a common expectation-formation process, for instance through careful and targeted central bank communication.

The remainder of the paper is organized as follows. Section 2 presents our dataset based on the Reuters survey as well as some key stylized facts of these data. The analysis of the determinants of differences in forecasting abilities across forecast institutions, distinguishing between geography and country-specific economic conditions, is presented in Section 3. Section 4 decomposes these forecast errors into systematic and unsystematic components, and tests whether country-specific economic conditions lead to a forecast bias. Section 5 summarizes the main findings and discusses policy implications.

2. Forecast errors of monetary policy and their effects on financial markets

We start by discussing the source of our expectations data and its quality, and outline a number of stylised facts (section 2.1), before proceeding to analyse to what extent the cross-sectional heterogeneity in expectations has effects on financial markets (section 2.2).

2.1 Data and Some Stylized Facts

The data on ECB policy expectations consist of about 4,500 observations. They comprise forecasts filed by financial institutions (rather than individual analysts) from 11 euro area and 13 other countries polled by Reuters prior to 97 ECB Governing Council meetings between January 1999 and January 2005. The panel encompasses on average about 30 respondents per Governing Council meeting in the first years and an average of about 60 toward the end of the sample (with a minimum of 18, a maximum of 64, and an average of 44 participants). Overall, the sample contains 120 institutions. The polled financial institutions vary somewhat from meeting to meeting, but a core group—comprising most large euro area financial institutions, those operating in the City of London, and a few in the US, Australia, and Japan—is represented for most meetings. About half the observations are from financial institutions with headquarters in euro area countries. Results of the poll are released by Reuters in one news item, after all responses have been obtained. Issues of herding, which have been shown to be relevant, e.g., in the context of security analysis (Welch 2000), should therefore play a relatively minor role in the current application.

As to the institutional background, the Governing Council usually meets twice per month, mostly on the first and third Thursday of each month. While monetary policy decisions were taken at both meetings before November 2001, they are since taken only at the first meeting each month.

Usually the Friday before each Governing Council meeting, Reuters asks a number of financial institutions to attach probabilities to different scenarios for the ECB's policy rate (the interest rate paid for the ECB's main refinancing operations). More specifically, forecasters assign probabilities that the ECB increases or decreases the policy rate by 50 or 25 basis points, or keeps it unchanged. In addition, the questionnaire asks for the institution's opinion on the probable timing of the next ECB policy change, its direction, and its size. Moreover, there are often (but not always) questions regarding the expected interest rate levels at the end of the current and following calendar year. Table 1 gives an example of the raw data collected by Reuters.

Table 1

While the questions changed somewhat over time and not all issues were touched upon prior to every meeting, the data allows us to construct various indicators that summarize well the expectations of forecasters: for instance, the expected ECB policy rate (*expected rate*) for the meeting ahead, the policy rate with the highest probability, and the expected speed of ECB decision-making as measured by the expected number of meetings until the next change in the policy rate. We will focus our analysis below on the *expected rate*.[1]

[1] The working paper version of the paper (Berger et al. 2006) also includes a discussion of possible time variations in our findings, as well as the results for the expected number of meetings until the next expected policy change (*meetings-to-change*) and the rate with the highest probability (*most likely rate*), both of which are qualitatively similar.

With a few exceptions, all responses are attributed, and the specific location where the Reuters questionnaire is answered is identified.[2] In many cases, the reporting location is the institution's headquarters. If headquarters and reporting location diverged, it is often the institution's subsidiary in Frankfurt or London that answers the questionnaire. We use this information to describe the geographic pattern of the forecasts.

How good and accurate is the expectations data based on these surveys? This is an important issue as the validity of the analysis in this paper crucially depends on the accuracy of the expectations data. As the survey is conducted a few days before each Governing Council meeting, it should reflect the latest information and expectations among the participants. Moreover, as the participants are largely from institutions for which monetary policy decisions are highly relevant, they are among the best informed about upcoming ECB decisions. Nevertheless, it cannot be ruled out that some participants may have an incentive to misreport their true expectations in the survey, for instance in order to stand out and attract more attention than others. However, it should be stressed that many of the survey participants are bound by their statements to their clients, such that it would be difficult for them to justify why they provide different forecasts to their clients than to the Reuters survey.

Nevertheless, we gauge the quality of our survey-based measure by comparing it to a market-based measure of monetary policy expectations, calculated from EONIA swap rates. The drawbacks of these market-based measures are that they give only the mean of the market expectations – rather than the whole distribution across agents as our survey-based measure does – and also that they are available only for a somewhat shorter time period. But still, a close match between the mean expectations from the survey and those from the market-based measure should indicate that the survey data provides a reasonably accurate measure of true expectations.

Figure 1

Figure 1 shows the two expectation errors over time, and indicates that the two indeed are highly correlated and also similar in magnitude. Indeed, the correlation coefficient between the two expectation errors is higher than 80%, and the average absolute expectation error is very similar with 6.2 basis points for the survey-based measure and 6.4 basis points for the market-based measure.

What do the data tell us? In a first attempt to gauge forecast accuracy, we compute the absolute forecast errors, that is, the absolute difference between the predicted and the ex-post action of the ECB at a particular Governing Council meeting. A first stylized fact is that the ability to anticipate ECB policy decisions differs markedly across forecast institutions. Figure 2 shows the average absolute forecast error by the financial institutions from the most accurate to the least accurate decile. The most accurate institutions have an average absolute forecast error of around 2 basis points, while the least accurate ones have an average absolute error of more than 10 basis points. Differences of up to 8 basis points are certainly significant in economic terms, both from a financial market perspective and from a policy point of view, reaching a level of about one-third of the typical ECB policy rate change of 25 basis points during the sample period. The differences in forecasting ability are quite notable compared to

[2] We discard all unattributed responses, as well as those of institutions that have participated less than seven times in the poll.

the observed average forecast error in the sample of about 6 basis points for the full sample period (see below Table 2).

Figure 2

Table 2 views forecasting performance through the country-lens, that is, we sort reporting institutions based on the country in which they are headquartered. The table reports the absolute forecast error by country, for all meetings and separately for these with or without a change in policy rates.

Table 2

The country-perspective suggests that geography matters. Forecasters based in some countries did better than others: while forecasters based in companies headquartered in Germany and Portugal show significantly lower absolute errors than the euro area average, forecasters in Austria, France, and Ireland show significantly higher errors. Also, as a rule, forecasters working with financial institutions located outside the euro area performed less well than their euro area counterparts.

Another interesting pattern uncovered by Table 2 is that, perhaps not unexpectedly, forecasters face particular difficulties correctly predicting interest rate changes. In fact, across the sample, the absolute forecast error in periods of changing policy rates is more than four times larger than in calmer periods. Moreover, there is some movement in the ranking of forecasting institutions by country, even though a clear pattern is hard to discern by simple descriptive statistics alone.

Despite these differences in forecast accuracy among institutions, generally the forecasts are unbiased and efficient. The very last columns in Table 2 provide information on standard tests allowing rejection of the hypothesis that the forecasts are biased or inefficient—both for the overall sample and, in most cases, by country. Figure 3 illustrates that the average forecast closely tracks the actual (post-meeting) ECB policy rate during most of the period, while individual forecasts varied more widely.

Figure 3

While there is evidence that forecast performance differs across countries, Table 2 also indicates that there is substantially more variation within than across countries. The table reports the standard deviation of the mean absolute forecast error across institutions within each country. This standard deviation ranges from 0.06 to 0.14, with an average of 0.08. Calculating the standard deviation of the average national forecast errors leads to a considerably lower figure, which stands at 0.02. This suggests that the search for determinants of forecasting performance needs to go beyond the country dimension.

2.2 Effects on financial markets

An important question is whether such heterogeneity has repercussions on financial markets. If survey participants truthfully reveal their expectations for the upcoming interest rate decisions and take according market positions, we should observe increased market volatility in response to the announcement of monetary policy decisions in the presence of larger heterogeneity, as a more substantial rebalancing of market positions may be required in such a case. This would provide strong evidence in favour of the above argument that the survey data are of high quality, and at the same time suggest that heterogeneity in expectations, in particular if it is systematic, has effects on financial markets.

Table 3 reports results of corresponding tests. We calculate and explain the volatility of the spot €/$ exchange rate following the announcement of the interest rate decision, based on tick-by-tick real time quotes recorded on Reuters. Volatility is measured by the standard deviation of the recorded exchange rates, starting with the first tick after the release of the monetary policy decision at 13:45:00 CET, over a window of 10 minutes (left panel) or 20 minutes (right panel).[3] The volatilities for each meeting day are then regressed on the absolute size of the surprise, defined by the median response of the Reuters respondents, and the heterogeneity in expectations, given by the standard deviation of the responses in the Reuters poll.

These regressions must take into account that market volatility on the day of the Governing Council meeting, the heterogeneity in expectations as well as the size of the surprise could possibly be jointly driven by some omitted variable. For instance, a surge in general macroeconomic uncertainty may make market expectations more heterogeneous and increase financial market volatility at the same time. To account for this possibility, Table 3 shows three different measures of market volatility. The first simply calculates volatility over the respective time window on the day of the meeting (σ_t), and is as such prone to the above critique. Alternatively, we have calculated how market volatility in the time windows increases relative to the volatility in other, unrelated, time windows. The second column of each panel in Table 3 relates the volatility on the day of meetings to the volatility at the same time on the previous day (σ_t/σ_{t-1}), whereas the third column relates it to the same time window a week in advance (σ_t/σ_{t-5}). Results are robust to all three measures. However, an advantage of the latter two measures is that their coefficients can be more easily interpreted.

Table 3

The table shows that market volatility does indeed depend on uncertainty. As expected, the absolute size of the surprise is not statistically significant in these models, indicating that the surprise component that is *common* to all market participants is priced in immediately after the announcement, and as such is not part of our time window that starts only with the first tick after the announcement. At the same time, however, heterogeneity in market expectations matters, as more heterogeneity increases volatility in response to the announcement of decisions. The measured effects are not only statistically significant, they are also large. The measure of heterogeneity ranges from 0 to 0.22. The results in Table 3 indicate that for the meeting with the most heterogeneous market expectations, market volatility has increased relative to the previous day by a factor of around 3, and relative to the previous week by a factor of around 2.

The results summarized so far invite further discussion. There is a substantial degree of heterogeneity in market expectations, which furthermore affects market behaviour. In what follows, we will therefore attempt to understand the determinants of market heterogeneity

[3] The data is available to us starting with the Governing Council meeting on September 14th, 2000, thus covering 66 meeting days.

better, decomposing the findings along two dimensions in order to study whether this heterogeneity is systematic or random.

One dimension pertains to the geographical pattern. This is open to a number of interpretations. The good forecasting performance of institutions headquartered in, say, Germany could be due to their relative closeness to Frankfurt. Alternatively, it might have to do with advantages from informational agglomeration. Another relates to theory. If institutions based their forecasts of ECB behaviour (mostly) on economic theory, one would expect that their predictions could, one way or the other, be explained by concepts like the Taylor rule. If true, this would allow filtering out the unsystematic errors made by forecasters—and focusing on the systematic, model-based error instead.

3. The Role of Geography and Macro Conditions

To learn more about the pattern underlying the forecast errors, we move to an econometric approach. Our main objective is to understand differences in the forecast accuracy (in the sense of a smaller absolute forecast error) across financial institutions in anticipating ECB monetary policy decisions. Hence we investigate whether the absolute value of the forecast error varies across institutions, and whether this difference can be explained with factors reflecting geography and country-specific macroeconomic conditions. While the directional differences in forecast errors are not central to the question of accuracy, we will shed some light on this issue in Section 4, where we analyze whether differences in forecasting performance are related to different weights forecast institutions give to various macroeconomic variables when forming their policy expectations.

Since we are interested in overall forecast accuracy, the variable of interest is the absolute observed forecast error e made by each forecast institution i for meeting t:

$$|e_{it}| = |r_{it}^{e} - r_t|,$$

where r_t is the policy rate after the Governing Council meeting t and r_{it}^{e} are institution i's published expectations for the policy rate. The model takes the general form

$$|e_{it}| = \alpha + \boldsymbol{\beta}\mathbf{y}_{it} + \gamma_t + u_{it} \quad (1)$$

where α is a constant, $\mathbf{y}_{it}$ is a vector of explanatory variables that are either institution-specific or institution- and meeting-specific, $\boldsymbol{\beta}$ is the matching coefficient vector, γ_t represents meeting-fixed-effects, and u_{it} is a residual following standard assumptions.[4]

Including meeting-fixed-effects in the model is a very flexible tool to robustly control for a number of potentially influential unobservables. Time variations in forecast accuracy can either arise in our dataset (e.g., due to changes in the composition of the survey panel) or they can reflect time variations in the predictability of ECB decisions. For instance, the change from bimonthly to monthly meetings in late 2001 or the clarification of the ECB's strategy in

[4] Note that we want to explain the difference in the magnitude of the forecast error made by each forecasting institution. Accordingly, the model contains a number of purely institution-varying explanatory variables, which precludes the inclusion of an institution-fixed effect.

May 2003 could affect how well the average ECB watcher can forecast upcoming decisions. Including meeting-fixed-effects will help to account for all of these factors in an efficient way.[5]

To learn more about the pattern underlying the forecast errors, we define a comprehensive set of explanatory variables along the dimensions of location, macro conditions, and controls for the type of the forecasting institution. The variables are defined as follows:

Location	
Headquarters in Frankfurt:	dummy variable; one for all financial institutions filing their forecasts of ECB policy rates from Frankfurt in the majority of cases
Subsidiary in Frankfurt:	dummy variable; marks financial institutions that report at least once from Frankfurt, but from other locations in the majority of cases
Headquarter or subsidiary in London:	dummy variable; identifies financial institutions reporting at least once from London that are not already included in the first two categories

Macro conditions	
Inflation differential:	difference between national and euro area HICP inflation rates, based on the latest information available at the time of the Reuters poll; set to zero for all financial institutions that are headquartered outside the euro area

Institution controls	
Type of institution:	set of 3 dummy variables, defined for i) investment banks, ii) commercial banks, iii) other institutions
Experience in ECB watching:	step dummy, for the number of participations of an institution in the survey; =1 as of the 13th, =2 as of the 30th participation in the survey (corresponding to the 33 and 66 percentiles in the distribution)

Table 4 presents the corresponding OLS estimates, with standard errors allowing for heteroskedasticity. The results show the importance of information and location. Institutions that are either headquartered in Frankfurt or operate a subsidiary there tend to have significantly lower absolute forecast errors than others.[6] This might be because it is helpful to observe the ECB's action up-close and on an everyday basis or because of informal information spillovers in a financial centre. Another explanation for the importance of Frankfurt as location might be previous experience: to the extent that the ECB's actions

[5] A different way to analyze time variations is to break down the estimation period into different sub-periods. However, there is no compelling evidence that the relationships we find change significantly across time periods.
[6] The headquarters identified by the reporting criterion are not always the *chartered* headquarters of the respective financial institution. However, we can safely assume that the dominant location from which Reuters polls are filed corresponds to the location of the section that performs the ECB-watching tasks. More often than not, this section will also conduct a substantial fraction of the euro area business of the respective financial institution.

followed similar patters as the Bundesbank before 1999, local analysts might have found the human capital they had accumulated watching the Bundesbank still helpful after 1999.

In addition, financial expertise in general and, possibly, positive agglomeration effects with regard to information processing might be at work as well. This notion is supported by the finding that also forecasters located in London are somewhat more accurate, even though the magnitude of the point estimate and its statistical significance are lower than for Frankfurt-based forecasters.

Table 4

Macro conditions also matter for forecasting performance. We find that inflation differentials affect accuracy, with forecasters in countries with high inflation relative to the euro area performing better, and those in countries with low relative inflation having a larger absolute forecast error. This could suggest the presence of forecasting biases, and we will return to this question below. More importantly, however, the fact that there is a significant impact of local economic conditions on forecasting accuracy at all suggests that forecasters have a national perspective, possibly because national information is more salient. While producing predictions on actions by the euro area's common central bank, forecasters continue to be informed by developments of their respective host countries. That is, they are tempted to take data at the national level as indicators of area-wide developments rather than aggregating properly weighted national data to form an opinion on developments in the area as a whole.

It is worthwhile noting that the magnitude of the absolute forecast errors linked to geographic and macro factors is non-negligible. For example, financial institutions with their headquarter or a subsidiary in Frankfurt show an *average* absolute forecast error that is lower by 1 to 1.5 basis points than their peers, compared to the 8 basis points difference overall between the best and worst forecasters reported in Figure 2. With respect to the macroeconomic variables, a one percentage point larger inflation rate differential leads to forecast errors that are smaller by a third of a basis point on average.

Are these results driven by the inclusion of a large number of forecasters located outside the euro area? The second set of columns of Table 4 reports estimates for the subsample of euro area forecasters, and shows that all results are robust.

The institutional controls contained in the model suggest that forecast accuracy differs across types of institutions. More experienced participants in the survey are somewhat more accurate, although this finding is only weakly significant. Furthermore, investment banks and commercial banks outperform the institutions classified as "Other". To test whether our findings are robust to the institutional composition of the sample, we drop all forecasters classified in the category "Other Institutions" from the sample. The third set of columns of Table 4 shows the corresponding results, which remain qualitatively the same.

4. Decomposition of the Forecast Errors

There is a remarkable (if not puzzling) degree of heterogeneity in the forecast errors. While we would expect geography and clustering around informational hubs not to matter a lot in tightly integrated and efficient financial markets, apparently they do. Moreover, even professional forecasters show some degree of confusion about the relative importance of national relative to euro area macroeconomic developments. Two questions arise from these findings. First, can the heterogeneity across forecasters be traced to differences in forecasting models, i.e. is it *systematic*? Second, is forecast accuracy affected by inflation differentials due to a bias, whereby forecasters in countries with above-average inflation tend to expect relatively higher interest rates, and vice versa?

4.1 Are the differences in forecast accuracy systematic?

Economic theory suggests that central banks might follow a monetary policy rule. This could be because they have optimally selected a particular rule, for instance one of the Taylor-variety, from a set of simple rules or because they follow an optimal monetary policy strategy that resembles such a rule. In both cases financial markets have reason to mimic the bank's supposed behaviour when formulating forecasting models for policy rates.

This argument can be used to filter out the systematic error from the observed (or overall) errors made by forecasting institutions. Above we defined the observed error e_{it} made by forecaster i at before meeting t as $e_{it} = r_{it}^e - r_t$, with r_t being the post-meeting policy rate and r_{it}^e i's published forecast. The question is how these expectations are formed?

A plausible assumption is that i's forecast of the policy rate will include a systematic component, based on economic theory. In addition, more often than not, the forecast might contain an unsystematic component representing the judgmental component, i.e. non-tangibles in the information set of the forecasting institution's analytical staff. Alternatively, this residual could reflect differences in the institutions' forecasts of relevant euro area variables. This would imply

$$r_{it}^e = \hat{r}_{it}^e + \hat{u}_{it},$$

with $\hat{r}_{it}^e$ representing the systematic and $\hat{u}_{it}$ the unsystematic component in i's forecast. While we do not directly observe i's forecasting model, a reasonable first attempt is perhaps to approximate its systematic component by a standard backward-looking Taylor-type rule.[7] Our model specification is

$$r_{it}^e = \alpha_i + \beta_{ir} r_{t-1} + \sum_{k=1}^{n} \beta_{ik} x_{kt} + \beta_{i\pi} \tilde{\pi}_{it} + \hat{u}_{it} \qquad (2)$$

where the β_i are forecaster-specific coefficients weighting the relevant variables, such as the previous ECB policy rate (r_{t-1}), relevant macro variables at the euro area level (x_{kt}), as well as the national inflation differential, defined as $\tilde{\pi}_{it} = \pi_{it} - \pi_{euro\ area,\, t}$. For the euro area macro variables x_{kt}, we obtain real time information from Bloomberg on year-on-year changes in consumer prices (HICP), industrial production, M3, and the level of consumer confidence as available at the time of the Reuters poll, i.e. on the Friday prior to meeting t.

[7] Sauer and Sturm (2007) show that using real time data (as we do in the empirical application), especially when including confidence indicators, can help maximizing the "fit" of Taylor-type rules for the ECB more than the inclusion of forward-looking terms.

This allows us to decompose the observed error in bank *i*'s expectations on the policy rate

$$\begin{aligned} e_{it} &= r^e_{it} - r_t \\ &= \hat{\alpha}_i + \hat{\beta}_{ir} r_{t-1} + \sum_{k=1}^{n} \hat{\beta}_{ik} x_{kt} + \hat{\beta}_{i\pi} \tilde{\pi}_{it} - r_t + \hat{u}_{it} \\ &= s_{it} + \hat{u}_{it} \end{aligned} \tag{3}$$

where s_{it} stands for the systematic and $\hat{u}_{it}$ for the unsystematic component in the observed (or overall) forecast error e_{it}. As before, we are mainly interested in the overall size of the forecast error—that is, we look at the absolute value. Moreover, in what follows, we will focus on the absolute systematic error of institution *i*, $|s_{it}|$, neglecting the unsystematic component in its forecasting performance.

Table 5 provides information about the systematic component of the observed error. A first finding is that the geographic country-pattern behind $|s_{it}|$ is fairly similar to the observed error discussed in Section 3. For instance, German institutions perform significantly better than the euro area average, while their Austrian and French counterparts significantly underperform. Moreover, the systematic forecasts, too, are generally unbiased and efficient.

Table 5

Following the analysis in the previous section, we proceed by explaining the systematic error component through our set of explanatory variables focusing on geography and macro conditions. The variables are as described above.

Table 6

The result emerging from Table 6 is that most of what has been said about the overall (or observed) absolute forecast error in Section 3 extends to the analysis of its Taylor-type rule based systematic component. In other words, differences in modelling the ECB's policy decisions underlie the observed differences in forecast accuracy. As with $|e_{it}|$, we find that $|s_{it}|$ is significantly lower for forecasters that report from *headquarters in Frankfurt* or a *subsidiary in Frankfurt*, or reside in London. At the same time, *inflation differentials* also matter significantly. All in all, the results are very close to those reported in Table 4 above, with the only exception of our proxy for experienced ECB watchers, which is no longer significant.

A similar picture emerges when we take a more general perspective on performance by again comparing top and bottom performers. In most cases, the best performers overall identified in Section 2.1. are at the same time also the best performers with respect to the systematic forecast error, and similarly, the worst performers overall generally show the worst performance when it comes to the systematic forecast. We conclude that it is indeed the systematic component that matters most in the overall forecast accuracy.

4.2 Do inflation differentials bias expectations?

So far, this paper has studied the determinants of *absolute* forecast errors, in order to assess the determinants of forecasting accuracy, and in particular the role of geography. However, the relevance of inflation differentials for accuracy raises questions about the underlying factors. In particular, it would be important to know whether national inflation *biases* the forecasts of ECB monetary policy decisions, in the sense that forecasters in countries with above-average inflation tend to expect relatively higher interest rates, and vice versa.

A first answer to this question can be provided through the estimated Taylor-type rules. As a matter of fact, there is evidence that national inflation differentials have a bearing on the interest rate forecasts along this dimension. Table 7 shows estimates for a pooled Taylor-type rule for the entire sample of forecasters. We report short-run coefficients, as these are more meaningful for forecasts of the imminent policy decisions than the long-run coefficients that are typically analyzed in the literature on central bank reaction functions. As discussed above, the model relates the expected rate to a number of explanatory variables.

Table 7

The results resemble standard estimated central bank reaction functions—that is, as a whole, forecasters follow a fairly conventional approach. On the nominal side, expectations for the ECB policy rate increase with *HICP inflation*, while the growth rate of *M3* is not considered relevant. On the real side, changes in *industrial production* and the level of *consumer confidence* are positively related to the expected rate. At the same time, the high estimated coefficients for *previous rate*, i.e., the policy rate prevailing before the meeting for which expectations are formed, suggest a high degree of interest rate smoothing and as such a rather protracted reaction to changes in the other model variables. Most importantly, however, there is a role for inflation differentials (on top of euro area HICP inflation): with increasing differentials, forecasters tend to expect significantly higher interest rates.

Another way to test whether inflation differentials generate a bias in expectations is to test whether the forecast error e_{it} (as opposed to the absolute forecast error) is a function of inflation differentials. As we do not have any prior as to how the other determinants of forecast accuracy of model (1) could bias expectations, we estimate a model with a full set of meeting-fixed effects as well as forecaster-fixed effects:

$$e_{it} = \alpha + \beta_{\pi}\widetilde{\pi}_{it} + \gamma_t + \delta_i + u_{it} \qquad (4)$$

Table 8 presents the corresponding results, for the forecast error e_{it} as dependent variable in the first row, and for the systematic forecast error s_{it} in the second row. As expected, relatively high inflation leads to larger forecast errors, as forecasters in countries with above-average inflation tend to expect relatively higher interest rates, and vice versa.[8]

Table 8

[8] This result and the impact of inflation differentials on the *absolute* forecast error in model (1) can be reconciled in different ways. Model (4) contains meeting and forecaster fixed effects, which control for the average forecast error per meeting and per forecaster. Model (1) controls for meeting and forecaster fixed effects in a different way, which can generate the negative effect of inflation differentials on accuracy. For instance, a forecasting institution stricken with a dovish bias in its interest rate forecast (captured by the δ_i in (4)) might experience a gain in forecasting accuracy when reacting to higher-than-average local inflation with an (otherwise not advisable) increase in its interest rate prediction.

5. Conclusions

EMU has implied the assignment of monetary policy making for 12 countries with varying histories of inflation, policy strategies, and economic environments to the ECB. Monetary policy is now conducted taking a euro area-wide perspective, but it operates in a multi-country, multi-cultural and multi-lingual context. This raises a number of issues: How does this heterogeneity of conditions and backgrounds in member countries affect the ability of economic agents to understand and anticipate monetary policy by the ECB? Is there convergence in the views how the ECB conducts monetary policy? Or are the differences in the ability to anticipate the ECB's decisions indicative of more permanent information asymmetries related to geographic proximity and country-specific factors?

Using a novel database on the forecasts of ECB policy decisions of 120 financial institutions in 24 countries since 1999, we find some marked differences in their ability to understand and anticipate policy decisions by the ECB. The paper shows that these differences have repercussions on financial markets, as heterogeneity in market expectations affects the volatility of asset prices in response to the announcement of monetary policy decisions.

Dissecting the differences in forecast ability, the paper finds that a substantial part is systematically explained by geography and country-specific economic conditions. We find that financial institutions that are based in Frankfurt, or have a subsidiary in Frankfurt, perform substantially better in predicting ECB policy decisions. A similar informational advantage also appears to be at play for institutions based in the City of London. This suggests that information asymmetries and agglomeration effects play a role in the ability of forecast institutions to anticipate monetary policy in the euro area. This finding is in line with earlier literature that analyses the reasons for the existence of international financial centres, and which argues that face-to-face contacts among financial market actors are facilitated by proximity, and are an essential factor in knowledge production and thus performance.

Country-specific economic conditions are also relevant, with the accuracy of forecasts depending on the levels of inflation in the institutions' host countries. Forecasters in countries with higher (lower) than euro area inflation tend to produce more hawkish (dovish) forecasts, suggesting that the national inflation environment biases forecasters' views of ECB monetary policy. Finally, we find that most of the heterogeneity in forecasting performance can be related to different models of ECB behaviour.

Our results have important policy implications. Expectations are a crucial factor in the transmission of monetary policy. And a central bank operating in a heterogeneous environment such as the ECB needs to be aware of differences in the ability of economic agents to understand and anticipate monetary policy—differences that appear to be significant in the case of the euro area. Furthermore, a systematically different perception of the ECB's monetary policy could prompt country-specific political controversy about the ECB's role.

The paper suggests that euro area financial markets have yet to converge on a homogeneous view of the ECB, to overcome locational and national biases, and to adopt a common expectation-formation process. Although some informational frictions and asymmetries or agglomeration effects may be a permanent feature of financial activity in any region, there seems to be scope for continuous guidance of this convergence process by a careful and targeted communication policy of the central bank.

There are multiple avenues for further research. A first extension would be to explore likely links between ECB communication and expectations on ECB policy. For instance, it could be asked whether communication in the form of speeches or interviews, perhaps targeted at particular audiences within EMU, is helpful in reducing systematic heterogeneity of expectations. A related question is whether certain forms of ECB communication are more likely to focus the attention of regional audiences than others. A second area for future research is to broaden the focus in two ways. First, it would be interesting to extend the forecast horizon; second, understanding the extent to which heterogeneity in expectations and perceptions of monetary policy is present not only among financial market participants, but also among the general public is of crucial importance for central banks.

References

Angeloni, I. and M. Ehrmann (2003), Monetary Transmission in the Euro Area: Early Evidence. *Economic Policy* 37, 469-492.

Bae, K.-H., R. Stulz and H. Tan (2009). Do Local Analysts Know More? A Cross-Country Study of the Performance of Local Analysts and Foreign Analysts. Forthcoming, *Journal of Financial Economics*.

Berger, H., M. Ehrmann and M. Fratzscher (2006), Forecasting ECB Monetary Policy, Accuracy Is (Still) a Matter of Geography, ECB Working Paper No. 578.

Cecchetti, S., (2001). Legal structure, financial structure and the monetary policy transmission mechanism, in: Deutsche Bundesbank (ed.), *The Monetary Transmission Process: Recent developments and lessons for Europe*, Palgrave, 170-194

Clark, G. (2002). London in the European Financial Services Industry: Locational Advantage and Product Complementarities, *Journal of Economic Geography*, 2(4), 433-53.

Cook, G., N. Pandit, J. Beaverstock, P. Taylor and K. Pain (2007). The Role of Location in Knowledge Creation and Diffusion: Evidence of Centripetal and Centrifugal Forces in the City of London Financial Services Agglomeration. *Environment and Planning* 39(6): 1325-1345.

Coval, J. D. and Moskowitz, T. J. (1999). Home bias at home: local equity preference in domestic portfolios. *Journal of Finance* 54 (6): 2045 – 2073.

Coval, J. D. and Moskowitz, T. J. (2001). The geography of investment: informed trading and asset pricing. *Journal of Political Economy* 109(4): 811 – 841.

Dornbusch, R., Favero, C.A. and F. Giavazzi (1998). The immediate challenges for the European Central Bank, *Economic Policy* 26, 17-64

Dvorak, T., (2005). Do Domestic Investors Have an Information Advantage? Evidence from Indonesia. *Journal of Finance* 60(2): 817 – 839.

Ehrmann, M. and M. Fratzscher (2005). Equal Size, Equal Role? Interest Rate Interdependence Between the Euro Area and the United States. *Economic Journal* 115: 930-950.

Faulconbridge, J. (2004). London and Frankfurt in Europe's Evolving Financial Centre Network, *Area*, 36(3): 235-44.

Froot, K. A., and J. Stein, (1991). Exchange Rates and Foreign Direct Investment: An Imperfect Capital Markets Approach. *Quarterly Journal of Economics* 106: 1191-1217.

Gordon, R.H. and A.L. Bovenberg, (1996). Why Is Capital So Immobile Internationally? Possible Explanations and Implications for Capital Income Taxation. *American Economic Review* 86(5): 1057 – 1075.

Grote, M. (2004), The Rise (and Fall?) of the Financial Centre Frankfurt—Locations of Foreign Banks in Germany 1949-2003, Typescript, University of Frankfurt.

Hau, H. (2001), Location Matters: An Examination of Trading Profits, *Journal of Finance*, 56(5): 1959-83.

Krugman, P. (1998), What's New About the New Economic Geography? *Oxford Review of Economic Policy*, 14(2), 7-17.

Malloy, C. (2005), The Geography of Equity Analysis, *Journal of Finance*, 60(2): 719-55.

Mihov, I., (2001), Monetary policy implementation and the transmission in the European Monetary Union, *Economic Policy*, 33, 369-402

Mojon B. and G. Peersman (2003), 'A VAR Description of the Effects of Monetary Policy in the Individual Countries of the Euro Area', in: Angeloni, I., Kashyap, A. and Mojon, B. (eds.), *Monetary Policy Transmission in the Euro Area*, Cambridge University Press, 56-74.

Portes, R., H. Rey and Y. Oh, (2001). Information and Capital Flows: The Determinants of Transactions in Financial Assets, *European Economic Review* 45: 783–96.

Sauer, S. and J.-E. Sturm (2007), Using Taylor Rules to Understand European Central Bank Monetary Policy. *German Economic Review* 8(3): 375-398.

Strauss-Kahn, V. and X. Vives (2005). Why and Where Do Headquarters Move? CEPR Discussion Paper, 5070, May.

Thrift, N. (1994), On the Social and Cultural Determinants of International Financial Centres: The Case of the City of London, in: S. Corbridge, N. Thrift, and R. Martin (eds.), *Money, Power and Space*, Oxford: Blackwell, 327-355.

Tschoegl, A. (2000), International Banking Centers, Geography, and Foreign Banks, *Financial Markets, Institutions and Instruments*, 9(1), 1-32.

Welch, I. (2000). Herding Among Security Analysts. *Journal of Financial Economics* 58(3), 369-396.

Figure 1: Survey-based versus market-based forecast errors of ECB policy decisions

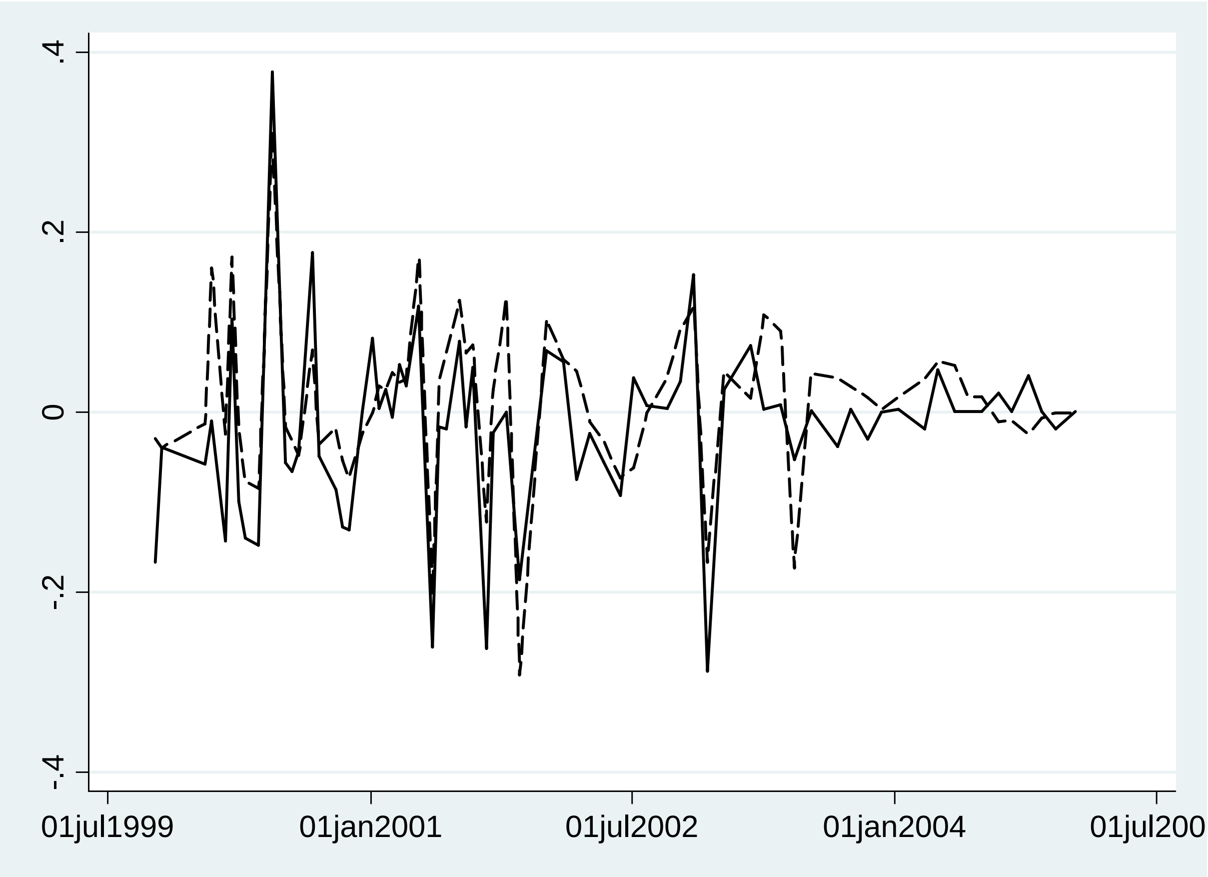

Note: The figure shows the mean forecast errors of the Reuters survey (dashed line) and the EONIA-based measure (solid line), in percent.

Figure 2: Absolute average forecast errors, by decile

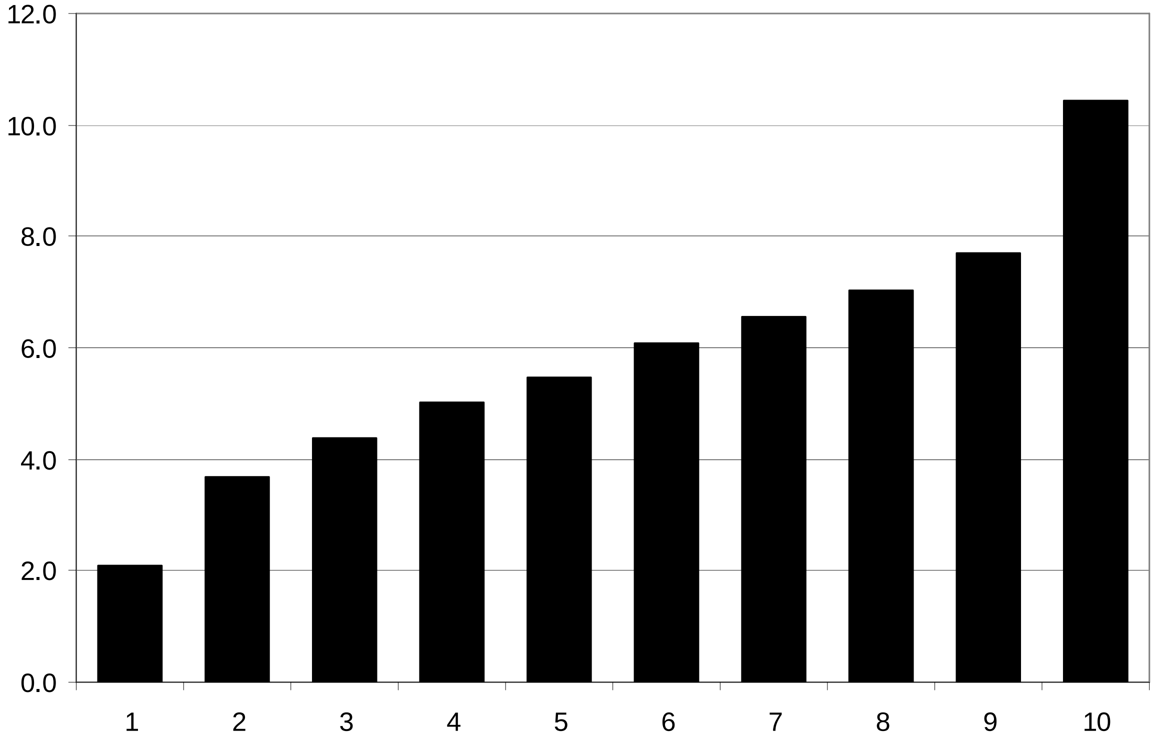

Note: The figure shows the average absolute forecast error (in basis points) of financial institutions, ranging from those 10% of institutions with the lowest forecast errors in decile 1 to those 10% with the highest error in decile 10.

Figure 3: Actual and expected ECB policy rates

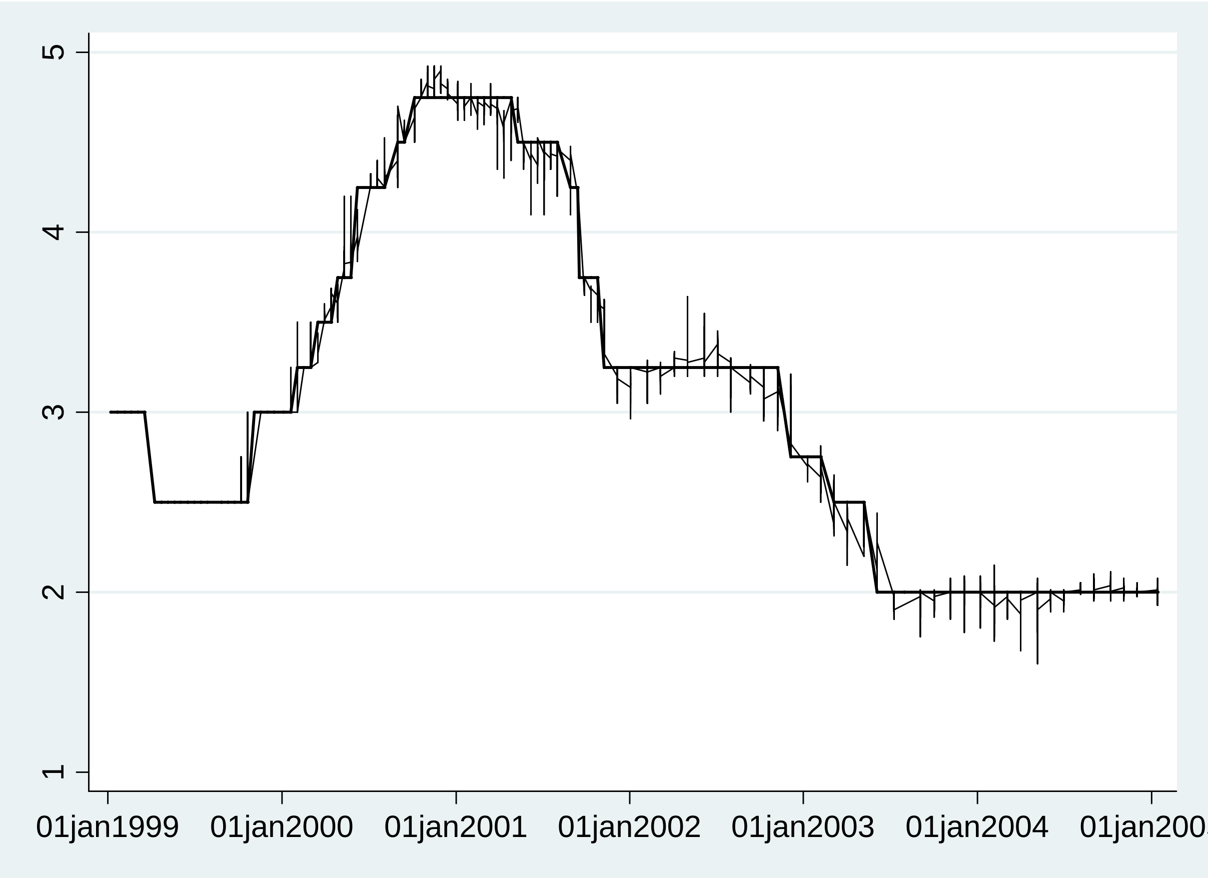

Note: For each Governing Council meeting, the vertical lines show the minimum and maximum of the expected policy rates in the Reuters survey. The connecting lines between these vertical lines indicate the mean expectation for each Governing Council meeting. The thick line depicts the actual policy rate as decided in the respective meeting.

Table 1: An example of the Reuters raw data (January 2002)

Contributor	City/Country	50P rise	25P Rise	No change	25P Cut	50P Cut	Next change when?	By how much?	Rates at end 2002?	Rates at end 2003?	Low – level?	Low when?
Invesco Asset Mgt	Frankfurt	0	0	40	25	35	Feb	-50	2.00	3.00	2.00	Sum 02
Standard Chartered	London	0	0	40	40	20	Feb	-50	2.00	3.00	2.00	-
American Express	London	0	0	60	30	10	Feb	-25	2.00	2.50	2.00	Q3-02
Exane	Paris	0	0	70	20	10	Feb	-25	2.50	4.00	2.00	Q2-02
Natexis Banques Populaires	Paris	0	0	10	65	25	.	.	2.50	3.00	2.50	Mar-02
….	….	….	….	….	….	….	….	….	….	…	…	….
HSBC	London	0	0	60	25	15	Feb	-50	2.75	.	2.50	Apr/May
Dresdner Bank	Frankfurt	0	0	60	20	20	March	-50	3.50	.	2.75	Mar-02
Commerzbank	Frankfurt	0	0	65	30	5	Feb/Mar	-25	3.00	4.25	2.75	Q2-02
RBS Financial Markets	London	0	0	67	25	8	Feb	-50	3.25	4.25	2.75	Feb
Sal Oppenheim	Cologne	0	0	70	20	10	Feb	-25	3.50	4.00	2.75	Mid-02
Halifax	London	0	0	70	25	5	.	-25	2.75	4.00	2.75	May-02
Stone & McCarthy	London	0	0	70	25	5	Feb	-25	3.50	4.25	2.75	End Q1- 02
CitigroupSSSB	London	0	0	75	20	5	Feb	-25	3.25	4.00	2.75	Q2-02
Nordea Group	Helsinki	0	0	80	15	5	Feb	-50	3.25	.	2.75	Q1/02
Alpha Bank	Athens	0	0	80	20	0	Feb	-25	3.25	3.75	2.75	Q1-02
Rabobank	Utrect	0	0	85	10	5	Feb	-25	3.25	3.50	2.75	Apr
Hypovereinsbank	Munich	0	1	85	13	1	Feb	-50	3.25	.	2.75	.
Standard & Poors MMS	London	0	0	90	9	1	Feb	-25	2.75	4.00	2.75	Q2-02
Banco Santander	Lisbon	0	0	50	50	0	Feb	-25	3.25	4.00	3.00	Q1-02
Fortis Bank	Amsterdam	0	5	60	25	10	Feb	-25	3.00	3.50	3.00	Q1/Q2
West LB	Duesseldorf	0	0	70	25	5	Feb	-25	3.25	4.00	3.00	Q1
Sanwa	London	0	0	75	25	0	Feb	-25	3.25	4.00	3.00	Q1-02
AM Generali Finanz	Cologne	0	0	90	10	0	Feb	-25	3.75	4.50	3.00	Feb-02
Bankgesellschaft	Berlin	0	0	90	10	0	Feb	-25	4.00	.	3.00	Feb-02
Investors Bank & Trust	Boston	0	1	98	1	0	March	-25	3.50	4.25	3.00	March
Nomura International	Frankfurt	0	0	100	0	0	March	-25	3.00	4.00	3.00	March

Note: The table provides an example of the raw data obtained from the Reuters polls, for a shortened sample of the survey conducted prior to the Governing Council meeting in January 2002. Entries in the third to seventh column refer to probabilities (in percent).

Table 2: Absolute forecast errors by country (in %)

	# obs	mean	diff.	st. dev.	meetings with change	meetings with no change	unbiased-ness test p-value	efficiency test p-value
Austria	90	0.087	***	0.098	0.125	0.083	0.000	0.000
Belgium	8	0.100		0.139	0.188	0.048	0.981	--
Denmark	147	0.053		0.081	0.211	0.030	0.156	0.895
Finland	98	0.060		0.081	0.188	0.039	0.612	0.508
France	314	0.070	**	0.083	0.169	0.058	0.000	0.000
United Kingdom	1,426	0.062		0.075	0.177	0.046	0.000	0.000
Germany	1,290	0.053	***	0.075	0.182	0.035	0.000	0.000
Ireland	72	0.073	**	0.060	0.148	0.061	0.170	0.280
Italy	150	0.046		0.085	0.195	0.026	0.050	0.136
Netherlands	119	0.066		0.077	0.177	0.053	0.220	0.551
Portugal	118	0.042	***	0.071	0.175	0.025	0.023	0.000
Sweden	71	0.063		0.060	0.156	0.054	0.520	0.941
USA	96	0.063		0.089	0.216	0.041	0.046	0.296
Others	243	0.069	**	0.073	0.178	0.054	0.029	0.675
Euro area	2,350	0.057		0.078	0.179	0.041	0.000	0.000
Non-euro area	1,892	0.062	*	0.076	0.180	0.046	0.000	0.000
All countries	4,242	0.060		0.077	0.179	0.043	0.000	0.000

Notes: -- indicates missing values; "diff." shows significance level of a t-test of whether the country-specific mean is different from the mean for all countries. ***, **, and * indicate significance at the 99%, 95%, and 90% levels, respectively. The unbiasedness test is based on the following equation: $r_t = \alpha + \beta r_t^e + \varepsilon_t$, where r_t is the policy rate after the Council meeting t and r_t^e are the expectations, with the Wald test of the joint hypothesis that H_0: $\alpha \neq 0$ and $\beta \neq 1$. Under the efficiency test, expectations are efficient if forecast errors cannot be predicted systematically on the basis of past policy decisions r_{t-p} : $r_t - r_t^e = \zeta + \sum_{p=1}^{P} \psi_p r_{t-p} + \varepsilon_t$, with the lag length chosen as P=6. The efficiency hypothesis to be tested is $\psi_1 = \psi_2 = ... = \psi_P \neq 0$. Note that both tests are based on country-specific, rather than institution-specific estimations. P-values of above 0.90 indicate that biasedness or inefficiency are accepted at the 90% level.For a more detailed discussion, see Ehrmann and Fratzscher (2005).

Table 3: Explaining the volatility in €/$ exchange rates following the announcement of ECB monetary policy decisions

	10 minute window					
	σ_t		σ_t/σ_{t-1}		σ_t/σ_{t-5}	
	coef.	*std. err.*	coef.	*std. err.*	coef.	*std. err*
absolute surprise	0.000	*0.000*	1.777	*2.590*	4.004	*3.473*
heterogeneity	0.002 ***	*0.001*	15.076 ***	*5.035*	10.156 **	*4.126*
# observations	66		65		65	
R^2	0.205		0.221		0.152	

	20 minute window					
	σ_t		σ_t/σ_{t-1}		σ_t/σ_{t-5}	
	coef.	*std. err.*	coef.	*std. err.*	coef.	*std. err.*
absolute surprise	0.000	*0.001*	4.886	*4.440*	1.918	*4.160*
heterogeneity	0.003 ***	*0.001*	15.741 ***	*3.716*	10.044 ***	*3.682*
# observations	66		65		65	
R^2	0.157		0.346		0.134	

Notes: The table shows the results of a regression that explains the volatility of the spot €/$ exchange rate following the announcement of the ECB's interest rate decision at 13:45 CET. Volatility is measured by the standard deviation of the recorded exchange rates, starting with the first tick after the release, over a window of 10 minutes (left panel) or 20 minutes (right panel). The regression model is $\sigma_t = \alpha + \beta\,\bar{s}_t + \gamma\,\psi_t + \varepsilon_t$, and explains the volatilities for each meeting day t by the absolute size of the surprise ($\bar{s}_t$, defined by the median response in the Reuters poll), and the heterogeneity in expectations (ψ_t, given by the standard deviation of the responses in the Reuters poll). σ_t/σ_{t-1} and σ_t/σ_{t-5} denote volatility relative to volatility over the same time window 1 and 5 business days before, respectively. The sample period starts with the Governing Council meeting on September 14, 2000. ***, **, and * indicate significance at the 99%, 95%, and 90% levels, respectively.

Table 4: Explaining absolute forecast errors: geography and macro conditions

	All forecasters		Euro area forecasters		Banks	
	coef.	*std. err.*	coef.	*std. err.*	coef.	*std. err.*
Distance:						
Headquarters in Frankfurt	-0.010 ***	*0.002*	-0.010 ***	*0.002*	-0.011 ***	*0.002*
Subsidiary in Frankfurt	-0.015 ***	*0.003*	-0.015 ***	*0.003*	-0.011 ***	*0.003*
HQ or Subsidiary in London	-0.004 *	*0.002*	-0.007 **	*0.003*	0.000	*0.002*
Macro conditions:						
Inflation Differential	-0.330 **	*0.132*	-0.407 ***	*0.145*	-0.357 **	*0.150*
Institution controls:						
Investment Bank	-0.006 **	*0.003*	-0.019 ***	*0.005*		
Commercial Bank	-0.010 ***	*0.003*	-0.020 ***	*0.005*	0.003	0.002
Experience in ECB Watching	-0.002 *	*0.001*	-0.001	*0.002*	-0.003 *	*0.001*
# observations	4242		2350		3624	

Notes: The table shows results from regression model (1), $|e_{it}| = \alpha + \boldsymbol{\beta}\mathbf{y}_{it} + \gamma_t + u_{it}$, which explains individual absolute forecast errors by location, macro conditions and institution controls, as well as a full set of Governing Council meeting dummies. Standard errors allow for heteroskedasticity. ***, **, and * indicate significance at the 99%, 95%, and 90% levels, respectively.

Table 5: Absolute *systematic* forecast errors by country (in %)

	# obs	mean	diff.	st. dev.	meetings with change	meetings with no change
Austria	90	0.107	***	0.093	0.216	0.095
Belgium	8	0.108		0.137	0.194	0.056
Denmark	147	0.066		0.090	0.270	0.037
Finland	98	0.076		0.093	0.236	0.050
France	314	0.085	***	0.082	0.239	0.067
United Kingdom	1,426	0.073		0.084	0.237	0.050
Germany	1,290	0.069	*	0.087	0.257	0.043
Ireland	72	0.072		0.066	0.159	0.058
Italy	150	0.066		0.085	0.246	0.042
Netherlands	119	0.073		0.089	0.275	0.050
Portugal	118	0.069		0.086	0.266	0.045
Sweden	71	0.065		0.074	0.219	0.051
USA	96	0.061		0.090	0.216	0.038
Others	243	0.078		0.085	0.240	0.055
Euro area	2,350	0.073		0.087	0.250	0.050
Non-euro area	1,892	0.072		0.084	0.237	0.049
All countries	4,242	0.073		0.086	0.244	0.050

Notes: The table shows summary statistics for $|s_{it}|$, the absolute systematic forecast error, averaged across institutions in a given country. s_{it} is defined as in equation (3), i.e. as the difference between forecasters' systematic forecast, estimated as $\hat{r}_{it}^{e} = \hat{\alpha}_i + \hat{\beta}_{ir} r_{t-1} + \sum_{k=1}^{n} \hat{\beta}_{ik} x_{kt} + \hat{\beta}_{i\pi} \tilde{\pi}_{it}$, and the actual policy rate. "--" indicates missing values; "diff." shows significance level of a t-test of whether the country-specific mean is different from the mean for all countries. ***, **, and * indicate significance at the 99%, 95%, and 90% levels, respectively.

Table 6: Explaining absolute systematic forecast errors: geography and macro conditions

	All forecasters		Euro area forecasters		Banks	
	coef.	*std. err.*	coef.	*std. err.*	coef.	*std. err.*
Distance:						
Headquarters in Frankfurt	-0.012 ***	*0.002*	-0.013 ***	*0.002*	-0.011 ***	*0.002*
Subsidiary in Frankfurt	-0.012 ***	*0.002*	-0.014 ***	*0.002*	-0.009 ***	*0.002*
HQ or Subsidiary in London	-0.004 ***	*0.001*	-0.008 ***	*0.002*	0.000	*0.001*
Macro conditions:						
Inflation Differential	-0.418 ***	*0.105*	-0.536 ***	*0.113*	-0.501 ***	*0.109*
Institution controls:						
Investment Bank	-0.004 **	*0.002*	-0.014 ***	*0.004*		
Commercial Bank	-0.006 ***	*0.002*	-0.015 ***	*0.004*	0.000	0.001
Experience in ECB Watching	0.000	*0.001*	0.000	*0.001*	0.000	*0.001*
# observations	4242		2350		3624	

Notes: The table shows results from the regression model $|s_{it}| = \alpha + \boldsymbol{\beta}\mathbf{y}_{it} + \gamma_t + u_{it}$, which explains individual absolute systematic forecast errors by location, macro conditions and institution controls, as well as a full set of Governing Council meeting dummies. Standard errors allow for heteroskedasticity. ***, **, and * indicate significance at the 99%, 95%, and 90% levels, respectively.

Table 7: Estimate for Taylor-type rules

	All forecasters			Euro area forecasters			Banks		
	coef.		*std. err.*	coef.		*std. err.*	coef.		*std. err.*
Previous Rate	0.915	***	*0.015*	0.917	***	*0.016*	0.918	***	*0.015*
M3	0.846		*0.822*	0.738		*0.835*	0.831		*0.807*
Industrial Production	0.671	**	*0.337*	0.698	**	*0.344*	0.723	**	*0.334*
Consumer Confidence	0.015	***	*0.002*	0.014	***	*0.002*	0.014	***	*0.002*
HICP Inflation	3.358	*	*1.916*	3.306	*	*1.969*	3.214	*	*1.890*
Relative Inflation	0.445	**	*0.210*	0.489	**	*0.207*	0.380	*	*0.228*

Notes: The table shows results from the regression model $r_{it}^{e} = \alpha_i + \beta_{ir} r_{t-1} + \sum_{k=1}^{n} \beta_{ik} x_{kt} + \beta_{i\pi} \tilde{\pi}_{it} + \hat{u}_{it}$, which explain individual expectations about ECB's policy rates with Tyalor-type rule parameters. Results from a pooled regression using all available forecasts. Standard errors allow for heteroskedasticity and clustering across panels. ***, **, * indicate significance at the 99%, 95%, and 90% level, respectively.

Table 8: Explaining forecast errors with inflation differentials

	All forecasters			Euro area forecasters			Banks		
	coef.		*std. err.*	coef.		*std. err.*	coef.		*std. err.*
Forecast error	0.515	**	*0.227*	0.536	**	*0.248*	0.431	*	*0.236*
Systematic error	0.267	*	*0.158*	0.463	***	*0.170*	0.083		*0.154*
# observations	4242			2350			3624		

Notes: The table shows results from regression model (4), $e_{it} = \alpha + \beta_{\pi}\tilde{\pi}_{it} + \gamma_t + \delta_i + u_{it}$, which explains individual forecast errors (and individual systematic forecast errors in the second row) with the national inflation differential and full sets of Governing Council meeting dummies and forecaster dummies. Standard errors allow for heteroskedasticity. ***, **, and * indicate significance at the 99%, 95%, and 90% levels, respectively.