

**How informative are macroeconomic
risk forecasts?
An examination of the Bank of England's
inflation forecasts**

Malte Knüppel
Guido Schulte Frankenfeld



Discussion Paper
Series 1: Economic Studies
No 14/2008

Discussion Papers represent the authors' personal opinions and do not necessarily reflect the views of the Deutsche Bundesbank or its staff.

Editorial Board:

Heinz Herrmann
Thilo Liebig
Karl-Heinz Tödter

Deutsche Bundesbank, Wilhelm-Epstein-Strasse 14, 60431 Frankfurt am Main,
Postfach 10 06 02, 60006 Frankfurt am Main

Tel +49 69 9566-1

Telex within Germany 41227, telex from abroad 414431

Please address all orders in writing to: Deutsche Bundesbank,
Press and Public Relations Division, at the above address or via fax +49 69 9566-3077

Internet <http://www.bundesbank.de>

Reproduction permitted only if source is stated.

ISBN 978-3-86558-437-3 (Printversion)

ISBN 978-3-86558-438-0 (Internetversion)

Abstract:

Macroeconomic risk assessments play an important role in the forecasts of many institutions. However, to the best of our knowledge their performance has not been investigated yet. In this work, we study the Bank of England's risk forecasts for inflation. We find that these forecasts do not contain the intended information. Rather, they either have no information content, or even an adverse information content. Our results imply that under mean squared error loss, it is better to use the Bank of England's mode forecasts than the Bank of England's mean forecasts.

Keywords: Forecast evaluation; risk forecasts; Bank of England inflation forecasts

JEL-Classification: E37, C12, C53

Non-technical summary

Many institutions providing macroeconomic forecasts add risk assessments to these forecasts. Not all of them are explicit about the definition of risk. However, those institutions giving a precise definition as e.g. the International Monetary Fund or the Bank of England state that an upward risk implies that the expected value of the forecasted variable lies above the value published as the central projection. Accordingly, a downward risk implies that the expected value of the forecasted variable lies below the value published as the central projection. The value published as the central projection by these institutions is the most likely single value, also called the mode forecast. The mean forecast, i.e. the expected value of the forecasted variable, thus results from the mode forecast and the forecasted risk.

Although many institutions publish risk forecasts, to the best of our knowledge the performance of these forecasts has not been studied yet. In this work, we try to close this gap by investigating the Bank of England's risk forecasts for inflation. Our findings indicate that these forecasts do not perform well.

If the risk forecasts are optimal they should improve the central projections. Moreover, they should not improve the mean forecasts. We find that the Bank of England's risk forecasts do not possess these properties. Apparently, it is possible to improve the mean forecasts by considering information contained in the risk forecasts, and it is not possible to improve the central projections by considering the risk forecasts. In the sample under study, for several forecast horizons it is even possible to improve the central projections by considering the opposite of the risk forecast. This finding is related to the fact that often forecasted upward risks are followed by the realization of downward risks or forecasted downward risks by the realization of upward risks.

A major reason for the poor performance of the risk forecasts is probably given by the difficulties to identify and to quantify risks in the determinants of future inflation, as e.g. oil prices and exchange rates.

Our results imply that the Bank of England's central projection for inflation, the mode forecast, is closer to expected inflation than the forecast for expected inflation, the mean forecast.

Nicht-technische Zusammenfassung

Viele Institutionen, die makroökonomische Prognosen erstellen, versehen ihre Prognosen mit Risikoeinschätzungen. Nicht alle geben explizit an, wie diese Risiken definiert sind. Aber jene Institutionen, die eine präzise Definition anführen, wie zum Beispiel der Internationale Währungsfonds und die Bank of England, erklären, dass das Vorliegen eines Aufwärtsrisikos bedeutet, dass der erwartete Wert der prognostizierten Variable über jenem Wert liegt, der als Basislinie prognostiziert wird. Entsprechend liegt ein Abwärtsrisiko vor, wenn der erwartete Wert der prognostizierten Variablen unter jenem Wert liegt, der als Basislinie prognostiziert wird. Der von diesen Institutionen als Basislinie veröffentlichte Wert ist der wahrscheinlichste Wert und wird auch als Modusprognose bezeichnet. Die Mittelwertprognose, also der erwartete Wert der prognostizierten Variablen, ergibt sich somit aus der Modusprognose und dem prognostizierten Risiko.

Obwohl viele Institutionen Risikoprognosen veröffentlichen, ist die Güte dieser Prognosen nach unserem Wissen noch nicht untersucht worden. In dieser Arbeit versuchen wir, diese Lücke zu schließen, indem wir die Risikoprognosen für Inflation der Bank of England auswerten. Unsere Ergebnisse deuten darauf hin, dass diese Prognosen keine guten Ergebnisse liefern.

Wenn die Risikoprognosen optimal sind, so sollten sie die Basislinie verbessern können. Andererseits sollten sie nicht in der Lage sein, die Mittelwertprognosen zu verbessern. Es stellt sich jedoch heraus, dass die Risikoprognosen für Inflation der Bank of England nicht diese genannten Eigenschaften besitzen. Es ist offenbar möglich die Mittelwertprognosen zu verbessern, indem man in den Risikoprognosen enthaltene Informationen berücksichtigt. Außerdem ist es nicht möglich, die Basislinie zu verbessern, in dem man die Risikoprognosen berücksichtigt. In der untersuchten Stichprobe kann bei einigen Prognosehorizonten sogar die Basislinie verbessert werden, indem man entgegengesetzte Risikoprognosen berücksichtigt. Diese Beobachtung hängt damit zusammen, dass häufig auf prognostizierte Aufwärtsrisiken die Realisation von Abwärtsrisiken oder auf prognostizierte Abwärtsrisiken die Realisation von Aufwärtsrisiken folgte.

Ein Hauptgrund für die schlechten Ergebnisse der Risikoprognosen dürfte sein, dass es sehr schwierig ist, Risiken für die Bestimmungsfaktoren der Inflation, wie

zum Beispiel Ölpreise und Wechselkurse, zu identifizieren und zu quantifizieren. Aus unseren Ergebnissen lässt sich schließen, dass die Basislinie der Bank of England, die Modusprognose, der erwarteten Inflation näherkommt als die Prognose der erwarteten Inflation, die Mittelwertprognose.

Contents

1	Introduction	1
2	The Data	3
3	Forecast Optimality	5
3.1	Mean Forecast Optimality and Risk	5
3.2	Bias of Mode, Mean and Risk Forecasts	8
3.3	Risk Forecast Optimality	9
3.4	Interpretation of the Test Results	14
4	Robustness Checks	16
4.1	Risk Forecast Optimality - A Qualitative Analysis	16
4.2	Stability Over Time	19
5	Possible Explanations for Poor Risk Forecasting Performance	22
5.1	Endogeneity of Outturns	22
5.2	Problems to Anticipate Risks in Determinants	24
6	Conclusion	25
A	Appendix: Mode and Mean Forecasts and Realizations	29

How Informative Are Macroeconomic Risk Forecasts?

An Examination of the Bank of England's Inflation Forecasts¹

1 Introduction

Many institutions providing macroeconomic forecasts add risk assessments to these forecasts. For example, the International Monetary Fund states in its World Economic Outlook from October 2007 that, considering the global GDP growth forecast, “the risks to the baseline forecast are distinctly to the downside. [...] The main sources of the increase in the downside risk since the July 2007 update come from deteriorating financial conditions and from the uncertain prospects for domestic demand in the United States and Europe” (p. 8). In its Inflation Report from February 2007, the Bank of England remarks that “The risks to inflation are weighted to the downside in the near term and to the upside in the medium term.” (pp. 45-46). Similar statements can be found in the publications of several other institutions, among them, for example, the European Central Bank, the Banco de Portugal and the Deutsche Bundesbank.² Thus, announcing forecast risks appears to play an important role in the communication of forecast results.

¹Authors: Malte Knüppel and Guido Schulte Frankenfeld, Deutsche Bundesbank, Research Centre, Wilhelm-Epstein-Straße 14, D-60431 Frankfurt am Main, Germany. Corresponding author's e-mail: malte.knueppel@bundesbank.de. This work represents the authors' personal opinion and does not necessarily reflect the views of the Deutsche Bundesbank or its staff. We would like to thank Jörg Breitung, Heinz Herrmann and Karl-Heinz Tödter as well as seminar participants at the Bank of England and the Deutsche Bundesbank for very helpful comments and discussions.

²For instance, the European Central Bank in its Monthly Bulletin from December 2007 declares that with respect to the inflation forecast, “Risks to this outlook are fully confirmed to lie on the upside. These risks include the possibility of further rises in oil and agricultural prices, as well as of unanticipated increases in administered prices and indirect taxes” (p. 55). The

The presence of forecast risks is caused by asymmetries of the respective forecast densities. Some institutions, for instance, speak of an upward risk to the forecast if the forecasted mean exceeds the forecasted mode of their density forecasts. Assessing the asymmetry of a forecast density, i.e. forecasting a phenomenon related to third moments is surely an extremely challenging task.

This might be illustrated by the fact that many institutions calculate forecast uncertainty, i.e. a phenomenon related to second moments and therefore in general easier to assess, based on past forecast errors. This is done due to the lack of models which can accomplish this task, as explained by Wallis (1989). However, if it is so difficult to forecast the uncertainty surrounding an institution's forecast appropriately, it is questionable whether risks can be forecasted in a reasonable manner. Given that so many institutions face the challenge of risk forecasting despite the difficulties to be encountered, it is important and interesting to find out how successful these risk forecasts are.

Of all the institutions mentioned, the Bank of England (henceforth BoE) features the largest published risk forecasting record. Moreover, in contrast to most other institutions, the BoE produces quarterly, and not only annual risk forecasts. Finally, it does not only publish qualitative, but also quantitative risk assessments. Actually, the BoE publishes density forecasts, from which point forecasts, uncertainty forecasts and risk forecasts can be derived. Therefore, our analysis focuses on the BoE's forecasts.

Since the BoE's record of macroeconomic density forecasts is relatively large, its forecasts are a highly favoured object of investigation in economics. The BoE's forecasts are studied inter alia by Dowd (2007), Wallis (2003), Wallis (2004), and Elder et al. (2005). Most studies are concerned with the accurateness of the point and uncertainty forecasts, but also the appropriateness of the entire forecast density is evaluated. Up to now, however, apparently no study has focused on the risk forecasts. So far, these forecasts have at best been evaluated in the context

Banco de Portugal claims that "As regards the projection for the inflation rate, risks appear to be broadly balanced in 2007 and slightly biased downwards in 2008, due to the risk of appreciation of the euro exchange rate" (p. 31) in its Economic Bulletin from summer 2007. The Deutsche Bundesbank in its Monthly Bulletin from December 2007 remarks that "[...] The price-dampening impact of the appreciation of the euro so far might also be stronger than expected and the euro could continue to appreciate. Taking everything together, however, the upside risks to future price developments predominate at the end of the forecasting horizon" (p. 29).

of investigations of the entire forecast density. For instance, Wallis (2003) states that “the excessive concern with upside risk was not justified over the period considered.” (p. 165). Yet, it remains to be analyzed how informative the BoE’s risk forecasts are in general. In this work, we attempt to assess the information content of the BoE’s risk forecasts for inflation.

This assessment is performed in the context of tests for forecast optimality. The risk forecasts are supposed not to contain information which can reduce the mean squared forecast error of the *mean* forecasts. However, for reasons that will become clear below, the risk forecasts are supposed to reduce the mean squared forecast error of the *mode* forecasts. Both hypotheses will be investigated in this study.

We briefly present the concepts underlying the BoE’s risk forecasts and the data in Section 2. In Section 3, the optimality of risk forecasts is investigated. Robustness checks are carried out in Section 4. Section 5 deals with possible explanations for the BoE’s risk forecasting performance. Section 6 concludes.

2 The Data

Since February 1996, the BoE publishes its inflation forecast in the form of a probability distribution, the “fan chart”. The BoE also publishes the location parameters mean, mode and median, as well as measures of dispersion and skewness. The mode forecast is considered the *central projection* of the BoE, whereas the mean forecast is mainly used to convey information about the prevailing forecast risk. Actually, during the entire subsequent analysis, it will be of no importance that the central projection is a mode forecast. It will only be important that the central projection is *not* the mean forecast.

Macroeconomic forecasts are typically based on certain assumptions concerning the future developments of variables which are considered exogenous with respect to the forecasting models used. Exchange rates, oil prices and foreign demand for domestic goods are typical examples for such variables in the context of macroeconomic models. According to Britton et al. (1998), the BoE’s mode forecast of inflation corresponds to the forecast obtained if the assumptions are based on the most likely future values of these variables. In contrast to that, the mean fore-

cast is based on the expected future values. If the density forecasts are symmetric, most likely and expected future values coincide. If, however, the density forecasts are asymmetric, most likely and expected future values differ. In the latter case, in general the mean and mode forecasts of inflation differ as well.

If the mean forecast exceeds the mode forecast, one speaks of an *upward risk* to the inflation forecast. A *downward risk* is present if the mean forecast is lower than the mode forecast. Thus, risks are present if the BoE's density forecast is asymmetric. For further details, see Britton et al. (1998). In Figure 1, two of the BoE's density forecasts with risks are displayed. For their calculations, we make use of the formulas given in Wallis (2004).

Our analysis uses the BoE's inflation forecasts based on the assumption that the official Bank rate, i.e. the interest rate paid on commercial bank reserves follows a path implied by market interest rates. In line with Elder et al. (2005), for the purpose of forecast evaluations we consider this assumption more adequate than the assumption of a constant official Bank rate.

The inflation forecasts considered in our analysis range from the first quarter of 1998 to the third quarter of 2007. Each of the BoE's quarterly projections covers the current and the subsequent 8 quarters. The data is displayed in Appendix A. Note that in some quarters, mean and mode forecast coincide, so that the forecast risk equals zero, i.e. there is no forecast risk or the forecast risks are balanced. Until 2003, the BoE forecasted the inflation of the all items retail prices index excluding mortgage interest payments (henceforth RPIX). Since 2004, it forecasts the inflation of the consumer price index (henceforth CPI).³

The BoE also publishes risk forecasts for GDP. We do not study these forecasts here, since the analysis of GDP risk forecasts would be more complicated due to the effects of data revisions. Such revisions play a substantial role for the assessment of the BoE's GDP forecasts, as noted by Elder et al. (2005).

³When outturns are compared with forecasts, this change has of course to be taken into account. For instance, an inflation forecast for the fourth quarter of 2004 has to be compared with CPI inflation data if the forecast was made in 2004. If the forecast was made before 2004, it must be compared with RPIX inflation data.




Figure 1: Two of the BoE's density forecasts for inflation. The forecast in the left panel implies an upward risk (mean > mode), the one in the right panel a downward risk (mean < mode).

3 Forecast Optimality

In this section, we are concerned with the partial optimality of forecasts, where partial optimality is defined as in Diebold & Lopez (1996). As mentioned by Diebold & Lopez (1996), the original concept of partial optimality refers to optimality conditional on the information set being used by the forecaster and goes back to Brown & Maital (1981). When we speak of optimality in this study, we always mean partial optimality with respect to the information set that is given by the independent variable(s) of a certain regression. Our tests for optimality assume a loss function being quadratic in the forecast error.

3.1 Mean Forecast Optimality and Risk

It is well known that using the mean of a density forecast as the point forecast minimizes the mean squared forecast error. In contrast to that, the loss function which is minimized by the mode forecast is a rather special all-or-nothing loss function as shown by Wallis (1999). Thus, the mean forecasts of the BoE should

Table 1: Root mean squared errors of mode and mean forecasts

h	0	1	2	3	4	5	6	7	8
n	39	38	37	36	35	34	33	32	31
RMSE of the forecasted...									
... mode	0.16	0.25	0.32	0.39	0.47	0.50	0.48	0.44	0.45
... mean	0.17	0.27	0.35	0.42	0.48	0.50	0.51	0.50	0.52
dev. in %	-5.7	-5.9	-7.1	-6.4	-2.2	0.0	-4.5	-12.7	-14.8

Note: h is the forecast horizon. n denotes the number of observations. dev. in % gives the deviation of the RMSE of the mode forecast from the RMSE of the mean forecast in %.

have a smaller mean squared forecast error than the mode forecasts.

The results for the root mean squared forecast error (RMSE) of the BoE's inflation forecasts displayed in Table 1, however, indicate the opposite. For 8 out of 9 forecast horizons, the mode forecasts yield a smaller RMSE than the mean forecasts.

This result is surprising. However, the observed negative deviations could be statistically insignificant. Tests of the hypothesis that the RMSEs of mode and mean forecasts differ would have to take into account that these forecasts are most likely generated by nested models. This implies that we cannot use the standard critical values of the test statistic proposed by Diebold & Mariano (1995), but have to make modifications as described in West (2006). Yet, it could even be that mode and mean forecasts come from identical models, and it is unclear how to handle this issue. Therefore, instead of testing the difference between RMSEs, we employ a test for forecast optimality.

If the forecasts are optimal, we should not be able to reject the hypothesis $H_0^a : (\alpha_h = 0, \beta_h = 1, \gamma_h = 0)$ in the regression

$$y_{t+h} = \alpha_h + \beta_h \mu_{t,t+h} + \gamma_h (\mu_{t,t+h} - m_{t,t+h}) + \varepsilon_{t,h} \quad (1)$$

where y_{t+h} denotes the inflation rate in period $t + h$, $\mu_{t,t+h}$ denotes the mean forecast made in t for $t + h$, $m_{t,t+h}$ is the corresponding mode forecast, $\varepsilon_{t,h}$ is a

zero-mean error term⁴ and $h = 0, 1, \dots, 8$ a positive integer denoting the forecast horizon. For $h = 0$, the forecast is actually a nowcast for the current quarter.

Note that $\mu_{t,t+h} - m_{t,t+h}$ is a measure of the *forecasted risk*. Also take into account that the mean and the mode forecast are based on the same information set. Thus, the forecasted risk should not be significant in this regression, since, conditional on the information set used, the forecasted mean $\mu_{t,t+h}$ is supposed to minimize the variance of $\varepsilon_{t,h}$.

In order to isolate the effect of the inclusion of the risk forecast, we also test the hypothesis $\gamma_h = 0$ separately. Moreover, we test the hypothesis of optimality $H_0^b : (\alpha'_h = 0, \beta'_h = 1)$ in the reduced regression $y_{t+h} = \alpha'_h + \beta'_h \mu_{t,t+h} + \varepsilon'_{t,h}$. H_0^a and H_0^b are tested with an F-test.

Table 2: Results of tests for partial optimality of mean forecasts

h	0	1	2	3	4	5	6	7	8
n	39	38	37	36	35	34	33	32	31
coefficient estimates									
α_h	0.19 (0.10)	0.40 (0.19)	0.73 (0.23)	1.25 (0.20)	2.26 (0.34)	2.89 (0.48)	2.91 (0.48)	2.70 (0.50)	2.71 (0.61)
β_h	0.92 (0.04)	0.84 (0.08)	0.71 (0.10)	0.48 (0.10)	0.01 (0.14)	-0.30 (0.20)	-0.29 (0.20)	-0.17 (0.21)	-0.16 (0.26)
γ_h	-1.58 (0.69)	-1.61 (0.57)	-1.74 (0.48)	-1.38 (0.39)	0.50 (0.77)	2.09 (0.74)	0.77 (0.64)	-0.37 (0.44)	-0.35 (0.52)
p -values									
H_0^b	0.22	0.06*	0.01***	0.00***	0.00***	0.00***	0.00***	0.00***	0.00***
H_0^a	0.06*	0.02**	0.00***	0.00***	0.00***	0.00***	0.00***	0.00***	0.00***
$\gamma_h = 0$	0.03**	0.01***	0.00***	0.00***	0.52	0.01***	0.24	0.41	0.50

Note: Standard errors are in parentheses. These are Newey-West (1987) standard errors. The bandwidth parameter is chosen based on the procedure proposed by Andrews (1991). *, (**, ***) denotes rejection of the null hypothesis at the 10 (5, 1) % significance level. n denotes the number of observations.

The results in Table 2 clearly indicate that using the forecasted risk significantly improves the forecasts for several horizons. For $h = 0, \dots, 3$ and $h = 5$, the null of $\gamma_h = 0$ can be rejected. In the light of the results obtained for the RMSEs, this finding is not too surprising. Since the mode forecasts have lower RMSEs than the

⁴In what follows, all error terms have the zero-mean property.

mean forecasts, and the mode forecasts are contained in the forecasted risk, the forecasted risk can help to improve the mean forecast.⁵

The null of forecast optimality can be rejected at lower significance levels with (1) than with the reduced regression at least for $h = 0, 1, 2$. However, even with the reduced regression the null can be rejected at a significance level of 1 % for $h \geq 2$.

3.2 Bias of Mode, Mean and Risk Forecasts

If tests reject the hypothesis of forecast optimality, this might be due to the presence of a bias in the forecasts. It is therefore interesting to test the hypothesis of $\alpha_h^\mu = 0$ in the equation

$$y_{t+h} - \mu_{t,t+h} = \alpha_h^\mu + \varepsilon_{t,h}^\mu$$

where α_h^μ is a constant and $\varepsilon_{t,h}^\mu$ is an error term. For the sake of completeness, we do the same for the mode forecasts, i.e. we test the hypothesis of $\alpha_h^m = 0$ in the equation

$$y_{t+h} - m_{t,t+h} = \alpha_h^m + \varepsilon_{t,h}^m$$

where α_h^m is a constant and $\varepsilon_{t,h}^m$ is an error term.

Finally, it might be interesting to know whether the BoE is mainly concerned with upward or downward risks to inflation. This question can be addressed by testing $\alpha_h^r = 0$ in the equation

$$\mu_{t,t+h} - m_{t,t+h} = \alpha_h^r + \varepsilon_{t,h}^r$$

where α_h^r is a constant, and the error term $\varepsilon_{t,h}^r$, like $\varepsilon_{t,h}^\mu$ and $\varepsilon_{t,h}^m$, can be serially correlated.⁶

⁵In principle, we could also test forecast optimality by directly using the forecasted mode $m_{t,t+h}$ instead of the forecasted risk $\mu_{t,t+h} - m_{t,t+h}$ in equation (1). Since, however, $\mu_{t,t+h}$ and $m_{t,t+h}$ are strongly correlated, this would result in multicollinearity.

⁶The presence of serial correlation in $\varepsilon_{t,h}^r$ is caused by the serial correlation in the forecasted risk $\mu_{t,t+h} - m_{t,t+h}$. This serial correlation can be explained as follows: If there is a risk to the forecast made in period t for the period $t + h + 1$, and this risk does not materialize until the forecast in $t + 1$ is made, the risk for the period $t + h + 1$ persists. Thus, $\mu_{t,t+h+1} - m_{t,t+h+1}$ and $\mu_{t+1,t+1+h} - m_{t+1,t+1+h}$ are correlated. Since in addition risks are typically correlated over horizons, i.e. $\mu_{t,t+h+1} - m_{t,t+h+1}$ is correlated with $\mu_{t,t+h} - m_{t,t+h}$, the forecast risks for a certain

The results in Table 3 show that neither the mode forecasts nor the mean forecasts have a significant bias.⁷ The estimates for α_h^μ and α_h^m are very similar, both implying that the BoE slightly underpredicted inflation for $h = 0, \dots, 6$ and slightly overpredicted inflation for $h = 7, 8$.

The estimates of α_h^r indicate that on average the BoE predicts upward risks to inflation. The average risks increase with the forecast horizon h . However, these average risks are very small, attaining at most 0.03 percentage points. Moreover, the presence of a non-zero average risk is not significant. Thus, we cannot reject the hypothesis that, on average, the BoE forecasts balanced risks, i.e. no risks.

Note that the latter result does not imply any kind of optimality, but is purely descriptive. If the true unconditional density of inflation was positively skewed, its mean would exceed its mode, and hence, upward risks should be forecasted on average. For the same reason, the absence of bias in the mode forecasts cannot be evaluated. Only for the mean forecast, unbiasedness is a property of optimal forecasts.

3.3 Risk Forecast Optimality

In Section 3.1, we have tested the optimality of mean forecasts with respect to the information contained in the risk forecasts. In case of optimality, the risk forecasts must not improve the mean forecasts. In the current section, our focus will be somewhat different. We want to test the optimality of risk forecasts in a more direct manner. In order to do so, we ask the question whether the central projection, i.e. the mode forecast can be improved by taking the risk forecast into account if agents have a quadratic loss function. If the risk forecasts have the intended information content, such an improvement should be observed.

In order to evaluate the risk forecasts, we need a measure of realized risks. The measure that corresponds to the risk forecasts of the BoE is given by the difference between realized inflation and the mode forecast. If the risk forecasts are optimal,

horizon h are serially correlated as well. However, the serial correlation of $\varepsilon_{t,h}^r$ turns out to be less pronounced than the serial correlation of $\varepsilon_{t,h}^\mu$ and $\varepsilon_{t,h}^m$.

⁷For the mean forecasts, this result is in line with, for instance, Clements (2004), Wallis (2004) and Elder et al. (2005).

Table 3: Results of regressions to test for bias

h	0	1	2	3	4	5	6	7	8
n	39	38	37	36	35	34	33	32	31
coefficient estimates									
α_h^μ	0.01 (0.03)	0.04 (0.05)	0.05 (0.07)	0.08 (0.10)	0.10 (0.12)	0.09 (0.14)	0.05 (0.15)	-0.04 (0.15)	-0.13 (0.16)
α_h^m	0.01 (0.03)	0.05 (0.05)	0.07 (0.06)	0.10 (0.09)	0.12 (0.12)	0.12 (0.13)	0.08 (0.13)	-0.01 (0.13)	-0.10 (0.13)
α_h^r	0.01 (0.01)	0.01 (0.01)	0.02 (0.01)	0.02 (0.01)	0.02 (0.01)	0.03 (0.02)	0.03 (0.02)	0.03 (0.04)	0.03 (0.04)
p -values									
$\alpha_h^\mu = 0$	0.78	0.46	0.45	0.41	0.43	0.51	0.72	0.80	0.43
$\alpha_h^m = 0$	0.61	0.31	0.29	0.27	0.30	0.37	0.56	0.93	0.44
$\alpha_h^r = 0$	0.36	0.26	0.21	0.14	0.12	0.13	0.31	0.46	0.51

Note: Standard errors are in parentheses. These are Newey-West (1987) standard errors. The bandwidth parameter is chosen based on the procedure proposed by Andrews (1991). *, (**, ***) denotes rejection of the null hypothesis at the 10 (5, 1) % significance level. n denotes the number of observations.

we have

$$y_{t+h} - m_{t,t+h} = \mu_{t,t+h} - m_{t,t+h} + \varepsilon_{t,h} \quad (2)$$

where $\varepsilon_{t,h}$ is an error term. The term $y_{t+h} - m_{t,t+h}$ is the forecast error with respect to the central projection, but in the context of the subsequent analysis it should be regarded as measure for the *realized risk*, as $\mu_{t,t+h} - m_{t,t+h}$ is the measure of the forecasted risk.

Note that we could rewrite equation (2) as $y_{t+h} = m_{t,t+h} + (\mu_{t,t+h} - m_{t,t+h}) + \varepsilon_{t,h}$. This equation states that in case of risk forecast optimality, the mode forecast must be shifted towards the mean by the amount $(\mu_{t,t+h} - m_{t,t+h})$. If the mode forecast was not shifted towards the mean by this amount, the resulting mean squared forecast error would exceed the variance of $\varepsilon_{t,h}$.

It is also important to note that we do not regress the true risk on the forecasted risk. The true risk is given by $E[Y_{t+h}] - M[Y_{t+h}]$, where $M[Y_{t+h}]$ denotes the mode of Y_{t+h} .⁸ Since we do not have an estimate for $M[Y_{t+h}]$, we cannot measure

⁸We adopt the common convention to use uppercase letters for random variables and lowercase letters for their realizations.

the true risk. It should be stressed that nothing in equation (2) relates the true mode $M[Y_{t+h}]$ to the forecasted mode $m_{t,t+h}$. This becomes especially clear if we suppose, for instance, that an institution simply publishes a constant c as its central projection. In this case, equation (2) would become $y_{t+h} - c = \mu_{t,t+h} - c + \varepsilon_{t,h}$. The forecasted risk would be given by $\mu_{t,t+h} - c$, and the realized risk by $y_{t+h} - c$. For the mentioned equation, it does not matter whether c is supposed to be the mode, median or any other point related to the distribution of Y_{t+h} . For optimal risk forecasts, it only matters that the forecasted risk $\mu_{t,t+h} - c$ equals the expected realized risk $E_t[Y_{t+h} - c]$.

Based on equation (2), it is possible to run Mincer-Zarnowitz regressions for the risk forecasts. That is, one can estimate the coefficients α_h and β_h of the equation

$$y_{t+h} - m_{t,t+h} = \alpha_h + \beta_h (\mu_{t,t+h} - m_{t,t+h}) + \varepsilon_{t,h} \quad (3)$$

and test whether $\alpha_h = 0$ and $\beta_h = 1$ hold. However, we will not consider the joint hypothesis $(\alpha_h = 0, \beta_h = 1)$, but test the hypothesis of the first ($\alpha_h = 0$) and of the second condition for optimality ($\beta_h = 1$) separately. This will turn out to deliver more insights than a joint test.

Moreover, we will test the hypothesis of qualitatively correct information content of the risk forecasts. This test is based on the assumption that $\alpha_h = 0$ holds and involves tests of two hypotheses. The first hypothesis is given by $\beta_h = 0$. If the risk forecasts have an information content, this hypothesis should be rejected. The second hypothesis is given by $\beta_h > 0$. If the risk forecasts have a qualitatively correct information content, this hypothesis should not be rejected.

With tests of two hypotheses, four cases can occur. In Table 4, we list these cases and their implications. They are ranked according to their desirability. Case I is clearly the most desirable case. Case II is less desirable, but not being able to reject $\beta_h = 0$ could not only be caused by lack of information content, but also, for example, by a too small sample size. Cases III and IV are the least desirable ones. In principle, they have identical implications, but in case III, the confirmation of the alternative hypothesis of a qualitatively adverse information content is weaker, since it occurs at a larger significance level.⁹

⁹Note that in case I, if at a given significance level $\beta_h = 0$ is rejected and $\beta_h > 0$ is not,

Table 4: Implications of test results for information content

case	$\beta_h = 0$	$\beta_h > 0$	implication for risk forecasts
I	rejected	not rejected	qualitatively correct information content confirmed
II	not rejected	not rejected	null of no information content not rejected
III	not rejected	rejected	qualitatively adverse information content confirmed
IV	rejected	rejected	qualitatively adverse information content confirmed

The estimation results are displayed in the upper panel of Table 5. While the α_h 's are all close to zero, the β_h 's are mostly negative. This indicates that, for a period with a forecasted upward risk, a downward risk is more likely to materialize.

The lower panel of Table 5 contains the p -values of tests of the various null hypotheses. Here and in the following, we use a significance level of 10 %. It turns out that the null of $\alpha_h = 0$ cannot be rejected for any forecast horizon h . However, the null of $\beta_h = 1$, i.e. of the second condition for optimality is rejected for all horizons except for $h = 4, 5, 6$ at very low significance levels. For $h = 4, 5, 6$, the largest standard errors of β_h are found, so that inference is relatively difficult for these horizons.

Since $\alpha_h = 0$ cannot be rejected, tests of the null hypothesis of qualitatively correct information content can be conducted. They yield rather disappointing results. Not for a single forecast horizon case I is found, i.e. for no forecast horizon the presence of qualitatively correct information content is confirmed. For the horizons $h = 0, 4, 5, 6, 8$ we cannot reject the null of no information content. For the other horizons, we even find a qualitatively adverse information content of the risk forecasts.

In summary, the risk forecasts appear to fulfill the first condition for optimality, but the intended information content, i.e. a qualitatively correct information content is not present in these forecasts. Rather they have either no information content, or even a qualitatively adverse information content.

In principle, the puzzling results with respect to the tests of the β_h 's could be related to the presence of a constant, albeit insignificant, in equation (3). Imagine

this implies that $\beta_h < 0$ is rejected. Therefore, in case I we cannot only state that the null of qualitatively correct information content cannot be rejected. We can even claim that, since the null of qualitatively adverse information content can be rejected, the alternative hypothesis of qualitatively correct information content is confirmed.

Table 5: Results of Mincer-Zarnowitz regressions

h	0	1	2	3	4	5	6	7	8
n	39	38	37	36	35	34	33	32	31
coefficient estimates									
α_h	0.02 (0.03)	0.06 (0.04)	0.09 (0.06)	0.13 (0.08)	0.14 (0.10)	0.12 (0.13)	0.09 (0.12)	0.01 (0.12)	-0.08 (0.12)
β_h	-0.77 (0.75)	-1.00 (0.67)	-1.35 (0.75)	-1.47 (0.73)	-0.59 (1.03)	0.13 (1.41)	-0.53 (1.08)	-0.82 (0.62)	-0.50 (0.56)
p -values									
$\alpha_h = 0$	0.49	0.18	0.17	0.12	0.20	0.37	0.46	0.93	0.50
$\beta_h = 1$	0.02**	0.01***	0.00***	0.00***	0.13	0.54	0.17	0.01***	0.01**
$\beta_h = 0$	0.31	0.15	0.08*	0.05*	0.57	0.93	0.63	0.19	0.38
$\beta_h > 0$	0.16	0.07*	0.04**	0.03**	0.29	0.54	0.31	0.10*	0.19
case	II	III	IV	IV	II	II	II	III	II

Note: Standard errors are in parentheses. These are Newey-West (1987) standard errors. The bandwidth parameter is chosen based on the procedure proposed by Andrews (1991). *, (**, ***) denotes rejection of the null hypothesis at the 10 (5, 1) % significance level. n denotes the number of observations.

that the risk forecasts for three periods are given by $\{0.5, -0.5, 1\}$, and the realizations are $\{0.3, 0.6, 0.4\}$. In this case, a forecasted upward risk would mostly be correctly associated with the realization of an upward risk. However, a regression according to equation (3) would yield $\alpha_h = 0.49$ and $\beta_h = -0.16$, where α_h (and also β_h) would not differ significantly from zero. Yet, a regression with the restriction $\alpha_h = 0$ would yield $\beta_h = 0.17$ which could be considered a more reasonable result. Therefore, we also estimate equation (3) with the restriction $\alpha_h = 0$. The results are displayed in the Table 6.

It turns out that the signs of the β_h 's are identical to those in the unrestricted estimation. The null of $\beta_h = 1$ is rejected for the same horizons as before. Again case I does not occur for any horizon, so there is no evidence for a qualitatively correct information content. However, in contrast to the estimations with a constant, now case II is found for all horizons except $h = 2$. For this horizon, there is evidence for qualitatively adverse information content. Yet, for several other horizons the p -values of the tests of $\beta_h > 0$ exceed 10 % by a small amount only, thereby again suggesting the possibility of a qualitatively adverse information content.

Table 6: Results of Mincer-Zarnowitz regressions with restriction $\alpha_h = 0$

h	0	1	2	3	4	5	6	7	8
n	39	38	37	36	35	34	33	32	31
coefficient estimates									
β_h	-0.69 (0.78)	-0.80 (0.82)	-1.08 (0.74)	-1.07 (0.85)	-0.17 (1.18)	0.50 (1.46)	-0.35 (1.13)	-0.80 (0.63)	-0.58 (0.50)
p -values									
$\beta_h = 1$	0.04**	0.03**	0.01***	0.02**	0.33	0.73	0.24	0.01***	0.00***
$\beta_h = 0$	0.38	0.33	0.16	0.22	0.89	0.73	0.76	0.21	0.25
$\beta_h > 0$	0.19	0.17	0.08*	0.11	0.44	0.63	0.38	0.11	0.13
case	II	II	III	II	II	II	II	II	II

Note: Standard errors are in parentheses. These are Newey-West (1987) standard errors. The bandwidth parameter is chosen based on the procedure proposed by Andrews (1991). *, (**, ***) denotes rejection of the null hypothesis at the 10 (5, 1) % significance level. n denotes the number of observations.

In summary, the analysis shows that the risk forecasts are not optimal. It is also highly probable that they lack qualitatively correct information content. Rather, they appear to have either no information content or even an adverse information content. The evidence for the latter phenomenon is not particularly strong, but due to the small sample size, we cannot discriminate more clearly between both alternatives.

3.4 Interpretation of the Test Results

It is important to understand that the results found in the previous section are not a consequence of the results found in Section 3.1. In that section, we found that the mean forecasts can be improved by considering the forecasted risk, and that the mode forecasts have lower RMSEs than the mean forecasts. These facts do not imply that the risk forecasts have no information content or an adverse information content, i.e. that the mode forecasts cannot be improved by considering the forecasted risk, or that the mode forecasts can only be improved by considering the reversed forecasted risk.

In order to clarify this issue, consider the densities and the mean and mode forecasts displayed in Figure 2. The true densities are symmetric around zero, so

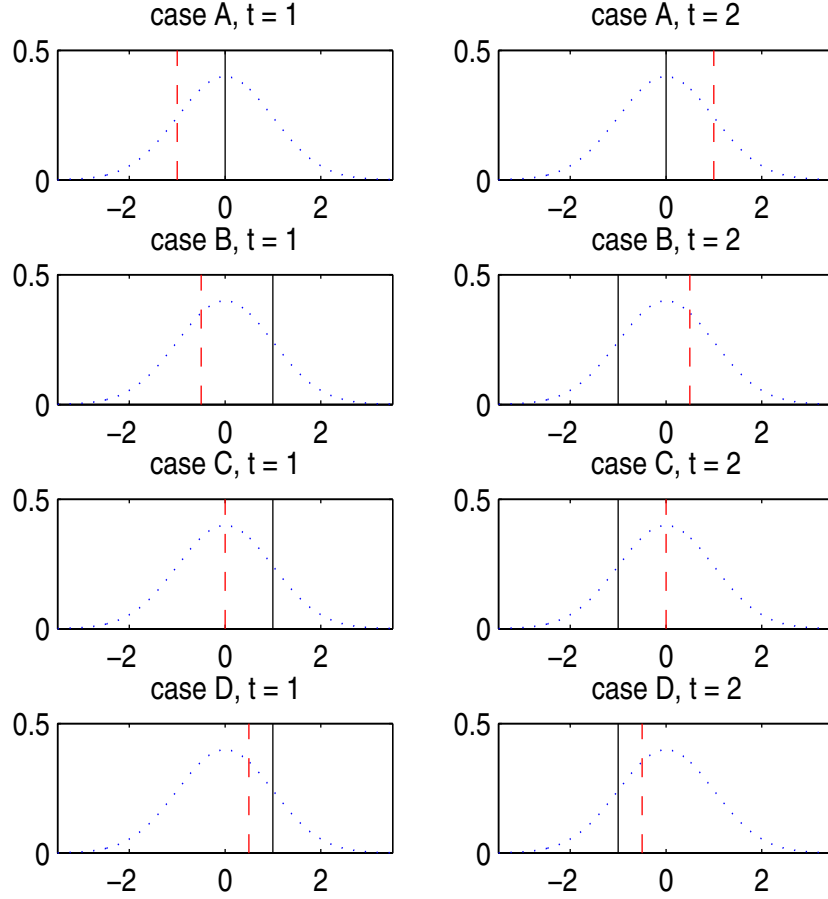


Figure 2: Four cases of mean and mode forecasts for two periods. The dotted lines are the true densities of Y_1 and Y_2 . The solid lines indicate the mean forecasts $\mu_{t-h,t}$, the dashed lines indicate the mode forecasts $m_{t-h,t}$.

that their mode and mean coincide at zero. We plot forecasts for two periods, $t = 1$ and $t = 2$ and distinguish between the four cases A, B, C and D.

In case A, the risk forecasts $\mu_{t-h,t} - m_{t-h,t}$ have the right sign and the correct size. With respect to our analysis, in this case mean forecasts and risk forecasts are optimal.¹⁰ In case B, the risk forecasts $\mu_{t-h,t} - m_{t-h,t}$ have the right sign, but they are exaggerated. That is, for $t = 1$, upward risks are present and are forecasted, and for $t = 2$, downward risk are present and are forecasted, but both risks are too large. In case C, the mode forecasts are actually equal to the mean

¹⁰Of course, the mode forecast does not coincide with the true mode, but, as mentioned above, this does not matter for our analyses.

of the true density, so the risk forecasts cannot improve the mode forecasts. In case D, the forecasted risks have the wrong signs. This could be interpreted as an exaggeration of risks already present in the mode forecast. For the mean forecasts, these risks are then exaggerated even further, leading to wrong signs of the risk forecasts.

In all four cases, mean and mode forecasts are unbiased. The same holds for the risk forecasts. In the three cases B, C and D, the mode forecasts have a smaller RMSE than the mean forecasts, because the mode forecasts are closer to the mean of the true density. Thus, for these three cases, we would obtain the results found in Section 3.1. However, case B cannot be reconciled with the results of Section 3.3, since case B would correspond to case I in Table 4. Yet, case I is never observed. In contrast to that, case C can be reconciled with case II, and case D can be reconciled with cases III and IV.

Thus, the results found in Section 3.1 and Section 3.3 suggest that the mode forecasts are closer to the mean of the true densities than the mean forecasts. Moreover, the mode forecasts might be improved by shifting them *away* from the forecasted mean. The latter possibility suggests that the probability of certain risks as e.g. exchange rate changes could already be overestimated in the mode forecast. In any case, for economic agents with a loss function being quadratic in the forecast error, the BoE's mean forecasts do not minimize these losses.

4 Robustness Checks

4.1 Risk Forecast Optimality - A Qualitative Analysis

In order to shed further light on the performance of the BoE's risk forecasts, we conduct an analysis of their qualitative performance. If the lack of information content or the adverse information content found above is mainly driven by outliers, i.e. mainly driven by very few periods in which large upward [downward] risks were forecasted, but large downward [upward] risks materialized, the qualitative performance of the risk forecasts could be expected to be considerably better than their quantitative performance.

Therefore, we regard all periods for which the BoE saw a forecast risk and

Table 7: Qualitative performance of risk forecasts

h	0	1	2	3	4	5	6	7	8	all
number of risk forecasts	19	19	19	18	18	26	25	26	24	194
share of failures in %	58	74	63	72	67	58	72	65	54	64
share of successes in %	42	26	37	28	33	42	28	35	46	36
failure-success ratio	1.4	2.8	1.7	2.6	2.0	1.4	2.6	1.9	1.2	1.8

compare the directions of these risks (downward or upward) with the outturns (below or above the mode forecast). If an upward [downward] risk was forecasted, and the outturn was above [below] the mode forecast, we count this as a successful risk forecast. Otherwise, the risk forecast is counted as a failure. In Table 7, we show the share of failures and of successes. Both shares add up to 1. We also report the ratio of failures to successes.

Obviously, the qualitative risk forecasts do not perform well either. For all forecast horizons, the failure-success ratio exceeds 1, ranging from 1.2 for $h = 8$ to 2.8 for $h = 1$. Considering all risk forecasts, the failure-success ratio attains a value of 1.8. Thus, there is no evidence of correct information content of the qualitative risk forecasts. Rather, the results indicate the presence of an adverse information content.

In order to test hypotheses about information content, one could think of setting up a (2×2) contingency table for each horizon, and then use a chi-square or a related test. The categories of the table would be forecasted upward and downward risk, and materialized upward or downward risk, i.e. outturns above or below the mode forecast. However, the tests mentioned do not take the time-series context of the data into account. This issue is emphasized by Christoffersen (1998) and Pesaran & Timmermann (2006). The inference would be distorted by serial correlation of the data.

To overcome this problem, Pesaran & Timmermann (2006) propose to transform the data to sequences of 1's and 0's, to run standard regressions and to use heteroscedasticity- and autocorrelation-consistent covariance estimators for infer-

ence. So we recode our data according to

$$\tilde{y}_{t,t+h}^q = \begin{cases} 1 & \text{if } y_{t+h} - m_{t,t+h} > 0 \\ 0 & \text{if } y_{t+h} - m_{t,t+h} < 0 \\ na & \text{if } y_{t+h} - m_{t,t+h} = 0 \end{cases}, \quad \tilde{x}_{t,t+h}^q = \begin{cases} 1 & \text{if } \mu_{t,t+h} - m_{t,t+h} > 0 \\ 0 & \text{if } \mu_{t,t+h} - m_{t,t+h} < 0 \\ na & \text{if } \mu_{t,t+h} - m_{t,t+h} = 0 \end{cases}$$

where na denotes a missing value.

In order to construct variables $y_{\tau,\tau+h}^q$ and $x_{\tau,\tau+h}^q$ without missing values, we use the transformations

$$\begin{aligned} \mathbf{y}_h^q &= \mathbf{A}_h \tilde{\mathbf{y}}_h^q \\ \mathbf{x}_h^q &= \mathbf{A}_h \tilde{\mathbf{x}}_h^q \end{aligned}$$

where the vector $\tilde{\mathbf{y}}_h^q$ is given by $\tilde{\mathbf{y}}_h^q = (\tilde{y}_{1,1+h}^q, \tilde{y}_{2,2+h}^q, \dots, \tilde{y}_{n,n+h}^q)'$, the vector $\tilde{\mathbf{x}}_h^q$ is given by $\tilde{\mathbf{x}}_h^q = (\tilde{x}_{1,1+h}^q, \tilde{x}_{2,2+h}^q, \dots, \tilde{x}_{n,n+h}^q)'$ and \mathbf{A}_h is a known $(m \times n)$ selection matrix consisting of 1's and 0's, and with $m \leq n$. \mathbf{A}_h is chosen such that the vectors $\mathbf{y}_h^q = (y_{1,1+h}^q, y_{2,2+h}^q, \dots, y_{m,m+h}^q)'$ and $\mathbf{x}_h^q = (x_{1,1+h}^q, x_{2,2+h}^q, \dots, x_{m,m+h}^q)'$ do not contain missing values. If there were no missing values, \mathbf{A}_h would be an $(n \times n)$ identity matrix. If, for example, there was no risk forecast for the first forecast of horizon h , i.e. in the case $y_{1+h} - m_{1,1+h} = 0$, m would equal $n - 1$ and \mathbf{A}_h would be given by the $(m \times n)$ matrix $\mathbf{A}_h = [\mathbf{0}_m \quad \mathbf{I}_m]$, where $\mathbf{0}_m$ denotes an $(m \times 1)$ vector of 0's and \mathbf{I}_m denotes the $(m \times m)$ identity matrix.

Having constructed the qualitative variables $y_{\tau,\tau+h}^q$ and $x_{\tau,\tau+h}^q$ with $\tau = 1, 2, \dots, m$, we run the regression

$$y_{\tau,\tau+h}^q = \alpha_h + \beta_h x_{\tau,\tau+h}^q + \varepsilon_{\tau,h} \quad (4)$$

where $\varepsilon_{\tau,h}$ is an error term. In this regression only the value of β_h is of interest. A value of 1 corresponds to a qualitatively optimal¹¹ information content, a value of 0 to no information content, and a value below 0 to an adverse information content. We do not test the hypothesis $\beta_h = 1$, because due to the construction of $y_{\tau,\tau+h}^q$ and $x_{\tau,\tau+h}^q$, $|\beta_h| \leq 1$ holds. Hence, as emphasized by Harding & Pagan (2006), testing $\beta_h = 1$ would mean testing on the boundary of the parameter

¹¹Qualitative optimality is achieved by the minimization of a quadratic loss function, where the argument of the loss function is a qualitative variable.

Table 8: Results of regressions with qualitative variables

h	0	1	2	3	4	5	6	7	8
m	19	19	19	18	18	26	25	26	24
coefficient estimates									
β_h	-0.15 (0.22)	-0.52 (0.20)	-0.36 (0.22)	-0.58 (0.13)	-0.33 (0.25)	-0.13 (0.23)	-0.40 (0.17)	-0.30 (0.16)	-0.08 (0.23)
p -values									
$\beta_h = 0$	0.50	0.02**	0.12	0.00***	0.20	0.59	0.02**	0.07*	0.72
$\beta_h > 0$	0.25	0.01***	0.06	0.00***	0.10	0.29	0.01**	0.04**	0.36
case	II	IV	II	IV	II	II	IV	IV	II

Note: Standard errors are in parentheses. These are Newey-West (1987) standard errors. The bandwidth parameter is chosen based on the procedure proposed by Andrews (1991). *, (**, ***) denotes rejection of the null hypothesis at the 10 (5, 1) % significance level. m denotes the number of observations.

space, leading to a non-standard distribution of the test statistic.

The estimation results are displayed in Table 8. For all forecast horizons, the estimated β_h 's are negative. For the horizons $h = 0, 2, 4, 5, 8$ we cannot reject the null of no information content. For the other horizons, the presence of qualitatively adverse information content is confirmed.¹²

To sum up, the qualitative analysis confirms the results of the quantitative analysis. There are no signs of qualitatively correct information content of the risk forecasts, and there might even exist a qualitatively adverse information content.

4.2 Stability Over Time

The BoE publishes density forecasts since February 1996. At that time, to the best of our knowledge the BoE was the only institution to publish macroeconomic projections of this kind, so that the BoE could not draw on others institutions' experiences. It is therefore quite likely that the BoE has gone through a learning process. If this process was still ongoing in February 1998 and later, the results

¹²Of course, the deletion of certain observations in the process of constructing $y_{\tau, \tau+h}^q$ and $x_{\tau, \tau+h}^q$ can be expected to cause breaks in the autocorrelation structure of these variables. The calculation of standard errors might suffer from this problem. Therefore, we vary the bandwidth parameters of the Newey-West procedure between 0 and 6, finding that the standard errors remain broadly unchanged.

of our analysis might be affected by it. In this case, it would be more appropriate to start the analysis at that point in time where the learning process came to an end.

Looking at the forecasts, there indeed appear to be signs of a learning process. Until May 2001, in every projection risks were forecasted for at least 8 out of 9 forecast horizons. However, from August 2001 onwards, 19 out of 25 projections had at most 4 out of 9 horizons for which risks were forecasted. In these 19 projections, risks were forecasted for $h = 5, \dots, 8$, i.e. the longer term, but not for $h = 0, \dots, 4$, i.e. the shorter term. Thus, the BoE seems to have become more cautious with respect to its risk forecasts, especially in the shorter term. Yet, it is not completely clear if August 2001 should be considered as the breakpoint.

Until May 2001, the share of forecasts with risks was 94 %. For forecasts from August 2001 onwards, this share equals only 40 %. The difference between these shares amounts to 55 %. If we considered May 2002 instead of August 2001 as the breakpoint, this difference would reach its largest possible value of 58 %.¹³ We therefore assume that the switch to a more cautious risk forecasting regime might have occurred in August 2001 or in May 2002. Increased caution could of course translate into changes of the risk forecasting performance.

It would therefore be interesting to split the data sample in August 2001 and May 2002 and to investigate whether the information content of the risk forecasts differs among the two subsamples. Unfortunately, for many horizons there are only very few risk forecasts differing from zero after these breakpoints, the lowest number being three if May 2002 is considered as the breakpoint. An analysis based on individual forecast horizons would thus not deliver many useful insights. Hence, we decide to pool all forecast horizons and to perform a panel analysis.

In the static panel model, we consider the forecast publication dates as the time variable and the forecast horizon as the group variable. Thus, the panel model emerging from equation (3) is given by

$$y_{t+h} - m_{t,t+h} = \alpha + \beta(\mu_{t,t+h} - m_{t,t+h}) + u_{t,h}, \quad (5)$$

¹³Until February 2002, the share of forecasts with risks was 92 %. For forecasts from May 2002 onwards, this share equals 34 %.

where $u_{t,h}$ denotes the error component. This random effects specification is supported by a Hausman test. Strictly speaking, we set $\alpha_h = \alpha + \eta_h$, where η_h is the specific disturbance. With assuming a remainder disturbance $\varepsilon_{t,h}$, we define the error component as $u_{t,h} = \eta_h + \varepsilon_{t,h}$. The intercept α is treated as the mean of the random effects. We estimate model (5) by Generalized Least Squares (henceforth GLS), assuming heteroskedasticity and imposing a panel-specific first-order autoregressive (henceforth AR(1)) process for the error term. Furthermore, cross-sectional correlation is explicitly taken into account. The latter option, however, requires the panel to be strictly balanced. Therefore, observations after August 2005 are deleted, resulting in a panel containing 279 entries rather than 315 for the unbalanced panel. Yet, with a glance at the data we remark that only three forecasts with risks are excluded.

GLS estimations are performed on the complete balanced panel data and on four subsamples. The four subsamples are determined by the two possible break-points. We set α to zero since it turns out not to differ significantly from this value in the full-sample estimation.¹⁴

The resulting estimates displayed in Table 9 show stark differences depending on the breakpoint date used. This is due to the fact that the risk forecasts published in November 2001 and February 2002 were very successful. If we consider the breakpoint to be August 2001, the risk forecasts have improved. While there was an adverse information content before the breakpoint, the null of no information content cannot be rejected after the breakpoint.¹⁵ However, the risk forecasts still do not have a qualitatively correct information content, with β being smaller than zero. If May 2002 is considered as the breakpoint, the risk forecasts have deteriorated. The second condition for optimality, i.e. $\beta = 1$ is rejected in all samples.

Thus, there is no evidence that increased caution with respect to the risk forecasts has led to a qualitatively correct information content of these forecasts.

¹⁴The estimate of α is 0.020, its standard error being 0.024.

¹⁵One could suppose that standard errors might be distorted by the use of a panel-specific AR(1)-process for the error term, since it would be more appropriate to employ moving-average processes with orders related to the respective forecast horizons. However, LM tests find evidence for serial correlation of the residuals only for the forecast horizons $h = 3, 4$. In any case, no estimation routine being able to handle cross-sectional correlation and moving-average or higher-order AR(k)-processes is known to us.

Table 9: Static panel estimation results

	full sample	old regime until May 2001	old regime until Feb 2002	new regime since Aug 2001	new regime since May 2002
n	279	126	153	153	126
s.r.f.	59 %	94 %	92 %	47 %	40 %
coefficient estimates					
β	-0.39 (0.24)	-0.70 (0.20)	-0.33 (0.23)	-0.06 (0.33)	-1.72 (0.49)
p -values					
$\beta = 1$	0.00***	0.00***	0.00***	0.00***	0.00***
$\beta = 0$	0.11	0.00***	0.16	0.86	0.00***
$\beta > 0$	0.05*	0.00***	0.08*	0.43	0.00***
case	III	IV	III	II	IV

Note: GLS estimation with AR(1)-process for error terms, and cross-sectional correlation. Standard error are in parentheses. *,(**, ***) denotes rejection of the null hypothesis at the 10 (5, 1) % significance level. n denotes the number of observations. s.r.f. denotes the share of risk forecasts. The s.r.f.'s differ from those mentioned in the text due to the different end date of the sample (August 2005 instead of August 2007).

For the whole sample, there is again weak evidence of a qualitatively adverse information content.

5 Possible Explanations for Poor Risk Forecasting Performance

5.1 Endogeneity of Outturns

If inflation reacts to the BoE's forecasts, then it could of course happen that estimations of equations like (3) yield misleading results. Consider, for example, the case of forecasted upward risks to inflation. In this case, economic agents could anticipate a risk of rising interest rates. In response to this risk, economic activity could be dampened, leading to lower demand and, consequently, to lower inflation. Then, even if the forecasted upward risk to inflation materializes, inflation could

still be lower than forecasted. In this case, the risk forecasts could even have an adverse information content. The same would happen if the BoE actually set interest rates according to its risk forecasts.

If there is a transmission channel from forecasts to realizations, then the analysis conducted above could easily come to wrong conclusions. However, the prevailing opinion in economics is that inflation can only be influenced with a lag by monetary policymakers. Actually, this is the reason why central banks are concerned with forecasting. They know that their current decisions will not affect the economy instantaneously, but in the future.

Taking this fact into account, it should be clear that, if a transmission channel from forecasts to realizations exists, its importance should increase with the forecasting horizon. Therefore, we should be able to assess the importance of this channel by regressing the β_h 's of our regressions on the respective h 's. If the transmission channel is important, the coefficient with respect to h should be negative, because for short horizons, the β_h 's should correctly measure the information content of the risk forecasts and should therefore be relatively large. For larger horizons, the transmission channel would become more important, lowering the values of the β_h 's.¹⁶

Therefore, we estimate the equation

$$\beta_h = a + bh + \varepsilon_h$$

where ε_h is an error term. We estimate this equation for the β_h 's of the standard Mincer-Zarnowitz regression (3), of the Mincer-Zarnowitz regression with $\alpha_h = 0$ and of the equation with qualitative variables (4). The results of these regressions are displayed in Table 10.

None of the b 's is negative, indicating that the transmission channel from forecasts to realizations does not play a major role. In any case, the negative signs of the β_h 's for short forecasting horizons could not have been explained by this

¹⁶Of course, there could also be another reason why the β_h 's decline with the forecasting horizon. The difficulty of forecasting for longer horizons tends to lower the β_h 's for longer horizons in standard Mincer-Zarnowitz regressions. However, if the β_h 's do not decrease as h increases, this can be interpreted as unimportance of the transmission channel even if there is another reason why the β_h 's should decrease.

Table 10: Results of regression of slope coefficients on forecast horizon

β_h 's from	b	
(3)	0.08	(0.21)
(3), $\alpha_h = 0$	0.06	(0.41)
(4)	0.02	(0.35)

Note: Standard errors are in parentheses.

channel.

As endogeneity of outturns is the only explanation we could have offered for a possibly adverse information content of risk forecasts, the evidence for adverse information content found above remains puzzling.

5.2 Problems to Anticipate Risks in Determinants

Risks to inflation or other aggregates are commonly identified via risks to variables that determine these aggregates. For example, an upward risk to inflation might be caused by an upward risk to oil prices, to the value added tax rate or by a risk of depreciation of the domestic currency. Thus, in order to correctly forecast the risks to inflation, one has to forecast the risks to these determinants. Actually, the process of risk forecasting might be thought of as a three-step process, where in the first step, one has to identify those determinants which are subject to forecast risks. In the second step, one has to quantify these risks, and in the third step, their impact on the aggregate has to be calculated.

While for the third step, models are in general available, the first two steps appear especially demanding. The first step requires to identify variables whose most likely future paths (represented by the mode forecast) differ from their expected future paths (represented by the mean forecast). This might be possible for fiscal variables like the value added tax rate, where one could imagine that a certain rate is likely, but that an alternative rate is discussed by the government at the time the forecast is made. For variables like oil prices and exchange rates, however, this task is extremely challenging. Even if the identification of risks is successful, the quantification appears equally difficult. But if the identified risks to the determinants are opposing risks to the aggregate, an incorrect quantification can easily lead to a qualitatively flawed risk forecast for the aggregate.

Without further knowledge about the assumptions made about risks of determinants in the BoE's forecasts, we cannot verify whether the difficulties in identifying and quantifying these risks are decisive for the performance of the risk forecasts for inflation. However, the BoE gives narrative support to this supposition. In the Inflation Report from August 2000 the BoE, referring to inflationary developments from 1997 to 1999, states that "In general, the modal inflation forecast has been closer to actual outturns than the mean projection. This is because the MPC judged the risks to the central projection to be on the upside, largely because of the risk that the sterling exchange rate might depreciate sharply. Up to 2000 Q2 this did not occur; indeed, the exchange rate tended to be higher than the central assumption." (pp. 63-64).

Although we cannot quantify the effect of errors in assumptions about risks in determinants, the reasoning given above and the cited statement by the BoE lead us to consider this effect as the major reason for the BoE's poor risk forecasting performance. If this is the major reason, then the appearance of an adverse information content for some of the forecast horizons would be related to the sample size of the analysis. With a larger sample size, one would then expect stronger evidence for the hypothesis of no information content of the risk forecasts.

6 Conclusion

Macroeconomic risk assessments play an important role in the forecasts of many institutions. However, to the best of our knowledge their performance has not been investigated yet. In this work, we study the BoE's risk forecasts for inflation. We find that these forecasts do not contain the intended information. Rather, they either have no information content, or maybe even an adverse information content. The poor performance of the risk forecasts is related to the fact that the mode forecasts have smaller RMSEs than the mean forecasts. Our results imply that economic agents with a loss function being quadratic in the forecast error should not use the mean forecasts. Instead, in order to attain the smallest expected loss, they should either use the mode forecasts and ignore the risk forecast. Or they even should, starting from the mode forecast, move into the opposite direction of the risk forecast. In any case, they can expect a lower loss when they use the

BoE's mode forecast instead of the mean forecast.

We find that the poor risk forecasting performance cannot be explained by outliers. We also find that in 2001 or 2002, the BoE changed its risk forecast pattern, and has apparently acted with more caution since then. Risks have only been forecasted relatively rarely since this change. However, we find that the increase in caution does not translate into a qualitatively correct information content of the risk forecasts.

The most convincing reason for the BoE's poor risk forecasting performance seems to be given by the difficulty to identify and quantify the forecast risks in the determinants of inflation. This, however, can only explain the lack of information content. We cannot offer conclusive explanations for a possibly adverse information content.

If our results are representative for macroeconomic risk forecasts, they call into question the common practice of adding risk assessments to forecasts. The only obvious rationale for publishing risk forecasts would then be the forecast institutions' desire for shaping the expectations of economic agents. However, this can only work properly as long as economic agents believe that the risk forecasts have the intended information content.

Yet, it might also be that risk forecasts are actually neither meant to deliver correct results nor to shape expectations, but rather aim to reach a "consensus" among those responsible for the forecast. By means of risk forecasts, it is possible to integrate disagreeing views on future developments into a single forecast. If a minority of an institution's forecasting committee does not agree with the view of the majority, the view of the majority could be published as the central projection, while the view of the minority could be represented in the form of risk assessments. In this case, our results imply that, for the sake of forecast optimality, it would be better to ignore the minority's view.

References

- Andrews, D. W. K. (1991), ‘Heteroskedasticity and Autocorrelation Consistent Covariance Matrix Estimation’, *Econometrica* **59**(3), 817–858.
- Britton, E., Fisher, P. & Whitley, J. (1998), ‘The inflation report projections: Understanding the fan chart’, *Bank of England Quarterly Bulletin* pp. 30–37.
- Brown, B. W. & Maital, S. (1981), ‘What do Economists Know? An Empirical Study of Experts’ Expectations’, *Econometrica* **49**(2), 491–504.
- Christoffersen, P. F. (1998), ‘Evaluating interval forecasts’, *International Economic Review* **39**(4), 841–862.
- Clements, M. P. (2004), ‘Evaluating the Bank of England density forecasts of inflation’, *The Economic Journal*, **114**, 844–866.
- Diebold, F. X. & Lopez, J. A. (1996), Forecast Evaluation and Combination, in G. Maddala & C. Rao, eds, ‘Forecast Evaluation and Combination’, North-Holland, Amsterdam, pp. 241–268.
- Diebold, F. X. & Mariano, R. S. (1995), ‘Comparing Predictive Accuracy’, *Journal of Business and Economic Statistics* **13**, 253–265.
- Dowd, K. (2007), ‘Too good to be true? The (In)credibility of the UK inflation fan charts’, *Journal of Macroeconomics* **29**, 91–102.
- Elder, R., Kapetanios, G., Taylor, T. & Yates, T. (2005), ‘Assessing the MPC’s fan charts’, *Bank of England Quarterly Bulletin* pp. 326–348.
- Harding, D. & Pagan, A. (2006), ‘Synchronization of cycles’, *Journal of Econometrics* **132**, 59–79.
- Newey, W. K. & West, K. D. (1987), ‘A Simple, Positive Semi-Definite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix’, *Econometrica* **55**(3), 703–708.
- Pesaran, M. H. & Timmermann, A. (2006), ‘Testing dependence among serially correlated multi-category variables’, *IZA Discussion Paper Series* (2196).
- Wallis, K. F. (1989), ‘Macroeconomic Forecasting: A Survey’, *The Economic Journal* **99**, 28–61.
- Wallis, K. F. (1999), ‘Asymmetric density forecasts of inflation and the Bank of England’s fan chart’, *National Institute Economic Review* **167**, 106–112.

- Wallis, K. F. (2003), ‘Chi-squared tests of interval and density forecasts, and the Bank of England’s fan charts’, *International Journal of Forecasting* **19**, 165–175.
- Wallis, K. F. (2004), ‘An Assessment of Bank of England and National Institute Inflation Forecast Uncertainties’, *National Institute Economic Review* **189**, 64–71.
- West, K. D. (2006), Forecast Evaluation, *in* G. Elliott, C. Granger & A. Timmerman, eds, ‘Handbook of Economic Forecasting’, Elsevier, Amsterdam, pp. 100–134.

A Appendix: Mode and Mean Forecasts and Realizations

date		modes of inflation forecasts									realizations	
		$h=0$	$h=1$	$h=2$	$h=3$	$h=4$	$h=5$	$h=6$	$h=7$	$h=8$	RPIX	CPI
1998	Feb	2.60	2.63	2.43	2.43	2.45	2.44	2.57	2.70	2.87	2.59	
	May	2.83	2.35	2.35	2.40	2.37	2.31	2.29	2.33	2.44	2.94	
	Aug	2.51	2.56	2.68	2.82	2.86	2.77	2.68	2.57	2.52	2.55	
	Nov	2.54	2.54	2.69	2.73	2.62	2.65	2.67	2.73	2.79	2.53	
1999	Feb	2.47	2.52	2.55	2.60	2.55	2.58	2.62	2.66	2.92	2.53	
	May	2.48	2.41	2.37	2.22	2.25	2.27	2.24	2.37	2.42	2.30	
	Aug	2.31	2.27	2.10	1.98	1.85	1.85	1.97	2.08	2.29	2.17	
	Nov	2.20	2.13	2.06	2.02	1.83	1.72	1.78	2.12	2.41	2.16	
2000	Feb	1.93	1.97	1.93	2.01	2.28	2.40	2.40	2.37	2.28	2.09	
	May	1.88	1.93	2.09	2.19	2.44	2.47	2.44	2.39	2.39	2.07	
	Aug	2.38	2.28	2.27	2.39	2.48	2.61	2.65	2.66	2.66	2.13	
	Nov	2.36	2.33	2.22	2.19	2.19	2.20	2.40	2.51	2.62	2.11	
2001	Feb	1.94	1.93	1.89	1.90	2.11	2.21	2.33	2.52	2.73	1.87	
	May	1.90	1.90	1.92	1.91	1.95	2.05	2.19	2.45	2.62	2.26	
	Aug	2.31	2.17	2.17	1.91	1.96	2.11	2.28	2.34	2.34	2.38	
	Nov	2.00	2.03	1.85	2.06	2.05	2.10	2.15	2.21	2.31	1.95	
2002	Feb	2.14	1.87	1.95	2.09	2.16	2.13	2.15	2.21	2.16	2.37	
	May	2.02	2.08	2.22	2.19	2.01	2.00	2.10	2.10	2.23	1.86	
	Aug	1.84	2.25	2.25	2.26	2.30	2.27	2.31	2.38	2.47	1.98	
	Nov	2.64	2.71	2.71	2.72	2.42	2.43	2.44	2.53	2.56	2.61	
2003	Feb	2.78	2.90	2.99	2.79	2.69	2.62	2.57	2.46	2.50	2.89	
	May	3.09	2.90	2.64	2.39	2.35	2.40	2.41	2.52	2.55	2.91	
	Aug	2.85	2.58	2.31	2.32	2.29	2.28	2.36	2.48	2.61	2.85	
	Nov	2.72	2.55	2.63	2.65	2.59	2.50	2.37	2.38	2.42	2.60	
2004	Feb	1.34	1.60	1.60	1.71	1.77	1.68	1.71	1.76	1.87	2.30	1.25
	May	1.46	1.36	1.39	1.43	1.60	1.71	1.76	1.98	2.13	2.17	1.45
	Aug	1.18	1.22	1.32	1.38	1.53	1.53	1.64	1.96	2.01	2.11	1.24
	Nov	1.18	1.20	1.36	1.57	1.65	1.73	1.84	1.95	2.03	2.26	1.44
2005	Feb	1.54	1.58	1.65	1.73	1.80	1.90	2.00	2.09	2.15	2.20	1.75
	May	1.98	2.06	2.10	2.01	1.86	1.90	1.93	1.96	2.00	2.23	1.94
	Aug	2.16	2.30	2.19	2.01	1.91	1.81	1.94	2.13	2.18	2.39	2.45
	Nov	2.23	2.16	2.03	1.85	1.77	1.81	1.84	1.89	1.95	2.26	2.13
2006	Feb	1.96	1.94	1.97	2.00	2.03	2.04	2.05	2.04	2.03		1.92
	May	2.27	2.19	2.33	2.35	2.12	2.10	2.09	2.06	2.01		2.20
	Aug	2.32	2.71	2.76	2.56	2.37	2.19	2.20	2.11	2.05		2.39
	Nov	2.56	2.68	2.36	2.06	2.03	2.00	1.98	1.97	1.98		2.78
2007	Feb	2.90	2.56	2.24	2.02	1.79	1.83	1.88	1.93	2.00		2.88
	May	2.52	2.25	2.06	1.95	1.83	1.79	1.88	1.92	1.98		2.64
	Aug	2.07	2.10	2.10	2.05	2.03	2.04	1.99	1.99	2.00		1.75

Note: Until 2003 RPIX forecasts, from 2004 CPI forecasts

date		means of inflation forecasts									realizations	
		$h=0$	$h=1$	$h=2$	$h=3$	$h=4$	$h=5$	$h=6$	$h=7$	$h=8$	RPIX	CPI
1998	Feb	2.64	2.69	2.51	2.52	2.54	2.60	2.80	3.00	3.25	2.59	
	May	2.74	2.20	2.16	2.18	2.15	2.22	2.33	2.50	2.74	2.94	
	Aug	2.56	2.66	2.81	2.96	3.00	2.94	2.89	2.81	2.79	2.55	
	Nov	2.57	2.62	2.79	2.85	2.75	2.73	2.69	2.70	2.70	2.53	
1999	Feb	2.49	2.56	2.60	2.65	2.61	2.60	2.60	2.59	2.81	2.53	
	May	2.51	2.46	2.45	2.31	2.36	2.40	2.39	2.54	2.61	2.30	
	Aug	2.35	2.35	2.21	2.12	2.00	1.98	2.07	2.16	2.35	2.17	
	Nov	2.19	2.10	2.03	1.98	1.78	1.64	1.63	1.91	2.17	2.16	
2000	Feb	1.96	2.02	2.00	2.10	2.38	2.51	2.52	2.51	2.43	2.09	
	May	1.89	1.95	2.12	2.24	2.49	2.50	2.44	2.36	2.34	2.07	
	Aug	2.38	2.28	2.27	2.39	2.48	2.59	2.60	2.58	2.56	2.13	
	Nov	2.37	2.35	2.26	2.23	2.24	2.24	2.42	2.52	2.62	2.11	
2001	Feb	1.92	1.90	1.85	1.85	2.06	2.12	2.18	2.31	2.48	1.87	
	May	1.88	1.87	1.88	1.87	1.90	1.98	2.09	2.32	2.47	2.26	
	Aug	2.31	2.17	2.17	1.91	1.96	2.08	2.20	2.22	2.19	2.38	
	Nov	2.10	2.17	2.02	2.25	2.25	2.30	2.35	2.41	2.51	1.95	
2002	Feb	2.24	2.01	2.12	2.28	2.36	2.37	2.45	2.57	2.56	2.37	
	May	2.02	2.08	2.22	2.19	2.01	2.05	2.22	2.30	2.48	1.86	
	Aug	1.84	2.25	2.25	2.26	2.30	2.30	2.39	2.50	2.62	1.98	
	Nov	2.64	2.71	2.71	2.72	2.42	2.46	2.51	2.65	2.71	2.61	
2003	Feb	2.78	2.90	2.99	2.79	2.69	2.65	2.64	2.58	2.65	2.89	
	May	3.14	2.97	2.72	2.49	2.45	2.48	2.46	2.54	2.55	2.91	
	Aug	2.85	2.58	2.31	2.32	2.29	2.26	2.31	2.40	2.51	2.85	
	Nov	2.72	2.55	2.63	2.65	2.59	2.50	2.37	2.38	2.42	2.60	
2004	Feb	1.34	1.60	1.60	1.71	1.77	1.68	1.71	1.76	1.87	2.30	1.25
	May	1.46	1.36	1.39	1.43	1.60	1.71	1.76	1.98	2.13	2.17	1.45
	Aug	1.18	1.22	1.32	1.38	1.53	1.53	1.64	1.96	2.01	2.11	1.24
	Nov	1.13	1.13	1.28	1.47	1.55	1.61	1.69	1.77	1.83	2.26	1.44
2005	Feb	1.49	1.51	1.56	1.63	1.70	1.78	1.85	1.91	1.95	2.20	1.75
	May	1.98	2.06	2.10	2.01	1.86	1.90	1.93	1.96	2.00	2.23	1.94
	Aug	2.16	2.30	2.19	2.01	1.91	1.79	1.89	2.05	2.08	2.39	2.45
	Nov	2.23	2.16	2.03	1.85	1.77	1.81	1.84	1.89	1.95	2.26	2.13
2006	Feb	1.96	1.94	1.97	2.00	2.03	2.04	2.05	2.04	2.03		1.92
	May	2.27	2.19	2.33	2.35	2.12	2.10	2.09	2.06	2.01		2.20
	Aug	2.32	2.71	2.76	2.56	2.37	2.19	2.20	2.11	2.05		2.39
	Nov	2.56	2.68	2.36	2.06	2.03	2.00	1.98	1.97	1.98		2.78
2007	Feb	2.85	2.49	2.16	1.92	1.69	1.79	1.94	2.08	2.20		2.88
	May	2.52	2.25	2.06	1.95	1.83	1.84	1.99	2.08	2.18		2.64
	Aug	2.07	2.10	2.10	2.05	2.03	2.06	2.05	2.07	2.10		1.75

Note: Until 2003 RPIX forecasts, from 2004 CPI forecasts

The following Discussion Papers have been published since 2007:

Series 1: Economic Studies

01	2007	The effect of FDI on job separation	Sascha O. Becker Marc-Andreas Mündler
02	2007	Threshold dynamics of short-term interest rates: empirical evidence and implications for the term structure	Theofanis Archontakis Wolfgang Lemke
03	2007	Price setting in the euro area: some stylised facts from individual producer price data	Dias, Dossche, Gautier Hernando, Sabbatini Stahl, Vermeulen
04	2007	Unemployment and employment protection in a unionized economy with search frictions	Nikolai Stähler
05	2007	End-user order flow and exchange rate dynamics	S. Reitz, M. A. Schmidt M. P. Taylor
06	2007	Money-based interest rate rules: lessons from German data	C. Gerberding F. Seitz, A. Worms
07	2007	Moral hazard and bail-out in fiscal federations: evidence for the German Länder	Kirsten H. Heppke-Falk Guntram B. Wolff
08	2007	An assessment of the trends in international price competitiveness among EMU countries	Christoph Fischer
09	2007	Reconsidering the role of monetary indicators for euro area inflation from a Bayesian perspective using group inclusion probabilities	Michael Scharnagl Christian Schumacher
10	2007	A note on the coefficient of determination in regression models with infinite-variance variables	Jeong-Ryeol Kurz-Kim Mico Loretan

11	2007	Exchange rate dynamics in a target zone - a heterogeneous expectations approach	Christian Bauer Paul De Grauwe, Stefan Reitz
12	2007	Money and housing - evidence for the euro area and the US	Claus Greiber Ralph Setzer
13	2007	An affine macro-finance term structure model for the euro area	Wolfgang Lemke
14	2007	Does anticipation of government spending matter? Evidence from an expectation augmented VAR	Jörn Tenhofen Guntram B. Wolff
15	2007	On-the-job search and the cyclical dynamics of the labor market	Michael Krause Thomas Lubik
16	2007	Heterogeneous expectations, learning and European inflation dynamics	Anke Weber
17	2007	Does intra-firm bargaining matter for business cycle dynamics?	Michael Krause Thomas Lubik
18	2007	Uncertainty about perceived inflation target and monetary policy	Kosuke Aoki Takeshi Kimura
19	2007	The rationality and reliability of expectations reported by British households: micro evidence from the British household panel survey	James Mitchell Martin Weale
20	2007	Money in monetary policy design under uncertainty: the Two-Pillar Phillips Curve versus ECB-style cross-checking	Günter W. Beck Volker Wieland
21	2007	Corporate marginal tax rate, tax loss carryforwards and investment functions – empirical analysis using a large German panel data set	Fred Ramb

22	2007	Volatile multinationals? Evidence from the labor demand of German firms	Claudia M. Buch Alexander Lipponer
23	2007	International investment positions and exchange rate dynamics: a dynamic panel analysis	Michael Binder Christian J. Offermanns
24	2007	Testing for contemporary fiscal policy discretion with real time data	Ulf von Kalckreuth Guntram B. Wolff
25	2007	Quantifying risk and uncertainty in macroeconomic forecasts	Malte Knüppel Karl-Heinz Tödter
26	2007	Taxing deficits to restrain government spending and foster capital accumulation	Nikolai Stähler
27	2007	Spill-over effects of monetary policy – a progress report on interest rate convergence in Europe	Michael Flad
28	2007	The timing and magnitude of exchange rate overshooting	Hoffmann Sondergaard, Westelius
29	2007	The timeless perspective vs. discretion: theory and monetary policy implications for an open economy	Alfred V. Guender
30	2007	International cooperation on innovation: empirical evidence for German and Portuguese firms	Pedro Faria Tobias Schmidt
31	2007	Simple interest rate rules with a role for money	M. Scharnagl C. Gerberding, F. Seitz
32	2007	Does Benford's law hold in economic research and forecasting?	Stefan Günnel Karl-Heinz Tödter
33	2007	The welfare effects of inflation: a cost-benefit perspective	Karl-Heinz Tödter Bernhard Manzke

34	2007	Factor-MIDAS for now- and forecasting with ragged-edge data: a model comparison for German GDP	Massimiliano Marcellino Christian Schumacher
35	2007	Monetary policy and core inflation	Michele Lenza
01	2008	Can capacity constraints explain asymmetries of the business cycle?	Malte Knüppel
02	2008	Communication, decision-making and the optimal degree of transparency of monetary policy committees	Anke Weber
03	2008	The impact of thin-capitalization rules on multinationals' financing and investment decisions	Buettner, Overesch Schreiber, Wamser
04	2008	Comparing the DSGE model with the factor model: an out-of-sample forecasting experiment	Mu-Chun Wang
05	2008	Financial markets and the current account – emerging Europe versus emerging Asia	Sabine Herrmann Adalbert Winkler
06	2008	The German sub-national government bond market: evolution, yields and liquidity	Alexander Schulz Guntram B. Wolff
07	2008	Integration of financial markets and national price levels: the role of exchange rate volatility	Mathias Hoffmann Peter Tillmann
08	2008	Business cycle evidence on firm entry	Vivien Lewis
09	2008	Panel estimation of state dependent adjustment when the target is unobserved	Ulf von Kalckreuth
10	2008	Nonlinear oil price dynamics – a tale of heterogeneous speculators?	Stefan Reitz Ulf Slopek

11	2008	Financing constraints, firm level adjustment of capital and aggregate implications	Ulf von Kalckreuth
12	2008	Sovereign bond market integration: the euro, trading platforms and globalization	Alexander Schulz Guntram B. Wolff
13	2008	Great moderation at the firm level? Unconditional versus conditional output volatility	Claudia M. Buch Jörg Döpke Kerstin Stahn
14	2008	How informative are macroeconomic risk forecasts? An examination of the Bank of England's inflation forecasts	Malte Knüppel Guido Schulte Frankenfeld

Series 2: Banking and Financial Studies

01	2007	Granularity adjustment for Basel II	Michael B. Gordy Eva Lütkebohmert
02	2007	Efficient, profitable and safe banking: an oxymoron? Evidence from a panel VAR approach	Michael Koetter Daniel Porath
03	2007	Slippery slopes of stress: ordered failure events in German banking	Thomas Kick Michael Koetter
04	2007	Open-end real estate funds in Germany – genesis and crisis	C. E. Bannier F. Fecht, M. Tyrell
05	2007	Diversification and the banks’ risk-return-characteristics – evidence from loan portfolios of German banks	A. Behr, A. Kamp C. Memmel, A. Pfingsten
06	2007	How do banks adjust their capital ratios? Evidence from Germany	Christoph Memmel Peter Raupach
07	2007	Modelling dynamic portfolio risk using risk drivers of elliptical processes	Rafael Schmidt Christian Schmieder
08	2007	Time-varying contributions by the corporate bond and CDS markets to credit risk price discovery	Niko Dötz
09	2007	Banking consolidation and small business finance – empirical evidence for Germany	K. Marsch, C. Schmieder K. Forster-van Aerssen
10	2007	The quality of banking and regional growth	Hasan, Koetter, Wedow
11	2007	Welfare effects of financial integration	Fecht, Grüner, Hartmann
12	2007	The marketability of bank assets and managerial rents: implications for financial stability	Falko Fecht Wolf Wagner

13	2007	Asset correlations and credit portfolio risk – an empirical analysis	K. Düllmann, M. Scheicher C. Schmieder
14	2007	Relationship lending – empirical evidence for Germany	C. Memmel C. Schmieder, I. Stein
15	2007	Creditor concentration: an empirical investigation	S. Ongena, G. Tümer-Alkan N. von Westernhagen
16	2007	Endogenous credit derivatives and bank behaviour	Thilo Pausch
17	2007	Profitability of Western European banking systems: panel evidence on structural and cyclical determinants	Rainer Beckmann
18	2007	Estimating probabilities of default with support vector machines	W. K. Härdle R. A. Moro, D. Schäfer
01	2008	Analyzing the interest rate risk of banks using time series of accounting-based data: evidence from Germany	O. Entrop, C. Memmel M. Wilkens, A. Zeisler
02	2008	Bank mergers and the dynamics of deposit interest rates	Ben R. Craig Valeriya Dinger
03	2008	Monetary policy and bank distress: an integrated micro-macro approach	F. de Graeve T. Kick, M. Koetter
04	2008	Estimating asset correlations from stock prices or default rates – which method is superior?	K. Düllmann J. Küll, M. Kunisch
05	2008	Rollover risk in commercial paper markets and firms' debt maturity choice	Felix Thierfelder
06	2008	The success of bank mergers revisited – an assessment based on a matching strategy	Andreas Behr Frank Heid

07	2008	Which interest rate scenario is the worst one for a bank? Evidence from a tracking bank approach for German savings and cooperative banks	Christoph Memmel
08	2008	Market conditions, default risk and credit spreads	Dragon Yongjun Tang Hong Yan
09	2008	The pricing of correlated default risk: evidence from the credit derivatives market	Nikola Tarashev Haibin Zhu
10	2008	Determinants of European banks' engagement in loan securitization	Christina E. Bannier Dennis N. Hänsel
11	2008	Interaction of market and credit risk: an analysis of inter-risk correlation and risk aggregation	Klaus Böcker Martin Hillebrand
12	2008	A value at risk analysis of credit default swaps	B. Raunig, M. Scheicher
13	2008	Systemic bank risk in Brazil: an assessment of correlated market, credit, sovereign and inter-bank risk in an environment with stochastic volatilities and correlations	Theodore M. Barnhill, Jr. Marcos Rietti Souto
14	2008	Regulatory capital for market and credit risk interaction: is current regulation always conservative?	T. Breuer, M. Jandačka K. Rheinberger, M. Summer
15	2008	The implications of latent technology regimes for competition and efficiency in banking	Michael Koetter Tigran Poghosyan
16	2008	The impact of downward rating momentum on credit portfolio risk	André Güttler Peter Raupach

Visiting researcher at the Deutsche Bundesbank

The Deutsche Bundesbank in Frankfurt is looking for a visiting researcher. Among others under certain conditions visiting researchers have access to a wide range of data in the Bundesbank. They include micro data on firms and banks not available in the public. Visitors should prepare a research project during their stay at the Bundesbank. Candidates must hold a PhD and be engaged in the field of either macroeconomics and monetary economics, financial markets or international economics. Proposed research projects should be from these fields. The visiting term will be from 3 to 6 months. Salary is commensurate with experience.

Applicants are requested to send a CV, copies of recent papers, letters of reference and a proposal for a research project to:

Deutsche Bundesbank
Personalabteilung
Wilhelm-Epstein-Str. 14

60431 Frankfurt
GERMANY

