

Modern Forecasting Models in Action: Improving Macroeconomic Analyses at Central Banks*

Malin Adolfson,^a Michael K. Andersson,^a Jesper Lindé,^{a,b}
Mattias Villani,^{a,c} and Anders Vredin^a

^aSveriges Riksbank

^bCEPR

^cStockholm University

There are many indications that formal methods are not used to their full potential by central banks today. In this paper, using data from Sweden, we demonstrate how BVAR and DSGE models can be used to shed light on questions that policymakers deal with in practice. We compare the forecast performance of BVAR and DSGE models with the Riksbank's official, more subjective forecasts, both in terms of actual forecasts and root mean-squared errors. We also discuss how to combine model- and judgment-based forecasts, and show that the combined forecast performs well out of sample. In addition, we show the advantages of structural analysis and use the models for interpreting the recent development of the inflation rate through historical decompositions. Last, we discuss the monetary transmission mechanism in the models by comparing impulse-response functions.

JEL Codes: E52, E37, E47.

*We are grateful to Jan Alsterlind for providing the implicit-forward-rate data, to Josef Svensson for providing real-time data on GDP, and to Marianne Nessén, Jon Faust, Eric Leeper, and participants in the “Macroeconomics and Reality—25 Years Later” conference in Barcelona in April 2005, for comments on an earlier version of this paper. 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. Address for correspondence: Anders Vredin, Sveriges Riksbank, SE-103 37 Stockholm, Sweden. E-mail: anders.vredin@riksbank.se.

1. Introduction

Over the last two decades, there has been a growing interest among macroeconomic researchers in developing new tools for quantitative monetary policy analysis. In this paper, we examine how the new generation of formal models can be used in the policy process at an inflation-targeting central bank. We compare official forecasts published by Sveriges Riksbank (the central bank of Sweden) with forecasts from two structural models—a dynamic stochastic general equilibrium (DSGE) model and an identified vector autoregression (VAR) estimated with Bayesian methods. We also discuss how the formal models can be used for storytelling and for the analysis of alternative policy scenarios, which is an important matter for central banks.

It is rather unusual that formal models are contrasted to official central bank forecasts. Such a comparison is especially interesting given that the latter also include judgments and “extra-model” information, which are very hard to capture within a formal setting. Hence, this can give an assessment of how useful expert knowledge is and shed light on the “role of subjective forecasting” raised by Sims (2002). An evaluation of official forecasts and formal model forecasts has, to our knowledge, only been conducted for the Board of Governors of the Federal Reserve System (see Altig, Carlstrom, and Lansing 1995 and Sims 2002).¹ We will supplement this analysis by evaluating official and model forecasts for a small open economy (Sweden).

Smets and Wouters (2004) have shown that modern *closed-economy* DSGEs (with various nominal and real frictions) have forecasting properties well in line with more empirically oriented models such as standard and Bayesian VARs (BVARs). However, this paper evaluates forecasts from DSGE and BVAR models that include *open-economy* aspects and compares these with judgmental forecasts. Given the increased complexity of open-economy models, and the different monetary policy transmission mechanism where exchange rate movements are of importance, our exercise adds an

¹Before this paper was completed, Edge, Kiley, and Laforte (2006) compared the forecasting performance of the Federal Reserve with a closed-economy DSGE as well as with a theoretical reduced-form model.

extra element to the analysis in, e.g., Sims (2002) and Smets and Wouters (2004). In section 2, we show the actual inflation and interest rate forecasts from the various setups as well as the root mean-squared errors (RMSEs) of the forecasts. In addition, we explicitly examine two episodes where official and DSGE forecasts diverge to look further into the role of subjective forecasting.

Both the DSGE and BVAR models have now been used within the policy process at Sveriges Riksbank between one and two years. However, this does not provide us with enough observations to carry out an extensive forecast evaluation in genuine real time. Since the DSGE model uses data on as many as fifteen macroeconomic variables, we have therefore estimated the formal models using a revised data set, which makes a much longer evaluation period of the various forecasting methods possible. We discuss the role of real-time versus revised data in detail in section 2 below.

In addition, we demonstrate how formal models can shed light on practical policy questions. In section 3.1, we let the two models interpret the underlying reasons for the recent economic development, using a historical decomposition of the forecast errors in each of the two models. We also show that a VAR model can be used to clarify what has happened in the economy, as long as we are willing to impose some structure to identify the underlying shocks. However, when one is interested in predictions conditioned upon alternative policy scenarios, an idea about how monetary policy is designed and how it affects the economy is required. By comparing impulse-response functions in section 3.2, we show that the DSGE model, using structure from economic theory, provides a much more reasonable transmission mechanism of monetary policy than the BVAR model. This is a necessary requirement for producing conditional forecasts that are meaningful from the perspective of a central banker.

Finally, section 4 summarizes our views on the advantages of formal methods and the reasons why such methods have not been more influential at central banks.

2. Forecasting Performance

This section provides an evaluation of the inflation, interest rate, and GDP forecasts by the Riksbank, a small-open-economy DSGE

model, and a Bayesian VAR model. Both actual forecasts and root mean-squared errors are examined for the period 1999:Q1–2005:Q4.² To characterize the forecasting advantage of the Riksbank's judgmental forecasts, we also analyze two specific episodes where the different forecasting approaches diverge for CPI inflation and where we, a priori, expect the sector experts to have an informational advantage. Since the Riksbank's forecasts have until recently been intended to be conditioned on the assumption of a constant short-term interest rate, we also look at the interest rate forecasts from the formal models.³ These are compared with implicit-interest-rate forecasts calculated from (market-based) forward interest rates. Finally, we also look at the GDP forecasts from the different models, but since no genuine real-time data set has been compiled for all fifteen variables in the DSGE model, these results should be interpreted with some caution.⁴

²Official inflation forecasts from the Riksbank cannot be obtained on a quarterly basis before 1999:Q1. Moreover, the DSGE model needs a sufficient number of observations after the transition from a fixed exchange rate to the inflation-targeting regime in 1993. Since GDP forecasts are not available from the Riksbank before 2000:Q1, their precision is evaluated between 2000:Q1 and 2005:Q4.

³Between October 2005 and February 2007, the Riksbank produced forecasts conditioned upon implicit forward rates, instead of the constant-interest-rate assumption. Since February 2007, the Riksbank has published its preferred path for the future repo rate (i.e., the short-term interest rate controlled by the Riksbank).

⁴The absence of a real-time data set implies that the formal models possibly have an information advantage relative to the official forecasts since the BVAR and DSGE forecasts are based on ex post data on GDP. (Preliminary GDP data are available with a delay of around one quarter, but they are subsequently revised.) This is not a problem with the other variables in the models, since data on prices, interest rates, and exchange rates are available on a monthly basis and are not revised. On the other hand, the official forecasts have a small information advantage in some quarters, when the *Inflation Report* has been published toward the end of the quarter and, thus, can be based on data on prices, interest rates, and exchange rates from the early part of the same quarter. Since most of the variables in the BVAR are available in real time, and we use a Litterman (i.e., random-walk) prior for the lag polynomial, there are good reasons to believe that the BVAR results with the exception of GDP growth are not very sensitive to our decision to use revised data. The DSGE model, which is estimated on a data set that contains several additional real quantities, is probably more sensitive to the use of revised data.

The Riksbank publishes official forecasts in its quarterly *Inflation Report* publication. The forecasts are not the outcome of a single formal model but rather the result of a complex procedure with input both from many different kinds of models and judgments from sector experts and the Riksbank's executive board.⁵

The BVAR model contains quarterly data for the following seven variables: trade-weighted measures of foreign GDP growth in logs (y_f), CPI inflation (π_f) and a short-term interest rate (i_f), the corresponding domestic variables (y , π , and i), and the level of the real exchange rate defined as $q = 100(s + p_f - p)$, where p_f and p are the foreign and domestic CPI levels (in logs) and s is the (log) trade-weighted nominal exchange rate. More details on the BVAR are provided in appendix 1.

The DSGE model is an extension of the closed-economy models developed by Altig et al. (2003) and Christiano, Eichenbaum, and Evans (2005) to the small-open-economy setting in a way similar to that of Smets and Wouters (2002). Households consume and invest in baskets consisting of domestically produced goods and imported goods. We allow the imported goods to enter both aggregate consumption and aggregate investment. By including nominal rigidities in the importing and exporting sectors, we allow for short-run incomplete exchange rate pass-through to both import and export prices. The foreign economy is exogenously given by a VAR for foreign inflation, output, and the interest rate. The DSGE model is estimated with Bayesian methods using data on the following fifteen variables: GDP deflator inflation, real wage, consumption, investment, real exchange rate, short-run interest rate (repo rate), hours, GDP, exports, imports, CPI inflation, investment deflator inflation, foreign (i.e., trade-weighted) output, foreign inflation, and foreign interest rate. The DSGE model is identical to the one developed and estimated by Adolfson et al. (forthcoming), and a more

⁵The forecast process during the evaluation period was of a recursive and iterative nature, where the foreign and financial variables entered first. Given these forecasts, the Swedish real variables were predicted. Typically, the GDP forecast was an aggregation of the components of the GDP identity. The labor-market variables entered in a third step, where the productivity and unit-labor-cost variables were determined. Finally, predictions of CPI and core inflation—mainly based on forecasts of import prices, unit labor cost, and the output gap—ended the first forecast round.

detailed description of the model, priors used in the estimation, and full-sample estimation results can be found in that paper.

The DSGE model contains more variables than the BVAR, but this does not imply that the DSGE model necessarily has an advantage in terms of the forecasting performance, since the inclusion of more variables in the BVAR also implies that a considerably larger number of parameters needs to be estimated. Adolfson et al. (forthcoming) consider a BVAR with the same set of variables as in the DSGE model, and the forecasting performance of this BVAR is not very different from the BVAR used in this paper.⁶

One difficulty when it comes to comparing official inflation forecasts from the Riksbank with model forecasts (or forecasts made by other institutions) is that the Riksbank's forecasts have until recently been intended to be conditioned on the assumption that the short-term interest rate (more specifically, the Riksbank's instrument, the repo rate) remains constant throughout the forecasting period. However, it is not unreasonable to assume that the official forecast actually lies closer to an unconditional forecast, given its subjective nature. In practice, it is extremely difficult to ensure that the judgmental forecast has been conditioned on a constant-interest-rate path rather than on some more likely path.⁷ Beyond the forecast horizon, the implicit assumption also seems to have been that the interest rate gradually returns to a level determined by some interest rate equation. Even so, we see from figure 1 (shown on the next set of facing pages) that the interest rate level (annualized average of daily repo-rate observations within each quarter) has been rather stable during the sample period, which is the reason why a

⁶If anything, the introduction of additional variables leads to a reduction in forecasting performance. In particular, this appears to be the case for the nominal interest rate.

⁷The *Inflation Reports* from March and June 2005 contain discussions of the problems with constant-interest-rate forecasts, as well as two forecasts of inflation and GDP growth that are conditional on either a constant interest rate or the implied forward rate. Although there was a considerable difference between the forward rate and the constant interest rate on these occasions (a gradual increase to about 150 basis points at the two-year horizon), the forecasts for GDP and inflation were not very different. This suggests that the official constant-interest-rate forecasts were in fact close to unconditional projections. See also Adolfson et al. (2005) for difficulties with constant-interest-rate forecasts in a model-based environment.

conceivable constant-interest-rate assumption might not have been that unfavorable. Moreover, the impulse-response functions from interest rate changes to inflation are typically fairly small, with a substantial time delay, at least in the BVAR model (see figure 6 in section 3.2), which implies that a constant-interest-rate assumption at long horizons should have relatively small effects on the inflation forecast. We therefore believe that, taken together, our forecast comparisons are valid and our conclusions will not be significantly affected by the alleged constant-interest-rate assumption underlying the official forecasts.

2.1 Inflation Forecasts

Figure 1a presents the outcome of CPI inflation in Sweden for 1998–2005 (bold line) together with the Riksbank’s official forecasts (first row), the forecasts from the DSGE model (second row), and the BVAR model (third row) for 1999:Q1–2005:Q4.⁸ The visual impression is that the official forecasts and the model forecasts have somewhat different properties. The Riksbank’s forecasts appear to be rather “conservative”; most often they predict a very smooth development of inflation, and the changes in the forecast paths are relatively small between quarters. The BVAR model, on the other hand, seems to view the inflation process as more persistent, since forecast errors have a larger influence on subsequent forecasts.

The top panel in figure 2 shows the root mean-squared errors (RMSEs) for different forecast horizons (one to eight quarters ahead) of the yearly CPI-inflation forecasts. The DSGE and the official forecasts have about the same precision for inflation forecasts made up to a year ahead; however, at somewhat longer horizons (five to eight quarters ahead), the forecasts from the DSGE model have a better accuracy than the official inflation forecasts. The BVAR forecasts, on the other hand, perform very well up to six quarters ahead but are beaten by the DSGE’s inflation forecasts at longer horizons. It should also be noted that all three forecasting methods are

⁸The data in figure 1 refer to yearly inflation rates ($p_t - p_{t-4}$), just like the inflation series published in the *Inflation Report*. Since the DSGE and BVAR models are specified in terms of quarterly rates of change ($p_t - p_{t-1}$), the model forecasts are summed up to fourth differences.

Figure 1a. Sequential Forecasts of Yearly CPI Inflation, 1999:Q1–2005:Q4, from the Riksbank (First Row), the DSGE Model (Second Row), and the BVAR Model (Third Row)




Figure 1b. Sequential Forecasts of Repo Rate, 1999:Q1–2005:Q4, from the Riksbank (First Row), the DSGE Model (Second Row), and the BVAR Model (Third Row)




Figure 1c. Sequential Forecasts of Yearly GDP Growth, 1999:Q1–2005:Q4, from the Riksbank (First Row), the DSGE Model (Second Row), and the BVAR Model (Third Row)





Figure 2. Root Mean-Squared Error (RMSE) of Yearly CPI-Inflation Forecasts (Top) and Annualized Interest Rate Forecasts (Middle) 1999:Q1–2005:Q4, and Yearly GDP-Growth Forecasts (Bottom) 2000:Q1–2005:Q4



better than a naive forecast of constant yearly inflation (the No Change—Yearly line). We take the similar precision in the different inflation forecasts as evidence in favor of all three forecasting approaches. It is encouraging that the model forecasts are performing as well as the official forecasts also at shorter horizons. We will, however, return to the specific advantages and the role of subjective forecasting below.

Although Smets and Wouters (2004) have shown that the forecasting performance of closed-economy DSGE models compares quite favorably to more empirically oriented models such as vector autoregressive models, it is not evident that our DSGE model will have similar properties, given that the open-economy dimensions add complexity to the model. In particular, it is well known that uncovered interest rate parity (UIP) is rejected empirically, which may deteriorate the forecasting performance of an open-economy DSGE model. The UIP condition in the DSGE model has therefore been modified to allow for a negative correlation between the risk premium and expected exchange rate changes; see Adolfson et al. (forthcoming) for further details. Figure 2 reveals that our DSGE model has a forecasting performance for CPI inflation that is remarkably good. The DSGE model makes smaller forecast errors for inflation six to eight quarters ahead than both the BVAR and the official Riksbank forecasts. Part of this can be attributed to the modified UIP condition (see Adolfson et al. forthcoming), but we also believe that the fact that we are considering a stable regime with a known and fixed inflation target makes the theoretical structure imposed by the DSGE model particularly useful.

2.2 Interest Rate Forecasts

Figure 1b shows outcomes of the repo rate along with either the expected interest rate paths derived from forward interest rates (first row), following the method described by Svensson (1995), or along with the repo-rate forecasts from the DSGE model (second row) and the BVAR model (third row). Throughout our sample, forward interest rates have systematically overestimated the future interest rate level. In principle, this could be due to (possibly time-varying)

term and risk premia; i.e., forward rates do not provide direct estimates of expectations. But, in practice, we believe it to be more likely that the market (along with the Riksbank and the DSGE model) on average overestimated the inflation pressure and the need for higher nominal interest rates during this period. Turning to the interest rate forecasts from the formal models, we see that future interest rates have been systematically overestimated also in this case, even if the BVAR forecasts have been more moderate than the interest rate forecasts from the DSGE model. The main reason for the relatively large forecast errors of the DSGE model for the repo rate is that the model has tended to overestimate both the inflation pressure and the GDP-growth prospects during the period, which can be seen in figures 1a and 1c. This indicates that it is not the estimated policy rule that does not fit the data: had we used the true inflation and output outcomes when projecting with the rule, the forecast errors for the repo rate would have been greatly reduced.

In the middle panel of figure 2, we compare the RMSEs for the various repo-rate forecasts. It can be seen that forward interest rates, a naive constant-interest-rate forecast, and the BVAR model all have about the same precision for forecasts three to four quarters ahead, whereas the DSGE has somewhat worse accuracy. Naturally, this raises some questions about how monetary policy is described in the DSGE model, although the above reasoning suggests that the policy rule itself may not be the key problem. Forward rates have a better precision one to two quarters ahead, while the BVAR model makes much better forecasts for longer horizons than the other approaches. The fact that the BVAR model provides a more realistic picture than forward rates at longer horizons, which even seem to have a lower predictive power than the constant-interest-rate assumption, may be interpreted as arguments against inflation forecasts conditioned on forward rates. Nevertheless, this is an assumption that the Bank of England recently adopted, and Norges Bank and Sveriges Riksbank recently abandoned. It should be emphasized, however, that our results have been obtained from a sample where the short-term interest rate has been unusually low and stable in a historical context. Thus, it is possible that forward rates are more informative in periods when interest rates change more.

2.3 GDP Forecasts

Figure 1c shows the forecasts for the yearly GDP growth from the Riksbank, the DSGE, and the BVAR models against actual yearly GDP growth (last data vintage in all cases). Given that the Riksbank's forecasts are made in real time, whereas the formal models are estimated using a revised data set for GDP, a direct comparison between them is hard to make. The pattern between the DSGE and BVAR forecasts is relatively similar, however. This also shows up in the root mean-squared errors for the forecasts in figure 2. The models' accuracy in terms of predicting GDP growth is almost the same in this sample (2000:Q1–2005:Q4).

To increase the comparability of the forecasting performances, we compute the RMSEs for the Riksbank's GDP forecasts in real-time data, since the formal models would have an informational advantage if all these GDP forecasts were evaluated on the revised data set. The models appear to perform reasonably well also against the judgmental forecasts of the Riksbank (see figure 2), although it should be emphasized that such a comparison should be interpreted with caution.


2.4 The Role of Subjective Forecasting

To obtain more information about the properties of the various approaches to forecasting, it is interesting to take a closer look at some specific episodes where the pure model forecast and the official Riksbank forecast differ. This provides us with useful information about whether sector experts can provide a better understanding of recent influences on inflation (which might only be temporary). In figure 3, we compare the Riksbank's official inflation forecasts with the corresponding DSGE forecasts on four different occasions (i.e., figure 3 contains a subset of the information in figure 1).⁹

The first episode concerns forecasts made immediately before and after the sudden increase in inflation in 2001:Q2. One important factor behind this increase was a rise in food prices due to the mad-cow and foot-and-mouth diseases, although the inflation rate had

⁹For ease of exposition, we only focus on the Riksbank and DSGE forecasts here.

Figure 3. Riksbank and DSGE Forecasts of Yearly CPI Inflation Made on Four Specific Occasions: 2000:Q4, 2001:Q2, 2003:Q1, and 2003:Q3



started to increase already in 1999. In 2000:Q4, the Riksbank's official forecast implied a slowly increasing inflation rate over the next two years. The DSGE model suggested a much stronger increase in inflation. This may reflect that the model forecast attributed a larger weight to recent increases in inflation, while the subjective procedure, leading up to the Riksbank's official forecast, underestimated the persistence in changes in inflation. However, once the shock had become apparent, the subjective approach proved to be very useful. At that time, 2001:Q2, the sector experts expected the food price increase to involve a persistent shock to the price level but with small further effects on the yearly inflation rate, and the Riksbank's forecasts at 2001:Q2 were more in line with the actual outcome for the next few quarters. The DSGE model, on the other hand, treated the food price shock as any other inflation shock and overestimated its effects on inflation during both 2001 and 2002.

The second episode concerns inflation forecasts made in 2003:Q1. Cold, dry weather had brought about extreme increases in electricity prices during the winter of 2002/03. This was a temporary shock

to the price level, but it had persistent effects on yearly inflation, which became unusually low when electricity prices declined during the spring and summer. The Riksbank's official forecasts described the decline in inflation extremely well for the first two quarters but underestimated the effects on the longer horizons, possibly because it was difficult to separate the effects of changes in energy prices (the oil price also fluctuated heavily) from the downward pressure on inflation from other forces in the economy (e.g., increases in productivity). In contrast, the DSGE model underestimated the drop in inflation at first, but once the decline had started, it more correctly predicted that inflation would be very low for the next one to two years (cf. the forecasts from 2003:Q3 in figure 3).

These episodes nicely illustrate how formal statistical models and judgments by sector experts can complement each other. Subjective forecasts may sometimes be too myopic and pay too little attention to systematic inflation dynamics related to the business cycle or other historically important regularities. Model forecasts, on the other hand, cannot take sufficient account of specific unusual but observable events. At the same time, judgments from sector experts based on their detailed knowledge about the economy can be extremely useful—in particular, when unusual shocks have hit the economy.

2.5 Combined Forecasts

The previous subsections have contrasted model-based and judgment-based macroeconomic forecasts, and have made it clear that judgments from sector experts can be useful in the short run, especially when unusual disturbances to the economy occur. Being equipped with several different forecasts, our natural question is, what can be gained from combining them into a single overall forecast? Combining a set of purely model-based forecasts is rather straightforward, especially within the Bayesian framework, where the weights are given by posterior model probabilities (Draper 1995). When at least one of the forecasting models cannot be represented by a probability model for the observed data, we need to resort to other solutions.

Winkler (1981) proposes an alternative Bayesian approach that may be used to combine model-based and judgment-based forecasts.

We will form weights separately for each variable and forecast horizon. Winkler's procedure is described in detail in appendix 2. The method assumes that the forecast errors from the different forecasts at a specific time period follow a multivariate normal distribution with zero mean and covariance matrix Σ . It is further assumed that the forecast errors are independent over time. The optimal weights on the individual forecasts can then be shown to be a simple function of Σ^{-1} . This means that the weights depend not only on the relative precision of the forecasts but also on the correlation between forecast errors. It should be noted that while the weights sum to unity, some weights may be negative. Negative weights arise quite naturally, especially when the forecast errors are highly correlated (Winkler 1981), but for convenience in interpretation, we shall restrict all weights to be non-negative. The results do not change substantially if we allow for negative weights.

The forecast-error covariance matrix Σ needs to be estimated from the realized forecast errors available at the time of the formation of the combined forecast. This means that Σ needs to be estimated from a small number of observations. We use a prior distribution to stabilize the estimate of Σ (see appendix 2 for details). Before turning to the results, it should be noted that we have not conducted this weighting experiment for GDP due to the real-time problems related to this variable.

The assumption of unbiased forecast errors does not seem to hold for the DSGE and implicit-forward-rate forecasts of the interest rate (see figure 7 in appendix 2), which may have consequences for the combined forecasts. However, the combined forecast with weights inversely related to the univariate mean-squared forecast errors (MSEs), which thus include any potential bias, yields similar results in terms of its accuracy (see figure 2).

From figure 2, we also see that the combined CPI-inflation forecast performs very well at all forecast horizons. The excellent performance of the combined forecast at the first-quarter horizon is particularly noteworthy. We want to stress that we only use those forecast errors that were actually available at the time of the forecast. This means that the RMSE evaluation at, e.g., the eight-quarter horizon only uses weights up to 2003:Q4 (the sample ends in 2005:Q4).

Turning to the short-term interest rate, we see once more that combining forecasts is a good idea. The RMSE of the combined forecast is low for all forecast horizons, only slightly beaten by the implicit forward rate at the first two horizons and the BVAR forecast at the longer horizons (see the middle panel of figure 2).

3. Advantages of Structural Analysis


3.1 *Historical Decompositions*

In section 2, we analyzed the usefulness of DSGE and BVAR models for forecasting purposes. But policymakers are also very interested in understanding the factors that have brought the economy to where it is right now. This necessitates a structural model that can disentangle the underlying causes for the recent economic development. In the DSGE model, all shocks are given an economic interpretation, while the BVAR model requires additional identifying restrictions. As an example of the use of models for structural analysis, we let the two models interpret the low inflation rate in Sweden during 2003–05, by decomposing the model projections into the various shocks that have driven the development of inflation and output.

In figure 4, we report the actual outcome of GDP growth, inflation, and the repo rate together with projections from the BVAR model. The first row of figure 4 reports forecasts made in 2003:Q1 under the assumption that no shocks would hit the Swedish economy during 2003–05. It can be seen that parts of the increase in GDP growth and the decreases in inflation and the interest rate were expected, but actual inflation turned out to be a great deal lower than anticipated by the BVAR model. From the second row of figure 4—where we have added the “foreign” shocks, identified by the BVAR model *ex post* (dashed line), to the BVAR model’s no-shock (expected) scenario (dotted line)—we can see that the sudden drop in inflation during 2003 and the lower GDP growth during the first quarters of 2003 were mainly due to foreign shocks hitting the economy.¹⁰ The last row of figure 4 shows the effects of “domestic”

¹⁰The foreign shocks are identified through the assumption that foreign GDP growth, foreign inflation, and the foreign interest rate are strictly

Figure 4. Actual Outcomes (—) and Predictions for 2003:Q2–2005:Q4 from the BVAR Model without Any Shocks (···) and with Only Subsets of the Shocks Active during the Forecasting Period (---)



shocks, which are simply identified residually as the parts of the forecast errors that are not accounted for by foreign shocks. From the last row, we see that with only domestic shocks, the BVAR model overestimates inflation during 2003, although overall macroeconomic growth seems to be well captured. For 2004 and 2005, the picture is somewhat different, and during these years it is clear from figure 4

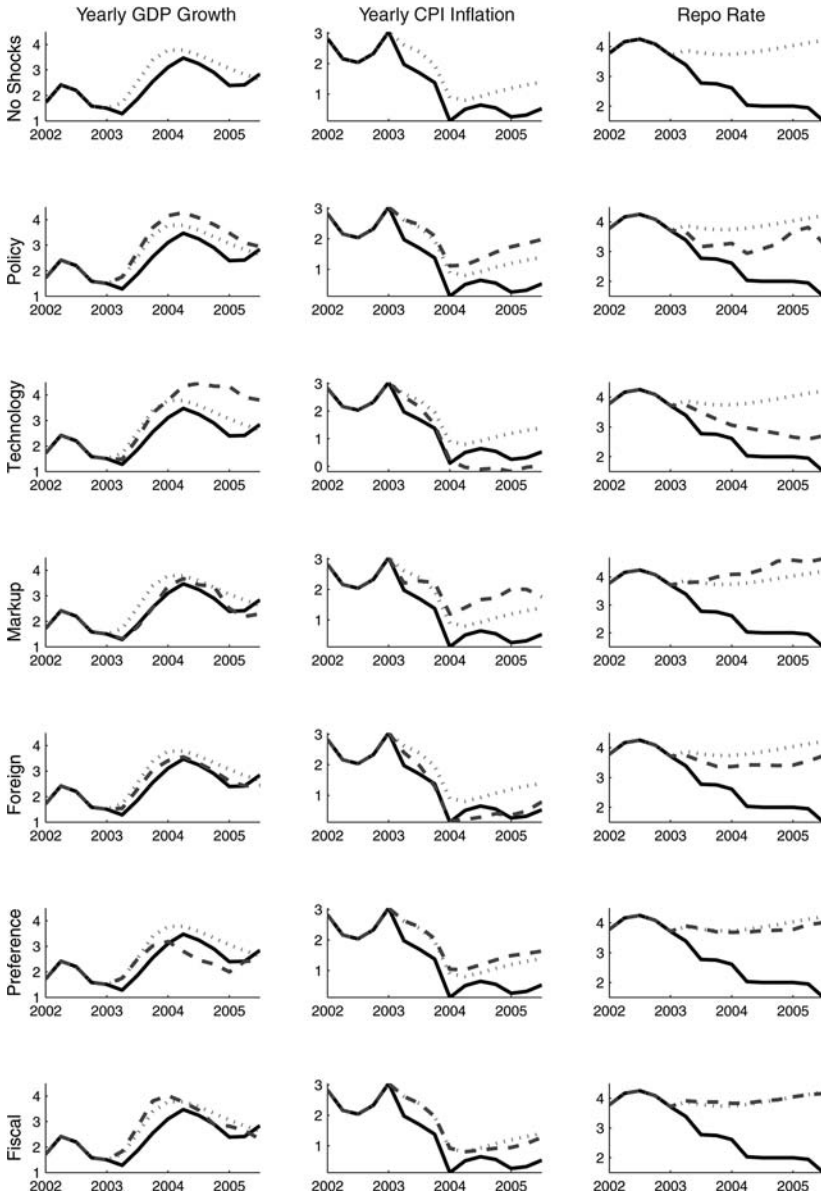
exogenous. Formally, this implies, among other things, that the foreign variables are ordered before all domestic variables in the Choleski decomposition. In addition, the forecast-error decompositions in figure 4 are based on the assumption that the real exchange rate is ordered last in the Choleski decomposition. The results are not affected much, however, if we instead assume that real-exchange-rate shocks are treated as foreign. The results are available upon request. No attempt is made to identify individual foreign shocks.

that the domestic shocks have been a more important source for the forecast errors in general and the low inflation rate in particular. In principle, it is also possible to get even more information about the shocks from the BVAR model, as long as we are willing to make additional identifying assumptions.¹¹

In figure 5, we instead decompose the forecast errors from the DSGE model. The first row shows the outcome of GDP growth, inflation, and the interest rate along with the predictions from the DSGE model under the assumption that no shocks are hitting the economy (i.e., the corresponding information to the first row of figure 4, which was based on the BVAR model). Both models overestimated inflation and the interest rate, but the DSGE model also overrated GDP growth to a somewhat larger extent. The other rows in figure 5 show the “ex post forecasts” from the DSGE model when we add the model’s estimates of different kinds of shocks during 2003–05 (dashed lines) to the original forecasts from 2003:Q1 (dotted lines): monetary policy shocks (second row), technology shocks (third row), markup shocks (fourth row), foreign shocks (fifth row), preference shocks (sixth row), and fiscal policy shocks (last row). It can be seen that when the estimated technology shocks and foreign shocks are individually taken into account, the “ex post forecasts” of inflation from the DSGE model are rather close to the outcome. The DSGE model thus supports the finding from the BVAR model that many of the forecast errors during 2003 were due to foreign shocks. In 2004 and 2005, when the BVAR model suggested that foreign shocks were less important, the DSGE model attributes a large part of the low inflation to both foreign shocks and domestic technology shocks. Interestingly, the model suggests that increased competition (i.e., lower markups) is not an important factor for directly understanding the low-inflation outcome. However, it is, of course, possible that the increased degree of openness (i.e., “globalization”) has stimulated the favorable development in total factor productivity. It is also clear from figure 5 that fiscal policy shocks have played a very

¹¹Other identifying restrictions may involve, e.g., restrictions on long-run impulses as in King et al. (1991). Jacobson et al. (2001) present some results based on such restrictions from a VAR model using similar data as this paper. Alternatively, restrictions may be imposed directly on the impulse-response functions as suggested by Canova and de Nicoló (2002) and Uhlig (2005).

Figure 5. Actual Values (—) and Predictions for 2003:Q2–2005:Q4 from the DSGE Model without Any Shocks (···) and with Only Subsets of the Shocks Active during the Forecasting Period (---)



limited role during this period and that monetary policy has, in fact, been expansionary. According to the DSGE model, the Riksbank reduced the repo rate more than the usual amount during this period to prevent inflation from falling too far below the target.

We find the results from these exercises very promising. Not only can the BVAR and DSGE models make forecasts that have, on average, an equal or better precision than the Riksbank's official, more subjective forecasts, but they can also, *ex post*, decompose the forecast errors in ways that are informative for policymakers and advisers.

3.2 The Effects of Monetary Policy Shocks


In the previous sections, we have shown that the BVAR model is a fine forecasting tool, but to some extent it can also explain what has happened in the economy, as long as we are willing to place some identifying assumptions on the shocks. However, to answer questions about how monetary policy is designed and how it influences the economy, we need to add further structure. Working with an identified model is especially important in a central bank environment where experiments such as predictions conditioned upon alternative interest rate paths are carried out. Therefore, we study impulse-response functions to see how the links between the interest rate setting and inflation outcomes differ between the BVAR and DSGE models.

Figure 6 displays the effects on output growth, CPI inflation, the repo rate, and the real exchange rate to a one-standard-deviation interest rate shock in the BVAR model (first column) and the DSGE model (last column).¹²

Qualitatively, there are some similarities between the impulse responses in the two models. For example, an increase in the interest rate appreciates the real exchange rate. However, there are also

¹²In the BVAR model, this is implemented through exogenous shocks to the interest rate in a Choleski decomposition where the interest rate is ordered after all other variables except the real exchange rate. Other nonrecursive identifying restrictions (e.g., allowing the central bank to react to changes in the real exchange rate within the period, but not to the two GDP variables) gave similar results. In the DSGE model, exogenous shocks are added to the central bank's reaction function (i.e., a monetary policy shock).

Figure 6. Posterior Median Impulse-Response Functions to a One-Standard-Deviation Interest Rate Shock for the Domestic Variables in the BVAR (First Column) and DSGE Model (Second Column) with 68 Percent and 95 Percent Probability Bands



discrepancies between the BVAR and the DSGE models. Although an increase in the interest rate causes a decline in output growth and inflation, consistent with typical prejudices, the responses in the BVAR are not significant in contrast to those in the DSGE model. Moreover, the DSGE model fulfills long-run nominal neutrality (i.e., monetary policy can only affect prices and not real quantities in the long run), whereas there is no such restriction in the BVAR. Further, the quantitative differences are very large. The DSGE model supports conventional wisdom: if the interest rate is unexpectedly increased by 0.35 percentage point (and then gradually reduced), the maximum effect on inflation is around 0.15 percentage point and is

recorded after about $1-1\frac{1}{2}$ years. The effects in the BVAR model are much smaller and typically insignificant.

Although there are reasons to expect that monetary policy shocks are more credibly identified in the DSGE model, the differences in impulse-response functions and forecasting properties between the BVAR and the DSGE models may create difficulties when using the models in the policy process, even if each individual model compares well with the subjective forecasts and rests on solid methodological grounds. Given the small impact of interest rate changes in the BVAR model, inflation projections conditional upon alternative interest rate paths would not differ to any considerable extent from the BVAR's unconditional forecasts. In contrast, the results in figure 6 suggest that the DSGE's conditional forecasts will change a great deal compared with the unconditional forecasts if the former are generated by injecting monetary policy shocks (which have much larger effects in the DSGE model). The impression of alternative policy assumptions will thus be very different in the two models.

4. Concluding Remarks

The theme of this paper is that modern macroeconomic tools like BVAR and DSGE models deserve to be used more in real-time forecasting and for policy advice at central banks. We have shown that it is possible to construct and use BVAR and DSGE models that make about as good inflation forecasts as the much more complicated judgmental procedure typically employed by central banks. In our view, central banks should use formal models—VARs and DSGEs—as benchmarks for forecasts and policy advice, and to summarize the implications of the continuous flow of new information about the state of the economy to which central bank economists are exposed.

We want to emphasize that we do not view our results as arguments against the use of judgments in monetary policy analysis. The key point here is not to dispute judgments versus formal models but, rather, to determine how to coherently combine the BVAR and DSGE models with beliefs about the current conditions. Our results suggest that it would be beneficial to incorporate judgments into the formal models, so that the forecasts reflect both

judgments and historical regularities in the data.¹³ We have analyzed a Bayesian weighting scheme based on the forecast errors of the different methods to combine the judgmental and model forecasts. The results are promising: the root mean-squared errors of the combined forecasts of the CPI and the interest rate are consistently low on all evaluated forecast horizons. As an alternative, short-run judgments by sector experts could be directly incorporated into the models by exploiting the methodology suggested in Waggoner and Zha (1999).

There are a number of gains in using formal models in the policy analysis. They make it possible to decompose forecast errors and provide a tool for characterizing the uncertainty involved in statements about the future development in the economy. Formal models make it possible to quantify the imprecision and uncertainties involved in the forecasting process. Formal models also serve as a learning mechanism, where lessons about the complex interdependencies in the economy can be accumulated. Our results suggest, for example, that subjective forecasts may be too myopic and not take sufficient account of important historical regularities in the data. Our present version of the DSGE model, on the other hand, may reflect problems with interest rate determination in financial markets, i.e., the empirical failure of the expectation hypothesis and the UIP condition; see the discussion in, e.g., Faust (2005).

Naturally, there are also limitations to the use of the current generation of modern macroeconomic models for policy purposes. Policymakers are often interested in details about the state of the current economy. Formal models cannot possibly cover all details within a tractable consistent framework. Neither can sector experts, but their insights into details often lead policymakers to rely on advice and forecasts from experts rather than from models. Another problem is that there are gaps between different models. Different models give quite different forecasts and imply different policy recommendations. Researchers are not typically bothered by this, as long as the models are considered to be good. Policymakers are, of course, bothered.

¹³Svensson (2005) offers a theoretical analysis of the links between judgments and monetary policy. One way of including judgments in the formal models is to approach this in a Bayesian manner. However, it is less clear how to translate the provided form of judgment into a usable prior distribution.

For some policy purposes, we therefore think that it makes sense to weight various models according to their empirical performance. We have discussed a procedure for combining forecasts that may also be used when some of the forecasts do not come from a well-specified formal model, which is typically the case in policy work. Another way of bridging the gap between formal models is to use Bayesian prior distributions that incorporate identical prior information on features that are common to the models, such as the steady state of the system (Villani 2005) or impulse-response functions (Del Negro and Schorfheide 2004).

Appendix 1. The BVAR Model

The BVAR model contains quarterly data on the following seven variables: trade-weighted measures of foreign GDP growth (y_f), CPI inflation (π_f) and the three-month interest rate (i_f), the corresponding domestic variables (y , π , and i , where i is the repo rate), and the level of the real exchange rate defined as $q = 100(s + p_f - p)$, where p_f and p are the foreign and domestic CPI levels (in logs) and s is the (log of the) trade-weighted nominal exchange rate.

The BVAR model used in this paper is of the form

$$\Pi(L)(x_t - \Psi d_t) = A\varepsilon_t, \quad (1)$$

where $x = (y_f, \pi_f, i_f, y, \pi, r, q)'$ is an n -dimensional vector of time series, $\Pi(L) = I_n - \Pi_1 L - \dots - \Pi_k L^k$, and L is the usual back-shift operator with the property $Lx_t = x_{t-1}$. The structural disturbances $\varepsilon_t \sim N_n(0, I_n)$, $t = 1, \dots, T$, are assumed to be independent across time. We impose restrictions on $\Pi(L)$ such that the foreign economy is exogenous. A is the lower-triangular (Choleski) contemporaneous-impact matrix, such that the covariance matrix Σ of the reduced-form disturbances decomposes as $\Sigma = AA'$. We also have experimented with nonrecursive identifying restrictions, in which case the equations are normalized with the Waggoner-Zha rule (Waggoner and Zha 2003b), and the Gibbs sampling algorithm in Waggoner and Zha (2003a) is used to sample from the posterior distribution. The deterministic component is $d_t = (1, d_{MP,t})'$, where

$$d_{MP,t} = \begin{cases} 1 & \text{if } t < 1993:\text{Q1} \\ 0 & \text{if } t \geq 1993:\text{Q1} \end{cases}$$

is a shift dummy to model the abandonment of the fixed exchange rate and the introduction of an explicit inflation target in 1993:Q1. Since the data are modeled on a quarterly frequency, we use $k = 4$ lags in the analysis. Larger lag lengths gave essentially the same results, with a slight increase in parameter uncertainty.

The somewhat nonstandard parameterization of the VAR model in (1) is nonlinear in its parameters but has the advantage that the unconditional mean, or steady state, of the process is directly specified by Ψ as $E_0(x_t) = \Psi d_t$. This allows us to incorporate prior beliefs directly on the steady state of the system, e.g., the information that steady-state inflation is likely to be close to the Riksbank's inflation target. To formulate a prior on Ψ , note that the specification of d_t implies the following parameterization of the steady state:

$$E_0(x_t) = \begin{cases} \psi_1 + \psi_2 & \text{if } t < 1993:Q1 \\ \psi_1 & \text{if } t \geq 1993:Q1, \end{cases}$$

where ψ_i is the i -th column of Ψ . The elements in Ψ are assumed to be independent and normally distributed a priori. The 95 percent prior probability intervals are given in table 1.

The prior proposed by Litterman (1986) will be used on the dynamic coefficients in Π , with the default values on the hyperparameters in the priors suggested by Doan (1992): overall tightness is set to 0.2, cross-equation tightness to 0.5, and a harmonic lag decay with a hyperparameter equal to 1. See Litterman (1986) and Doan (1992) for details. Litterman's prior was designed for data in levels and has the effect of shrinking the process toward the univariate random-walk model. Therefore, we set the prior mean on the first own lag to 0 for all variables in growth rates. The two interest rates and the real exchange rate are assigned a prior that centers on the $AR(1)$ process with a dynamic coefficient equal to 0.9. The usual

Table 1. Ninety-Five Percent Prior Probability Intervals of Ψ

	y_f	π_f	r_f	y	π	r	q
ψ_1	(2, 3)	(1.5, 2.5)	(4.5, 5.5)	(2, 2.5)	(1.7, 2.3)	(4, 4.5)	(-1, 1)
ψ_2	(-1, 1)	(1.5, 2.5)	(1.5, 2.5)	(-1, 1)	(4.3, 5.7)	(3, 5.5)	(-9, 9)

random-walk prior is not used here, as it is inconsistent with having a prior on the steady state. Finally, the usual non-informative prior $|\Sigma|^{-(n+1)/2}$ is used for Σ .

The posterior distribution of the model's parameters and the forecast distribution of the seven endogenous variables were computed numerically by sampling from the posterior distribution with the Gibbs sampling algorithm in Villani (2005).

Appendix 2. Combining Judgmental and Model Forecasts

Suppose that we have available forecasts, at a given forecast horizon, from k different forecasting methods over T different time periods. Let \hat{x}_{jt} denote the j -th method's forecast of a variable x_t , and $e_{jt} = \hat{x}_{jt} - x_t$ the corresponding forecast error, where $j = 1, \dots, k$ and $t = 1, \dots, T$. The question here is how to merge these k forecasts into a single combined forecast. Following Winkler (1981), we shall assume that the vector of forecast errors from the k methods, $e_t = (e_{1t}, \dots, e_{kt})'$, can be modeled as independent draws from a multivariate normal distribution with zero mean and covariance matrix Σ . This implies that $\hat{x}_t \sim N_k(x_t u, \Sigma)$, where $\hat{x}_t = (\hat{x}_{1t}, \dots, \hat{x}_{kt})'$ is the vector of forecasts of x_t from the k forecasting methods, and $u = (1, \dots, 1)'$. We use an uninformative (uniform) prior on x_t and an inverted Wishart density for Σ a priori: $\Sigma \sim IW(\Sigma_0, v)$, where $\Sigma_0 = E(\Sigma)$ and $v \geq k$ is the degrees-of-freedom parameter. The prior on Σ is important, as historical forecast errors are limited and an estimate of Σ is typically unreliable. This is particularly important when the correlations in Σ are large, which is often the case with forecast errors from competing methods. As v increases, the prior becomes increasingly concentrated around Σ_0 . The specification of Σ_0 and v is discussed below.

The posterior mean of the true value x_t , which is the natural combined forecast for a Bayesian, can now be shown to be a linear combination of the individual forecasts (Winkler 1981):

$$E(x_t | \hat{x}_t) = \sum_{j=1}^k w_{jt} \hat{x}_{jt}, \quad (2)$$

where the weights of the forecasting methods are given by

$$w'_t = (w_{1t}, \dots, w_{kt}) = \frac{u' \tilde{\Sigma}_t^{-1}}{u' \tilde{\Sigma}_t^{-1} u}, \quad (3)$$

and the posterior estimate of Σ is

$$\tilde{\Sigma}_t = E(\Sigma | e_1, \dots, e_t) = \frac{v}{t+v} \Sigma_0 + \frac{t}{t+v} \hat{\Sigma}_t,$$

where $\hat{\Sigma}_t$ is the usual unbiased estimator of a covariance matrix.

The weights w_t sum to unity at all dates t , but they need not be positive. Negative weights may result quite naturally, as explained in Winkler (1981), especially when the forecasts are positively correlated across methods.

Strictly speaking, the weighting scheme in (3) is only known to be the Bayesian solution under the assumption that forecast errors are independent and unbiased. The independence assumption is likely to be violated for forecast errors beyond the first horizon. For simplicity, we will continue to assume independent forecast errors at all forecast horizons. An alternative approach would be to stick to the weighting scheme in (3) but with a more sophisticated $\hat{\Sigma}_t$ estimate that accounts for autocorrelation, e.g., the Newey-West estimator. This procedure is unlikely to be a Bayesian solution, however, and also suffers from the drawback that the Newey-West estimator is likely to be unstable when the history of available forecast errors is short. The second assumption behind (3) is that forecasts are unbiased. This does not seem to be supported for the implicit forward rate or the DSGE's interest rate forecast at longer horizons, and can potentially have a large effect on the combined forecast. Therefore, we also look at an ad hoc method for combining forecasts with weights inversely proportional to the mean-squared errors from past forecasts. Note that this method ignores the fact that forecast errors of different methods are typically correlated.

We need to determine Σ_0 and v in the inverted Wishart prior for Σ . We will use the following parameterization of the prior mean of Σ :

$$\Sigma_0 = \sigma_0^2 \begin{pmatrix} 1 & \rho & \rho \\ \rho & 1 & \rho \\ \rho & \rho & 1 \end{pmatrix},$$

which leaves σ_0 , ρ , and v to be specified. It seems fair to expect all forecasting methods to produce fairly correlated forecasts, so that ρ is comparatively large and also increases with the forecast horizon (most methods will produce long-run forecasts that are quite close to their steady-state value, and the steady states in the different methods should not be too different). Moreover, σ_0 should also increase with the forecast horizon. We will assume that both ρ and σ_0 increase linearly with the forecast horizon (σ_0 ranges from 0.5 to 1, and ρ equals 0.5 at the first forecast horizon and 0.92 at the eighth horizon). Finally, we need to pin down the overall precision in the prior—the degrees-of-freedom parameter, v . We set $v = 50$, which gives us a 95 percent prior probability interval for ρ at the first horizon equal to (0.4, 0.75). The results are robust to nondrastic variations in the prior.

Turning to the results, we show the bias for the different models' CPI inflation and interest rate forecasts in figure 7. The interest rate forecasts from the DSGE model and the implicit forward rate especially seem to be biased. Therefore, we also consider a weighting scheme that is inversely related to the univariate mean-squared

Figure 7. Forecast Bias






Figure 8. Sequential Weighting Schemes for Yearly CPI Inflation



forecast errors, which thus includes any potential bias. In figures 8 and 9, we compare the two weighting schemes. However, as seen from figure 2 in the main text, the accuracy of the combined forecast does not seem to be affected much by which scheme is used.

Figure 8 shows the sequential weighting schemes for CPI inflation at different horizons. Note that the weights are equal on all three forecasts at the beginning of the evaluation period, where there are not enough realized forecast errors to estimate Σ . The weights for CPI inflation in the two different weighting schemes are similar on the first- and second-quarter horizons, at least in the latter part of the evaluation period. At longer horizons, there are fewer forecast errors for constructing the weights and larger biases in the forecasts. This causes the two weighting schemes to differ much more than at shorter horizons.

Figure 9. Sequential Weighting Schemes for the Interest Rate



For the interest rate, the sequential forecast weights are once more stable at the two shortest horizons compared with the longer ones; see figure 9. From figure 9, it can also be seen that the superiority of the implicit forward rate at the first- and second-quarter horizons is immediately picked up by both weighting schemes.

References

Adolfson, M., S. Laséen, J. Lindé, and M. Villani. 2005. “Are Constant Interest Rate Forecasts Modest Policy Interventions? Evidence from a Dynamic Open-Economy Model.” *International Finance* 8 (3): 509–44.

- . Forthcoming. “Evaluating an Estimated New Keynesian Small Open Economy Model.” *Journal of Economic Dynamics and Control*.
- Altig, D. E., C. T. Carlstrom, and K. J. Lansing. 1995. “Computable General Equilibrium Models and Monetary Policy Advice.” *Journal of Money, Credit and Banking* 27 (4): 1472–93.
- Altig, D., L. Christiano, M. Eichenbaum, and J. Lindé. 2003. “The Role of Monetary Policy in the Propagation of Technology Shocks.” Manuscript, Northwestern University.
- Canova, F., and G. de Nicoló. 2002. “Monetary Disturbances Matter for Business Fluctuations in the G-7.” *Journal of Monetary Economics* 49 (6): 1131–59.
- Christiano, L. J., M. Eichenbaum, and C. L. Evans. 2005. “Nominal Rigidities and the Dynamic Effects of a Shock to Monetary Policy.” *Journal of Political Economy* 113 (1): 1–45.
- Del Negro, M., and F. Schorfheide. 2004. “Priors from General Equilibrium Models for VARs.” *International Economic Review* 45 (2): 643–73.
- Doan, T. A. 1992. *RATS User’s Manual*. Version 4. Evanston, IL: Estima.
- Draper, D. 1995. “Assessment and Propagation of Model Uncertainty.” *Journal of the Royal Statistical Society* 57: 45–97.
- Edge, R., M. Kiley, and J.-P. Laforte. 2006. “A Comparison of Forecast Performance Between Federal Reserve Staff Forecasts, Simple Reduced-Form Models, and a DSGE Model.” Manuscript, Board of Governors of the Federal Reserve System.
- Faust, J. 2005. “Is Applied Monetary Policy Analysis Hard?” Manuscript, Board of Governors of the Federal Reserve System.
- Jacobson T., P. Jansson, A. Vredin, and A. Warne. 2001. “Monetary Policy Analysis and Inflation Targeting in a Small Open Economy: A VAR Approach.” *Journal of Applied Econometrics* 16 (4): 487–520.
- King, R. G., C. I. Plosser, J. H. Stock, and M. W. Watson. 1991. “Stochastic Trends and Economic Fluctuations.” *American Economic Review* 81 (4): 819–40.
- Litterman, R. B. 1986. “Forecasting with Bayesian Vector Autoregressions—Five Years of Experience.” *Journal of Business and Economic Statistics* 4 (1): 25–38.

- Sims, C. A. 2002. "The Role of Models and Probabilities in the Monetary Policy Process." *Brookings Papers on Economic Activity* 2: 1–62.
- Smets, F., and R. Wouters. 2002. "Openness, Imperfect Exchange Rate Pass-Through and Monetary Policy." *Journal of Monetary Economics* 49 (5): 947–81.
- . 2004. "Forecasting with a Bayesian DSGE Model: An Application to the Euro Area." *Journal of Common Market Studies* 42 (4): 841–67.
- Svensson, L. E. O. 1995. "Estimating Forward Interest Rates with the Extended Nelson and Siegel Method." *Sveriges Riksbank Quarterly Review* 3: 13–26.
- . 2005. "Monetary Policy with Judgment: Forecast Targeting." *International Journal of Central Banking* 1 (1): 1–54.
- Uhlig, H. 2005. "What Are the Effects of Monetary Policy on Output? Results from an Agnostic Identification Procedure." *Journal of Monetary Economics* 52 (2): 381–419.
- Villani, M. 2005. "Inference in Vector Autoregressive Models with an Informative Prior on the Steady State." Sveriges Riksbank Working Paper No. 181.
- Waggoner, D. F., and Zha, T. 1999. "Conditional Forecasts in Dynamic Multivariate Models." *Review of Economics and Statistics* 81 (4): 639–51.
- . 2003a. "A Gibbs Sampler for Structural Vector Autoregressions." *Journal of Economic Dynamics and Control* 28 (2): 349–66.
- . 2003b. "Likelihood-Preserving Normalization in Multiple Equation Models." *Journal of Econometrics* 114 (2): 329–47.
- Winkler, R. L. 1981. "Combining Probability Distributions from Dependent Sources." *Management Science* 27 (4): 479–88.