SVERIGES RIKSBANK
WORKING PAPER SERIES

247

SVERIGES
RIKSBANK

Density-Conditional Forecasts in Dynamic Multivariate Models

Michael K. Andersson, Stefan Palmqvist and Daniel F. Waggoner

SEPTEMBER 2010

WORKING PAPERS ARE OBTAINABLE FROM

Sveriges Riksbank • Information Riksbank • SE-103 37 Stockholm
Fax international: +46 8 787 05 26
Telephone international: +46 8 787 01 00
E-mail: info@riksbank.se

Density-Conditional Forecasts in Dynamic Multivariate Models*

Michael K. Andersson†, Stefan Palmqvist‡, and Daniel F. Waggoner§

Sveriges Riksbank Working Paper Series
No. 247

September 2010

Abstract

When generating conditional forecasts in dynamic models it is common to impose the conditions as restrictions on future structural shocks. However, these conditional forecasts often ignore that there may be uncertainty about the future development of the restricted variables. Our paper therefore proposes a generalization such that the conditions can be given as the full distribution of the restricted variables. We demonstrate, in two empirical applications, that ignoring the uncertainty about the conditions implies that the distributions of the unrestricted variables are too narrow.
Keywords: Central Bank, Market Expectation, Restrictions, Uncertainty.
JEL-codes: C 53, E 37, and E 52

*We thank Mattias Villani for helpful suggestions and David Kjellberg for providing the implied forward rates' forecast errors. The views expressed in this paper are solely the responsibility of the authors and should not be interpreted as reflecting the views of the Executive Board of Sveriges Riksbank, the Federal Reserve Bank of Atlanta, or the Federal Reserve System.

†Sveriges Riksbank, Monetary Policy Department, SE-103 37 Stockholm, Sweden. E-mail: Michael.Andersson@riksbank.se.

‡Sveriges Riksbank, Monetary Policy Department, SE-103 37 Stockholm, Sweden. E-mail: Stefan.Palmqvist@riksbank.se.

§Research Department, Federal Reserve Bank of Atlanta, 1000 Peachtree Street N.E., Atlanta, GA 30309-4470. E-mail: dwaggoner@frbatlanta.org

1 Introduction

Inflation targeting central banks commonly publish forecasts of most of the relevant macro-variables in their inflation or monetary policy reports. These forecasts are sometimes published as numbers in a table, where the numbers represent the central tendency of the forecast, and sometimes as graphs containing not only the central tendency but also a "fan chart" representing the uncertainty surrounding the projection.

We use the phrase "most of the relevant macro-variables" since not all central banks publish forecasts of their policy instrument, the short-term nominal interest rate. The Reserve Bank of New Zealand, Norges Bank, the Riksbank, Sedlabanki Islands, and the Czech National Bank publish their own forecasts of the policy instrument, whereas others, like the Bank of England and the Central Bank of Brasil, condition their forecasts on the assumption that the interest rate follows the market expectation, as measured by implied forward rates.[1] Yet others condition their forecasts on the assumption of an unchanged interest rate. We can think of these differences in terms of unconditional vs. conditional forecasts. Central banks that publish their own forecasts of the interest rate can be thought of as publishing unconditional forecasts, whereas the others can be thought of as publishing conditional forecasts, albeit with different conditioning assumptions (market expectation or constant interest rate).

In this paper we address the question of how to generate model-consistent uncertainty bands surrounding a conditional forecast. If we start with the unconditional forecasts we can think of them in terms of a model, like a vector autoregression (VAR). One then simply estimates the VAR and then produce a forecast. This forecast (given an assumption of normality) is characterized by the first and second moments. So, it is pretty simple to generate model-consistent fan charts when the purpose is to produce an unconditional forecast. In a Bayesian setting the forecast distribution is retrieved by direct simulation methods, see for example Thompson and Miller (1986). But how does one go about when the objective is to produce a conditional forecast? One possibility that is nowadays common is to use the procedure of Waggoner and Zha (1999). Their procedure provides a tool to generate forecasts for some variables conditional upon restricitions on other variables' forecasts. Jarocinski and Smets (2008) use this procedure to generate such conditional forecasts, including uncertainty bands for the unrestricted variables.

Our proposed procedure gives an estimate of the full forecast distribution, including second moments as measured by the conditional variances. When we condition the forecast on interest rates following the market expectation there is no uncertainty about the interest rates, that is, the forecasts can be thought of as "hard-conditional" forecasts. In other words, the policy rate is assumed to *exactly* follow the central tendency of what the market expects. The fan charts can therefore be interpreted as how uncertain the forecasts for the non-conditioned variables are given that the interest rate exactly follows what the

[1] For more on this, see Svensson (2008).

market expects.

However, using the market expectation as a predictor of future interest rates generates substantial forecast errors, see for example Adolfson *et al.* (2007). Some banks therefore also report a fan chart for the implied forward rates, indicating that the future policy rate is likely to deviate from what the market expects. The uncertainty surrounding the market expectation as a predictor of the future policy rate is however ignored when producing the conditional forecasts for inflation and GDP with the Waggoner and Zha procedure.

As the market expectation of future interest rates is an uncertain predictor of future interest rates, the conditional variances of the Waggoner and Zha procedure generates uncertainty intervals for unrestricted variables, like inflation and GDP in the example, that are too narrow. In this paper we therefore propose a generalization of the Waggoner and Zha procedure that allows us to incorporate the uncertainty surrounding restricted variables, thus giving more accurate fan charts for the unrestricted variables.

The paper is organized as follows. Section 2 presents our workhorse model, a SVAR, lays out the notation and shows how density restrictions can be imposed. The section also shows how the Waggoner and Zha procedure for both hard and soft conditions relates to our procedure. Section 3 contains two examples that illustrate the procedure. The first example illustrates how the density restrictions can be imposed in order to generate a forecast that is conditioned on the full distribution of the market expectation of future monetary policy. The second example presents how uncertainty surrounding foreign variables affect the uncertainty around the domestic variables. Finally, Section 4 concludes.

2 Forecasting in a Dynamic Model

This section describes a simple dynamic model and shows how it can be used to generate unconditional and conditional forecasts. We start with the unconditional forecasts and then move on to the density-conditional forecasts. We then show how the hard-conditional forecasts of Waggoner and Zha (1999) obtain as a special case of our procedure. Waggoner and Zha (1999) also suggest a soft-conditional procedure, based on a conditional Gibbs-sampler. We show that the soft-conditional forecasts also can be formulated within our setup and use this to develop more efficient simulation techniques.

2.1 Unconditional Forecasts

The procedure will be presented in terms of a structural vector autoregressive model (SVAR), which can be written as

$$\begin{aligned} \sum_{\ell=0}^{L} \underset{(1\cdot n)}{y'_{t-\ell}} \underset{(n\cdot n)}{A_\ell} &= \underset{(1\cdot n)}{d} + \underset{(1\cdot n)}{\varepsilon'_t}, \text{ where} \\ \varepsilon_t &\sim N(0, I_n). \end{aligned} \tag{1}$$

In this setting, y_t is a n-dimensional vector of observable variables at time t. A_l is a $(n \times n)$ matrix of dynamic parameters, L is the maximum lag-length, and d is the model's intercepts. The forecasts from model (1) for horizons $t+1$ to $t+h$, which we denote as $y'_{t+1,t+h} = [y'_{t+1}, \cdots, y'_{t+h}]$, can be written as

$$y'_{t+1,t+h} = b'_{t+1,t+h} + \varepsilon'_{t+1,t+h} M, \tag{2}$$

where $b_{t+1,t+h}$ and M are given in Appendix A and $\varepsilon'_{t+1,t+h} = [\varepsilon'_{t+1}, \cdots, \varepsilon'_{t+h}]$. Both $b_{t+1,t+h}$ and M depend on the model parameters, A_ℓ for $1 \leq \ell \leq L$ and d.

2.2 Density-conditional Forecasts

In this paper, we are not primarily interested in where the restrictions come from, but they may for instance result from another model's forecast or expert judgments. It suffices to know that they exist as a distributional forecast of future observable variables. Our procedure is flexible in that we can handle restrictions on some variables at some horizons, or on all variables and horizons, or on linear functions of the variables in the model. An example of such a linear function is annual growth-rates in a quarterly model or real wages in a model with nominal wages and inflation.

Formally, we assume that we want to condition upon p of the forecasts $Cy_{t+1,t+h}$, where C is a $(p \times nh)$ matrix where each of the rows point at one restriction.[2] The restricted variables' forecasts then follow the distribution

$$Cy_{t+1,t+h} \sim N\left(f_{t+1,t+h}, \Omega_f\right), \tag{3}$$

where $f_{t+1,t+h}$ and Ω_f represent the mean and covariance matrix of the restrictions.

While getting the central tendency of the restrictions, $f_{t+1,t+h}$, typically is straight-forward, it may be somewhat trickier to get the covariance matrix of the restrictions, Ω_f. If we have past forecast errors for the restrictions we may use these to estimate the full covariance matrix. If we have no information about some or all of the elements in Ω_f we may use the model to fill in the missing elements. Using equation (2), we note that the unrestricted distribution of the restricted variables is

$$Cy_{t+1,t+h} \sim N\left(Cb_{t+1,t+h}, DD'\right), \tag{4}$$

where $D = CM'$.[3] If we have no information about Ω_f, the matrix DD' may be used. If we only have information about the diagonal elements of Ω_f, through RMSE-evaluations for instance, we can decompose the unrestricted covariance matrix into

$$DD' = \Xi \odot \left(\sigma\sigma'\right),$$

[2] Note that the restrictions may involve several variables and/or several horizons. If the restriction concerns one variable at one horizon C will be a row-vector with one element set to unity and the rest to zero. If there are restrictions on all variables and all horizons C will be the identity matrix.

[3] See Appendix B for further information.

where Ξ is the correlation matrix of the forecast errors, σ is the vector of standard deviations, and $\odot$ denotes the Hadamard product, which is element by element multiplication. If the only information we have about the covariance matrix of the restrictions are the RMSE:s of each of the restrictions, denoted σ_f, we can use the unrestricted correlation matrix of the forecast errors to form

$$\Omega_f = \Xi \odot \left(\sigma_f \sigma_f'\right). \tag{5}$$

The procedure is thus flexible in that we can use the information we have about some or all of the elements in Ω_f and let the model fill in the missing pieces.

Given the structure of the model (1) and the restrictions we want to impose, the distribution of the structural disturbances necessary to satisfy the density restrictions is

$$\tilde{\varepsilon}_{t+1,t+h} \sim N\left(\mu_{\tilde{\varepsilon}}, \Sigma_{\tilde{\varepsilon}}\right), \tag{6}$$

where $\mu_{\tilde{\varepsilon}}$ and $\Sigma_{\tilde{\varepsilon}}$ are given by

$$\begin{aligned} \mu_{\tilde{\varepsilon}} &= D^* f_{t+1,t+h} - D^* C b_{t+1,t+h}, & (7) \\ \Sigma_{\tilde{\varepsilon}} &= D^* \Omega_f (D^*)' + \widehat{D}' \widehat{D}, & (8) \end{aligned}$$

D^* denotes the generalized inverse of D, and the rows of $\widehat{D}$ form a orthonormal basis for the null space of D. The shocks are such that the density restrictions are satisfied for the restricted variables, and for the unrestricted variables, the model's own structural shocks fill in. Both the mean and variance depend on density restrictions and the model structure. The corresponding forecast distribution is

$$y_{t+1,t+h} \sim N\left(\mu_y, \Omega_y\right), \tag{9}$$

where

$$\begin{aligned} \mu_y &= M' D^* f_{t+1,t+h} + M' \widehat{D}' \widehat{D} (M')^{-1} b_{t+1,t+h}, \\ \Omega_y &= M' \left(D^* \Omega_f (D^*)' + \widehat{D}' \widehat{D} \right) M. \end{aligned}$$

See Appendix B for details on the distributions for $\tilde{\varepsilon}_{t+1,t+h}$ and $y_{t+1,t+h}$. In the examples below we present how the density-conditional procedure may be used with a Bayesian VAR.

2.3 Hard- and Soft-conditional Forecasts

When Ω_f is zero, our density restrictions reduce to the hard-conditions case of Waggoner and Zha (1999). The details of this are given in Appendix C. More interestingly, the soft-conditions of Waggoner and Zha (1999) can also be represented in our framework, but with the density restriction being a truncated normal as opposed to a normal.[4] This is important because, using these

[4] We wish to thank Junior Maih of Norges Bank for making this point.

ideas, we can sample directly from the soft-conditions using a Gibbs sampler as opposed to the rejection method proposed in Waggoner and Zha (1999).

A multivariate truncated normal distribution can be characterized by location and scale parameters and its region of support. We denote these by μ_{TN}, Ω_{TN}, and $\mathfrak{R}$, respectively. The parameters μ_{TN} and Ω_{TN} denote the mean and variance of the underlying normal distribution and will not, in general, be the mean and variance of the truncated normal. While the support could be quite general, in the examples it will be a rectilinear region. Thus, in the one-dimensional case $\mathfrak{R}$ will be an interval, in the two-dimensional case it will be a rectangle with sides parallel to the coordinate axes, and so forth. We will denote the truncated normal distribution by

$$N_T(\mu_{TN}, \Sigma_{TN}, \mathfrak{R}).$$

On its support, the density function of the truncated normal will be proportional to the density of the underlying normal.

The soft-conditions in Waggoner and Zha (1999) can be represented in our framework by using the density restriction

$$Cy_{t+1,t+h} \sim N_T\left(Cb_{t+1,t+h}, DD', \mathfrak{R}\right),$$

where

$$\mathfrak{R} = \{Cy_{t+1,t+h} \mid \underline{c} \leq Cy_{t+1,t+k} \leq \bar{c}\}$$

for some p-dimensional vectors $\underline{c}$ and $\bar{c}$. This density restriction should be interpreted as follows. The underlying normal distribution is the same as that coming from (2) so that the shape of the density is unchanged. However, the i^{th} coordinate of $Cy_{t+1,t+h}$ is restricted to lie between the i^{th} coordinates of $\underline{c}$ and $\bar{c}$. Geweke (1995) proposed an efficient Gibbs sampler for simulating from the truncated normal distribution. In appendix D, we will briefly describe this Gibbs sampler.

3 Examples

This section presents the density-conditional procedure at work in two examples. In the first example we show how to condition a forecast on the market expectation of future policy rates in a Bayesian VAR. We then study the forecast distributions of the Swedish variables when the foreign variables are restricted. The restrictions are given as central tendencies only and as full densities, respectively.

3.1 Density-conditional Forecasts in a Small Open Economy VAR

The first two examples compare forecast distributions when some variables are hard-conditioned on a central tendency to when they are density-conditioned. More precisely, we set up a VAR consisting of three foreign variables (trade weighted GDP, CPI, and short interest rate), three domestic (Swedish) variables (GDP, CPI, and a short interest rate), and a trade weighted real exchange rate. The GDP and CPI variables are measured in first log-differences, the interest rates in levels (annualized return), and the real exchange rate in log-level. The reduced form VAR can be written as

$$\begin{aligned}(y_t - x_t\Phi) &= \sum_{\ell=1}^{4}(y_{t-\ell} - x_t\Phi)\, B_\ell + u_t, \text{ or in mean-adjusted form} \\ \tilde{y}_t &= \sum_{\ell=1}^{4}\tilde{y}_{t-\ell} B_\ell + u_t. \end{aligned} \tag{10}$$

y_t is a row vector of the seven variables (that are to be modelled) at time t and x_t a row vector of a non-modelled (deterministic) variable influencing y_t. The only deterministic variable here is a constant (which is interpreted as the mean or steady state of the processes). B_ℓ and Φ are parameter matrices of dimension (7×7) and (1×7) respectively and u_t is a random error term. The model is estimated on data covering 1993-2006. The reason for selecting 1993 as the first year in the sample is the introduction of a floating exchange rate. The structural form of (10) is obtained through a recursive identification scheme.

The model (10) can be compactly expressed as

$$\tilde{y}_t' = z_t'\Pi + u_t'$$

with $z_t' = \left[\tilde{y}_{t-1}', \ldots,\ \tilde{y}_{t-4}'\right]$, $\Pi' = [B_1, \ldots,\ B_4]$. The random error term, u_t, is normally distributed with mean zero and covariance matrix Σ. The VAR-model is estimated by Bayesian methods, see for instance Kadiyala and Karlsson (1997) for more information on Bayesian VAR estimation.

Prior distribution. In the examples in this paper we use a Normal-Diffuse prior on the parameters in the model, that is, $vec\,(\Pi) \sim N\,(\beta_0, \Psi_0)$ and $p\,(\Sigma) \propto |\Sigma|^{-(L+1)/2}$. The prior for Π is a Minnesota type prior, see Doan *et al.* (1984). The mean, β_0, is zero except for the elements corresponding to the diagonal of the first lag coefficient. These elements will be set to 0.9 for variables in levels and 0.0 for elements in first differences.[5] The matrix Ψ_0 is diagonal with prior standard deviations set to

$$\begin{array}{rl} \frac{\lambda_1}{l^{\lambda_3}} & \text{own lags, } l = 1, ..., 4 \\ \frac{\lambda_1 \lambda_2}{l^{\lambda_3}} \frac{s_i}{s_j} & \text{lags of variable } j \text{ in equation } i \\ \lambda_4 & \text{influence from domestic on foreign variables} \end{array}, \tag{11}$$

where s_i denotes the residual standard deviation for equation i from the OLS fit of univariate AR-specifications. The overall tightness, λ_1 in equation (11), is set to 0.2, the cross variables tightness λ_2 is set to 0.5 and the lag decay hyper parameter λ_3 equals unity. The hyper parameter λ_4 handles that the rest of the world is exogenous to Sweden and is set to a very small number. The prior

[5] The reason for using a value of 0.9 instead of the original 1 is to ensure stationarity, and thus, the existence of the mean. This is crucial when working with the mean-adjusted form.

distribution for the steady state Φ is multivariate normal with mean μ_Φ and covariance matrix Ω_Φ, where

$$\begin{aligned} \mu'_\Phi &= \begin{bmatrix} 0.63 & 0.50 & 5.00 & 0.56 & 0.50 & 4.25 & 3.90 \end{bmatrix}, \\ \sqrt{diag\left(\Omega_\Phi\right)'} &= \begin{bmatrix} 0.064 & 0.064 & 0.255 & 0.032 & 0.006 & 0.128 & 0.051 \end{bmatrix}, \end{aligned}$$

and the covariances in Ω_Φ are all set to zero. The same specification is used by Adolfson *et al.* (2007). Villani (2008) derives the full conditional posterior distribution for the mean-adjusted VAR-model.

Posterior distribution. The posterior distribution is obtained by simulation. Because the distribution of each of the parameters, Π, Φ and Σ, conditional on the other other parameters, is tractable, a Gibbs-sampler can be used to simulate. Following Waggoner and Zha (1999) we augment the data with forecasts from the previous draw in the Gibbs-sampler when simulating the model's parameters. The conditional posterior distribution is thus obtained by the following algorithm. The notation $\Pi^{(i)}$, $\Phi^{(i)}$, and $\Sigma^{(i)}$ denotes the i^{th} draw of the parameters from the augumented posterior and $y^{(i)}_{t+1,t+h}$ denotes a draw from the forecast distribution conditional on the draw of the i^{th} parameters.

1. draw the residual variance $\Sigma^{(i)}$ conditional upon the data augmented by $y^{(i-1)}_{t+1,t+h}$, $\Pi^{(i-1)}$, and $\Phi^{(i-1)}$.

2. draw the dynamic parameters $\Pi^{(i)}$ conditional upon the data augmented by $y^{(i-1)}_{t+1,t+h}$, $\Sigma^{(i)}$, and $\Phi^{(i-1)}$.

3. draw $\Phi^{(i)}$ conditional upon the data augmented by $y^{(i-1)}_{t+1,t+h}$, $\Sigma^{(i)}$, and $\Pi^{(i)}$

Usually forecasts, conditional on the dynamic coefficients and residual variance, are are constructed by making draws of the structural shocks, which are mean zero and independent. In our case, the forecasts must take into account the density restrictions. Also, in this example, Π and Σ are reduced form parameters and must be transformed using the recursive identification scheme[6] to the structural parameters before applying the formulas given by (2) and (6).

4. mean-adjust the restrictions by $\tilde{f}^{(i)}_{t+1,t+h} = f_{t+1,t+h} - x_{t+1,t+h}\Phi^{(i)}$

5. draw structural disturbances $\tilde{\varepsilon}^{(i)}_{t+1,t+h}$ according to (6).

6. compute $\tilde{y}^{(i)}_{t+1,t+h}$ using equation (2)

7. compute $y^{(i)}_{t+1,t+h} = \tilde{y}^{(i)}_{t+1,t+h} + x_{t+1,t+h}\Phi^{(i)}$.

[6] This holds only for exactly identified specifications of the VAR. If over-identifying restrictions are used, the VAR must be estimated on structural form, see Rubio-Ramirez *et al* (2010).

Our first two examples take advantage of the variables as they where known when the Riksbank presented its first monetary policy report 2007. At that point in time National Accounts data up to and including the third quarter of 2006 was known, and the interest rate was known for the full year of 2006. The model forecasts reach up to nine quarters ahead, thus covering the years 2007-2008. We explore our procedure in two different examples, one where we impose restrictions on the domestic policy rate, and one where we impose restrictions on foreign variables. In both examples we illustrate that there may be excess variance shrinkage on the unrestricted variables. This is due to the fact that we ignore the uncertainty surrounding the restricted variables. Maih (2010) takes this one step further and argues that there may still be excess variance shrinkage even if there is no uncertainty about the restricted variables.

3.1.1 Forecasts Conditioned on Market Expectation of Interest Rates

As a baseline scenario it may be interesting to start with an unconditional forecast. These unconditional forecasts are shown in the top panel of Figure 1. The left chart shows the model's forecast for the Riksbank's policy rate, the repo rate, and the right chart shows the corresponding forecast for CPI-inflation. The other five variables in the model are omitted from the figure for expositional purposes. Each forecast is given by a central tendency (the median) and a fan chart representing the 50 and 90 per cent probability intervals. We also include the subsequent outcomes of the repo rate and CPI-inflation (which are of course not used when estimating the model or generating the forecasts).

Suppose now that we are interested in conditioning the forecast on the market expectation of the policy rate, as measured by implied forward rates. This is a conditioning assumption that was previously used by Norges Bank and the Riksbank, and it is used today at the Bank of England and the Central Bank of Brazil. At the point of time we are studying, the market expected a somewhat higher policy rate than the model's unconditional forecast. We therefore use the procedure of Waggoner and Zha, which draws shocks that guarantee that the policy rate follows the market expectation, and substitutes these shocks into the model. The results of this experiment are shown in the middle panel of Figure 1. We note that there is no uncertainty surrounding the future policy rate. This is because the Waggoner and Zha procedure ignores such uncertainty and therefore is referred to as a hard-conditional forecast. The 90 percent probability interval for CPI-inflation is up to 0.7 percentage points narrower than the corresponding interval for the unconditional forecast. We note a similar shrinkage in the forecast distributions of, for example, domestic GDP-growth. This variance shrinkage comes from the fact that the procedure places restrictions on some of the model's shocks, thus reducing the number of shocks that can vary freely.

We may also note that the subsequent policy rates actually deviated from what the market expected. It is therefore somewhat odd to treat the market expectation as an assumption that will hold *exactly*. In the left part of the bottom panel of Figure 1 we show the fan charts surrounding the market expectation of future policy rates.[7] So, there is considerable uncertainty surrounding the market expectation of future policy rates. To incorporate this uncertainty in the density-conditional forecasts we compute the covariance matrix of the risk-premia adjusted implied forward rate's forecast errors at different horizons, which provide the matrix Ω_f.[8,9] The fan chart surrounding the interest rate in the left part of the bottom panel of Figure 1 is narrower than the fan chart surrounding the unrestricted forecast in the top panel of the same figure. As a consequence, the resulting 90 per cent probability interval for CPI-inflation is up to 0.2 percentage points narrower than the unconditional forecast intervals. Compared with the hard-conditional forecasts in the middle panel of Figure 1 the 90 per cent probability intervals are up to 0.6 percentage points wider. There is thus a notable reduction in the uncertainty surrounding the unrestricted variables when using hard conditions.

In a recent paper, Österholm (2009) suggests a procedure to incorporate judgments in fan charts. His procedure amounts to first generating a few alternative scenarios, which then are combined to get a main scenario with fan charts. The different alternative scenarios are constructed by the means of imposing hard conditions in a VAR-model.[10] Since the fan charts surrounding each of those alternative scenarios suffer from excess variance shrinkage, this analysis would benefit from using densisty restrictions.

3.1.2 Domestic Forecasts Conditioned on the Rest of the World

This example investigates when the forecasts for the foreign block of the model are restricted. As restrictions we use the Riksbank's forecasts for the quarterly growth rates of the trade-weighted foreign GDP and CPI, and the trade-weighted foreign interest rate as the centre of the distributions.[11] In this case we have no information about the variances or covariances of the forecast errors, that is, the elements of Ω_f, so we let the model fill in on these entries. The covariance matrix for the restrictions is thus given by DD', see equation (4).

Figure 2 presents forecasts for foreign and Swedish GDP and CPI. The other three variables in the model are omitted from the figure for expositional purposes. The model is conditioned on the foreign development using the hard-conditional procedure (see row 1 in the figure) and our proposed density-conditional procedure (see row 3 in the figure). The resulting forecasts for the domestic variables are given in row 2 (hard-conditioned upon the foreign vari-

[7] The width of these fan charts are taken from Monetary Policy Report 2008:3, Sveriges Riksbank.

[8] There is, of course, no uncertainty around the repo rate for the fourth quarter of 2006, since it was known when the Riksbank published its first Monetary Policy Report of 2007. The variance of the one-step ahead forecast for the repo rate is thus set to zero.

[9] This covariance matrix differs from the covariance matrix that obtains from the unrestricted forecasts in two ways. The first difference is that the variance terms are smaller. The second difference is that the correlations between forecast errors at different horizons are smaller. In our example, almost all of the shrinkage in the probability intervals for unrestricted variables stems from the variance terms in Ω_f.

[10] Österholm does not use the Waggoner-Zha procedure. Instead he merely uses shock i to get a restricted forecast for variable i.

[11] These forecasts are taken from Monetary Policy Report 2007:1, Sveriges Riksbank.

ables) and row 4 (density-conditioned upon the foreign development). When the foreign variables are density-conditioned, with the same mean as in the hard-condition case and the restricted covariance matrix taken from the model, the uncertainty about the Swedish variables naturally increases. The width of the 50 per cent probability interval for the Swedish CPI forecast widens by 0.3 percentage points at the one-step horizon and by 0.6 percentage point at the nine-step horizon when using the density-conditions compared to the hard-conditions. The 50 per cent probability interval for the domestic GDP forecast is 0.1 percentage point wider one step ahead and 0.3 percentage points wider nine steps ahead. The corresponding numbers for the 90 per cent probability intervals are 0.6 and 1.5 percentage points for CPI one and nine steps ahead and 0.3 and 0.9 percentage points for GDP one and nine steps ahead. These results show that there may be a sizeable reduction in the uncertainty of the unrestricted variables when using the hard-conditional procedure. Consequently, to make reliable inference on the forecasts a density-conditional approach is preferable.

4 Conclusions

We have introduced a new method to generate model-based conditional forecasts, where the conditions take the form of densities. While the examples explored in this paper take the conditional density to be normal, that is, we condition on the first and second moments, our procedure allows us to condition on other distributions as well. In addition to the normal distribution, the details of this procedure are also exemplified for the truncated normal distribution.

We demonstrate, in two applications, that ignoring the uncertainty around the conditioned variables may give probability intervals for the unrestricted variables that are too narrow. Our method can therefore solve a practical shortcoming at central banks who base their forecasts on the market expectation of the future path of policy rates or on constant interest rates. Since both those assumptions are uncertain predictors of the future policy rates, our procedure better handles such conditional forecasts. Moreover, our procedure can also be used in a sequential forecasting scheme where exogenous variables are forecasted first and the endogenous variables are then predicted, conditioned upon those exogenous variables. Using density-restrictions may also be a way to check the consistency of a judgemental forecast, since the full distribution of shocks that replicate a judgmental forecast can be solved for and investigated.

References

Adolfson M, Andersson MK, Lindé J, Villani M, Vredin A. 2007., Modern Models in Action: Improving Analyses at Central Banks. *International Journal of Central Banking* **3**:111-144.

Doan T, Litterman RB, Sims CA. 1984. Forecasting and Conditional Projection Using Realistic Prior Distributions. *Econometrics Reviews* **3**: 1-100.

Geweke J. 1995., Bayesian Inference for Linear Models Subject to Linear Inequality Constraints. Working Paper 552. Federal Reserve Bank of Minneapolis, July 1995.

Jarocinski M, Smets F. 2008. House prices and the stance of monetary policy. Federal Reserve Bank of St. Louis Review **90**: 339-65

Kadiyala KR, Karlsson S. 1997. Numerical Methods for Estimation and Inference in Bayesian VAR-models. *Journal of Applied Econometrics* **12**: 99-132.

Maih J. 2010. Conditional Forecasts in DSGE Models. Norges Bank Working Paper No. 2010/07.

Rubio-Ramirez JF, Waggoner DF, Zha TA. 2010. Structural Vector Autoregressions: Theory of Identification and Algorithms for Inference. *Review of Economic Studies* **77**: 665-696.

Svensson LEO. 2008. Transparency under Flexible Inflation Targeting: Experiences and Challenges. Paper prepared for the Riksbank's conference "Refining Monetary Policy: Transparency and Real Stability," Stockholm, September 5-6, 2008.

Sveriges Riksbank. 2007. Monetary Policy Report 2007:1.

Sveriges Riksbank. 2008. Monetary Policy Report 2008:3.

Thomson PA, Miller RB. 1986. Sampling the Future: a Bayesian Approach to Forecasting from Univariate Models. *Journal of Business and Economic Statistics* **4**: 427-436.

Villani, M. 2009. Steady State Priors for Vector Autoregressions. *Journal of Applied Econometrics* **24**: 630-650.

Waggoner DF, Zha TA. 1999. Conditional Forecasts In Dynamic Multivariate Models. *The Review of Economics and Statistics* **81**: 639-651.

Österholm P. 2009. Incorporating Judgement in Fan Charts. *Scandinavian Journal of Economics* **111**: 387-415.

Appendix A: Forecasts from a VAR

The workhorse model in this paper is a VAR with the following structural form

$$\begin{aligned}
\sum_{\ell=0}^{L} \underset{(1\cdot n)}{y'_{t-\ell}} \underset{(n\cdot n)}{A_\ell} &= \underset{(1\cdot n)}{d} + \underset{(1\cdot n)}{\varepsilon'_t}, \text{ where} \qquad (12)\\
\varepsilon_t &\sim N(0, I_n).
\end{aligned}$$

The model (12) has the corresponding reduced form

$$\begin{aligned}
y'_t &= \underset{(1\cdot n)}{c} + \sum_{\ell=1}^{L} \underset{(1\cdot n)}{y'_{t-\ell}} \underset{(n\cdot n)}{B_\ell} + \underset{(1\cdot n)}{\varepsilon'_t A_0^{-1}}, \text{ where}\\
c &= dA_0^{-1}, \text{ and}\\
B_\ell &= -A_\ell A_0^{-1}.
\end{aligned}$$

From the model, the forecasts may be computed as

$$\begin{aligned}
y'_{t+h} &= \underbrace{cK_{h-1} + \sum_{\ell=1}^{L} y'_{t+1-\ell} N_h^l}_{b'_{t+h}} + \sum_{j=1}^{h} \varepsilon'_{t+j} M_{h-j}, \text{ where} \qquad (13)\\
K_0 &= I_n,\\
K_i &= I_n + \sum_{j=1}^{i} K_{i-j} B_j,\\
N_1^\ell &= B_\ell,\\
N_i^\ell &= \sum_{j=1}^{i-1} N_{i-j}^\ell B_j + B_{i+\ell-1},\\
M_0 &= A_0^{-1},\\
M_i &= \sum_{j=1}^{i} M_{i-j} B_j, \text{ and}\\
B_j &= 0 \ \forall \ j > L.
\end{aligned}$$

Equation (13) provides a convenient way to characterize the forecasts. Define also the following matrices

$$\underset{(1\cdot nh)}{y'_{t+1,t+h}} = \left[\begin{array}{cccc} y'_{t+1} & y'_{t+2} & \dots & y'_{t+h} \end{array}\right],$$

$$\underset{(1\cdot nh)}{b'_{t+1,t+h}} = \left[\begin{array}{cccc} b'_{t+1} & b'_{t+2} & \dots & b'_{t+h} \end{array}\right],$$

$$\underset{(1\cdot nh)}{\varepsilon'_{t+1,t+h}} = \left[\begin{array}{cccc} \varepsilon'_{t+1} & \varepsilon'_{t+2} & \dots & \varepsilon'_{t+h} \end{array}\right], \text{ and}$$

$$\underset{(nh\times nh)}{M} = \begin{bmatrix} M_0 & M_1 & M_2 & \cdots & M_{h-1} \\ 0 & M_0 & M_1 & \cdots & M_{h-2} \\ 0 & 0 & M_0 & \cdots & M_{h-3} \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & 0 & \cdots & M_0 \end{bmatrix}.$$

The forecasts for horizons $t+1$ to $t+h$ can then be written as

$$y'_{t+1,t+h} = b'_{t+1,t+h} + \varepsilon'_{t+1,t+h}M.$$

Appendix B: Density-Conditional Forecasts

The density-restrictions are imposed in a similar fashion as in Waggoner and Zha (1999), but we allow the restrictions to be given as densities rather than just central tendencies. Suppose that we want some linear combination of forecasts of observable variables (say GDP, inflation, etc.) to have a normal distribution. This distribution may be used to form a density-condition on a forecast

$$Cy_{t+1,t+h} \sim N\left(f_{t+1,t+h}, \Omega_f\right), \tag{14}$$

where C is a $(p \times nh)$ matrix of rank p. Unrestricted, y will have the distribution given by the model (2)

$$y_{t+1,t+h} \sim N\left(b_{t+1,t+h}, M'M\right). \tag{15}$$

Let $D = CM'$. Now the unrestricted distribution for the restricted variables is

$$Cy_{t+1,t+h} \sim N\left(Cb_{t+1,t+h}, DD'\right).$$

Choose $\widehat{D}$ to be any $(nh-p) \times nh$ matrix such that the rows of $\widehat{D}$ form an orthornormal basis for the null space of D.[12] This implies that

$$\widehat{D}\widehat{D}' = I_{nh-p},\ D\widehat{D}' = \mathbf{0}_{p,nh-p}, \text{ and } \widehat{D}D' = \mathbf{0}_{nh-p,p},$$

where $\mathbf{0}_{i,j}$ denotes the $i \times j$ zero matrix. If

$$z = \begin{bmatrix} z_1 \\ z_2 \end{bmatrix} = \begin{bmatrix} D \\ \widehat{D} \end{bmatrix} (M')^{-1} y_{t+1,t+h} = \begin{bmatrix} Cy_{t+1,t+h} \\ \widehat{D}(M')^{-1} y_{t+1,t+h} \end{bmatrix},$$

then the distribution of z will be

$$N\left(\begin{bmatrix} Cb_{t+1,t+h} \\ \widehat{D}\left(M'\right)^{-1} b_{t+1,t+h} \end{bmatrix}, \begin{bmatrix} DD' & \mathbf{0}_{p,nh-p} \\ \mathbf{0}_{nh-p,p} & I_{nh-p} \end{bmatrix}\right).$$

Equation (14) is equivalent to requiring $z_1 \sim N\left(f_{t+1,t+h}, \Omega_f\right)$. Because z_2 is independent of z_1, we can simply impose this. So, under the restriction given by (14), the restricted distribution of z will be

$$N\left(\begin{bmatrix} f_{t+1,t+h} \\ \widehat{D}\left(M'\right)^{-1} b_{t+1,t+h} \end{bmatrix}, \begin{bmatrix} \Omega_f & \mathbf{0}_{p,nh-p} \\ \mathbf{0}_{nh-p,p} & I_{nh-p} \end{bmatrix}\right).$$

The variable $y_{t+1,t+h}$ can be recovered via

$$y_{t+1,t+h} = M' \begin{bmatrix} D \\ \widehat{D} \end{bmatrix}^{-1} z. \tag{16}$$

Because D has full row rank, if the generalized inverse of D is D^*, then $DD^* = I_p$. Thus

$$\begin{bmatrix} D \\ \widehat{D} \end{bmatrix}^{-1} = \begin{bmatrix} D^* & \widehat{D}' \end{bmatrix}.$$

[12] In Matlab, $\widehat{D}$ can be obtained via the command null$(D)'$.

This implies that the restricted distribution of the $y_{t+1,t+h}$ is $N\left(\mu_y, \Omega_y\right)$, where

$$\mu_y = M' D^* f_{t+1,t+h} + M' \widehat{D}' \widehat{D} (M')^{-1} b_{t+1,t+h}, \text{ and} \tag{17}$$

$$\Omega_y = M' \left(D^* \Omega_f (D^*)' + \widehat{D}' \widehat{D} \right) M. \tag{18}$$

Since $y_{t+1,t+h} = b_{t+1,t+h} + M' \varepsilon_{t+1,t+h}$, we can back out the distribution of the restricted structural shocks. The restricted distribution of $\varepsilon_{t+1,t+h}$, denoted by $\tilde{\varepsilon}_{t+1,t+h}$, is $N\left(\mu_{\tilde{\varepsilon}}, \Sigma_{\tilde{\varepsilon}}\right)$, where

$$\mu_{\tilde{\varepsilon}} = D^* f_{t+1,t+h} - D^* C b_{t+1,t+h}, \text{ and} \tag{19}$$

$$\Sigma_{\tilde{\varepsilon}} = D^* \Omega_f (D^*)' + \widehat{D}' \widehat{D}, \tag{20}$$

where we have used the fact that $\widehat{D}' \widehat{D} + D^* D = I_{nh}$.

This analysis can be extended to restrictions involving distributions other than normal. For example, suppose that we want $C y_{t+1,t+h}$ to have a truncated normal distribution, so that

$$C y_{t+1,t+h} \sim N_T \left(f_{t+1,t+h}, \Omega_f, \mathfrak{R} \right). \tag{21}$$

Following the above analysis, we see that the restricted distributions of z_1 and z_2 are independent and given by

$$z_1 \sim N_T \left(f_{t+1,t+h}, \Omega_f, \mathfrak{R} \right), \text{ and}$$

$$z_2 \sim N \left(\widehat{D} \left(M' \right)^{-1} b_{t+1,t+h}, \widehat{D} \widehat{D}' \right).$$

As before, the variable $y_{t+1,t+h}$ can be recovered by (16), which can be written as

$$y_{t+1,t+h} = M' D^* z_1 + M' \widehat{D}' z_2.$$

Similarly, $\tilde{\varepsilon}_{t+1,t+h}$ can be written as

$$\tilde{\varepsilon}_{t+1,t+h} = D^* z_1 + \widehat{D}' z_2 - (M')^{-1} b_{t+1,t+h}.$$

Because a linear transformation of a truncated normal is a truncated normal, both $y_{t+1,t+h}$ and $\tilde{\varepsilon}_{t+1,t+h}$ can be represented as a sum of a truncated normal and a normal. Following Geweke (1995), in the Appendix D we show how to simulate from a truncated normal when the truncated region is rectilinear.

Appendix C: Hard Conditions

In Waggoner and Zha (1999), hard-conditions on the structural shocks were investigated. There, restrictions were of the form $R' \varepsilon_{t+1,t+h} = r$ and the main result was that the restricted structral shocks were normal with mean μ_R and variance Σ_R where

$$\mu_R = R (R'R)^{-1} r$$

$$\Sigma_R = I_{nh} - R (R'R)^{-1} R'.$$

We show that $\mu_{\tilde{\varepsilon}} = \mu_R$ and $\Sigma_{\tilde{\varepsilon}} = \Sigma_R$ when $\Omega_f = 0$. The restriction $\Omega_f = 0$ gives $Cy_{t+1,t+h} = f_{t+1,t+h}$. So, $f_{t+1,t+h} = Cy_{t+1,t+h} = Cb_{t+1,t+h} + CM'\varepsilon_{t+1,t+h}$, which implies $R' = CM' = D$ and $r = f_{t+1,t+h} - Cb_{t+1,t+h}$. The results will follow from the singular value decomposition of D.

Let $D = UdV'$, where U is an $p \times p$ orthogonal matrix, d is a $p \times p$ invertible diagonal matrix, and V is a $nh \times p$ matrix with orthonormal columns. Because the rows of $\widehat{D}$ form an orthonormal basis for the null space of D, it follows that $[V \ \widehat{D}']$ is an orthogonal matrix and thus $I_{nh} - VV' = \widehat{D}'\widehat{D}$. Also, $D^* = Vd^{-1}U'$. Thus,

$$\begin{aligned}
\mu_R &= R(R'R)^{-1}r = D'(DD')^{-1}(f_{t+1,t+h} - Cb_{t+1,t+h}) \\
&= VdU'(UdV'VdU')^{-1}(f_{t+1,t+h} - Cb_{t+1,t+h}) \\
&= Vd^{-1}U'(f_{t+1,t+h} - Cb_{t+1,t+h}) \\
&= D^*(f_{t+1,t+h} - Cb_{t+1,t+h}) = \mu_{\tilde{\varepsilon}},
\end{aligned}$$

and

$$\begin{aligned}
\Sigma_R &= I_{nh} - R(R'R)^{-1}R' = I_{nh} - D'(DD')^{-1}D \\
&= I_{nh} - VdU'(UdV'VdU')^{-1}UdV' \\
&= I_{nh} - VV' \\
&= \widehat{D}'\widehat{D} = \Sigma_{\tilde{\varepsilon}}.
\end{aligned}$$

■

Appendix D: A Gibbs Sampler for the Truncated Normal Distribution

We first develop some notation. Let v be a p-vector and V be a $p \times p$ matrix. As usual, let v_i denote the i^{th} element of v, let V_i denote the i^{th} column of V, and let $V_{i,i}$ denote the element in the i^{th} row and column of V. Let v^{-i} denote the $(p-1)$-vector obtained by removing the i^{th} element from v and let V^{-i} denote the $(p-1) \times (p-1)$ matrix obtained by removing the i^{th} row and column from V. Thus V_i^{-i} will be the $(p-1)$-vector obtained by removing the i^{th} element from the i^{th} column of V.

Let x be a p-dimensional truncated normal random varible with

$$x \sim N_T(\mu, \Sigma, \mathfrak{R}),$$

where

$$\mathfrak{R} = \{x \mid \underline{c} \leq x \leq \bar{c}\}.$$

The distribution of x_i, conditional on x^{-i}, will be a one dimensional truncated normal random variable with

$$x_i \mid x^{-i} \sim N_T(\tilde{\mu}_i, \tilde{\sigma}_i^2, \mathfrak{R}_i),$$

where

$$\begin{aligned}\tilde{\mu}_i &= \mu_i + (\Sigma_i^{-i})'\Sigma^{-i}(x^{-i} - \mu^{-i}) \\ \tilde{\sigma}_i^2 &= \Sigma_{i,i} - (\Sigma_i^{-i})'\Sigma^{-i}\Sigma_i^{-i}\end{aligned}$$

and

$$\Re_i = \{x_i \mid \underline{c}_i \le x_i \le \bar{c}_i\}.$$

Drawing directly from a one dimensional truncated normal is straightforward. If $\phi(\cdot, \tilde{\mu}_i, \tilde{\sigma}_i^2)$ denotes the cumulative density function for the normal distribution with mean $\tilde{\mu}_i$ and variance $\tilde{\sigma}_i^2$, then

$$\phi^{-1}\left(\phi(\underline{c}_i) + u\left(\phi(\bar{c}_i) - \phi(\underline{c}_i)\right)\right), \tag{22}$$

where u is drawn from the uniform distribution on $(0, 1)$, will be a variate from the required truncated normal distribution. Since the inverse cumulative normal density is a standard function in Matlab and many other packages, this technique is easy to implement and will give accurate results as long $\underline{c}_i$ is not too far out in the *upper* tail or $\bar{c}_i$ is not too far out in the *lower* tail. However, more accurate and efficient simulation techniques are available for the one-dimensional truncated normal distribution. See Geweke (1995) for a description of these techniques.

Using these ideas, a Gibbs sampler can be constructed to make draws of x by successively making draws of $x_i \mid x^{-1}$. Our experience is that as long as the variance matrix Σ is not too close to being singular, the Gibbs sampler will be efficient.

Figure 1: Forecasts Conditioned on Different Policy Rate Assumptions

Unconditional Forecasts

Policy Rate

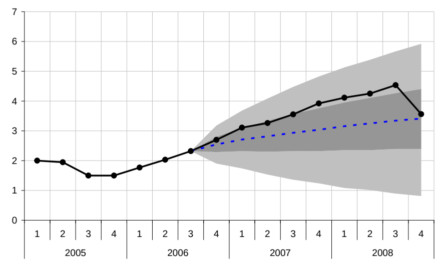

CPI

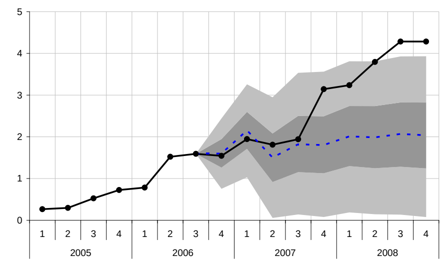

Hard-conditional Forecasts (Waggoner-Zha)

Policy Rate

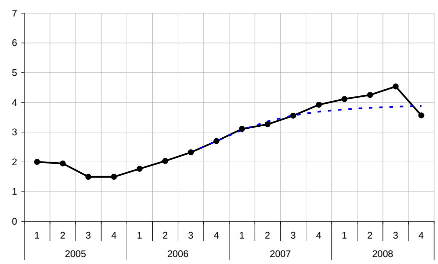

CPI

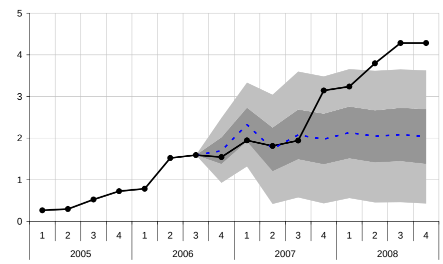

Density-conditional Forecasts

Policy Rate

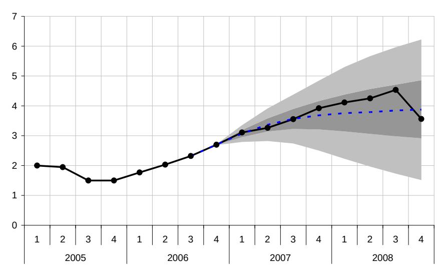

CPI

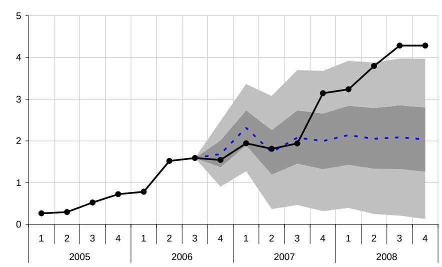

Note: The solid lines with dots represent the outcomes available as of August 2010. The dashed line represents the median of the forecast distribution and the shaded areas the 50 and 90 per cent probability intervals.

Figure 2: Forecasts Conditioned on Foreign Variables

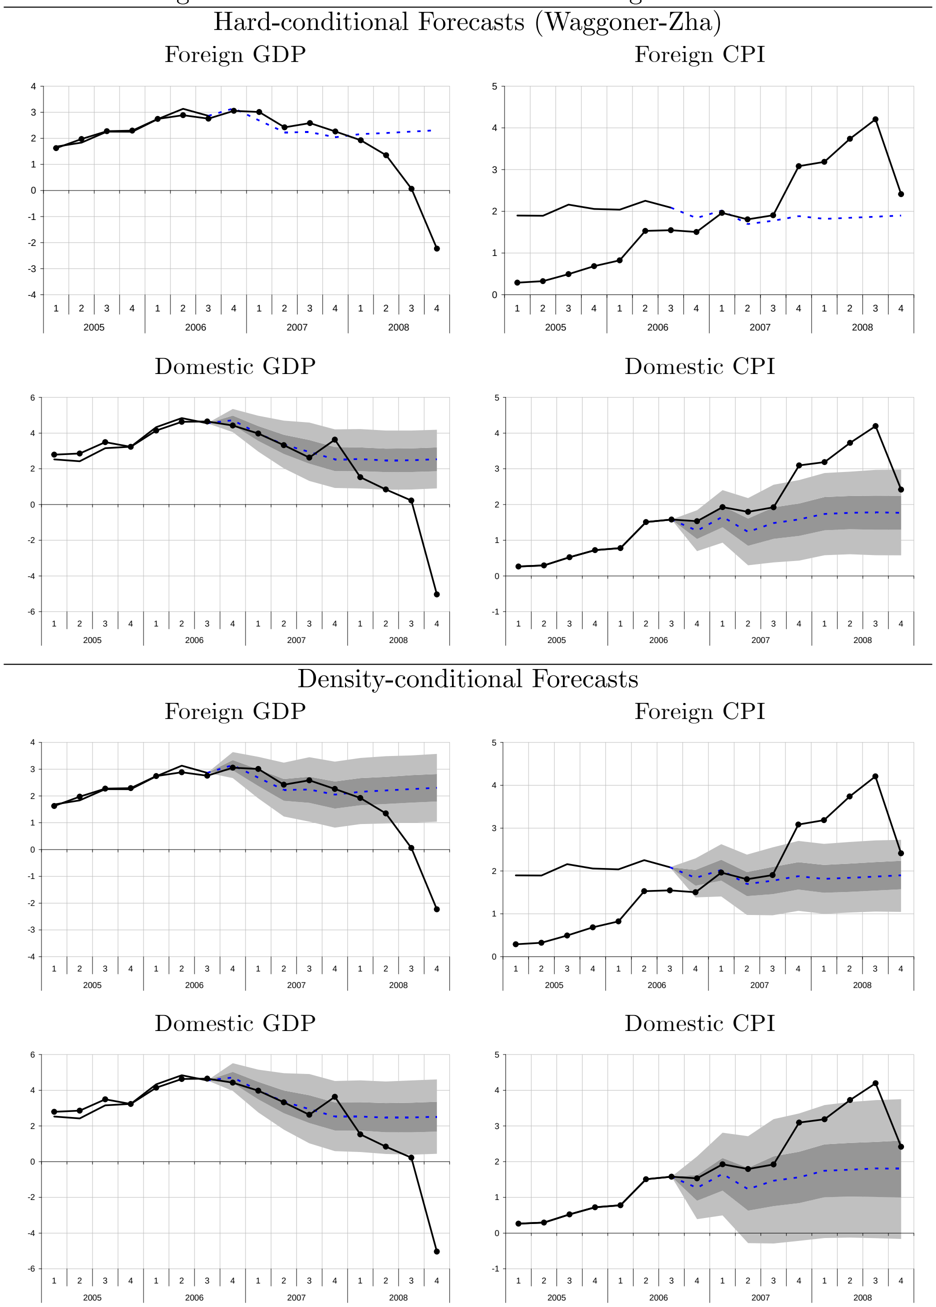

Note: The solid line in the figures represents outcomes available at the time of the forecasts, the dashed line the median of the forecast distribution and the grey areas the 50 and 90 per cent probability intervals respectively. The solid lines with dots are the outcomes available as of August 2010.

Earlier Working Papers:

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

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

ISSN 1402-9103 E-Print, Stockholm 2010