

Working Paper Series

**A Statistical Forecasting Method
for Inflation Forecasting**

Ippei Fujiwara and Maiko Koga***

Working Paper 02-5

July 2002

Research and Statistics Department

Bank of Japan

C.P.O BOX 203 TOKYO

100-8630 JAPAN

*e-mail: ippei.fujiwara@boj.or.jp

** e-mail: maiko.koga@boj.or.jp

Views expressed in Working Paper Series are those of the authors and do not necessarily reflect those of the Bank of Japan or Research and Statistics Department.

A Statistical Forecasting Method for Inflation Forecasting*

Ippeï Fujiwara[†] and Maiko Koga[‡]

July, 2002

Abstract

Typically, when using econometric techniques to forecast economic variables, estimation is carried out on a forecasting model that is built upon some assumed economic structure, based upon *a priori* knowledge and economic principles. However, such techniques cannot avoid running into the possibility of misspecification, which will occur should there be some error in the assumptions underlying this economic structure. Even when diagnostic tests have been easily cleared, a small change in the way this structure is set up can induce large differences in the forecast value. In other words, the researcher's subjective choices in setting up the model can have a substantial influence on the estimated forecast.

In this paper, in which we concentrate upon inflation forecasting, we present a statistical forecasting method (SFM) that stresses statistical relationships among time series data, and that makes no structural assumptions, other than to set up the underlying variables. When putting

*In producing this paper we have had the benefit of valuable comments from Munehisa Kasuya and Hideaki Shimizu (both of the Research and Statistics Department, Bank of Japan), as well as from many other staffs of the Bank. In carrying out the analysis, we received considerable support from Toshinao Yoshiba (Institute for Monetary and Economic Studies, Bank of Japan) and Toshiro Nomura (Graduate School of Economics, Keio University, currently at Finacial Research Center, Nomura Securites). We would like to offer them our sincere thanks. In addition, in conducting estimation of the Reduced Rank VAR model and in calculation of the Pesaran-Timmermann statistics, we made use of the GAUSS program created by Dr. George Kapetianos (Bank of England, currently at Queen Mary college, University of London), with some modifications of our own. It remains only to emphasize that the opinions expressed in this paper are our own and do not reflect the public position of the Bank of Japan or of the Research and Statistics Department.

[†]Economic Research Division, Research and Statistics Department, Bank of Japan (E-mail: ippei.fujiwara@boj.or.jp)

[‡]Economic Research Division, Research and Statistics Department, Bank of Japan (E-mail: maiko.koga@boj.or.jp)

together a forecast, this SFM first builds a number of VAR models from combinations of the underlying variables; it then automatically ranks these, based upon their performance. Furthermore, it has the additional property that it produces forecasts not merely by looking at the movements of the forecast themselves over time, but by taking into account the uncertainty in both the model and the forecast value captured in the forecast distribution (and illustrated in the fan charts). We also carry out analysis that looks just at the question of whether future inflation will move upwards or downwards, attempting to produce a qualitative forecast of this movement.

Use of this SFM, in addition to establishing a more objective setting and enabling us to produce forecasts which take uncertainty into account, also gives better results when forecasting these qualitative movements in inflation. Although a further extended comparison between forecast and actual values is still required to confirm the practical value of this SFM, at this juncture we can state the following: not only does the SFM offer useful forecasting information that cannot be extracted when using just a single structural-type estimating model, but it can also play a valuable role in providing a cross-check for forecasts produced using such structural-type models.

Keywords: Inflation, Forecast, Reduced Rank VAR, Nonparametric test of predictive performance

JEL Classification: C32, C35, C53, E31

1 Introduction

Typically, when using econometric techniques to forecast economic variables, estimation is carried out on a forecasting model that is built upon some assumed economic structure, based upon *a priori* knowledge and economic principles. However, such techniques cannot avoid running into the possibility of misspecification, which will occur should there be some error in the assumptions underlying this economic structure. Even when diagnostic tests have been easily cleared, a small change in the way this structure is set up can induce large differences in the forecast value. In other words, the researcher's subjective choices in setting up the model can have a substantial influence on the estimated forecast.¹

Given this, rather than using a model that relies on some specific hypothesis above, it may be considered desirable to employ a forecasting technique that is as objective as possible and based purely on statistical relationships additionally.

Of course, however objective the methodology is, it is not feasible to completely remove the chance that some misspecification occurs in writing the model. For this reason, in order to produce forecasts which reflect observed reality as accurately as possible, it is necessary to make full use of all the information we possess concerning the uncertainty inherent in both the model and in the forecast value itself. In order to achieve this, in making our forecasts we do not look merely at the fluctuations observed in the forecast value of some specific model, but instead take account of the uncertainty inherent in the forecast errors and in the choice of model itself, admirably captured in the forecast distribution (illustrated in the fan charts).

In full awareness of the above problems, the object of this paper is to use our forecasting of the inflation rate as an illustration² with which we present our Statistical Forecasting Method (SFM hereafter). This SFM is one example of a forecasting technique which scores highly on objectivity, and which takes appropriate account of uncertainty. Specifically, we produce forecasts using

¹Krolzig and Hendry (2001) introduce a process for selecting a model objectively based on the Gets (General to Specific) concept. However, we do not cover this selection methodology in this paper.

²Various papers dealing with inflation forecasting have been published by the Bank of Japan's Research and Statistics department, following sessions such as the "Workshop on Inflation Forecasting Errors" (Sept. 2000) [cf. for example, Ban and Saito (2001), Kitagawa and Kawasaki (2001), Kasuya and Shinki (2001), and Fukuda and Keida (2001) among others]. Although the analysis in this paper is close to that of Kitagawa and Kawasaki (2001), we differ from them in that we produce the forecast distributions (fan charts) which take account of the uncertainty inherent in the forecast errors and the model selection.

our SFM in accordance with the following process.

To begin with we construct a VAR forecasting model for the inflation rate. Even for a reduced-form VAR model, which is highly objective, it is not possible to do away entirely with all traces of subjectivity.³ The forecast result will inevitably depend to some extent upon the choice of data, the combination of adopted variables, and the period for which estimation is carried out. Recognizing that some bias may emerge as a result of our choice of data, or of some particular combination of variables, we follow Stock and Watson (1999) and Pesaran and Timmermann (1995) and carry out estimations of the model for as many different combinations of variables as possible, in hope that by thus constructing a large number of VAR forecasting models we may improve our results. As our estimating methodology, we employ a VAR model that takes account of rank restrictions amongst the variables, in other words a Reduced Rank VAR model (RR-VAR model⁴ hereafter). As described above, in estimating VAR models based on various different combinations of variables, it is not unreasonable to assume that some quantity of redundant information will be included in the matrix of parameters, since several of the variable combinations are likely to display similar tendencies. In order to deal with this we employ rank restrictions on the parameter matrix and are thus able to carry out estimation using a more parsimonious⁵ model, which is likely to produce improved results. On this point, by comparing the RR-VAR model with the standard VAR model, Camba-Mendez, Kapetanios, Smith, and Weale (1999) have demonstrated the performance of the former to be superior than the latter for forecasting purposes.

Having reached this point, we then construct the forecast distribution, or fan chart, which takes account of the uncertainty in both the model and the forecast value. We produce three different constructions of the forecast distribution: (1) the top model distribution; (2) the nonparametric distribution; (3) the mixed distribution. (1) is the distribution we get when we make use of the forecast error from the model which performs best out of all of the forecasting models that we constructed. (2) is the estimate of future uncertainty that we get when we depict a dispersion of the forecast values (point estimates) obtained from each of our forecasting models. For (3) we take matters a step

³There are models other than VAR models that place their main emphasis on statistical relationships; we might for example have chosen to adopt a method such as the Neural Network Approach. However, in the analysis which follows we have restricted ourselves to the VAR framework, attempting to produce inflation forecasts that take account of the uncertainty inherent in the forecast errors and the model selection.

⁴For details, see Velu, Reinsel, and Wichern (1986), and Lütkepohl (1991).

⁵i.e., a model that obeys the principles of parsimony.

further: each of the forecast values obtained from each of the forecasting models is assumed to be normally distributed according to parameters extracted from the forecast errors of the respective model; then, using the “forecast error reciprocals” from each of the distributions as the base for constructing weights, we compose a distribution which is able to capture uncertainty.⁶ To put it differently: we construct three different forecast distributions, where (1) is based on parametric methodology, (2) on nonparametric methodology, and (3) on semiparametric methodology.⁷ Lastly, by arranging these various constructed distributions into time-series, we illustrate the extent of uncertainty in three fan charts.

In addition to the quantitative forecasting question “Roughly what percentage rate of inflation do we expect to see?”, we also consider the question “Is future inflation more likely to move upwards or downwards?”, and we construct a qualitative forecast (in other words, a simple up or down forecast) in order to examine this.

2 Forecasting techniques and results

Forecasting in our SFM is carried out via three processes: (1) forecasts with the RR-VAR models; (2) construction of forecast distributions and fan charts; (3) the qualitative up-down directional forecast that makes use of the Pesaran-Timmermann test (for an outline of techniques (1) and (2), cf. Chart 1). More detail on each of these processes follows.

2.1 Forecasting with the RR-VAR model

Here, we first of all construct a large number of RR-VAR models from different combinations of stationary data, and use these to calculate a correspondingly large number of forecast values over several periods. Next, we arrange the

⁶Model uncertainty enters in three possible ways: (i) the average difference between the real values and those forecast by the model (the forecast error); (ii) misspecification of the estimating model; (iii) parameter uncertainty (i.e., how the parameters obtained via estimation are distributed). When we carry out stochastic simulations for the parameters, although it is possible to consider (iii), it is highly computationally demanding to do so. Thus for our forecast distribution number (3) we in fact think only about the uncertainty that arises from a mixture of concepts (i) and (ii).

⁷The mixed distribution technique may be thought of as one form of semiparametric method (cf. Powell (1994) for details), and it may also be considered to follow the thinking of the Forecast Combination method (for which see Clements and Hendry (1998), Diebold (1998), and Granger and Newbold (1986)).

forecasting models in order, based on their in-sample dynamic forecast⁸ performance over the past two years (where performance is judged according to Root Mean-Squared Error - RMSE hereafter), and we select the model that performs best for each forecasting period (1 period ahead,..., 4 periods ahead). Stringing together the out-of-sample dynamic forecasts obtained from these best-performing models, we construct the “movement over time of the forecast value obtained from the best performing forecasting model.” In this way we manage to account for the possibility that, when we compare the performance for each object forecasting period across models, the combination of variables that produces the best forecasts for, for example, 1 period ahead (i.e., one quarter from now) may be different from the combination of variables that works best when forecasting 4 periods ahead (1 year from now).

Turning to the details, we see that there are seven specific steps involved in calculating the forecast values. These are as follows:

(a) Choice of data series: Having selected the appropriate series of representative macro-economic variables, we carry out the necessary seasonal adjustment⁹ and take logs (cf. Chart 2 for data selection and seasonal adjustment procedures).

(b) Ensuring data stationarity: We perform unit-root tests¹⁰ on each data series, taking differences until stationarity is ascertained. Making use of their means and standard deviations, we then standardize¹¹ these now stationary data series. Having discovered that the CPI, which we are looking to forecast, follows an I(2) process, we employ the twice-differenced series, namely “change in the inflation rate.”¹²

(c) Choice of explanatory variables: From the data series obtained in (b), we select the arbitrary number of explanatory variables. Specifically: From N data series we select \underline{m} to M different combinations of explanatory

⁸Clements and Hendry (1998) observe the following: when for example carrying out a 4 periods ahead forecast, using only data which precedes the object forecasting period by at least 4 periods to estimate the model directly can produce forecasts whose performance is superior to dynamic forecasts. However we do not use this technique here.

⁹Monthly variables are reconstituted, after seasonal adjustment, in quarterly form.

¹⁰We use the ADF (augmented Dickey-Fuller) test here.

¹¹We transform them into data series with means of 0 and standard deviations of 1.

¹²Since our data is in logs, the second difference gives the (magnitude of the) change in the rate of change. Having taken differences to ensure stationarity, the combination of the differenced explained variables and the differenced explanatory variables may become difficult to explain from the standpoint of economic theory.

variables which are assumed to be included in each model. The total number of combinations of variables turns out to be $K = \sum_{m=\underline{m}}^M \frac{N!}{m!(N-m)!}$.

In this paper there are 16 data series from which we select combinations of 2 to 4 explanatory variables. Thus we produce a total of exactly 2,500 different models.

(d) Estimation of the RR-VAR models: We carry out estimation using 3 to 5 dimensional RR-VAR models, made up of the CPI which we are looking to forecast and the explanatory variables selected in (c). Where \mathbf{y}_t is a multivariate stationary time series made up of the group of variables included in some given model k , the RR-VAR model is set up as follows. Note that the lag length (L), which is to take one to four, is determined where the AIC (Akaike Information Criterion) is the smallest. The example below illustrates the case when the estimation period for the parameters is 83/Q1 - 99/Q4.

$$\begin{aligned} \mathbf{y}_t &= \sum_{l=1}^L \mathbf{B}_l \mathbf{y}_{t-l} + \boldsymbol{\eta}_t \\ &= \mathbf{B} \mathbf{x}_t + \boldsymbol{\eta}_t \\ &= \boldsymbol{\alpha} \boldsymbol{\beta}' \mathbf{x}_t + \boldsymbol{\eta}_t, \quad t = 83/Q1, \dots, 99/Q4. \end{aligned}$$

Here \mathbf{B}_l is an (m, m) matrix. $\boldsymbol{\alpha}$ and $\boldsymbol{\beta}_l$ are (m, r^*) matrices ($r^* \leq m$) when they are of rank r^* . Also, $\boldsymbol{\beta} = (\boldsymbol{\beta}'_1, \dots, \boldsymbol{\beta}'_L)'$, $\mathbf{x}_t = (\mathbf{y}'_{t-1}, \dots, \mathbf{y}'_{t-L})'$, and $\mathbf{B} = \boldsymbol{\alpha} \boldsymbol{\beta}'$.

In carrying out estimation using the RR-VAR model, we first estimate the matrix of parameters that we get using a standard VAR model; then, taking the rank r^* of this matrix as a base, we reduce the matrix into two more parsimonious matrices without any loss of information (i.e., having estimated the rank of the parameter matrix \mathbf{B} , we then reduce this into the form of the two matrices $\boldsymbol{\alpha}$ and $\boldsymbol{\beta}$).

We follow Bartlett (1947) in estimating the rank. Bartlett's rank test is presented below in the form of a likelihood ratio test.

$$H_0 : \text{rank}(\mathbf{B}) = r^* \quad \text{against} \quad H_1 : \text{rank}(\mathbf{B}) = r > r^*,$$

$$l_c(H_0) - l_c(H_1) = \frac{T}{2} \sum_{i=r^*+1}^m \ln(1 + \hat{\lambda}_i^2) \xrightarrow{d} \chi^2[(m - r^*), (mL - r^*)].$$

l_c are the concentrated log likelihoods,¹³ while $\hat{\lambda}_i$ are the eigenvectors of the matrix $T^{-1}\Sigma_\eta^{-1/2}\mathbf{Y}\mathbf{X}'(\mathbf{X}\mathbf{X}')^{-1}\mathbf{X}\mathbf{Y}'\Sigma_\eta^{-1/2}$, where $\mathbf{Y} = [\mathbf{y}_1, \mathbf{y}_2, \dots, \mathbf{y}_T]$, $\mathbf{X} = [\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_T]$. $\Sigma_\eta = E(\boldsymbol{\eta}\boldsymbol{\eta}')$ denotes the error variance-covariance matrix, estimated by the quasi-maximum likelihood method.¹⁴

Under the null hypothesis H_0 , the concentrated log likelihood ratio $l_c(H_0) - l_c(H_1)$ is known to converge to the chi-squared distribution described above, and comparing it with the significance points (corresponding to the adopted significance level) for the chi-squared distribution determines the rank r^* .

Having calculated the rank in this way, we then use this to estimate the RR-VAR parameter matrices (in other words $\boldsymbol{\alpha}$ and $\boldsymbol{\beta}'$) as¹⁵

$$\hat{\boldsymbol{\alpha}} = \Sigma_\eta^{1/2}\hat{\mathbf{V}}, \quad \hat{\boldsymbol{\beta}}' = \hat{\mathbf{V}}'\Sigma_\eta^{-1/2}\mathbf{Y}\mathbf{X}'(\mathbf{X}\mathbf{X}')^{-1}.$$

Here $\hat{\boldsymbol{v}}_r$ is the standardized eigenvector corresponding to the r th largest eigenvalue $\hat{\lambda}_r$ (where r is the rank calculated above), and $\hat{\mathbf{V}}$ describes the matrix

$$\hat{\mathbf{V}} = [\hat{\boldsymbol{v}}_1, \hat{\boldsymbol{v}}_2, \dots, \hat{\boldsymbol{v}}_{r^*}].$$

(e) Calculation of the in-sample forecasts: For each of our 2500 models we carry out a series of in-sample dynamic forecasts,¹⁶ ranging from 1 period ahead forecasts to 4 periods ahead forecasts. The evaluation period for gauging forecast performance stretches across two years of past data (98/Q1 - 99/Q4). In other words, we produce forecasts recursively, shifting the starting period for the dynamic forecast each time, and obtaining values for the 1 period

¹³In cases when the log likelihood is made up of several parameters, this is concentrated so as to be able to express just the parameters of interest. Specifically, we concentrate the log likelihood function by sequential replacement of the first order condition in the log likelihood function for all parameters other than the parameters of interest.

¹⁴i.e., We calculate $T^{-1}\mathbf{Y}\mathbf{Y}' - T^{-1}\mathbf{Y}\mathbf{X}'(T^{-1}\mathbf{X}\mathbf{X}')^{-1}(T^{-1}\mathbf{Y}\mathbf{X})'$.

¹⁵For more detail see Velu, Reinsel, and Wichern (1986), or Lütkepohl (1991).

¹⁶Here, having ranked models based on their in-sample dynamic forecast performance, we produce out-of-sample dynamic forecasts. Since our interest lies ultimately in the out-of-sample forecasts, there is an argument to suggest that we should also rank models based on their out-of-sample forecast performance (e.g., Clark (2000)). However, for evaluations based on out-of-sample forecasts, since the estimation period for performance evaluation is set to be a certain number of periods (here 8 periods) prior to the first forecast period, we become unable to use all the information potentially available to us when carrying out the estimation. Here, therefore, we choose to evaluate forecast performance as we do, based on the conjecture that the out-of-sample performance will be superior if we conduct our estimations using all of the information available to us.

ahead forecast ($\hat{y}_{\tau|\tau-1}$) to the 4 periods ahead forecast ($\hat{y}_{\tau|\tau-4}$), where τ is taken across the 8 periods from 98/Q1 to 99/Q4 (see Chart 1-(2)).¹⁷ Since we have 2500 models and we calculate forecast values for 1 to 4 periods ahead for each of the 8 past periods, we produce a total of 80,000 forecasts ($2500 \times 4 \times 8$).

$$\hat{y}_{\tau|\tau-h} = [\alpha\beta']^h \mathbf{x}_{\tau-h}, \quad \tau = 98/Q1, \dots, 99/Q4, \quad h = 1, \dots, 4.$$

(f) Calculation of the RMSE: We calculate the RMSE,¹⁸ which indicate the relative performance of our respective inflation forecasts. However, since for the CPI our estimate is of the change in the inflation rate,¹⁹ we produce forecasts of the inflation rate itself by adding as appropriate the estimated changes in the inflation rate to the inflation rate from the last estimating period of the model. The RMSE are calculated based on the differences between the actual values and the forecasts produced in this way.

In other words, denoting the actual value of the inflation rate in period τ as R_{τ}^{π} , and the forecast value as $\hat{y}_{\tau|\tau-h}^{\pi}$, we get

$$RMSE_h = T^{-1} \sum_{\tau=1}^T \epsilon_{\tau,h}^2, \quad \tau = 98/Q1, \dots, 99/Q4, \quad h = 1, \dots, 4,$$

$$\epsilon_{\tau,h} = R_{\tau}^{\pi} - \hat{y}_{\tau|\tau-h}^{\pi}.$$

Carrying out the above calculations for every model, we obtain $\{RMSE_{h,k}\}_{k=1}^{2500}$.

(g) Calculation of the out-of-sample forecasts: Having obtained the $RMSE_{h,k}$, we use these to select, for each respective forecasting period, the

¹⁷Note that the estimation period itself remains unchanged throughout.

¹⁸As pointed out in Clements and Hendry (1998), there is potentially a problem here since we do not consider possible serial correlation in the forecast errors when calculating the RMSE. There are two possible ways that can be considered for dealing with this: (i) we could set the forecasting periods so that there is no overlap between the forecasts (i.e., we could set the data frequency so that, for example, it is quarterly for a 1 period ahead forecasting model or yearly for a 4 periods ahead forecasting model); or (ii) we could use instrumental variables. However, whichever of these methods we use, we run into problems, such as a loss of degrees of freedom or the difficulty of choosing the right instruments. For this reason, although the RMSE represents extremely important information for constructing the forecast distributions, we do not revise them to correct for the possibility that they are excessive because of serial correlation in the forecast errors.

¹⁹Strictly speaking, since as we pointed out in 2.1.(b) the estimated forecast values are normalized, we carry out these calculations after first reversing the normalization procedure.

best-performing model (i.e., the model with the smallest RMSE). Having transformed into inflation rates the out-of-sample dynamic forecasts ($\hat{\mathbf{y}}_{\tau|\tau-h}$) obtained from these best-performing models, we produce a time series for our object forecast period (i.e., 4 periods into the future) by linking these estimates together and constructing the “movement over time of the forecast value obtained from the best-performing models.”²⁰

$$\hat{\mathbf{y}}_{\tau|\tau-h} = [\boldsymbol{\alpha}\boldsymbol{\beta}']^h \mathbf{x}_{\tau-h}, \quad \tau = 00/Q1, \dots, 00/Q4, \quad h = 1, \dots, 4.$$

2.2 Construction of the forecast distribution and fan charts

The movements over time in the forecast values obtained from the best-performing models and calculated in the preceding section may be considered, in themselves, to be highly useful as forecasts based on a technique that is as objective as possible. However, as discussed earlier, we are interested here in capturing the uncertainty inherent in the forecast, and this is well accounted for in the forecast distribution. Thus we now demonstrate how we produce this forecast distribution, making use of the out-of-sample forecast values and the forecast errors.

As shown in Charts 1-(3a), (3b), and (3c), the forecast distribution captures in distributional form the uncertainty inherent in the forecast value. Setting the desired significance level for the forecast distribution at each point in time, and then linking the resulting significance points as a time series give us our fan charts (Chart 1-(4)).

Here, we depict first of all (1) the “top model distribution,” which uses information from the best-performing models alone (their forecast values and forecast errors). Next, we depict (2) the “nonparametric distribution,” which makes use of the forecast values from the full set of models. Lastly, we produce (3) the “mixed distribution,” which makes use of the forecast values and the forecasts errors from the full set of models. In contrast with the “top model distribution,” which reflects only the information gleaned from the best-performing model, the “nonparametric distribution” and the “mixed distribution” capture uncertainty about the future: specifically, the former does so through a dispersion of forecast values taken from the full set of models (i.e., it addresses the possibility of model misspecification); while the latter take into account in addition the forecasting error inherent in each forecast value (i.e., it addresses the issues of both forecasting errors and model misspecification).

²⁰This is depicted as the central line (the line showing movement of the mean over time) in the fan chart (Chart 3) for the “top model distribution” (see below for further explanation).

2.2.1 Top model distribution

First, for the best-performing models selected above, we create, for each forecast period, a normal distribution with the forecast value as its mean and the RMSE as its standard deviation. We then produce fan charts with starting periods of 2000/Q1 and 2001/Q1,²¹ by joining together as time series the significance points attained for each 10% increment in the one-sided significance level for this forecast distribution.

With regard to forecasts based on this top model distribution, when we consider that there may not have been such large differences between the RMSE of the top-ranking models, it may well be that we are relying too heavily on just the information gleaned from one model, picked out from its peers on the basis of infinitesimal differences. Putting it another way, when we take into account the possibility that the model may have been misspecified, it brings into question the appropriateness of evaluating our inflation rate forecasts in light of only the forecasts and the expression of forecast uncertainty taken from the top model distribution alone.

As econometric principles, sharpening up an existing model or searching out the best possible model may be thought important. However here, where interest lies in the construction of a practical forecasting model, we are searching to find a forecasting technique that takes into account the forecasts obtained from a number of models. Thus in what follows we no longer restrict attention to the best-performing model, and instead focus on building forecast distributions that make full use of the information contained in the forecasts and concomitant forecasting errors from the other models.

2.2.2 Nonparametric distribution

Having produced forecasts from a variety of models for the inflation rate of one particular period, it seems reasonable to assume that, should it be the case that these forecast values are roughly the same whichever forecasting model is used, then the actual value will also lie within this neighborhood. For this reason, we regard a dispersion of the forecast values obtained from a

²¹The estimation periods for the parameters are as follows: for fan charts whose starting period is 2000/Q1, the estimation period runs from 1983 through 1999; while for fan charts starting in 2001/Q1, it runs from 1983 through 2000. Likewise, the forecast performance evaluation periods are 1998 and 1999, and 1999 and 2000 respectively. The forecast distribution is drawn based on the RMSE calculated during the evaluation period (i.e., the prior two year period specified above), and in cases (such as 98-99) when movements in the actual inflation rate for the relevant period are not easily explained by the model (in other words, there is a large error), the spread of the forecast distribution tends to become wider. We can attribute the wider spread of the fan chart in 2000, compared with that in 2001, to this.

number of models as an indicator of the uncertainty that attends a forecast of the future.

Turning to the risk of misspecification, it becomes important to make full use of the information gained from several models and not just look at the information contained in one particular model.

In addition, because we made the assumption of normality when looking at the uncertainty in the top model distribution, it was not possible to represent any skewness in the distribution. With the nonparametric distribution, it becomes possible to represent asymmetry in the distribution of the inflation rate with regard to its movement in an upward versus a downward direction.

In accordance with this line of thought, and following the two steps outlined below, in this subsection we first of all carry out the estimation of the nonparametric distribution, which makes use only of the information contained in the dispersion of forecast values.

(a) Estimation of the nonparametric distribution: The aim of producing a fan chart with these nonparametric distributions²² is to construct forecast distributions using just the information contained in the dispersion of 2500 forecast values²³ we have for each period. What we are producing here is not a forecast distribution built assuming some particular shape of probability distribution and using parameters calculated from the forecast values themselves, such as their means and variances. Instead we elicit a probability distribution from the frequency distribution that derives from all 2500 forecast values, thus constructing our distribution in a nonparametric manner (i.e., neither estimating any parameters, nor placing any *a priori* restrictions upon the shape of the probability distribution). The nonparametric distribution is expressed using the Kernel density estimator below. $\kappa(u)$ is the Kernel weighting function (here a Gaussian Kernel), ϕ is the smoothing parameter that determines bandwidth, and the formula for the probability density func-

²²For detailed treatments of the nonparametric methodology in general, see Silverman (1986), Härdle and Linton (1994), or Pagan and Ullah (1999).

²³We could consider building the forecast distribution not from the full set of sample values (2500 values), but from a subset of, say, the top 5% in terms of their performance (125 values). Since such a forecast distribution could prove extremely useful, we would like to examine, in the future, issues such as the most suitable sample size and so on, through an extensive comparison with actual values. Of course, since here our primary point of focus is the introduction of techniques, and since the only strategy that we can think of for approaching the question of how to choose the sample size is an extensive comparison with actual values to be carried out later, in our production below of the forecast distributions, including the mixed distributions, we use the whole sample.

tion (*pdf*) is written as follows²⁴:

$$pdf(q) = \frac{1}{K\phi} \sum_{k=1}^K \kappa\left(\frac{q - y_k^\pi}{\phi}\right),$$

$$\kappa(u) = \frac{1}{\sqrt{2\pi}} e^{-\frac{1}{2}u^2},$$

$$\phi = 0.9K^{-\frac{1}{5}} \min\{S, A/1.34\}.$$

Here, $y^\pi = \{y_1^\pi, \dots, y_k^\pi, \dots, y_K^\pi\}$ is the set of forecast values, K is the number of forecast values (2500 of them), S is the standard deviation of y^π , A is the interquartile range of y^π , and q denotes the estimation point of the density function.

(b) Construction of the fan charts: For the forecast distributions estimated for each period in (a), we set the significance level, which we vary in one-sided 10% increments, and by joining together as time series the respective significance points attained for each significance level, we produce a fan chart.

The forecast distributions produced in this way are illustrated in Chart 4. In contrast to the top model distribution, because the fan chart is constructed with reference to the forecast values from all 2500 models, it is not overly affected by the movements in the explanatory variables of one particular model (i.e., the best-performing model).

It is also noteworthy that, if we look at for example the forecasts for mid 2000, the upper half of the forecast distribution can be seen to have a long tail. In other words, the mode is situated below the mean, reflecting the fact that, restricting attention to frequency, there are a large number of forecasts predicting inflation rates lower than the mean.

²⁴According to Silverman (1986), there are two methods for calculating the bandwidth ϕ : he refers to the method on which the expression used in the text is based as the “subjective method”; the alternative “objective method” sets out to minimize the “cross validation” criterion (an MSE estimator that measures the difference between the actual value and the estimate obtained with a given ϕ). Although the latter method gets rid of some arbitrariness, in cases such as in this paper, where we confirm the spread of the distribution in the fan charts, it is not desirable to have a bandwidth that varies greatly depending on the forecasting period. For this reason, placing emphasis on the goal of getting a rough idea of the shape of the distribution, we use the former method here.

2.2.3 The mixed distribution

In the technique based on the nonparametric distribution, we looked to obtain the forecast uncertainty from the dispersion of forecast values. Of course in this method we take no account at all of the uncertainty inherent in each individual forecast value, not of the information gleaned about the relative performance of models in the in-sample dynamic forecast. Thus it is possible that we are not giving due weight to the forecast values from the top-performing models. For this reason, in the development of the mixed distribution that follows, we produce a forecast distribution that takes appropriate account of such information.

To begin with, we construct a normal distribution for each forecast value based on its RMSE. Then we take a weighted linear combination of these normal distributions, with the respective weights in proportion to the inverse RMSE of each forecast distribution.²⁵ Although admittedly using the inverse RMSE as weights is somewhat ad hoc, this is based on the notion that those models whose performance has been good over the past two years may be considered to provide relatively high forecast accuracy in the future as well.²⁶

²⁵Although the application of the weights is different, distributions that capture uncertainty by compounding other distributions in this way can be found in the field of finance. For an example of existing research in the area, we would suggest Melick and Thomas (1997).

²⁶This is indeed just one example of a number of possible ways of applying weights. Creating the same type of forecast distribution of the period for which we have actual values, we also tried the grid search method, in which weights are chosen to bring the mean of the forecast distribution closer to the actual value. However, even with this method, we are ultimately unable to confirm whether or not we are extracting the appropriate information from the forecast values. Among examples of the semiparametric methods, rather than carrying out estimation of the distribution using weights fixed in advance as above, it is common to estimate the weights, the mean, and the variance, by using the EM algorithm, or something similar, to maximize log likelihoods based on a linear combination of distributions. However, in this method we are unable to make use of the information contained in the RMSE.

In addition to the above, there are further alternative methods, as introduced in Clements and Hendry (1998): (i) the “regression method” in which, weights are calculated by regressing the actual values on a number of the forecast values such that the forecast errors from the combination forecasts is minimized; (ii) the “variance-covariance approach,” in which the weights are calculated such that the forecast error variance is minimized. However, in the case of the former, not only are we unable to secure a large enough sample to calculate the requisite parameters for 2500 forecast values, but there are also a number of other definitional problems that emerge, as summed up in Diebold (1998): whether to include a constant term, whether to estimate time-variant weights, and whether to allow for serial correlation in the errors in the constructed forecasts. For the latter method, it is unclear whether, with a small sample, we would be able to obtain reliable estimates of the variances and covariances.

For these reasons, we choose to adopt the inverse RMSE as weights here, as a straightfor-

Specifically, the mixed distribution is constructed following the three steps outlined below.

(a) Construction of a normal distribution for each forecast value:

For a forecast h periods ahead using model k , we construct a normal distribution with the forecast value $\mu_{h,k}$ as its mean, and a standard deviation $\sigma_{h,k} = RMSE_{h,k}$. The probability density function $f_{h,k}$ can then be written as follows:

$$f_{h,k}(z_{h,k}; \mu_{h,k}, \sigma_{h,k}) = \frac{1}{\sqrt{2\pi}\sigma_{h,k}} e^{-\frac{(z_{h,k} - \mu_{h,k})^2}{2\sigma_{h,k}^2}},$$

$$\int f_{h,k}(z_{h,k}) dz_{h,k} = 1, \quad 0 \leq f_{h,k}(z_{h,k}).$$

We then convert the probability density function into a cumulative distribution function:

$$F_{h,k}(z_{h,k}; \mu_{h,k}, \sigma_{h,k}) = \int_{-\infty}^{z_{h,k}} f_{h,k}(\zeta_{h,k}; \mu_{h,k}, \sigma_{h,k}) d\zeta_{h,k}, \quad 0 \leq F_{h,k}(z_{h,k}) \leq 1.$$

(b) Constructing a weighted linear combination of cumulative distributions:

The cumulative distribution function for h periods ahead is expressed as F_h , a weighted linear combination of K cumulative distribution functions $F_{h,k}$. The weights $w_{h,k}$ are based on the inverse RMSE and are constructed so as to sum to unity (thus the area under the weighted linear combination of probability density function f_h is also 1).

$$F_h(z_h; \mu_h, \sigma_h) = \sum_{k=1}^K w_{h,k} F_{h,k}(z_{h,k}; \mu_{h,k}, \sigma_{h,k}), \quad 0 \leq F_h(z_h) \leq 1,$$

$$\sum_{k=1}^K w_{h,k} = 1, \quad 0 \leq w_{h,k} \leq 1,$$

$$f_h(z_h; \mu_h, \sigma_h) = \sum_{k=1}^K w_{h,k} f_{h,k}(z_{h,k}; \mu_{h,k}, \sigma_{h,k}), \quad 0 \leq f_h(z_h),$$

$$\int f_h(z_h) dz_h = 1.$$

ward and uninvolved method.

(c) Construction of the fan chart: For the weighted linear combination of cumulative density function estimated for each period in (b), we set the significance level, which we vary in one-sided 10% increments, and we calculate the resultant significance points. Joining these together as time series, we produce a fan chart.

Comparing the fan chart constructed from the mixed distribution (Chart 5) with that constructed from the nonparametric distribution (Chart 4), we observe that although the movement over time of the means of the forecast distributions are very similar for each period, the spread of the distribution is more pronounced in the former (i.e., there is more uncertainty). We can hypothesize two reasons for this. First, (i) while the nonparametric distribution focuses on “pin point” forecast values from each model, the mixed distribution by contrast focuses on the distribution centered on the forecast value (i.e., a spread), calculated from each model. In addition we may point out: (ii) while in the nonparametric distribution all the forecast values are given the same weight, so that forecasts from models that performed well during the period over which RMSE were calculated are treated no differently from others, in the mixed distribution on the other hand, greater importance is attached to forecast values (and their concomitant distributions) taken from models that performed relatively well over the same period. The result of this is that, in cases when these better-performing forecast values lie apart from the majority of the other forecasts, they may cause the spread of the forecast distribution to widen.²⁷

Although we are ultimately judging only from relative performances during 2000 and 2001, the mixed appears to offer the most insightful guide when we are considering at roughly in what range the forecast value will fall.

2.3 Qualitative up-down forecast using the Pesaran-Timmermann test

The aim of the forecast distribution is to obtain a measure of the forecast value and the size of the associated uncertainty (i.e., a quantitative forecast). However interest may lie in the question of “whether next period’s inflation rate will be higher or lower than that observed in this period,” and when thinking about the future inflation rate, this kind of forecast of the direction of change (i.e., a qualitative forecast) may also be considered important. Thus we look here to produce a simple up-down forecast of the direction of the

²⁷For a discussion of issues underlying the wider base of the forecast distribution in 2000 compared to 2001, see note 21.

inflation rate, using the large number of VAR models that we constructed in section 2.1. We evaluate the respective performances of the qualitative forecasts that emerge from each model using the Pesaran-Timmermann test, but the qualitative forecast that makes use of all the models is produced by following the two steps outlined below.

(a) Calculation of the Pesaran-Timmermann statistic for each forecast value: The Pesaran-Timmermann test examines whether the directional movements of the real and forecast values are in step with one another, or to put it another way, it checks how well rises and falls in the forecast value follow actual rises and falls: the larger the Pesaran-Timmermann statistic, the better the match.²⁸

The way the test works can be simply explained via the following example. Forecasting of coin tosses is carried out according to some model. Whenever heads comes up (or is forecast to come up) the value of +1 is assigned, while whenever tails comes up (or is forecast to come up) the value of -1 is assigned. In each instance the resulting values from the actual coin toss and from the forecast from the model are multiplied together. At this point, if, for example, the mean of this product is close to unity, it suggests that the performance of the forecasting model is high. The Pesaran-Timmermann statistic is a statistic that indicates qualitative forecasting performance based on this line of reasoning. Substituting a rise (+1) or a fall (-1) in the inflation rate for the heads or tails of the coin toss, we can apply the Pesaran-Timmermann test to our inflation forecasts.

More specifically, denoting the in-sample dynamic forecast calculated in 2.1 (e) (i.e., the forecast value of the magnitude of the change in the inflation rate) as \hat{y}_t ,²⁹ and the actual value recorded for the change in the inflation rate as R_t (i.e., $\hat{y}_t = E(R_t|\Omega_{t-1})$, where Ω_{t-1} sums up all the information available at $t - 1$), we first of all define the following quantities:

²⁸It takes no account, however, of the differences between the respective magnitudes of rises and falls.

²⁹As explained in note 19, strictly speaking, the estimated value has been normalized, so that we need first to reverse this process.

$$\begin{aligned}
X_t &= 1 \text{ if } R_t > 0 \\
&= 0 \text{ otherwise,} \\
Y_t &= 1 \text{ if } \hat{y}_t > 0 \\
&= 0 \text{ otherwise,} \\
Z_t &= 1 \text{ if } R_t \hat{y}_t > 0 \\
&= 0 \text{ otherwise.}
\end{aligned}$$

Now the Pesaran-Timmermann statistic S_n defined below follows, asymptotically, a standard normal distribution, under the null hypothesis that R_t and \hat{y}_t are independent random variables.³⁰

$$\begin{aligned}
S_n &= \frac{\hat{P} - \hat{P}_*}{[\hat{V}(\hat{P}) - \hat{V}(\hat{P}_*)]^{\frac{1}{2}}} \xrightarrow{d} N(0,1), \\
\hat{P}_X &= \sum_{t=1}^n X_t/n, \quad \hat{P}_Y = \sum_{t=1}^n Y_t/n, \quad \hat{P} = \sum_{t=1}^n Z_t/n, \\
\hat{P}_* &= \hat{P}_X \hat{P}_Y + (1 - \hat{P}_X)(1 - \hat{P}_Y).
\end{aligned}$$

Here, $\hat{V}(\hat{P})$ and $\hat{V}(\hat{P}_*)$ denote the variances of \hat{P} and \hat{P}_* respectively and defined as follows:

$$\begin{aligned}
\hat{V}(\hat{P}) &= n^{-1} \hat{P}_*(1 - \hat{P}_*), \\
\hat{V}(\hat{P}_*) &= n^{-1} (2\hat{P}_Y - 1)^2 \hat{P}_X (1 - \hat{P}_X) + n^{-1} (2\hat{P}_X - 1)^2 \hat{P}_Y (1 - \hat{P}_Y) \\
&\quad + 4n^{-2} \hat{P}_Y \hat{P}_X (1 - \hat{P}_Y)(1 - \hat{P}_X).
\end{aligned}$$

Having calculated the Pesaran-Timmermann statistic for each model from the forecast values and real values for the change in the inflation rate over the object forecasting period (1 to 4 periods ahead), we go on to select those models that reject the null hypothesis described above at the 5% significance level³¹ (i.e., those models for which the Pesaran-Timmermann statistic lies in the right-side critical region).

³⁰For more detail, see Pesaran and Timmermann (1995). We should note that, although the test is based on asymptotic theory, the sample size in the example in this paper is small.

³¹Models that clear a left-sided 5% significance test may contain important information regarding, for example, inverse correlation. However, since here we are looking to pick out those models whose directional movements are in step with actual movements in the inflation rate, we only select those models that clear a right-sided 5% significance test.

(b) Construction of the qualitative up-down Forecast: Using the models selected in (a) above, we now calculate out-of-sample dynamic forecasts for the object forecasting period. Construction follows that of the fan charts above in that forecasting of the 4 periods from 2000/Q1-Q4 is carried out using data from an estimation period that goes up to 1999/Q4; while the estimation period for forecasting of the 4 periods 2001/Q1-Q4 goes up to 2000/Q4.

Having produced these forecasts, we then examine, for each period, the proportion of models that produce positive forecasts of the change in the inflation rate relative to those that produce negative forecasts. We note down which is greater with the requisite sign.

The results are shown in Chart 6, where, comparing the respective signs of the forecasts versus the realized values, we see that rises or falls in the inflation rate are correctly predicted in 5 out of 8 periods. In particular, it is worth noting the consistent accuracy of the qualitative forecast in periods in which there is a conspicuous difference in the proportion of models registering either negative or positive changes³² (i.e., 2000/Q1, Q3, Q4; 2001/Q2).

While the relative success of this qualitative up-down forecast clearly still requires some verification through many more examples in terms of its practical applicability, it may nevertheless be deemed useful as a supplement to the information provided by the forecast distribution and associated fan charts.

3 Conclusion

In this paper, using forecasting of the inflation rate as an illustration, we present our SFM as an example of a technique for producing forecasts that not only scores highly on objectivity, but also takes account of uncertainty. The construction process can be outlined as follows. First of all, we build a large number of VAR forecasting models; we then arrange these in order of forecasting performance (in sample), select the model that performs best for each of our objective forecasting periods, and carry out (out-of-sample) forecasting accordingly. This technique is then supplemented by the production of fan charts that consider the uncertainty embodied in the forecast errors and indeed inherent in the choice of model itself. We also suggest a way of producing a qualitative up-down forecast.

The three forecast distributions that we produce (the top model distribu-

³²In a related point, for the 3 periods in which the forecast was out, it is worth mentioning that the proportions of models predicting positive versus those predicting negative values were almost the same.

tion, the nonparametric distribution, and the mixed distribution) offer information that may be deemed valuable when thinking quantitatively about the extent of the uncertainty involved in inflation forecasting. In addition, our qualitative up-down forecast roughly follows rises and falls in the actual inflation rate. Although we still need to confirm the practical value of our SFM by carrying out further comparison between forecast and actual values, it is nevertheless fair to say that this SFM both offers valuable forecasting information that it would be impossible to extract from a single structural forecasting model, and that it can also play a useful role in providing a cross-check on the forecasts produced using such structural models.

Without a more comprehensive set of forecast and real values, we cannot judge which of our three forecast distributions is to be preferred.³³ For this reason, at the present time, our SFM comes into its own when we are looking to evaluate the extent to which a given inflation forecast from some structural model may actually be realized in future. In such a case, we can put to work both the fan charts, produced using three different techniques, and the results of the qualitative up-down forecast, and in approaching the task of evaluation from several directions at once, demonstrate the value of our SFM as a tool for practical forecasting.

³³In other words, we are unable to rank in any simple way the parametric, nonparametric, and semiparametric methodologies.

References

- [1] Ban, K., and Saito, M. (2001), “Makuro keiryō moderu niyoru infure-ritsu yosoku gosa no bunseki (Analysis on the inflation rate forecasting errors in a macroeconometric model),” *Bank of Japan Research and Statistics Department Working Paper* 01-12 (in Japanese).
- [2] Bartlett, M.S. (1947), “Multivariate Analysis,” *Journal of the Royal Statistical Society, Series B*, 9, 176-197.
- [3] Camba-Mendez, G., Kapetanios, G., Smith, R.J., and Weale, M.R. (1999), “Tests of Rank in Reduced Rank Regression Models,” *National Institute of Economic and Social Research Discussion Paper*, No.150.
- [4] Clark, T.E. (2000), “Can Out-of-Sample Forecast Comparisons Help Prevent Overfitting?” *Federal Reserve Bank of Kansas City Research Working Papers*, 00-05.
- [5] Clements, M.P., and Hendry, D.F. (1998), *Forecasting Economic Time Series*, Cambridge: Cambridge University Press.
- [6] Diebold, F.X. (1998), *Elements of Forecasting*, Cincinnati: South-Western College Publishing.
- [7] Fukuda, S., and Keida, M. (2001), “Infure yosoku niokeru jisho bunseki no tembo -Phillips kyokusen no nihon niokeru yosokuryoku wo chushin ni (Prospects for empirical analysis on inflation forecasts: the predictive power of Phillips curves in Japan),” *Bank of Japan Research and Statistics Department Working Paper* 01-21 (in Japanese).
- [8] Granger, C.W.J., and Newbold, P. (1986), *Forecasting Economic Time Series*, San Diego: Academic Press.
- [9] Härdle, W., and Linton, O. (1994), “Applied Nonparametric Methods,” in Engle, R.F., and McFadden, L. (eds.), *The Handbook of Econometrics*, Vol.IV, New York, North-Holland, 2295-2339.
- [10] Kasuya, M., and Shinki, K. (2001), “Bukka hendo no tenkanten yosoku nitsuite (Forecasting turning points in price movements),” *Bank of Japan Research and Statistics Department Working Paper* 01-20 (in Japanese).
- [11] Kitagawa, G., and Kawasaki, Y. (2001), “Jikeiretsu moderu niyoru infure-ritsu yosoku gosa no bunseki (Analysis on the inflation rate forecasting

- errors from a time series model),” *Bank of Japan Research and Statistics Department Working Paper 01-13* (in Japanese).
- [12] Krolzig, H.M., and Hendry, D.F. (2001), “Computer automation of general-to-specific model selection procedures,” *Journal of Economic Dynamics and Control*, 25, 831-866.
- [13] Lütkepohl, H. (1991), *Introduction to Multiple Time Series Analysis*, New York: Springer-Verlag.
- [14] Melick, W.R., and Thomas, C.P. (1997), “Recovering an Assets Implied PDF from Option Prices: An Application to Crude Oil during the Gulf Crisis,” *Journal of Financial and Quantitative Analysis*, 32(1), 91-115.
- [15] Pagan, A.R., and Ullah, A. (1999), *Nonparametric Econometrics*, Cambridge: Cambridge University Press.
- [16] Pesaran, M.H., and Timmermann, A. (1992), “A Simple Nonparametric test of Predictive Performance,” *Journal of Business and Economic Statistics*, 10(4), 461-465.
- [17] Pesaran, M.H., and Timmermann, A. (1995), “Predictability of Stock Returns: Robustness and Economic Significance,” *Journal of Finance*, 50, 1201-1228.
- [18] Powell, J.L. (1994), “Estimation of Semiparametric Models,” in Engle, R.F., and McFadden, L. (eds.), *The Handbook of Econometrics*, Vol.IV, New York, North-Holland, 2443-2521.
- [19] Silverman, B.W. (1986), *Density Estimation for Statistics and Data Analysis*, London: Chapman and Hall.
- [20] Stock, J.H., and Watson, M.W. (1999), “Forecasting Inflation,” *Journal of Monetary Economics*, 44, 293-335.
- [21] Velu, R.P., Reinsel, G.C., and Wichern, D.W. (1986), “Reduced Rank Models for Multiple Time Series,” *Biometrika*, 73, 105-118.

An Outline of the Forecasting Method

(1) Forecasts from the RR-VAR model (1)	Data: different combinations of the data in Chart 2
	Estimation period: 1983/Q1 ~ 1999/Q4 / 2000/Q4

(2) Forecasts from the RR-VAR model (2)

→ 2 period ahead forecast model
→ 1 period ahead forecast model

t t+1 t+2 t+3 t+4

(3a) Top Model Distribution

forecast value from the best-performing model

(3b) Nonparametric Distribution

forecast values from all 2500 models

(3c) Mixed Distribution

forecast values from all 2500 models

(4) Fan Chart Construction

dashed line: fan chart
fine curved line: forecast distribution for each period
thick line: mean of the forecast distribution
black shaded area: adopted critical region

t=0 t=1 t=2

Data Series used in Performing Estimation

	Data	Method of adjustment
0	Consumer Price Index (excluding fresh food, "package tours to overseas," personal computer)	Seasonally adjusted; consumption tax adjusted
1	Domestic wholesale price index	consumption tax adjusted
2	Import price index	
3	Industrial production index (value-added weights)	Seasonally adjusted
4	Index of tertiary industry activity	Seasonally adjusted
5	SNA private consumption (in real terms)	Seasonally adjusted
6	SNA non-residential investment (in real terms)	Seasonally adjusted
7	SNA public investment (in real terms)	Seasonally adjusted
8	SNA exports of goods and services (in real terms)	Seasonally adjusted
9	Rate of unemployed in labor force	Seasonally adjusted
10	Tankan diffusion index of employment conditions	
11	Hourly contractual cash earnings (= contractual cash earnings/ total number of working hours)	Seasonally adjusted
12	Monetary aggregate (M2+CDs) (average outstanding)	Seasonally adjusted
13	TOPIX (average closing)	
14	Government bond yield (10-year)	
15	Nominal effective exchange rate	
16	Import penetration of consumption goods (imports of consumption goods / aggregate supply of consumption goods)	

Notes:

Seasonal adjustment of the CPI is carried out after exclusion of (a) fresh food, (b) package tours to overseas, and (c) personal computers.

This is because each item:

(a) is subject to large seasonal variability depending on climate change;

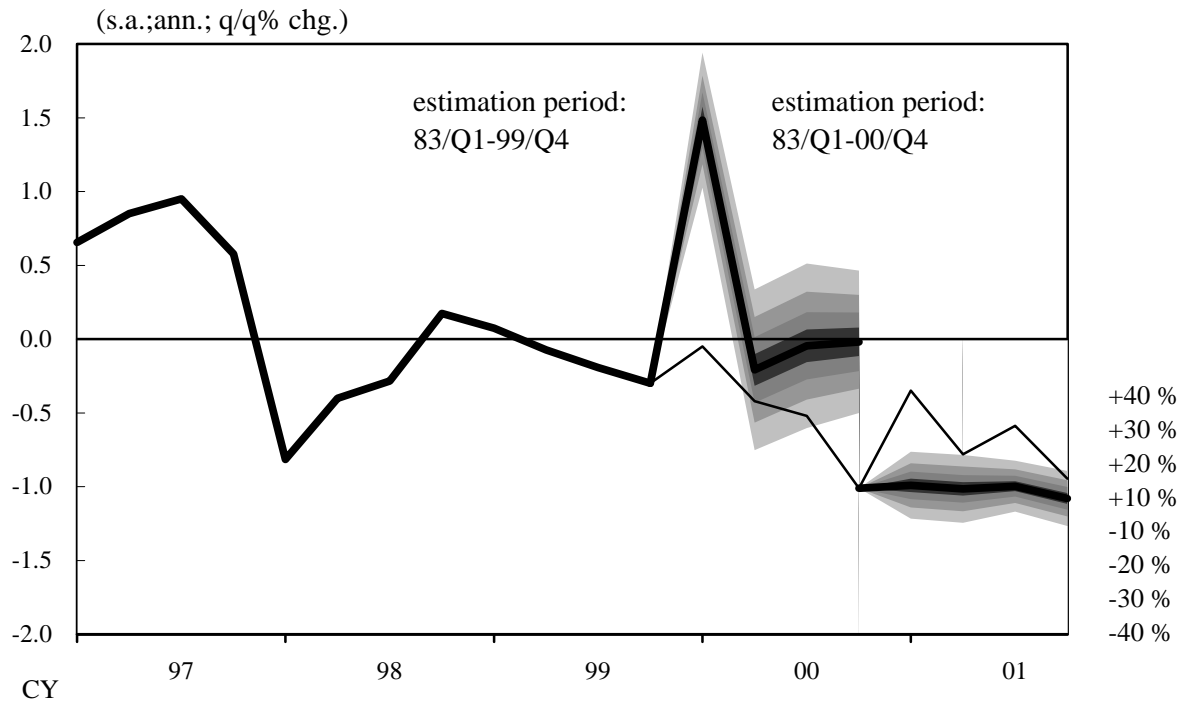
(b) is subject to large seasonal fluctuations, however, since this is newly adopted from the 2000 base, sample size of this item is insufficient for seasonal adjustment;

(c) has exhibited downward trend movement since 2000, because of the new adoption

from the 2000 base and as a result these items have been factors causing disturbance in the overall CPI time series.

Sources: Cabinet office, "National Accounts"; Ministry of Health, Labor, and Welfare, "Monthly Labor Survey"; Ministry of the Economy, Trade, and Industry, "Indices of Industrial Production," "Indices of Industrial Domestic Shipments and Imports," "Indices of Tertiary Industry Activity"; Bank of Japan "Balance of Payment Monthly," "Wholesale price index," "Tankan", Short-term Economic Survey of Enterprises in Japan, "Money Supply," etc.

Fan Charts based on the Top Model Distribution

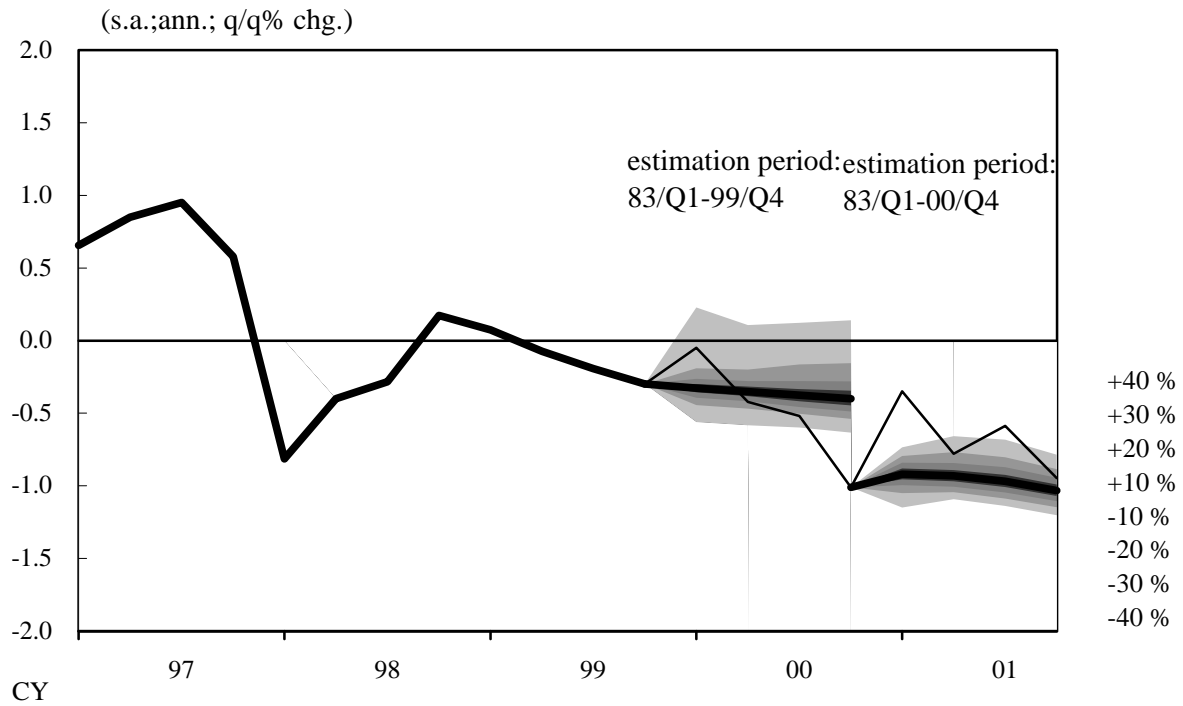


Note:

The thinner line depicts the actual value, and the thicker line the mean of the forecast distribution.

The fan charts illustrate 10% increments in the width of the band around the mean using gradations in the shading.

Fan Charts based on the Nonparametric Distribution

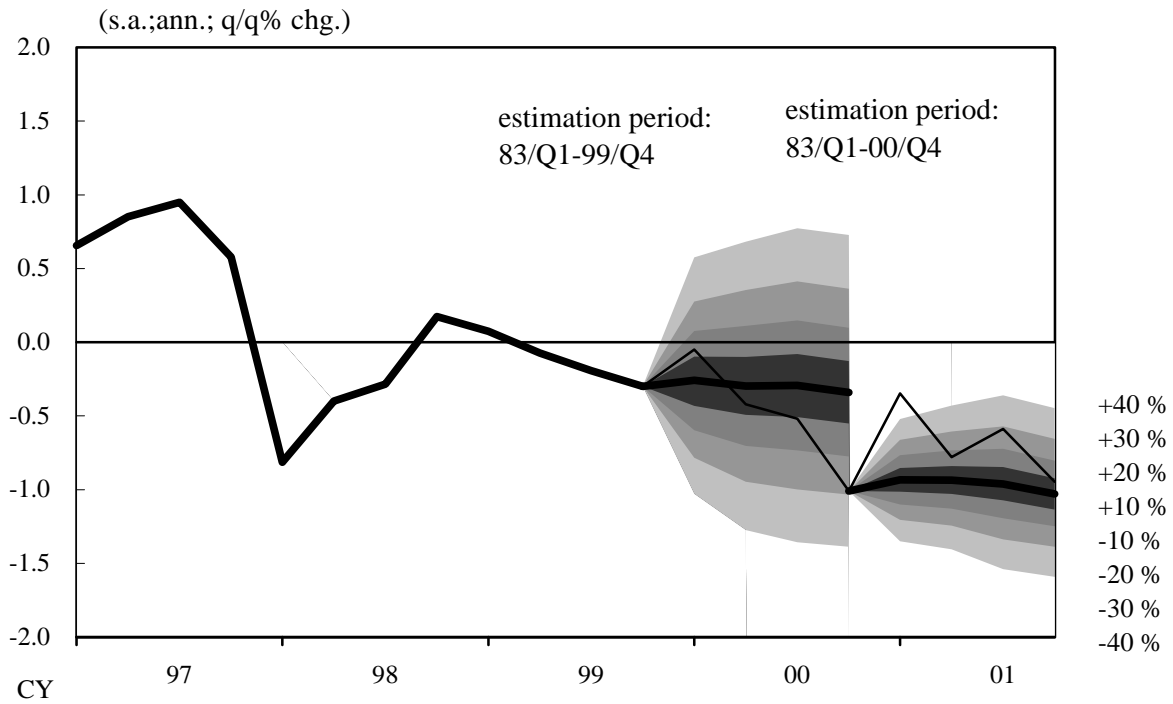


Note:

The thinner line depicts the actual value, and the thicker line the mean of the forecast distribution.

The fan charts illustrate 10% increments in the width of the band around the mean using gradations in the shading.

Fan Charts based on the Mixed Distribution



Note:

The thinner line depicts the actual value, and the thicker line the mean of the forecast distribution.

The fan charts illustrate 10% increments in the width of the band around the mean using gradations in the shading.

A Qualitative Up-Down Forecast based on
the Pesaran-Timmermann Test

		00/1Q	2Q	3Q	4Q	01/1Q	2Q	3Q	4Q
(1)	+	66%	57%	29%	12%	53%	12%	42%	51%
(2)	-	34%	43%	71%	88%	47%	88%	58%	49%
(3)	forecast value	+	+	-	-	+	-	-	+
(4)	actual value	+	-	-	-	+	-	+	-

Note:

Out of the group of models for which the Pesaran-Timmermann statistic lay in the right-side 5% critical region, (1) describes the proportion for which the forecast value was positive (i.e., the inflation rate rose) in that respective period; while (2) describes the proportion for which the forecast value was negative (i.e. the inflation rate fell). By illustrating the case when a larger proportion of models have positive than have negative values with a plus sign, and the reverse case with a minus sign, we express in (3) whether a relatively larger number of the forecasts of the inflation rate from the selected models suggest a rise or a fall. In (4), for the respective period we look at the sign of the change in the actual value of the inflation rate from the value recorded in the previous period: we illustrate a rise with a plus sign, and a fall with a minus sign.