SVERIGES RIKSBANK
WORKING PAPER SERIES

328

Adjusting for Information Content when Comparing Forecast Performance

Michael K. Andersson, Ted Aranki and André Reslow

August 2016

WORKING PAPERS ARE OBTAINABLE FROM

www.riksbank.se/en/research
Sveriges Riksbank • SE-103 37 Stockholm
Fax international: +46 8 21 05 31
Telephone international: +46 8 787 00 00

Adjusting for Information Content when Comparing Forecast Performance*

Michael K. Andersson† Ted Araranki‡ André Reslow§

August 2016

Sveriges Riksbank Working Paper Series

No. 328

Abstract

Cross institutional forecast evaluations may be severely distorted by the fact that forecasts are made at different points in time, and thus with different amount of information. This paper proposes a method to account for these differences. The method computes the timing effect and the forecaster's ability simultaneously. Monte Carlo simulation demonstrate that evaluations that do not adjust for the differences in information content may be misleading. In addition, the method is applied on a real-world data set of 10 Swedish forecasters for the period 1999–2015. The results show that the ranking of the forecasters is affected by the proposed adjustment.

Keywords: Forecast error; Forecast comparison; Publication time; Evaluation; Error component model; Panel data

JEL classification: C23; C53; E37

*The authors thank seminar participants at Sveriges Riksbank and the National Institute of Economic Research. The views expressed in this paper are solely the responsibility of the authors and should not be interpreted as reflecting the views of Finansinspektionen or Sveriges Riksbank.

†Finansinspektionen, Sweden. E-mail: michael.andersson@fi.se

‡Sveriges Riksbank, Sweden. E-mail: ted.aranki@riksbank.se

§Sveriges Riksbank, Sweden. E-mail: andre.reslow@riksbank.se (corresponding author)

1 Introduction

Many agents in the economy, including economic policy makers, regularly publish forecasts of the development of the economy. Since important economic and political decisions usually are based on forecasts it is crucial that these predictions are as accurate as possible. Therefore, evaluations are regularly undertaken to assess the performance. Usually evaluations compare the forecast at a specific point in time with the outcome, or using a statistical measure based on the two quantities. One such measure is the mean absolute errors (MAE), see for example Diebold (2007) and Gneiting (2011). Albeit informative, such evaluations are in general insufficient. A large forecast error, for a particular observation, can be a consequence of a shock that could not have been foreseen. Comparing different agencies is a common way to handle this problem. For example, central banks, international organizations and other institutions regularly publish evaluations of their forecasts and compare them with other forecasters (see e.g., Sveriges Riksbank (2016); Bank of England (2015); European Central Bank (2013); Vogel (2007); Timmermann (2007)). More in-depth analyses of forecasts are presented at a less frequent basis. For instance Blix et al. (2001) compare different forecaster's performance of key Swedish macroeconomic variables. Andersson et al. (2007) estimate the Riksbank's accuracy in relation to that of the National Institute of Economic Research and simple econometric specifications. Davies and Lahiri (1995) and Bauer et al. (2003) compare the agents of the Blue Chip Survey and Boero et al. (2008) assess the Bank of England Survey of External Forecasters. Goh and Lawrence (2006) compares the accuracy of New Zealand forecaster's RMSE and their average relative rank. Cabanillas and Terzi (2012) present an assessment of the European Commission's track record and compares the forecast errors of GDP growth to those of other international institutions.

Comparing forecasters is appealing but not necessary straightforward. One problem with comparisons is that forecasts are published at different points in time. In practice, this implies that different forecasters have different amount of information when they prepare their forecasts. To make a fair comparison, this must be accounted for. One attempt to reduce this problem is to compare forecasts produced at almost the same point in time. This approach is far from flawless since one month of information may be important, e.g. if a quarter of National Account figure is published that particular month. Another problem is that this approach utilizes only a subset of the available forecasts. We argue that a fair comparison between forecasters requires that the differences in information must be adjusted for in a quantitative way.

The contribution of this paper is that we propose a method to account for the fact that forecasts are made at different points in time, and thus with different amount of information, when comparing the accuracy. We represent the amount of information by timing, that is, the distance between the publication date and outcome date. In the model, the timing effect and the respective forecasters underlying performance (ability) are estimated simultaneously. The model rests on the assumption that part of a forecaster's average error can be explained by the timing. To our knowledge this paper is the first study that, in a quantitative way, takes into account the differences in timing when comparing the performance of forecasters. A forecaster that publishes its forecasts later than others can also be expected, ceteris paribus, to have a slightly better accuracy. A Monte Carlo simulation shows that this is the case and that the proposed model can correct for this. In an empirical application, the model is used to assess the forecasting performance of Swedish forecasters with respect to GDP growth and CPI inflation for the period 1999–2015. The ranking of the forecasters depends on weather information differences are accounted for or not. Overall, we find the model well-suited for comparing forecast performance.

The paper unfolds as follows: Section 2 presents the empirical framework and the econometric model and Section 3 holds a Monte Carlo simulation. Section 4 presents the empirical application and Section 5 concludes the paper.

2 Empirical Framework

This section discusses the effect of the information content, measured by time to outcome, and explains why this is an important determinant of the size of the forecast error. The forecast error is expected to become smaller the nearer (in time) to the outcome the prediction is made. We then present an econometric model that adjusts for differences in timing when estimating the ability of the different forecasters.

2.1 The Forecast Error

Let $\hat{y}_{it,h}$ be the forecast for the growth rate of variable y for year t, made by individual i, h months prior to the end of year t (the period the outcome is known). The error of a particular forecast is then defined by:

$$\nu_{it,h} = y_t - \hat{y}_{it,h} \tag{1}$$

This error is expected to decline as the horizon shrinks, because the information set grows over time. For a forecast of a variable expressed as an annual average, the errors decrease at a non-linear rate as quarterly or monthly outcomes of the particular prediction year become available. We exemplify this with quarterly data, by looking at how annual and quarterly growth rates are related. The annual change of a variable is a function of the quarterly changes during the last four quarters according to:[1]

$$\frac{y_{T,1} - y_{T-1,1}}{y_{T-1,1}} = \frac{y_{T,1}}{y_{T-1,1}} - 1 = \frac{y_{T,1}}{y_{T-1,4}} * \frac{y_{T-1,4}}{y_{T-1,3}} * \frac{y_{T-1,3}}{y_{T-1,2}} * \frac{y_{T-1,2}}{y_{T-1,1}} - 1 \tag{2}$$

[1] Similar calculations can easily be generalized to an arbitrary frequency.

where $y_{T,q}$ is the level of variable y in quarter q of year T. Let $g^4_{T,q}$ and $g^1_{T,q}$ be the annual and quarterly change for quarter q of year T. Equation (2) can then be expressed as:

$$1 + g^4_{T,1} = (1 + g^1_{T,1}) * (1 + g^1_{T-1,4}) * (1 + g^1_{T-1,3}) * (1 + g^1_{T-1,2}) \tag{3}$$

The annual change of quarters 2, 3 and 4 are calculated as equation (3). Thus, the yearly (average) growth for year T, $\Delta_4 y_T$, can be stated as the average of the annual change in the year's quarters:

$$\begin{aligned}\Delta_4 y_T = \frac{1}{4}\Big[&(1 + g^1_{T,1}) * (1 + g^1_{T-1,4}) * (1 + g^1_{T-1,3}) * (1 + g^1_{T-1,2}) + \\ &(1 + g^1_{T,2}) * (1 + g^1_{T,1}) * (1 + g^1_{T-1,4}) * (1 + g^1_{T-1,3}) + \\ &(1 + g^1_{T,3}) * (1 + g^1_{T,2}) * (1 + g^1_{T,1}) * (1 + g^1_{T-1,4}) + \\ &(1 + g^1_{T,4}) * (1 + g^1_{T,3}) * (1 + g^1_{T,2}) * (1 + g^1_{T,1})\Big] - 1\end{aligned} \tag{4}$$

Equation (4) shows that the annual growth rate of year T is a function of all quarterly growth rates from the second quarter of the previous year, $T - 1$, until the last quarter of the forecast year, T. The equation also indicates that the quarters weigh differently in terms of yearly growth. Table 1 present the weights.

Table 1: The relative importance of each quarter in annual average calculation

	Year: $T-1$				Year: T			
Quarter	1	2	3	4	1	2	3	4
Weight	0/16	1/16	2/16	3/16	4/16	3/16	2/16	1/16
Accumulated weight	0/16	1/16	3/16	6/16	10/16	13/16	15/16	16/16

Note: The table shows the weights of each quarter in the yearly calculation, together with the proportion of the full year value (accumulated) known at each point in time. See also equation (4).

Table 1 reveals that a forecast of the annual average of year T made in the first quarter the year prior to the outcome, $T - 1$, do not include any information of the outcome. One quarter later, i.e. second quarter in year $T - 1$, 1/16 of the outcome is known. As the first quarter of the forecast year becomes available, 10/16 of the yearly outcome is known to the forecaster. Note that these calculations are stylized and do not consider recurrent revisions of quarterly outcomes, which is a common feature of macroeconomic data.[2] The practical implication of the growing information set is that less of the outcome needs to be forecasted as time proceed. We call this the *outcome effect*.

In addition to the outcome effect, there exist other reasons to why the forecast errors are expected to decrease the closer to the outcome the forecast is made. For example, the degree of persistence in

[2] GDP data are usually revised back in time when a new outcome is published. Seasonal adjustments may also lead to revision of historical GDP observations.

the data. If we assume that the quarterly growth of variable y in period t is a function of the quarterly growth in the previous period we get the following relation:

$$y_t = c + \alpha y_{t-1} + \epsilon_t \tag{5}$$

For future periods $t+1, ..., t+H$, equation (5) is a simplified description of the expected future quarterly growth rates which the forecaster can use at each forecast occasion. The parameter α (persistence of y) determines how far into the future a new outcome will remain important. We call this effect the *persistence effect*.[3]

Given the timing effects (i.e., *outcome effect* and *persistence effect*), we can calculate how the expected forecast errors, in absolute terms, will decrease as the distance to the outcome decreases. In Figure 1 we plot a calculation of the expected forecast error as a function of the timing assuming α to be 0.3 and 0.6 respectively.[4] The size of the forecast error before any information of the outcome is known, i.e. the first quarter of the year prior to the outcome year $T-1$, has been normalized to 1. Figure 1 provides a sketch showing that the functional form of the expected size of the forecast error depends on the outcome effect and the persistence effect. This study uses this functional form to account for the expected decline in the errors.

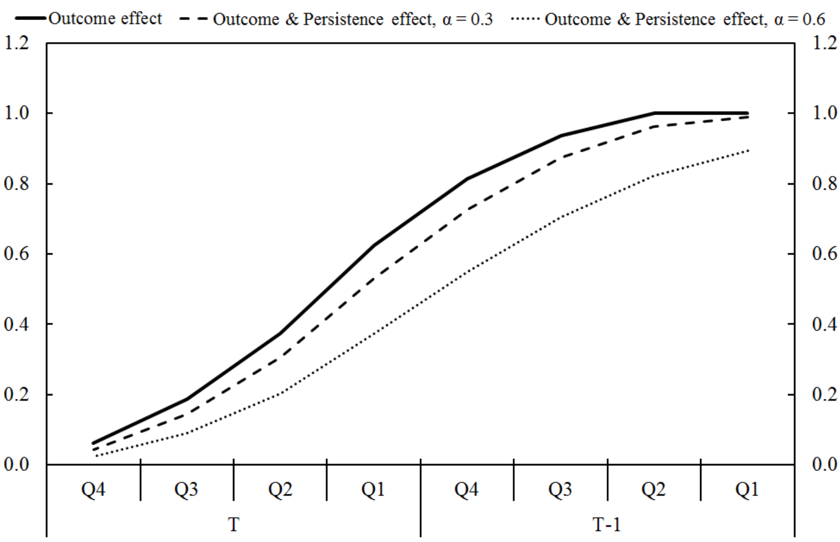

Figure 1: The impact of the timing on the size of the forecast error

Note: The black solid line indicates the extent to which the expected error depends on the outcome effect, the dashed and dotted lines indicates the extent to which the error depends on the outcome and persistence effects, where the persistence coefficient α is set to 0.3 and 0.6, respectively.

[3] Moreover, indicator series and forecasts given by other institutions may enhance the information set even more.

[4] As a reference, an estimation of α based on quarterly data for Swedish GDP during 1994–2015 gives the value 0.3.

2.2 The Econometric Model

Given the panel structure of data and following the previous section, the absolute forecast error, denoted by ϵ, can be explained by:[5]

$$\epsilon_{it,h} = \delta M_{it,h} + u_{it,h} \tag{6}$$

where $M_{it,h}$ is the timing effect, defined as $M_{it,h} = 1 - W_{it,h}$. The accumulated weight, $W_{it,h}$, increases as the parts of the prediction year become known.[6] As $W_{it,h}$ increases $M_{it,h}$ decreases and could be seen as a proxy for missing information at the time of the publication. The coefficient δ represents the marginal effect on the absolute forecast error of having less information. We model the absolute errors in equation (6) without an intercept, hence, at horizon zero the absolute forecast error is also zero since the full outcome is then known. The disturbances $u_{it,h}$ can be divided into two unobserved components and an idiosyncratic random disturbance term, according to:

$$u_{it,h} = \mu_i + \lambda_t + e_{it,h} \tag{7}$$

where the individual specific effect μ_i is interpreted as the (genuine) ability of forecaster i. This ability is assumed to be constant over time. λ_t is a time specific effect which refers to the year the forecast concerns. This effect is assumed to be common for all forecasters. $e_{it,h}$ is the idiosyncratic disturbance term. Rewriting equation (6) gives the following specification:

$$\epsilon_{it,h} = \delta M_{it,h} + \mu_i + \lambda_t + e_{it,h} \tag{8}$$

Hence, the method to adjust for information differences when comparing the forecast performance across institutions is based on the assumption that the (absolute) forecast errors can be divided up into different components: a component that is due to the amount of information available at the time of publication ($M_{it,h}$), a component that reflects the different forecasters' general forecasting performance/ability (μ_i) and a component that captures that different years can be more or less difficult to forecast (λ_t).[7]

In equation (8) we assume a common information effect for all years. Nevertheless, we will also control for time varying slope coefficients. One might argue that the fact that different years can be more or less difficult to forecast does not only affect the level of the forecast error but also the functional

[5] The method may also be adopted to other error measures. The measure should be selected on the base of the loss-function of the forecasters, see for example Patton (2015). However, comparing different evaluation measures is beyond the scope of this paper.

[6] In the estimation we calculate the weights based on monthly observations in contrast to the quarterly weights presented in Table 1. Thus, a forecast produced in December will have a different weight compared to a forecast produced in October. We also assume a publication lag of one month, which means that the December outcome becomes available in January.

[7] A somewhat similar decomposition for analyzing forecast errors in a panel data setting has been used by Davies and Lahiri (1995), Davies (2006) and Boero et al. (2008) with the difference that they do not adjust for timing.

form of the information effect (i.e., the slope). Hence, a more flexible specification becomes:

$$\epsilon_{it,h} = \delta_t M_{it,h} + \mu_i + \lambda_t + e_{it,h} \tag{9}$$

We estimate model (9) using the within (or fixed effect) estimator. The within estimator is preferred when the individual-specific factors are constant over time. Moreover, the within estimator allows for correlation between the fixed individual-specific effects and the other explanatory variables, without violating the assumption of no correlation between observable and unobservable determinants of the dependent variable. This is due to the fact that the within transformation wipes out the individual-specific effects, see Baltagi (2001). Since the individual-specific effects in this model, which are interpreted as each forecasters individual forecast ability, may be correlated with the other explanatory variables, the within estimator is appropriate.

3 Monte Carlo Simulation

3.1 Mode of Procedure and Method of Analysis

As described in previous section, the functional form of how the forecast errors decline over time is given by the *outcome effect* and the *persistence effect*. However, the sizes of these effects are unknown, but can be estimated in an econometric model such as equation (9). This model also estimates individual-specific effects that we argue represent the individual forecasting performances. In order to test if the model actually do capture the ability of the individuals, we conduct a Monte Carlo simulation. In the simulation we evaluate the model's comparison between institutions and contrast the model to a traditional mean absolute error (MAE) statistic. We use the following data generating process for the absolute forecast errors ($\epsilon_{it,h}$):

$$\epsilon_{it,h} = L_{it,h} + \mu_i + \varepsilon_{it,h} \tag{10}$$

where,

$$L_{it,h} = 1 - \Big(\sum_{j=h+1}^{24} \omega_j + \sum_{\substack{i=1,\dots,h \\ j=h,\dots,1}} \alpha^j \omega_i \Big) \tag{11}$$

$$\omega_h = \begin{cases} \frac{h}{12^2} & \text{if } h \leq 12 \\ \frac{24-h}{12^2} & \text{if } h > 12 \end{cases} \tag{12}$$

$$\varepsilon_{it,h} = |e_{it,h}| \tag{13}$$

$$e_{it,h} \sim N(0, \sigma^2) \tag{14}$$

$L_{it,h}$ represents the information effect h months before the complete outcome is known. $L_{it,h}$ is calculated as one minus the *outcome effect* and the *persistence effect.* ω_h is the monthly weights in the full outcome. In the simulations we assume that all years are equally hard to predict when generating data.

We model the ability using an individual-specific bias (μ_i) that we assign to each forecaster in the simulation. ε_{it} is an error term that is the absolute value of a normally distributed error with mean zero and variance σ^2.[8]

Furthermore, we create seven fictive forecast institutions, (A, B, $\cdots$, G). Each institute produces four forecasts each year according to a fixed publication calendar.[9] This gives eight forecasts from each institute for each outcome since we assume a maximum forecast horizon of 24 months prior to the outcome. In the set up, institute A has the shortest average horizon while institute G has the longest.[10] Thus, A has an advantage, due to most information available, at average when forecasting. Besides being affected by the forecast occasion, the errors are influenced by the individual specific biases. In the evaluation one would like the institution with the lowest bias to be ranked by the model as the best (number one). We construct two different assignments (cases) of ability:

1. $\mu_A = \mu_B = \cdots = \mu_G$: Equal ability

2. $\mu_A > \mu_B > \cdots > \mu_G$: Unequal ability, different direction to the timing effect

In the first case we expect the performance of all institutions to be equal, and in the second case we expect A to have the poorest forecast ability and G the best. In Case 2 we set the sequential differences between forecasters A–G to 0.02. Subsequently, the difference between the best and the worst forecaster is 0.12.[11]

The value of α is varied in the study and we also do two different evaluation samples. The first is a single-year evaluation ($N = 1$) and the second is based on 10 years. However, the evaluations are always computed with $S = 100,000$ Monte Carlo replicates. For each replicate, S_i, we rank the forecast performance according to the estimated ability in equation (9) and the MAE statistic. From the $100,000$ rankings we compute the proportion of each institution in each position.

Figure 2 presents the main findings from the Monte Carlo simulation. Our proposed model is labeled Adj. MAE and is compared to the MAE. The results in the figure are presented as the observed proportion of each forecaster in each ranking position.

[8] The expected mean of $\varepsilon_{it,h}$ is $\sigma\sqrt{\frac{2}{\pi}}$. In the main simulations we set σ to 0.3 but we also do the simulations using σ equal to 0.15.

[9] Publication months: A = {Mar, Jul., Oct., Dec.}, B = {Mar., Jun., Sep., Dec.}, C = {Mar., May., Sep., Nov.}, D = {Feb., May., Aug., Nov.}, E = {Feb., Apr., Aug., Oct.}, F = {Jan., Apr., Jun., Nov.}, G = {Jan., Mar., Jul., Sep.}.

[10] Average horizon in months: A=11, B=11.5, C=12, D=12.5, E=13, F=13.5, G=14.

[11] The values of μ_i and σ are inspired from Swedish GDP and CPI forecast error data.

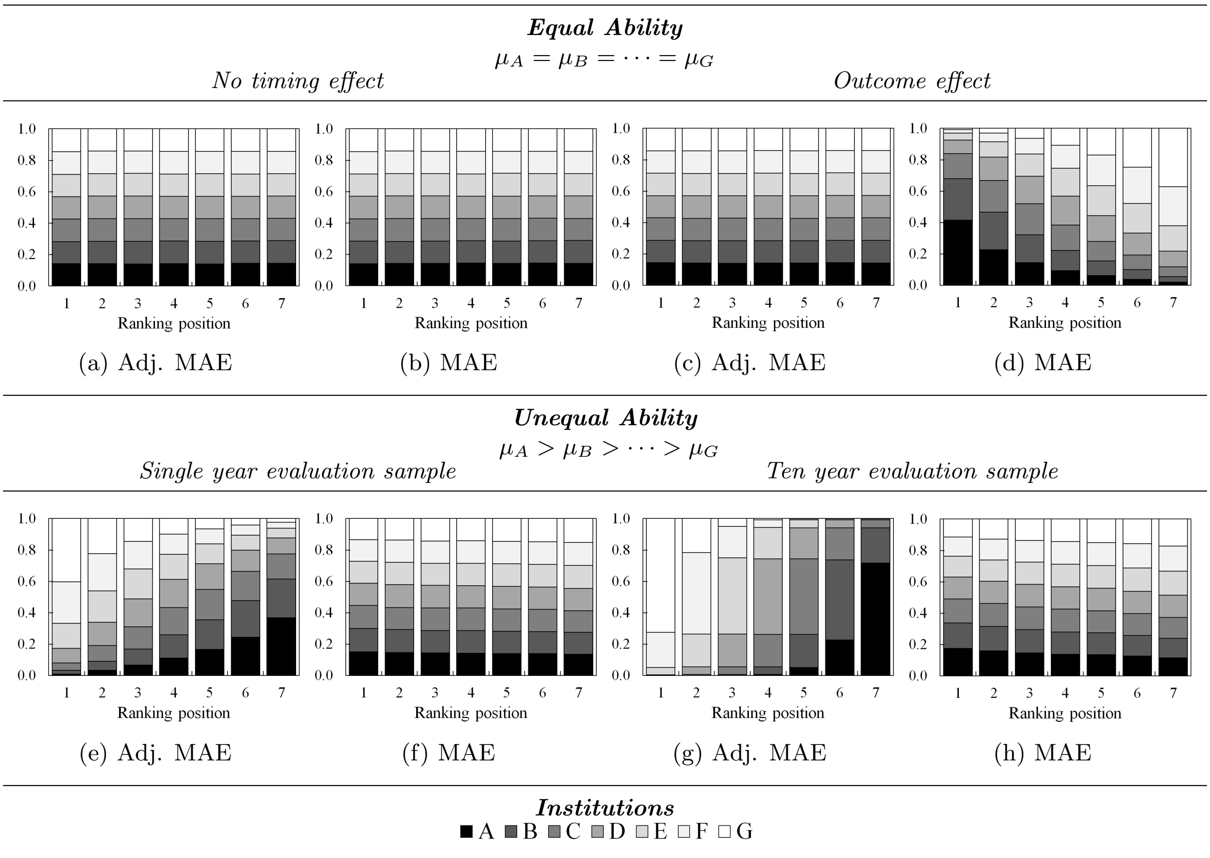

Figure 2: Monte Carlo simulations

Note: The graphs show the observed proportion of each forecaster in each ranking position according to our proposed model in equation (9) (Adj. MAE) and according to the MAE statistic. In the top panel we have equal ability, $\mu_A = \mu_B = \cdots = \mu_G$ and the sample size, N, is equal to one year. The graphs (a) and (b) show a simulation assuming no timing effect ($\alpha = 1$). Graphs (c) and (d) show a simulation where the timing effect is given by outcome effect only ($\alpha = 0$). The lower panel presents the case of a outcome effect and unequal ability according to $\mu_A > \mu_B > \cdots > \mu_G$, thus, A is expected to be ranked as number seven. The the single year evaluation graphs (e) and (f) have N equal to one while the ten year evaluation graphs (g) and (h) have N equal to ten. Results based on S = 100,000 replicates. All simulations have the standard deviation, σ, in the error term set to 0.3.

3.2 The Equal Ability Case

We start with the case with equal ability and no timing effect (case 1). This case is included because we do not want our procedure to distort the performance of the MAE statistic when no timing effect is present. The simulations are based on a single-year evaluation. The results are given in Figure 2a and 2b for the two statistics. The timing effect is eliminated by equating the persistence coefficient α to one. Since the abilities are set equal for all seven forecasters we expect each forecasters ranking to be uniform in each position. As given in Figure 2a and 2b both our model and the MAE statistic handle this case well. The rankings are evenly distributed for all forecasters.

In the second equal ability simulation we add a timing effect, consisting only of the outcome effect. The rank proportions are still expected to be uniformly distributed across the institutions. Figure 2c shows that our proposed model achieve the expected. However, the MAE fail to rank the forecasters properly in this case, see Figure 2d. The information advantage of forecaster A is interpreted as A being the best forecaster in about 40 per cent of the replicates. This is actually not surprising, but will distort an evaluation.

We have also computed a case of random ability. In this case each forecaster is given an ability from a uniform distribution between 0 and 0.16 in each replicate.[12] Since the forecasters share the same distribution, the forecasters still have the same expected ranking over the simulation. The result of the random ability case is thus similar to the results shown in case 1 with equal ability.

3.3 An Unequal Ability Case

The lower part of Figure 2 presents results when we do not assume equal ability (case 2). More precisely we now assume that the generated ability order of the forecasters are such that the forecaster with the shortest average horizon is the poorest, $\mu_A > \mu_B > \cdots > \mu_G$. In other words, the timing effect and the ability now pulls in different directions. In this case we expect forecaster G to be ranked as number one and A as number seven. Figure 2e and 2f show results for a single year, $N = 1$. Since the information effect and the ability works in different directions the MAE statistic concludes that the forecasters are approximately equally good, see Figure 2f.[13] However, our model does a better job in this case, see Figure 2e. Forecaster G is top ranked in about 40 per cent of the cases and forecaster A exhibits the poorest rank in approximately 40 per cent of the replicates.

If we instead of evaluating a single year undertake an evaluation based on 10 years, the performance of our model improves further. This is presented in Figure 2g. Forecaster G is now correctly ranked as best in 70 per cent of the replicates, forecaster F as second best in 50 per cent of the times and forecast A as being worst in over 70 per cent of the generated samples. The MAE statistic is still giving wrong rankings, see Figure 2h.

The ability may also work in the ”same direction” as the timing effect. Then the unadjusted MAE statistic would work better than in case 2, and slightly better than the adjusted statistic, since it has a bias towards a ranking that follows the direction of the timing effect. The adjusted statistic performance is in line with the results of case 2.

3.4 A Measure of Incorrect Rankings

The overall results from the Monte Carlo are robust also for different designs of the simulations. This is evident in Table 2, which present simulations based on different values of the residual variance, the persistence parameter, as well as the sample size. In the table we present a measure of incorrect rankings (MIR) for a comparison based on observed, $O_{i,r}$, minus expected, $E_{i,r}$, frequency for individual i in

[12] By setting the upper limit to 0.16 we get an average sequential difference of 0.02 and an average difference between the best and the worst forecaster of 0.12, as in case 2.

[13] This is of course not always the case. If we vary the size of the information effect and abilities the resulting proportions are also varied.

ranking r for method j according to the following:

$$MIR_j = \frac{1}{S}\sqrt{\frac{1}{49}\sum_{i=A}^{G}\sum_{r=1}^{7}(O_{i,r} - E_{i,r})^2(1+\rho_i)^2} \tag{15}$$

where ρ_i is the absolute distance between the expected, $E[r_i]$, and the observed ranking, r_i:

$$\rho_i = |E[r_i] - r_i|. \tag{16}$$

Thus, ρ works as a penalty parameter of having large deviations in the ranking position. In the first ability case we do not have an expected ranking position, $E[r_i]$, meaning that all ρ_i is set to zero, while in the second ability case we set ρ according to equation (16). The column *Adjusted* refers to MIR according to our proposed adjusted MAE statistic, while the column *Unadjusted* refers to MIR based on a traditional MAE statistic.

Table 2: Monte Carlo simulation: Design and incorrect rankings results

				Measure of Incorrect Rankings	
μ_i	N	σ	α	Adjusted	Unadjusted
Case 1: $\mu_A = \mu_B = \cdots = \mu_G$	1	0.30	0.0	0.00	0.08
			0.8	0.01	0.08
			1.0	0.00	0.00
		0.15	0.0	0.00	0.15
			0.8	0.01	0.14
			1.0	0.00	0.00
	10	0.30	0.0	0.00	0.20
			0.8	0.02	0.19
			1.0	0.00	0.00
		0.15	0.0	0.00	0.30
			0.8	0.04	0.29
			1.0	0.00	0.00
Case 2: $\mu_A > \mu_B > \cdots > \mu_G$	1	0.30	0.0	0.45	0.62
			0.8	0.46	0.61
			1.0	0.45	0.45
		0.15	0.0	0.36	0.64
			0.8	0.37	0.59
			1.0	0.36	0.36
	10	0.30	0.0	0.28	0.65
			0.8	0.29	0.58
			1.0	0.28	0.28
		0.15	0.0	0.10	0.68
			0.8	0.11	0.56
			1.0	0.10	0.10

Note: μ_i is the generated individual forecast ability. N is the number of years in the evaluation sample. σ is the standard deviation in the error term and α is the persistence parameter. Note that the value of 0.8 on monthly data approximately corresponds to the line with 0.6 on quarterly data in Figure 1. See equation (10) for the data generating process. The column *Adjusted* refers to the measure of incorrect rankings (MIR) values based on our proposed model in equation (9). The column *Unadjusted* refers to MIR based on a traditional mean absolute error statistic. MIR calculations according to equation (15), where a lower value indicates a smaller ranking error. Results based on 100,000 replications.

Overall, for both ability cases, the MIR values are lower for the *Adjusted* method, indicating a smaller ranking error, compared to the *Unadjusted* method. Only when α is one (i.e. no timing effect), the MIR value for both methods are the same. This is, however, in line with what one would expect. All in all, the simulations suggest that the model can handle different situations regarding to the timing effects, which is shown in Figure 2 and Table 2. The Monte Carlo simulations also suggest that our model can adjust for different publication dates and identify the ability of the forecasters. Furthermore, the model does not distort the estimated ability if no information effect is present. The MAE statistic on the other hand is severely affected in the Monte Carlo study by differences in information. We also note that our model performs much better when the sample size increase and that the model performance increases as the variance in the error term is reduced. Hence, we conclude that our model, overall, is preferred compared to an unadjusted statistic.

4 Empirical Application

We now use the model for a real-world application. We pay special attention to two things: The importance of the timing effect by itself and how it affects the ranking. This is done by looking at the data and through estimated quantities in our model.

4.1 Data and Descriptive Statistics

In the application we use data from 10 Swedish forecasters.[14] The forecasts are annual averages for Swedish GDP growth and CPI inflation and are made up to two years prior to the outcome, yielding a maximum forecast horizon of 24 months. The data covers the period 1999–2015. From the forecasts we calculate the absolute value of the forecast error based on the first available outcome. The absolute errors then give rise to two quantities: the MAE statistic and our timing adjusted ability statistic.

Table 3 presents the data by means of summary statistics for each institutions absolute forecast errors. The number of forecasts differs for the various institutions and so does the average forecast horizon.[15] The average horizon provides a first indication on how the adjustment will affect the evaluation, but is by no means exhaustive. In practice the horizon for each forecast will matter. The average horizon differs at most by (slightly more than) two months. Two months in this case is definitely enough for potential information disadvantages. The lower part of each panel in Table 3 demonstrates that the forecast errors vary more within each institutions set of errors than between the institutions. This is an indication that forecasters regularly adjust their forecasts, probably due to new outcomes and short-term forecast errors. The small variation between the institutions suggest a prominent “herd behavior” among the forecasters.[16]

[14] The evaluated institutions are the same as in Sveriges Riksbank’s yearly evaluation, see for example Sveriges Riksbank (2016): Swedish Ministry of Finance (MoF), HUI Research (HUI), National Institute of Economic Research (NIER), Swedish Trade Union Confederation (LO), Nordea (NDA), the Riksbank (RB), Scandinaviska Enskilda Banken (SEB), Svenska Handelsbanken (SHB), Confederation of Swedish Enterprises (SN), and Swedbank (SWED).

[15] Although there are some regularities regarding when various institutions publish their forecasts the publication calenders are not fixed and can vary between the years.

[16] The occurrence of *herd behavior* among institutions is not particularly remarkable. Forecasters study approximately the same information and have access to each others forecasts and analyses.

Table 3: Descriptive statistics for absolute forecast errors, 1999–2015

	Number of forecasts	Mean horizon	Mean value	Standard deviation	Min	Max	Mean value	Standard deviation	Min	Max
			Panel (a) GDP				*Panel (b) CPI*			
MoF	84	11.7	1.15	1.34	0.00	6.7	0.51	0.60	0.00	3.1
HUI	136	11.3	1.16	1.33	0.00	7.5	0.56	0.62	0.00	3.5
NIER	141	11.7	1.17	1.46	0.00	7.7	0.48	0.63	0.00	3.2
LO	69	12.1	1.25	1.40	0.00	7.2	0.51	0.52	0.00	2.3
NDA	125	12.6	1.13	1.22	0.00	6.9	0.59	0.65	0.00	2.9
RB	164	11.5	1.09	1.30	0.00	6.9	0.55	0.73	0.00	3.8
SEB	136	12.3	1.09	1.24	0.00	6.5	0.51	0.58	0.00	3.5
SHB	115	12.0	1.27	1.36	0.00	7.0	0.56	0.67	0.00	3.3
SN	121	12.1	1.23	1.40	0.00	7.5	0.59	0.68	0.00	3.3
SWED	92	13.5	1.35	1.44	0.00	7.7	0.64	0.70	0.00	3.1
Overall	N=1183	12.0	1.18	1.34	0.00	7.70	0.55	0.65	0.00	3.80
Between	n =10			0.08	1.09	1.35		0.05	0.48	0.64
Within	N/n=118.3			1.34	−0.17	7.71		0.64	−0.09	3.80

Note: The table describes the absolute errors from the forecasts (GDP in panel (a) and CPI in panel (b)) published by institutions over the period 1999–2015. In total there are 10 forecasters with a total of 1183 forecasts. The overall average of the absolute forecast error is 1.18 for GDP and 0.55 for CPI. *Overall* refers to the entire dataset. The *Between* variation shows the variation of the means of each institution (across time periods), and *Within* refers to the variation about the respective mean of each institution. Note that the between and the within standard deviation does not add up to the overall deviation since the two standard deviations are not calculated around the same mean value.

4.2 Estimated Timing Effect

This section discusses and estimates the information content effect in the data. Starting with the results reported in Panel A in Table 4. Panel A in Table 4 presents the estimation results of equation (8) for GDP and CPI, respectively. The models include a full set of fixed forecast year and individual-specific effects. As expected, the contribution from the distance between the forecast and the outcome on the absolute forecast error is positive and significant for both variables. The estimated coefficient for the timing suggests that the marginal effect on the absolute forecast error, of producing the forecast with full information instead of no information of the outcome, is around 1.4 percent for GDP and 1.1 percent for CPI.

We cannot reject that the estimated effect of the information content is common for all forecasters. This is tested by the joint hypothesis $\delta_{i=1} = \delta_{i=2} =, ..., \delta_{i=10} = 0$. Moreover, the forecast year effects are strongly significant for both variables. This means that some years have been harder to predict than others. This effect is included in the model to separate the difficulty of different years from the individual effect, which we call ability. We test differences in the ability of the various institutions with a joint F-test. The test finds no significant differences between the institutions accuracy for the assessed variables. This is because the variation between the forecasters is much smaller than the variation within each forecaster.

Figure 3 plots the observed absolute forecast errors for GDP and CPI for the horizons together with a regression line according to the estimations of timing from Panel A in Table 4. The regression lines represent the best fits of the whole sample of the data, and can be used to adjust the absolute forecast

Table 4: Estimation results

	Panel A						Panel B					
	GDP			CPI			GDP			CPI		
	Coef.	SE	*p*-val.	Coef.	SE	*p*-val.	Coef.	SE	*p*-val.	Coef.	SE	*p*-val.
Timing	1.425***	0.071	0.000	1.085***	0.036	0.000						
*Timing*Year*												
1999							0.722***	0.189	0.000	1.092***	0.114	0.000
2000							0.262*	0.139	0.060	0.259***	0.067	0.000
2001							2.212***	0.138	0.000	0.616***	0.114	0.000
2002							0.820***	0.101	0.000	0.570***	0.051	0.000
2003							1.208***	0.100	0.000	0.205*	0.118	0.083
2004							1.084***	0.114	0.000	1.179***	0.111	0.000
2005							0.022	0.067	0.738	1.230***	0.077	0.000
2006							1.558***	0.077	0.000	0.211***	0.053	0.000
2007							−0.329***	0.124	0.008	0.270***	0.078	0.001
2008							1.944***	0.119	0.000	1.063***	0.091	0.000
2009							7.563***	0.196	0.000	3.412***	0.154	0.000
2010							3.713***	0.192	0.000	0.316***	0.078	0.000
2011							0.549***	0.119	0.000	1.019***	0.124	0.000
2012							1.429***	0.147	0.000	1.448***	0.093	0.000
2013							0.211*	0.114	0.064	1.468***	0.089	0.000
2014							0.280***	0.081	0.001	1.732***	0.087	0.000
2015							0.194*	0.103	0.061	1.695***	0.113	0.000
Year effects	Yes***			Yes***			Yes***			Yes***		
Individual effects	Yes			Yes			Yes***			Yes**		
Individuals	10			10			10			10		
Observations	1183			1183			1183			1183		
Adj. R-squared	0.642			0.539			0.907			0.745		
RMSE	0.803			0.438			0.410			0.326		
F(16, 1140)							130.07			59.61		
Prob.>F							0.000			0.000		

Note: *** denotes statistically significant at the 1 per cent level, ** significant at the 5 per cent level and * significant at the 10 per cent level. The table shows selected estimates for 10 Swedish forecasters over the period 1999–2015 together with robust standard errors (SE). *Year* is the the forecast year and not the year of the forecast origin. In Panel A the reported estimates are based on equation (8), while Panel B show estimates from equation (9).

errors with respect to the timing effect. This is done by subtracting the estimated effect from individual absolute errors. The individual effect (μ_i) represent each forecasters mean deviation from the estimated line (after accounting for year effects). Thus, the estimated individual effect is appropriate to assess and compare the forecasting ability of the different institutions.

It is clear from Figure 3 that the variability in the errors increase with the forecast horizon. Although, the estimated marginal effects presented in Table 4 are still unbiased and consistent, the longer horizons will dominate the short horizons.[17] This is due to the fact that errors are generally further from the expected error (fitted line) in the long horizons and will therefore be adjusted more.

The increasing variability in the errors with the horizon is to a large extent given by the fact that the difficulty to forecast a variable can differ across years. This means that both the size of the error and the marginal timing effect can differ across years. To illustrate this we also estimate equation (9),

[17] The presence of heteroskedasticity should not alter the 'central' position of the OLS line. Least square estimates are still unbiased and consistent but not efficient. Thus, the estimated standard errors are biased, so the usual confidence intervals and test statistics are incorrect. The standard errors reported in Table 4 are, however, corrected for heteroskedasticity using White's robust standard errors, see White (1980). Clustering the standard errors on year or institutions does not alter the results.

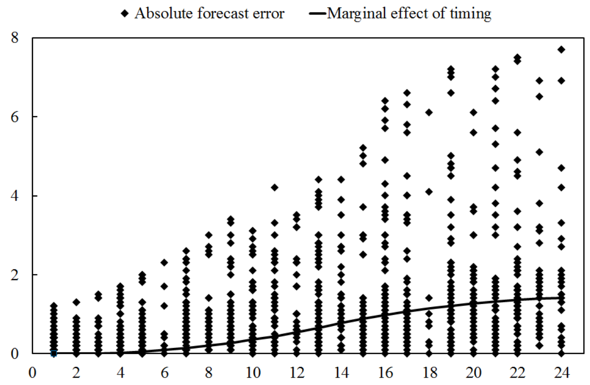

(a) GDP growth

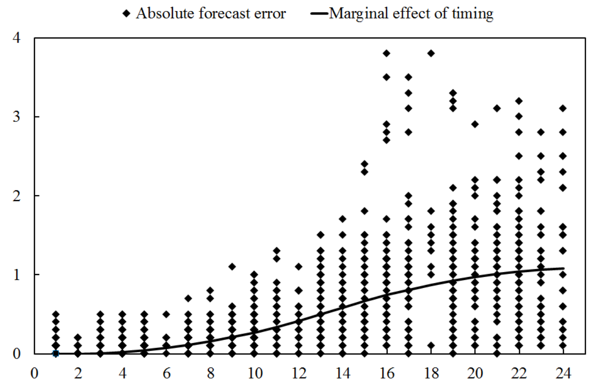

(b) CPI inflation

Figure 3: Absolute forecast errors and estimated average effect of timing

and thus allow for time-varying marginal timing effects (i.e. δ_t). Panel B of Table 4 clearly suggests that there are large and significant variations in the estimated timing effects across different years for both GDP and CPI. The hypothesis of homogeneous timing effects across forecast years is strongly rejected.[18] Subsequently, a model that accounts for heterogeneous timing effects across years is required to capture the behavior of the forecast errors correctly. One caveat with this specification is that it is more sensitive to the number of observations in each year. Our data set has about 60–70 observations per year.

All years, except for GDP in 2007, exhibit the expected timing effect. This is shown by positive estimates in Panel B in Table 4. The GDP timing coefficient for 2007 is negative, which is driven by data.[19] In an application one may think this is appropriate, whereas, it is also possible to restrict the coefficients to be non-negative.

4.3 Estimated Abilities and Rankings

For the flexible specification in Panel B in Table 4 the joint F-test on the individual-specific (ability) effects becomes significant. Thus, a more careful modeling of the forecast errors pick up systematic differences between different forecasters GDP and CPI predictions over time.[20]

The estimated individual effects from equation (9), which are interpreted as each institutions forecast ability (adjusted for the differences in timing when forecasting) can be used for a fair comparison of the performance. In Figure 4 we plot the estimated ability for each institution (for GDP and CPI respectively) for the entire period 1999–2015. The figure presents two measures. The black bars show the ability, that is the adjusted mean absolute error, estimated from the individual effects obtained from Panel B in

[18] A joint F-test, presented in Panel B of Table 4, of parameter constancy test of the hypothesis that the coefficients are identical across years is rejected at the 0.0001 level for both GDP and CPI.

[19] During 2007 the forecasters seem to been presented to GDP disinformation as time evolved. In practice, the outcomes during the first half of 2007 were stronger than the forecasts, which led to upward revisions. When the third quarter outcome was published, the first and second quarter outcomes where revised downwards. This had a non-negligible impact on the 2007 annual average, since these quarters together weigh 7/16. See Sveriges Riksbank (2008) for more information. As a consequence this led to the unusual pattern where early forecasts were more accurate than later ones.

[20] This joint test investigates the null hypothesis that all forecasters share the same ability against the alternative that at least one of the forecasters ability differs from the others.

Table 4. This ability is expressed as deviations from the all-institutions average, $\mu_i^* = \hat{\mu}_i - \frac{1}{n}\sum_{i=1}^{n}\hat{\mu}_i$. A positive bar ($\mu_i^*$ >0) corresponds to a larger adjusted forecast error than the average forecaster and thus a lower forecast ability, whereas a negative bar (μ_i^* <0) implies the opposite. The gray bars represent an MAE statistic, also stated as deviations from the all-institutions average MAE.

Figure 4, ordered by the black bars, shows that there exist differences in the forecasting ability between the assessed institutions. In pair-wise comparisons these differences are sometimes quite small but in other cases non-negligible. Furthermore, it is noted that the two measures give rise to somewhat different rankings between the institutions. Thus, an adjustment of the errors with respect to the differences in timing have implications on the comparison.

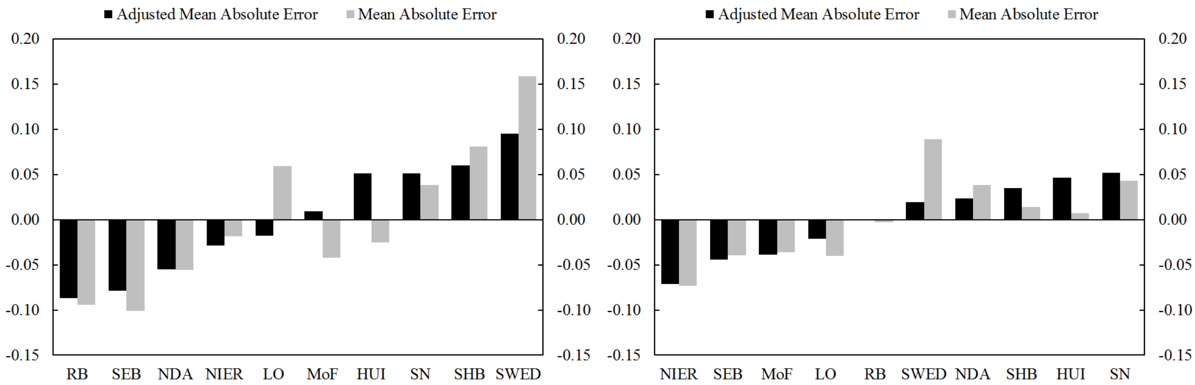

(a) GDP growth (b) CPI inflation

Figure 4: Estimated relative forecast ability

Note: Adjusted Mean Absolute Error refers to estimated ability according to our proposed model in equation (9). Mean Absolute Error refers to a MAE statistic. Both series are demeaned by the overall mean based on all institutions. Thus, a negative value tells us that the individual have a forecast performance above the average.

5 Conclusions

Different forecasters typically prepare and publish their forecasts at different points in time. This means that the forecasts are based on different amount of information. When the forecasters performance is assessed and compared this may severely distort the conclusions. This paper introduces a way to handle these problems by means of a simultaneous estimation of the forecasters' ability and the effect of different information content. We then use the ability measure to compare the forecasters. We argue that a forecast comparison that do not adjust for differences in timing may be misleading due to information advantage or disadvantage. This view is supported by Monte Carlo evidence. The Monte Carlo simulation also concludes that increasing the evaluation sample size is beneficial to the proposed method. The implication of this is that a single year evaluation is not sufficient to retrieve reliable estimates of the forecasters' ability. Finally, the method is applied to a real-world dataset of 10 Swedish forecasters for the sample period 1999–2015. These forecasters are ranked based on their ability to forecast Swedish GDP and CPI inflation. Our first finding is that the timing of the forecasts do matter for the forecast error. Thus, comparisons based on unadjusted mean absolute errors may give unfair results. We conclude that our proposed method is preferred over the traditional statistic when comparing the performance of forecasters.

References

Andersson, M. K., Karlsson, G., and Svensson, J. (2007). An evaluation of the Riksbank's forecasting performance. *Sveriges Riksbank Economic Review*, 3:59–75.

Baltagi, B. H. (2001). *Econometric Analysis of Panel Data.* John Wiley and Sons, 2 edition.

Bank of England (2015). Evaluating forecast performance. Independent Evaluation Office.

Bauer, A., Eisenbeis, R. A., Waggoner, D. F., and Zha, T. (2003). Forecast evaluation with cross-sectional data: The blue chip surveys. *Federal Reserve Bank of Atlanta Economic Review*, 88(2):17–31.

Blix, M., Wadefjord, J., Wienecke, U., and Ådahl, M. (2001). How good is the forecasting performance of major institutions? *Sveriges Riksbank Economic Review*, 3:38–68.

Boero, G., Smith, J., and Wallis, K. F. (2008). Evaluating a three-dimensional panel of point forecasts: The Bank of England Survey of External Forecasters. *International Journal of Forecasting*, (24):354–367.

Cabanillas, L. G. and Terzi, A. (2012). The accuracy of the European Commission's forecasts re-examined. European Economy – Economic Papers 476, Directorate General Economic and Financial Affairs (DG ECFIN), European Commission.

Davies, A. (2006). A framework for decomposing shocks and measuring volatilities derived from multi-dimensional panel data of survey forecasts. *International Journal of Forecasting*, (22):373−393.

Davies, A. and Lahiri, K. (1995). A new framework for analyzing survey forecast using three-dimensional panel data. *Journal of Econometrics*, (68):205–227.

Diebold, F. X. (2007). *Elements of Forecasting.* Thomson, South-Western, 4 edition.

European Central Bank (2013). An assessment of eurosystem staff Macroeconomic projections. Ecb monthly bulletin.

Gneiting, T. (2011). Making and Evaluating Point Forecasts. *Journal of the American Statistical Association*, 106(494):746–762.

Goh, K. L. and Lawrence, D. (2006). Treasury's Forecasting Performance: A Head-to-Head Comparison. New Zealand Treasury Working Paper 06/10.

Patton, A. J. (2015). Evaluating and Comparing Possibly Misspecified Forecasts. Working paper, March 2015 version.

Sveriges Riksbank (2008). Material for assessing monetary policy 2005–2007.

Sveriges Riksbank (2016). Account of Monetary Policy 2015.

Timmermann, A. (2007). An Evaluation of the World Economic Outlook Forecasts. IMF Staff Papers 1.

Vogel, L. (2007). How Do the OECD Growth Projections for the G7 Economies Perform? A Post-Mortem. OECD Working Papers 573.

White, H. (1980). A Heteroskedasticity-Consistent Covariance Matrix Estimator and a Direct Test for Heteroskadasticity. *Econometrica*, 48(4):817–838.

Earlier Working Papers:

For a complete list of Working Papers published by Sveriges Riksbank, see www.riksbank.se

Estimation of an Adaptive Stock Market Model with Heterogeneous Agents *by Henrik Amilon*	2005:177
Some Further Evidence on Interest-Rate Smoothing: The Role of Measurement Errors in the Output Gap *by Mikael Apel and Per Jansson*	2005:178
Bayesian Estimation of an Open Economy DSGE Model with Incomplete Pass-Through by Malin Adolfson, Stefan Laséen, Jesper Lindé and Mattias Villani	2005:179
Are Constant Interest Rate Forecasts Modest Interventions? Evidence from an Estimated Open Economy DSGE Model of the Euro Area *by Malin Adolfson, Stefan Laséen, Jesper Lindé and Mattias Villani*	2005:180
Inference in Vector Autoregressive Models with an Informative Prior on the Steady State *by Mattias Villani*	2005:181
Bank Mergers, Competition and Liquidity *by Elena Carletti, Philipp Hartmann and Giancarlo Spagnolo*	2005:182
Testing Near-Rationality using Detailed Survey Data *by Michael F. Bryan and Stefan Palmqvist*	2005:183
Exploring Interactions between Real Activity and the Financial Stance *by Tor Jacobson, Jesper Lindé and Kasper Roszbach*	2005:184
Two-Sided Network Effects, Bank Interchange Fees, and the Allocation of Fixed Costs *by Mats A. Bergman*	2005:185
Trade Deficits in the Baltic States: How Long Will the Party Last? *by Rudolfs Bems and Kristian Jönsson*	2005:186
Real Exchange Rate and Consumption Fluctuations follwing Trade Liberalization *by Kristian Jönsson*	2005:187
Modern Forecasting Models in Action: Improving Macroeconomic Analyses at Central Banks *by Malin Adolfson, Michael K. Andersson, Jesper Lindé, Mattias Villani and Anders Vredin*	2005:188
Bayesian Inference of General Linear Restrictions on the Cointegration Space *by Mattias Villani*	2005:189
Forecasting Performance of an Open Economy Dynamic Stochastic General Equilibrium Model *by Malin Adolfson, Stefan Laséen, Jesper Lindé and Mattias Villani*	2005:190
Forecast Combination and Model Averaging using Predictive Measures *by Jana Eklund and Sune Karlsson*	2005:191
Swedish Intervention and the Krona Float, 1993-2002 *by Owen F. Humpage and Javiera Ragnartz*	2006:192
A Simultaneous Model of the Swedish Krona, the US Dollar and the Euro *by Hans Lindblad and Peter Sellin*	2006:193
Testing Theories of Job Creation: Does Supply Create Its Own Demand? *by Mikael Carlsson, Stefan Eriksson and Nils Gottfries*	2006:194
Down or Out: Assessing The Welfare Costs of Household Investment Mistakes *by Laurent E. Calvet, John Y. Campbell and Paolo Sodini*	2006:195
Efficient Bayesian Inference for Multiple Change-Point and Mixture Innovation Models *by Paolo Giordani and Robert Kohn*	2006:196
Derivation and Estimation of a New Keynesian Phillips Curve in a Small Open Economy *by Karolina Holmberg*	2006:197
Technology Shocks and the Labour-Input Response: Evidence from Firm-Level Data *by Mikael Carlsson and Jon Smedsaas*	2006:198
Monetary Policy and Staggered Wage Bargaining when Prices are Sticky *by Mikael Carlsson and Andreas Westermark*	2006:199
The Swedish External Position and the Krona *by Philip R. Lane*	2006:200

Price Setting Transactions and the Role of Denominating Currency in FX Markets 2007:201
by Richard Friberg and Fredrik Wilander

The geography of asset holdings: Evidence from Sweden 2007:202
by Nicolas Coeurdacier and Philippe Martin

Evaluating An Estimated New Keynesian Small Open Economy Model 2007:203
by Malin Adolfson, Stefan Laséen, Jesper Lindé and Mattias Villani

The Use of Cash and the Size of the Shadow Economy in Sweden 2007:204
by Gabriela Guibourg and Björn Segendorf

Bank supervision Russian style: Evidence of conflicts between micro- and macro-prudential concerns 2007:205
by Sophie Claeys and Koen Schoors

Optimal Monetary Policy under Downward Nominal Wage Rigidity 2007:206
by Mikael Carlsson and Andreas Westermark

Financial Structure, Managerial Compensation and Monitoring 2007:207
by Vittoria Cerasi and Sonja Daltung

Financial Frictions, Investment and Tobin's q 2007:208
by Guido Lorenzoni and Karl Walentin

Sticky Information vs Sticky Prices: A Horse Race in a DSGE Framework 2007:209
by Mathias Trabandt

Acquisition versus greenfield: The impact of the mode of foreign bank entry on information and bank lending rates 2007:210
by Sophie Claeys and Christa Hainz

Nonparametric Regression Density Estimation Using Smoothly Varying Normal Mixtures 2007:211
by Mattias Villani, Robert Kohn and Paolo Giordani

The Costs of Paying – Private and Social Costs of Cash and Card 2007:212
by Mats Bergman, Gabriella Guibourg and Björn Segendorf

Using a New Open Economy Macroeconomics model to make real nominal exchange rate forecasts 2007:213
by Peter Sellin

Introducing Financial Frictions and Unemployment into a Small Open Economy Model 2007:214
by Lawrence J. Christiano, Mathias Trabandt and Karl Walentin

Earnings Inequality and the Equity Premium 2007:215
by Karl Walentin

Bayesian forecast combination for VAR models 2007:216
by Michael K. Andersson and Sune Karlsson

Do Central Banks React to House Prices? 2007:217
by Daria Finocchiaro and Virginia Queijo von Heideken

The Riksbank's Forecasting Performance 2007:218
by Michael K. Andersson, Gustav Karlsson and Josef Svensson

Macroeconomic Impact on Expected Default Freqency 2008:219
by Per Åsberg and Hovick Shahnazarian

Monetary Policy Regimes and the Volatility of Long-Term Interest Rates 2008:220
by Virginia Queijo von Heideken

Governing the Governors: A Clinical Study of Central Banks 2008:221
by Lars Frisell, Kasper Roszbach and Giancarlo Spagnolo

The Monetary Policy Decision-Making Process and the Term Structure of Interest Rates 2008:222
by Hans Dillén

How Important are Financial Frictions in the U S and the Euro Area 2008:223
by Virginia Queijo von Heideken

Block Kalman filtering for large-scale DSGE models 2008:224
by Ingvar Strid and Karl Walentin

Optimal Monetary Policy in an Operational Medium-Sized DSGE Model 2008:225
by Malin Adolfson, Stefan Laséen, Jesper Lindé and Lars E. O. Svensson

Firm Default and Aggregate Fluctuations 2008:226
by Tor Jacobson, Rikard Kindell, Jesper Lindé and Kasper Roszbach

Re-Evaluating Swedish Membership in EMU: Evidence from an Estimated Model — 2008:227
by Ulf Söderström

The Effect of Cash Flow on Investment: An Empirical Test of the Balance Sheet Channel — 2009:228
by Ola Melander

Expectation Driven Business Cycles with Limited Enforcement — 2009:229
by Karl Walentin

Effects of Organizational Change on Firm Productivity — 2009:230
by Christina Håkanson

Evaluating Microfoundations for Aggregate Price Rigidities: Evidence from Matched Firm-Level Data on Product Prices and Unit Labor Cost — 2009:231
by Mikael Carlsson and Oskar Nordström Skans

Monetary Policy Trade-Offs in an Estimated Open-Economy DSGE Model — 2009:232
by Malin Adolfson, Stefan Laséen, Jesper Lindé and Lars E. O. Svensson

Flexible Modeling of Conditional Distributions Using Smooth Mixtures of Asymmetric Student T Densities — 2009:233
by Feng Li, Mattias Villani and Robert Kohn

Forecasting Macroeconomic Time Series with Locally Adaptive Signal Extraction — 2009:234
by Paolo Giordani and Mattias Villani

Evaluating Monetary Policy — 2009:235
by Lars E. O. Svensson

Risk Premiums and Macroeconomic Dynamics in a Heterogeneous Agent Model — 2010:236
by Ferre De Graeve, Maarten Dossche, Marina Emiris, Henri Sneessens and Raf Wouters

Picking the Brains of MPC Members — 2010:237
by Mikael Apel, Carl Andreas Claussen and Petra Lennartsdotter

Involuntary Unemployment and the Business Cycle — 2010:238
by Lawrence J. Christiano, Mathias Trabandt and Karl Walentin

Housing collateral and the monetary transmission mechanism — 2010:239
by Karl Walentin and Peter Sellin

The Discursive Dilemma in Monetary Policy — 2010:240
by Carl Andreas Claussen and Øistein Røisland

Monetary Regime Change and Business Cycles — 2010:241
by Vasco Cúrdia and Daria Finocchiaro

Bayesian Inference in Structural Second-Price common Value Auctions — 2010:242
by Bertil Wegmann and Mattias Villani

Equilibrium asset prices and the wealth distribution with inattentive consumers — 2010:243
by Daria Finocchiaro

Identifying VARs through Heterogeneity: An Application to Bank Runs — 2010:244
by Ferre De Graeve and Alexei Karas

Modeling Conditional Densities Using Finite Smooth Mixtures — 2010:245
by Feng Li, Mattias Villani and Robert Kohn

The Output Gap, the Labor Wedge, and the Dynamic Behavior of Hours — 2010:246
by Luca Sala, Ulf Söderström and Antonella Trigari

Density-Conditional Forecasts in Dynamic Multivariate Models — 2010:247
by Michael K. Andersson, Stefan Palmqvist and Daniel F. Waggoner

Anticipated Alternative Policy-Rate Paths in Policy Simulations — 2010:248
by Stefan Laséen and Lars E. O. Svensson

MOSES: Model of Swedish Economic Studies — 2011:249
by Gunnar Bårdsen, Ard den Reijer, Patrik Jonasson and Ragnar Nymoen

The Effects of Endogenuos Firm Exit on Business Cycle Dynamics and Optimal Fiscal Policy — 2011:250
by Lauri Vilmi

Parameter Identification in a Estimated New Keynesian Open Economy Model — 2011:251
by Malin Adolfson and Jesper Lindé

Up for count? Central bank words and financial stress — 2011:252
by Marianna Blix Grimaldi

Wage Adjustment and Productivity Shocks *by Mikael Carlsson, Julián Messina and Oskar Nordström Skans*	2011:253
Stylized (Arte) Facts on Sectoral Inflation *by Ferre De Graeve and Karl Walentin*	2011:254
Hedging Labor Income Risk *by Sebastien Betermier, Thomas Jansson, Christine A. Parlour and Johan Walden*	2011:255
Taking the Twists into Account: Predicting Firm Bankruptcy Risk with Splines of Financial Ratios *by Paolo Giordani, Tor Jacobson, Erik von Schedvin and Mattias Villani*	2011:256
Collateralization, Bank Loan Rates and Monitoring: Evidence from a Natural Experiment *by Geraldo Cerqueiro, Steven Ongena and Kasper Roszbach*	2012:257
On the Non-Exclusivity of Loan Contracts: An Empirical Investigation *by Hans Degryse, Vasso Ioannidou and Erik von Schedvin*	2012:258
Labor-Market Frictions and Optimal Inflation *by Mikael Carlsson and Andreas Westermark*	2012:259
Output Gaps and Robust Monetary Policy Rules *by Roberto M. Billi*	2012:260
The Information Content of Central Bank Minutes *by Mikael Apel and Marianna Blix Grimaldi*	2012:261
The Cost of Consumer Payments in Sweden *by Björn Segendorf and Thomas Jansson*	2012:262
Trade Credit and the Propagation of Corporate Failure: An Empirical Analysis *by Tor Jacobson and Erik von Schedvin*	2012:263
Structural and Cyclical Forces in the Labor Market During the Great Recession: Cross-Country Evidence *by Luca Sala, Ulf Söderström and AntonellaTrigari*	2012:264
Pension Wealth and Household Savings in Europe: Evidence from SHARELIFE *by Rob Alessie, Viola Angelini and Peter van Santen*	2013:265
Long-Term Relationship Bargaining *by Andreas Westermark*	2013:266
Using Financial Markets To Estimate the Macro Effects of Monetary Policy: An Impact-Identified FAVAR* *by Stefan Pitschner*	2013:267
DYNAMIC MIXTURE-OF-EXPERTS MODELS FOR LONGITUDINAL AND DISCRETE-TIME SURVIVAL DATA *by Matias Quiroz and Mattias Villani*	2013:268
Conditional euro area sovereign default risk *by André Lucas, Bernd Schwaab and Xin Zhang*	2013:269
Nominal GDP Targeting and the Zero Lower Bound: Should We Abandon Inflation Targeting?* *by Roberto M. Billi*	2013:270
Un-truncating VARs* *by Ferre De Graeve and Andreas Westermark*	2013:271
Housing Choices and Labor Income Risk *by Thomas Jansson*	2013:272
Identifying Fiscal Inflation* *by Ferre De Graeve and Virginia Queijo von Heideken*	2013:273
On the Redistributive Effects of Inflation: an International Perspective* *by Paola Boel*	2013:274
Business Cycle Implications of Mortgage Spreads* *by Karl Walentin*	2013:275
Approximate dynamic programming with post-decision states as a solution method for dynamic economic models *by Isaiah Hull*	2013:276
A detrimental feedback loop: deleveraging and adverse selection *by Christoph Bertsch*	2013:277
Distortionary Fiscal Policy and Monetary Policy Goals *by Klaus Adam and Roberto M. Billi*	2013:278
Predicting the Spread of Financial Innovations: An Epidemiological Approach *by Isaiah Hull*	2013:279

SVERIGES
RIKSBANK

Sveriges Riksbank
Visiting address: Brunkebergs torg 11
Mail address: se-103 37 Stockholm

Website: www.riksbank.se
Telephone: +46 8 787 00 00, Fax: +46 8 21 05 31
E-mail: registratorn@riksbank.se