No. 8 | 2010

Staff Memo

Monetary policy analysis in practice - a conditional forecasting approach

Ida Wolden Bache, Leif Brubakk, Anne Sofie Jore, Junior Maih and Jon Nicolaisen
Norges Bank Monetary Policy

NB NORGES BANK

ISSN 1504-2596 (online only)

ISBN 978-82-7553-549-6 (online only)

Monetary policy analysis in practice - a conditional forecasting approach

Ida Wolden Bache, Leif Brubakk, Anne Sofie Jore, Junior Maih and Jon Nicolaisen*
Monetary Policy Wing, Norges Bank (Central Bank of Norway)

2 July 2010

Abstract

In this paper we provide a broad outline of the forecasting and policy analysis system adopted at Norges Bank.

1 Introduction

Monetary policy works mainly through affecting private agents' expectations. As a consequence, the effectiveness of monetary policy depends on the way the central bank communicates its future policy intentions. Norges Bank has gone further than most central banks in this respect. Since 2005, a central element in our communication strategy has been to publish our projections of the key policy rate along with the forecasts of other key variables such as inflation and the output gap. The uncertainty surrounding the mean projections is illustrated by fan charts. In addition to being transparent about our future policy intentions, we also aim to be precise about how these intentions are formed. The overall communication strategy and the main arguments for publishing the interest rate path are laid out in Holmsen et al. (2008). In this paper, we focus instead on the choice of analytical framework and the forecasting process at Norges Bank.

The forecasting system is organized around our core macroeconomic model, NEMO (Norwegian Economy Model). NEMO is a medium-scale, small open economy DSGE model similar in size and structure to the DSGE models developed recently by many other central banks. A distinguishing feature of our approach is that the interest rate projection is based on optimal policy in the sense of minimizing an intertemporal loss function that is consistent with the monetary policy mandate.[1] The medium- to long-term projections are largely model-based, but since all economic models are incomplete and simplified descriptions of

*This note was prepared as a background document for the Chief Economist Workshop on state-of-the-art modelling for central banks organized by the Bank of England 18-20 May 2010. We are grateful for comments from colleagues at Norges Bank. The views expressed in this paper are our own and do not necessarily represent the views of Norges Bank. Corresponding author: Leif Brubakk, leif.brubakk@norges-bank.no.

[1] To our knowledge, Norges Bank is the only central bank that has stated publicly that it uses optimal monetary policy as the normative benchmark for assessing the appropriateness of the interest rate path (see e.g., Holmsen et al. (2007) and Monetary Policy Report 2/10).

reality, some degree of judgement will always be needed. However, organizing the policy process around a single core model adds discipline to the process and helps ensure that the analyses are consistent over time. Furthermore, it would be very difficult to communicate a single policy rate projection constructed on the basis of projections from different models.

Given a balanced set of historical data, NEMO can be used to provide unconditional forecasts for any desired forecasting horizon. In the projection exercise at Norges Bank, however, we have adopted a conditional forecast approach. As shown by Maih (2010) it may be possible to improve the forecast performance of DSGE models by conditioning on e.g., financial market information or short-term forecasts from models that are able to exploit recent data and information from large datasets. Conditioning information may also come in the form of policymaker judgement that is not directly interpretable in terms of the DSGE model. Conditioning information is likely to be particularly important in the event of large disturbances that are not evident in the most recent national accounts data, as was the case in the autumn of 2008. The conditional forecasting approach allows us to exploit this information in a consistent manner without changing the structure of the model. An alternative to publishing model consistent conditional forecasts is to start with the pure unconditional model forecasts and then, ex post, adjust the projections in the direction suggested by off-model considerations and judgement. In our experience, however, both the internal consistency of the forecasts and the policy discussion is improved by producing conditional forecasts based on a single model.

The conditioning information used in NEMO consists of nowcasts and short-term forecasts provided by sector experts.[2] Detailed knowledge of developments in the different sectors of the economy is particularly useful for forecasting short-term developments. Each sector expert generally has one or more small econometric models that incorporate a broader information set than NEMO. The sector experts monitor a large amount of data from disparate sources, including information of a more qualitative nature. One example is the data from Norges Bank's regional network, which is a business survey that provides information about production, investment, prices and wages. The sector experts also have an understanding of how disaggregated bits of data feed into the preparation of the aggregate numbers that are published with a lag by the statistical agencies. This provides a starting point that is likely to be more accurate than can be obtained from a model based on e.g. quarterly data.

An additional tool for short-term forecasting is the System for Model Averaging (SAM). SAM is used to produce density forecasts for the current and the next few quarters by averaging forecasts from a large set of different models. Currently, the system provides forecasts for inflation and output growth, however, the goal is to extend the system to produce forecasts for the full set of observable variables in NEMO.

When implementing a system of conditional forecasting, several choices have to be made. First, the type of conditioning method employed in a DSGE model depends on whether the conditioning information is anticipated or not. As forward-looking agents

[2] For some variables (e.g., government spending, oil investment and foreign variables) we condition on off-model information for the whole forecasting horizon.

exploit available information that can improve their forecasts, anticipated events matter for their current decisions. Hence, when conditioning on leading information in DSGE models, an important question is to what extent private agents can be assumed to internalize this information. Our baseline forecasts are based on the assumption that the conditioning information is known to all agents in the model at the beginning of the forecast period.

A second issue is whether to treat the conditioning information as certain (referred to in the literature as 'hard' conditioning) or uncertain ('soft' conditioning). Most of the literature on conditional forecasting has focused on hard conditioning. So far, this has also been the approach taken at Norges Bank. However, Norges Bank, like other inflation targeting central banks, publish fan charts for key macroeconomic variables. Ideally, these probability bands should also reflect the probabilistic nature of the conditioning information. In the last part of this paper, we discuss the procedures for density-conditional forecasting for DSGE models developed in Maih (2010) to illustrate how uncertainty about the conditioning information can be incorporated in a formal manner to produce model consistent density forecasts (or fan charts).

2 The forecasting and policy analysis system

Norges Bank publishes forecasts for the key policy rate and other key macro variables three times per year in the Monetary Policy Report (MPR). Decisions concerning the interest rate are normally taken at the Executive Board's monetary policy meeting every sixth week. In conjunction with the publication of the MPR, the Executive Board decides on a strategy interval for the key policy rate that applies for the period up to the next MPR. The analyses and the monetary policy strategy prepared by the staff form the basis for the Executive Board's discussions of the monetary policy strategy.

The overall structure of the forecasting and policy analysis system is illustrated in figure 1. The medium-term projections and hence the policy advice are based on two premises in particular. The first is an assessment of the current economic situation and short-term forecasts up to four quarters ahead. One important input to the short-term forecasts is the System of Averaging Models (SAM) which provides point and density forecasts for GDP growth and inflation. Short-term forecasts for other key variables are based on current statistics, information from Norges Bank's regional network and simple econometric models. The final short-term forecasts are the result of an overall assessment based on both models and judgement. The second key premise is forecasts for exogenous variables – those that have to be determined outside our model. Examples include foreign variables, commodity prices and government spending. On the basis of these premises, we use our core macroeconomic model NEMO to produce a set of projections for key macroeconomic variables, including the key policy rate.

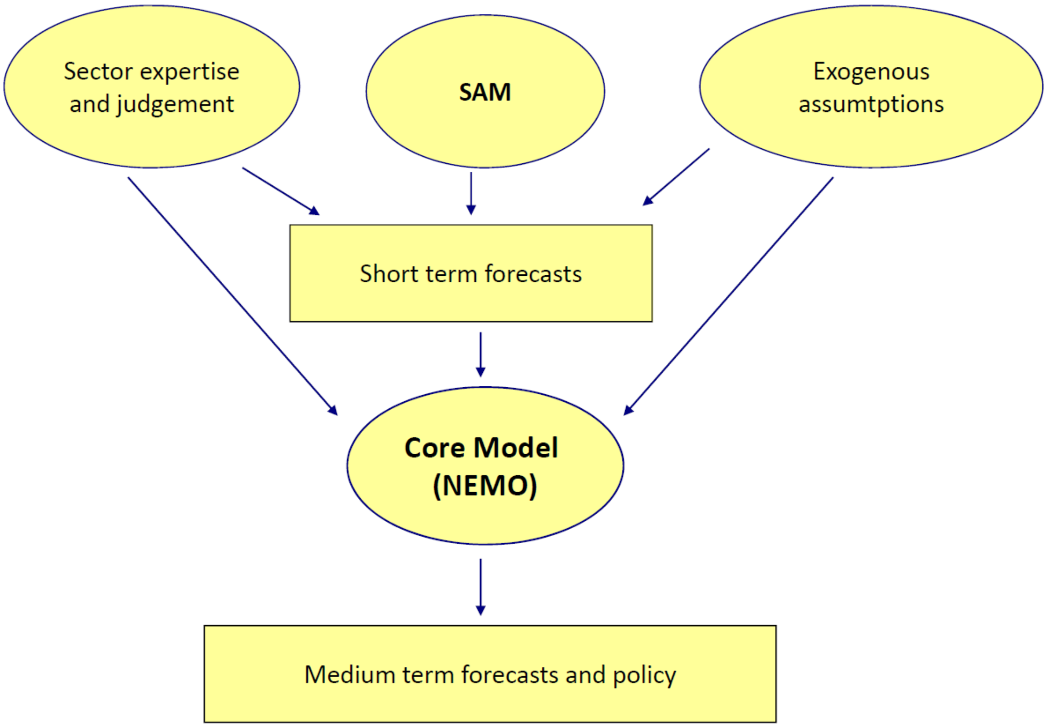

Figure 1: The forecasting and policy analysis system

2.1 System of Averaging Models

SAM is based on the idea of forecast combination.[3] Model or forecast combination has a long history. Timmermann (2006) highlights three main reasons why forecast combinations may produce better forecasts on average than methods based on the ex-ante best individual forecasting model. First, forecast combination can be motivated by a simple portfolio diversification (hedging) argument. A second rationale for combining forecasts is that there may be unknown instabilities (structural breaks) that sometimes favour one model over another. Some models may adapt to breaks quickly while others may have parameters that will only adjust slowly to structural breaks. By combining forecasts from different models, the decision maker may obtain forecasts that are more robust to these instabilities than if they had chosen a single model.[4] A third motivation is that forecast combination may be desirable if the models are misspecified in unknown ways. In this case combining forecasts may average out the biases, improving forecast accuracy. Hence, even if the combined forecast may not always be superior, model combination is preferable as it will ensure against selecting a bad model.

The SAM forecasts are density forecasts, that is, they give a statement about the probability distribution of the forecasts. One main objective when developing SAM, was

[3] See Bjørnland et al. (2008) for a description of the project to improve short-term forecasting at Norges Bank. The project benefitted greatly from discussions and cooperation with forecasters and researchers from other central banks. A number of central banks have developed forecasting systems based on the idea of forecast combination, see e.g., Kapetanios et al. (2008), Andersson & Löf (2007) and Andersson et al. (2007), Bloor (2009) and Coletti & Murchison (2002).

[4] See Jore et al. (2010) for an example using US data and Bache et al. (2009a) for an example using Norwegian data.

to better characterise the uncertainty surrounding Norges Bank's short term forecasts. With a characterisation of the probability distribution at hand, one can ask questions like: "What is the probability that inflation will exceed the inflation target the next four quarters?"

The density forecasts in SAM are purely model-based. For the smaller and least complex models the standard errors are estimated, while the forecast distribution for larger, more complex models are derived using simulation methods.[5]

2.1.1 The models in SAM

The current version of SAM produces forecasts up to one year ahead for GDP Mainland-Norway and CPI-ATE (consumer prices adjusted for taxes and without energy prices). The models in SAM vary both in terms of structure and in the information set they use. Currently, a total of 237 models are used to forecast GDP and 167 models are used to forecast CPI-ATE. Some models are used to forecast both variables.

The models are of a variety of different types, including autoregressive integrated moving average (ARIMA) models, vector autoregressive (VAR) models, Bayesian VAR (BVAR) models, factor models, a DSGE model and a macroeconometric (VEqCM) model.[6] For each type of model there are several variants with different specifications. E.g., SAM includes 36 specifications of a bivariate VAR with GDP and inflation, each specification with different detrending assumptions, lag lengths and/or estimation periods.

The forecasts are based on a large information set, including quarterly national accounts data, monthly data on manufacturing production, employment, retail sales, accommodation statistics and building start indicators, disaggregated CPI series, data on the term structure of interest rates, asset prices, monetary aggregates and household and business tendency surveys, including information from Norges Bank's own regional network.

2.1.2 The forecast combination scheme

The forecasts in SAM are combined in two steps. In the first step we group models that loosely share the same information set or model structure into distinct groups or "ensembles".[7] In the second step, we combine predictive densities from the ensembles in a "grand ensemble". The idea is that the lower the degree of information overlap, the more useful a combination of forecasts is likely to be.[8]

The 16 ensembles for forecasting GDP growth and the 10 ensembles for forecasting inflation are listed in tables 1 and 2, respectively. The second column in both tables provides a short description of the ensembles. The third column states the number of

[5] See Hall & Mitchell (2009) for an extensive exposition.

[6] For details about the models in SAM see Gerdrup et al. (2009).

[7] The concept of ensemble modelling comes from the weather forecasting literature. It involves predictive density construction from a large number of models and combination based on out-of-sample performance and time varying weights. See Bache et al. (2009b) and Garratt et al. (2009) for elaborations. In SAM, we combine predictive densities from "similar" models into ensembles.

[8] The method proposed by Winkler (1981) to take account of the covariance between the forecast errors is infeasible in our model suite because of near-singularity in the variance-covariance matrix.

models within each ensemble. For example, in the ensemble eRegN in table 1 we only have one model. At the other end of the spectrum, the ensemble eVAR3 contains 72 models. If we were using equal weights for all models in one step, the average forecast would be heavily influenced by groups of very similar models, e.g. large groups of only slightly different AR or VAR models. The forecasts from these models would then dominate forecasts from other types of models or information sets. This would reduce the benefit of forecast combination.

Table 1: Models for forecasting GDP Mainland-Norway

Ensemble	Description	No of models
eRegN	Regional network model	1
eTstruc	Term structure models	4
eMI	Monthly indicator models	2
eFM	Factor models	2
eEmod	Macro model (VECM)	1
eDSGE	Macro model (DSGE)	1
eBVAR	Bayesian VARs	10
eUniv	Univariate autoregressive models (ARs)	38
eVAR2	VARs with GDP and inflation	36
eVAR3	VARs with GDP, inflation and/or interest rate	72
eTNSG	Bivariate VARs with household surveys	6
eBuild	Bivariate VARs with building and construction	10
eOrd	Bivariate VARs with orders to manufacturing	4
eEmpl	Bivariate VARs with employment data	10
eMny	Bivariate VARs with money and credit	7
eBTS	Bivariate VARs with Business Tendency Survey	33
Sum		237

Table 2: Models for forecasting CPIATE

Ensemble	Description	No of models
eDisAgg	ARs for CPI-disaggregates	1
eMth	Monthly VARs	3
eFM	Factor models	3
eEmod	Macro model (VECM)	1
eDSGE	Macro model (DSGE)	1
eBVAR	Bayesian VARs	10
eUniv	Univariate autoregressive models (ARs)	39
eVAR2	VARs with GDP and inflation	36
eVAR3	VARs with inflation, GDP and/or interest rate	72
eMny	VAR with GDP and money	1
Sum		167

At both steps our preferred aggregation method is the linear opinion pool, in which the combined density is a linear combination of the densities from the individual models (in step 1) or ensembles (in step 2):

$$P\left(y_t\right) = \sum_{i=1}^{n} \omega_i P_i(y_t), \qquad \omega_i \in [0,1] \text{ and } \sum_{i=1}^{n} \omega_i = 1, \tag{1}$$

where $P(y_t)$ is the combined density, $P_i(y_t)$ is the density from model (or ensemble) i, and ω_i the corresponding weight. If the individual densities are normal, then the combined density will be mixed normal. The distribution may be multimodal, if one or more distributions are "very far" from the majority.[9]

Within each ensemble we use weights based on the logarithmic score (or log score) of the respective predictive densities. The log score is the probability that the realization of the variable would occur, given the density function specified by the model. The logarithmic scoring rule gives a high score to a density forecast that provides a high probability to the realized value of the forecasted variable.[10] Using log score weights ensures that the best model, i.e., the model with the ex ante highest probability of having generated the data, gets the highest weight. The log score weights can be calculated as:

$$\omega_i = \frac{\exp\left(\log(P_i(\underline{y}))\right)}{\sum_{j=1}^{n} \exp\left(\log(P_j(\underline{y}))\right)}, \tag{2}$$

where $\underline{y} = (y_1, ..., y_T)'$ and $\log(P_i(\underline{y})) = \sum_{t=1}^{T} \log(P_i(y_t))$.[11]

To combine the ensemble forecasts we use inverse Mean Squared Error (MSE) weights:

$$\omega_i = \frac{\frac{1}{MSE_i}}{\sum_{j=1}^{n} \frac{1}{MSE_j}}, \tag{3}$$

By constructing the ensembles in a sensible way, the ensemble forecasts should be approximately independent of each other. In that case, MSE weights will minimise a quadratic loss function based on point forecast errors. We denote the final forecast, combined in two steps, as the grand ensemble forecast. A more detailed exposition is provided in Gerdrup et al. (2009).

Figure 2 illustrates the performance of alternative weighting schemes for forecasting inflation. The models are first estimated up to 1998Q4, and forecasts for 5 quarters ahead are calculated. Then we expand the estimation window recursively in quasi real-time. The models are combined using univariate, horizon-specific weights. The four panels show the development of the Root Mean Squared Forecast Error (RMSFE) for horizons 1 to 4, respectively. The alternatives are "Unweighted average" (equal weights for all individual models), "Selection" (the best model is picked at each point in time) and "Grand ensemble" (combining in two steps as described above).

[9] The logarithmic opinion pool is an alternative aggregation method. See Bjørnland et al. (2010) for an application of different forecasting combination and selection methods to Norwegian GDP.

[10] Hall & Mitchell (2007) show that maximizing the log score is equivalent to minimising the Kullback-Leibler distance between the models and the true, but unknown density.

[11] Bjørnland et al. (2009) demonstrate that the log score for a model with normally distributed errors is a transformation of the Mean Squared Error (MSE), modified by the sample size and the unknown variance.

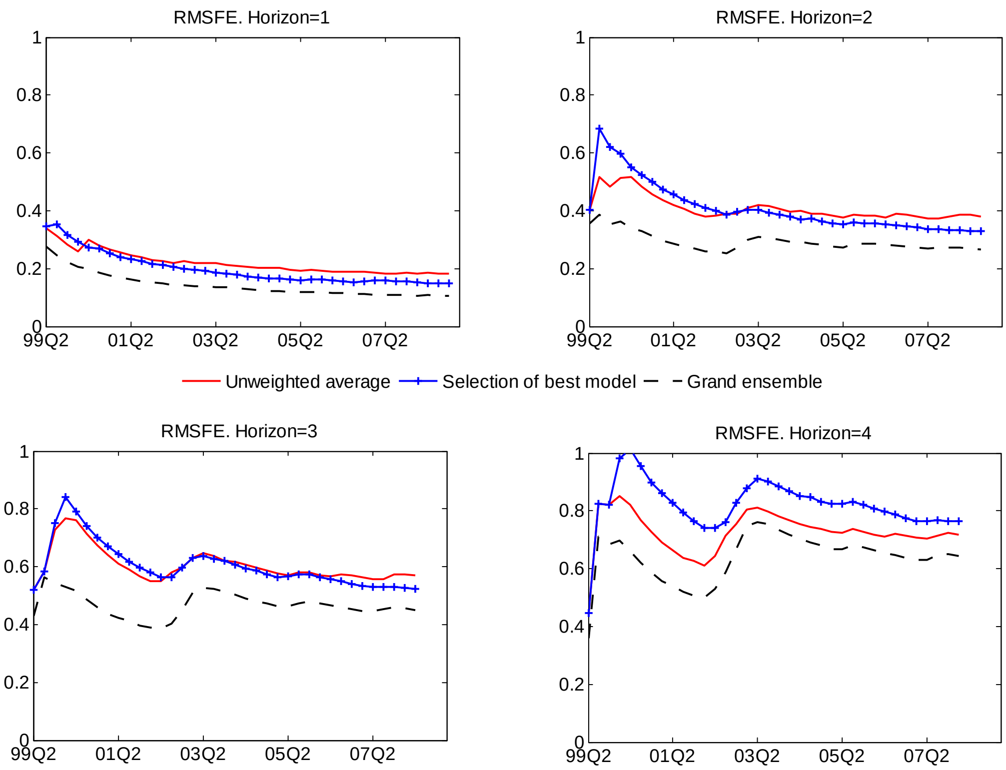

Figure 2: Recursive RMSFE for different model combination schemes. Inflation (CPI-ATE). Based on information available in April 2010.

In terms of point forecast accuracy, the grand ensemble performs somewhat better than the other weighting schemes. Trying to pick the best model at each point in time seems risky, as this alternative tends to perform worse than the combination alternatives for some horizons. This conclusion is supported by evaluating the density fit, based on log-score weights, for the three alternatives (not shown here). In figure 3, we show the individual forecast densities for the 10 ensembles used to forecast inflation for the forecast round in April 2010. As the horizon increases, the forecast densities are getting wider and the means from the different models are increasingly dispersed.

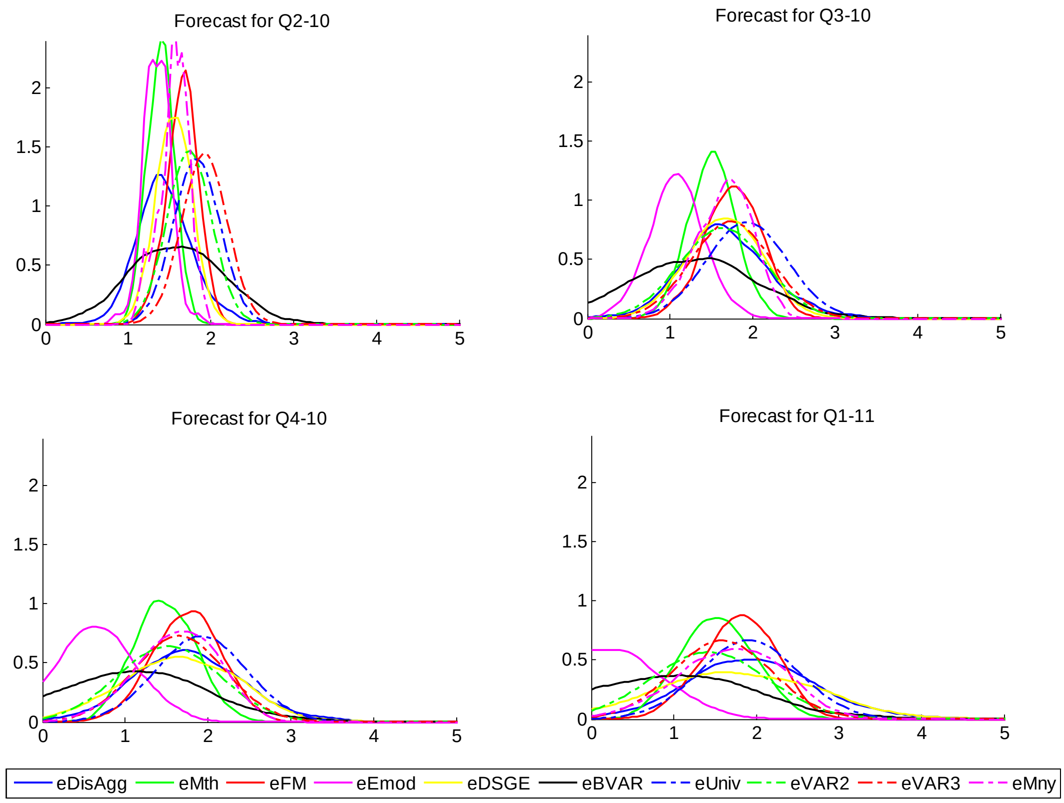

Figure 3: Density forecasts for inflation. All ensembles. Based on information available in April 2010.

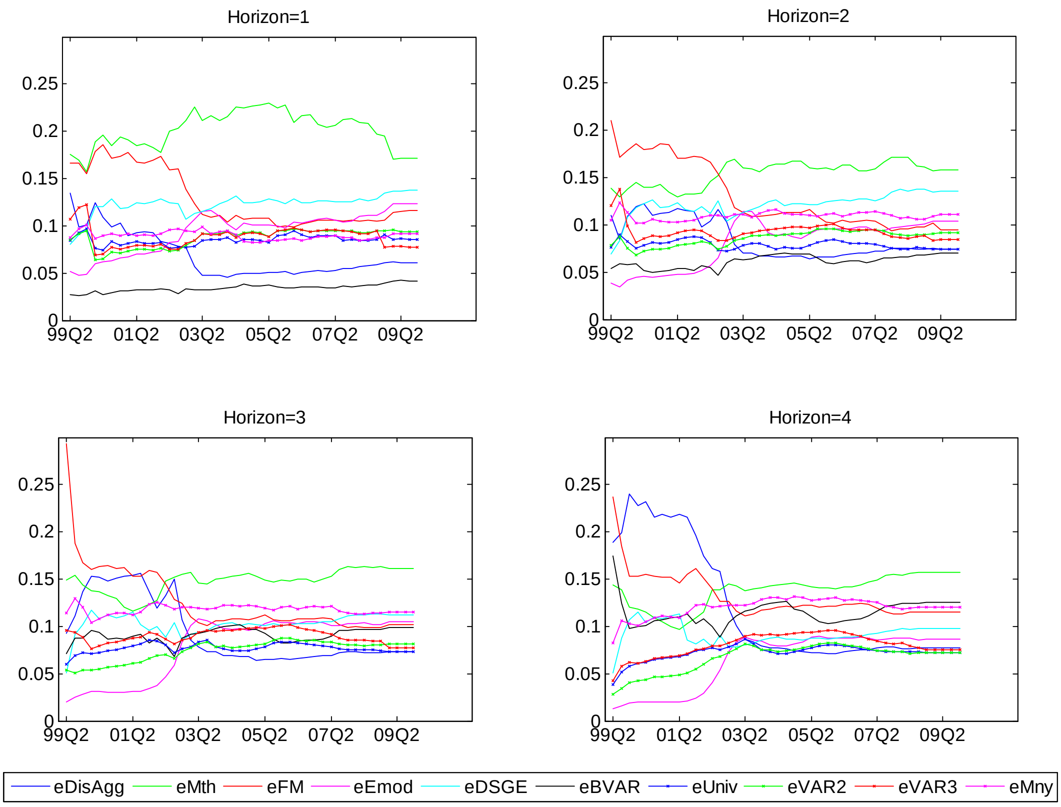

Figure 4: Time-varying (ex-post) weights for inflation forecasts for horizons 1 to 4. Based on information available in April 2010.

As described above, the ensemble densities are combined by using inverse MSE weights, calculated recursively. Figure 4 illustrates the developments of the weights for the inflation forecasts through the evaluation period.[12] For all horizons, models with monthly information get the highest weight. At the beginning of the evaluation period we have very few observations, and the weights are fluctuating considerably. But after a few years the weights are quite stable.

2.1.3 The final short term forecasts

Forecasts produced by SAM are free of judgement.[13] However, the final short-term forecasts for GDP and CPI-ATE that are used as starting values and conditioning assumptions in NEMO are in general subject to judgement. This judgement can e.g., be based on information not yet embedded in published statistics. Moreover, in some situations it is useful to study forecasts from selected individual models, or ensembles, in addition to forecasts from the grand ensemble. The final short term forecasts that are used as starting values and conditioning assumptions in NEMO are the responsibility of the sector experts. The sector expert forecasts are informed by SAM forecasts, forecasts from other models

[12] The dates on the horizontal axes correspond to the quarter when the forecasts are computed.

[13] Of cource, the selection of models, the choice of ensembles and the weigthing scheme will all influence the forecasts. But once these choices are made, the forecasts can be interpreted as being free of judgement.

and off-model judgement.

SAM forecasts are updated regularly and will be published as part of the background material for the Executive Board's monetary policy decisions. Figures 5 and 6 depict the fan-charts for inflation and GDP-growth, based on combined densities. The dotted, black lines are the point forecasts published in MPR 1/10 (March). The green dotted lines are the mean forecast from March calculated by SAM for the report. Judgement and sector expertise resulted in lower inflation forecasts in the MPR than the SAM-forecasts from March indicated, whereas the opposite was true for GDP-growth.

The current version of SAM only produces forecasts for two variables. The system is still under development, and our aim is to produce density forecasts for all the observable endogenous variables in NEMO. Currently, short term forecasts for key variables such as the components of demand, wage growth, labour market variables, foreign trade, the krone exchange rate and developments among our trading partners are made outside SAM, on the basis of current statistics and simple econometric models. A simple accounting framework ensures consistence between short-term forecasts for the different sectors of the economy.

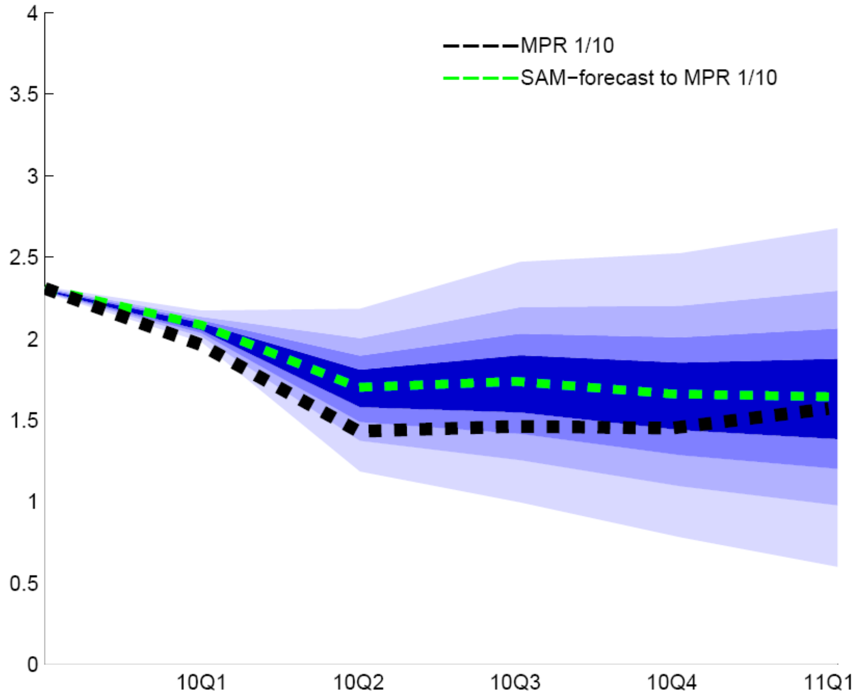

Figure 5: SAM forecast density and final short-term forecasts MPR 1/10. Inflation.

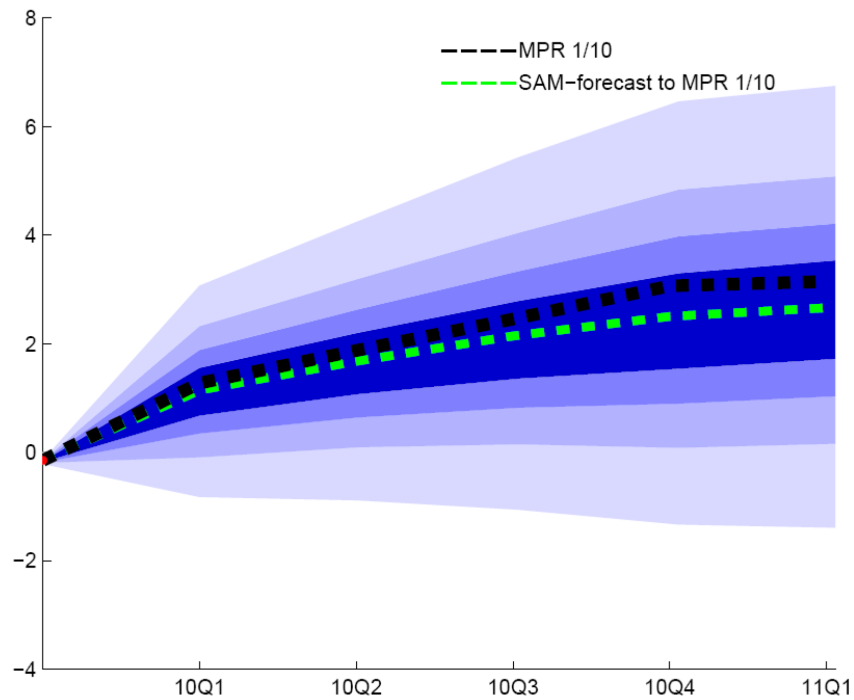

Figure 6: SAM forecast density and final short-term forecasts MPR 1/10. GDP growth.

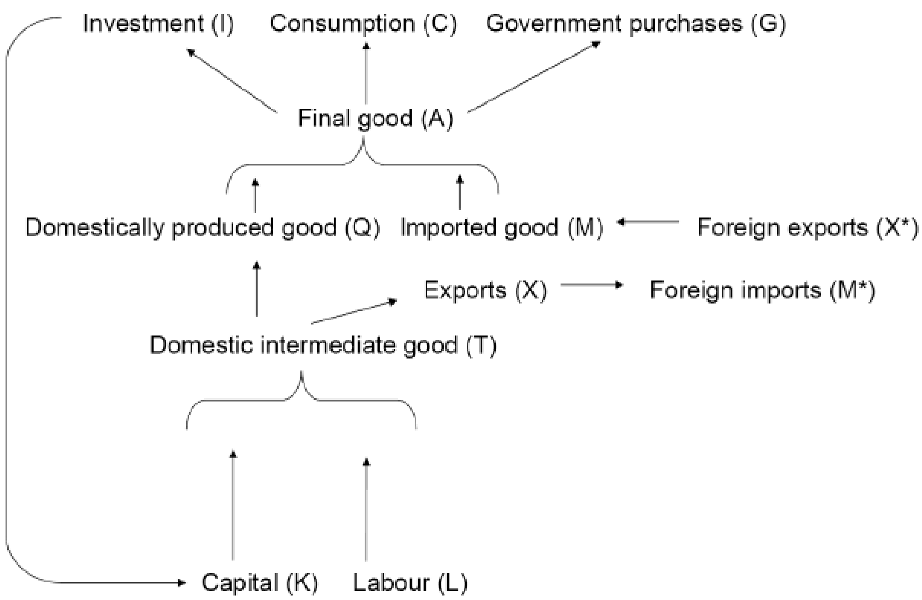

Figure 7: The production structure of NEMO

2.2 The Norwegian economy model (NEMO)

NEMO is Norges Bank's main forecasting and monetary policy analysis tool.[14] The model is based on international research and model development over the past 10–15 years and has many features in common with similar models in other central banks. NEMO has been under development since autumn 2004 and has been used as the core model since 2007.

A consistent theoretical framework makes it easier to interpret relationships and mechanisms in the model in light of economic theory. One advantage is that we can analyse the economic effects of changes of a more structural nature. In NEMO, developments in the Norwegian economy can be explained by changes in firms' technology, competitive conditions in product and labour markets, household preferences between consumption and leisure, and monetary policy. The structural framework makes it possible to provide a consistent, theoretical rationale for Norges Bank's projections. This distinguishes NEMO from the purely statistical models, which to a limited extent provide scope for economic 'story-telling'.

When constructing NEMO particular emphasis was given to developing a model that would be a useful decision-making tool in monetary policy. Therefore, it has been constructed with a view to being transparent and manageable. Output, price-setting, wage formation and all the main demand components are modelled and a distinction is made between domestic and imported inflation.

Figure 7 depicts the overall demand and supply structure of NEMO. The domestic economy has two production sectors, an intermediate goods sector and a final goods sector. Each intermediate good is produced by a single firm, using differentiated labour (L)

[14] A sketch of the model is given in the appendix. For a more detailed description of the model, see Brubakk et al. (2006).

and capital (K) services as inputs. The market for intermediate goods is characterized by monopolistic competition. The intermediate good (T) can be exported (M^*) or sold domestically (Q) to the final goods sector. The monopolistically competitive intermediate good firms set prices as a mark-up over marginal costs. Since we abstract from the possibility of arbitrage across countries, intermediate good firms can set different prices at home and abroad. Furthermore, we assume that it is costly for intermediate firms to change their prices. Prices are set in the currency of the buyer (local currency pricing). The specification of the price adjustment costs is consistent with Rotemberg (1982). This assumption implies a 'hybrid' Phillips curve that includes both expected future inflation and lagged inflation. Intermediate firms choose hours, capital[15], investment, the capital utilization rate and prices to maximize the present discounted value of cash-flows, taking into account the law of motion for capital, and demand both at home and abroad. Firms in the perfectly competitive final goods sector combine domestically produced (Q) and imported intermediate goods (M) into an aggregate good (A) that can be used for private consumption (C), investment (I), government spending (G).[16]

There are two types of households in the economy 'spenders' (or liquidity constrained households) and 'savers'. The spenders simply consume their disposable income. The remaining households, the savers, have access to domestic and foreign capital markets, and base their consumption decisions on an intertemporal optimization problem. Each household is the monopolistic supplier of a differentiated labour input. The household sets the nominal wage subject to the labour demand of intermediate goods firms and subject to quadratic costs of nominal wage adjustment. This assumption implies a 'hybrid' Phillips curve for wages. The model is closed by assuming that domestic households pay a premium on the foreign interest rate when they borrow in foreign bonds. The premium is increasing in the aggregate level of foreign debt in the domestic economy. The model evolves around a balanced growth path, where the growth rate is determined by exogenous technological growth. For simplicity, the fiscal authority is assumed to run a balanced budget each period, financed by lump-sum taxes. The small open economy assumption implies that the foreign economy is fully exogenous from the point of view of domestic agents. Hence, economic developments in Norway have no effects on its trading partners.

A key aspect of the model relates to the assumption regarding monetary policy. In the literature, there are two common ways to model monetary policy: either by a simple interest rate rule, or by 'optimal' monetary policy, in the sense of minimising an (ad-hoc) loss function. A popular approach is to use a generalised Taylor rule of the following type:

$$r_t^* = \lambda_r r_{t-1}^* + (1 - \lambda_r) \left[\lambda_\pi \pi_t + \lambda_{\Delta\pi} \Delta\pi_t + \lambda_y y_t + \lambda_{\Delta y} \Delta y_t \right], \tag{4}$$

where r_t^* refers to the key policy rate, π_t is the inflation rate and y_t denotes the output

[15] Capital is firm-specific, but since all firms are identical and there is no price dispersion this assumption does not affect the linearised dynamics of the model.

[16] We model the mainland economy, that is, the total economy excluding the oil sector. However, although oil production is not modeled, we include (exogenously) oil investments on the demand side, affecting mainland industries.

gap.[17] This type of rule has been shown to perform reasonably well in a variety of models, and is the most commonly used way to specify monetary policy in forecasting models. Among the other central banks publishing interest rate forecasts, the Reserve Bank of New Zealand and the Riksbank model the interest rate path using such rules.[18]

Modelling monetary policy by an interest rate rule like (4) has the advantage of being simple, intuitive and easy to implement. It has, however, at least two disadvantages. First, it does not address the time-inconsistency problem explicitly. Second, due to its simplicity it is not "optimal" in the sense of fully minimising a loss function.

When computing 'optimal' policy projections, one needs to specify a loss function. The wording of the Bank's mandate indicates that both inflation and the output gap are obvious candidates. However, minimising a loss function with the inflation gap and the output gap as the only arguments often leads to quite aggressive interest rate responses to shocks and may therefore look unacceptable to the policymakers. Hence, a natural extension to the standard set-up is to add an interest rate smoothing term in the loss function, i.e.:

$$E_t \sum_{i=0}^{\infty} \beta^i \left[\pi_{t+i}^2 + \omega_y y_{t+i}^2 + \omega_{\Delta r} \left(\Delta r_{t+i}^* \right)^2 \right], \tag{5}$$

where β is the discount factor of the central bank. The central bank minimises the loss function subject to the log-linearised first-order conditions of the private sector and the exogenous shock processes.[19]

Several authors have argued that simple instrument rules could be more robust to model misspecification than the loss-function based approach. Interest rate paths based on simple rules could thus serve as useful cross-checks. In NEMO we operationalise this cross-check by including deviations from a simple instrument rule as an argument in the loss function:

$$E_t \sum_{i=0}^{\infty} \beta^i \left[\pi_{t+i}^2 + \omega_y y_{t+i}^2 + \omega_{\Delta r} \left(\Delta r_{t+i}^* \right)^2 + \omega_s \left(r_t^* - r_t^{*,SIMPLE} \right)^2 \right], \tag{6}$$

where $r_t^{*,SIMPLE}$ is the interest rate implied by a simple instrument rule.[20]

NEMO has been estimated using Bayesian techniques on data for the Norwegian mainland economy for the period 1981–2007. The model is estimated under two different assumptions regarding monetary policy: a simple instrument rule and optimal policy under timeless perspective commitment. The results are reported in Bache et al. (2010).[21] The results show that the in-sample fit of the model with optimal policy is superior to the model with a simple instrument rule. However, in terms of forecasting accuracy, which is our favoured measure of model fit, the models perform about equally well. The forecast performance of the DSGE model is superior to that of an unrestricted VAR based on the

[17] The variables are measured as log-deviations from their steady-state values. The output gap is defined as the deviation of output from a permanent technology trend.

[18] See Hampton (2002) and the Riksbank's Monetary Policy Report 1/2007.

[19] The optimal policy projections depend on the assumptions about the degree of commitment. See Alstadheim et al. (2010) for a discussion.

[20] See Alstadheim et al. (2010) for details and a more thorough discussion of the monetary policy analysis at Norges Bank. See also box in MPR 2/10.

[21] The estimated model differs slightly from the version of NEMO used in the actual forecasting process.

same set observables, but slightly inferior to those of a BVAR.

3 Conditional forecasting

Conditional forecasting has a long history at policy institutions like central banks, and has also received attention in the recent academic literature.[22] There are many arguments in favour of using conditional information to guide the core model forecasts. First, most data are published with a considerable time lag and arrive at different frequencies. For example, preliminary Norwegian national accounts data for quarter t are published almost two months after the end of that quarter. However, a significant amount of relevant information in the form e.g., monthly observations of manufacturing production and retail sales, will be available long before the national accounts data are published. Conditioning on this information should provide an advantage in the short end of the forecasting horizon, but could also translate into a more persistent advantage.[23] Second, our structural DSGE models are likely to be misspecified along one or several dimensions. Conditioning on forecast information from other sources (e.g. other models, market-based indicators or off-model judgement) for variables of the core model where misspecification is believed to be of particular concern, potentially improves the overall forecast performance of the model. Third, in the event of huge and unexpected shocks, like experienced e.g. during the recent financial turmoil, models that lack the flexibility to adapt are very likely to deliver poor (short-term) forecasts. Thus, conditioning on information obtained from alternative sources may significantly reduce the uncertainty in the endogenous variables and thereby improve the forecasting performance of a DSGE model without necessarily having to change its structure.

Technically, conditional forecasting involves adding a sequence of structural shocks to the model over the forecasting period so that the model exactly reproduces the conditioning information. For the exogenous variables there is a unique sequence of shocks that will reproduce any given path.[24] Moreover, if the number of conditioning variables is equal to the number of structural shocks, the combination of shocks is still unique. If, however, the number of shocks exceeds the number of conditioning variables, a choice has to be made regarding which combination of disturbances to include. This choice can be made on the basis of judgement or on the basis of an optimality criterion as proposed by Waggoner & Zha (1999) in the context of VARs. This approach involves selecting the combination of shocks with the smallest variance that is consistent with the conditioning information. Thus the conditional forecasts will represent the most likely outcomes from the perspective of the model, given the conditioning assumptions.

[22] See e.g., Doan et al. (1984), Waggoner & Zha (1999) and Robertson et al. (2005) for applications to VARs and Adolfson et al. (2005), Christoffel et al. (2007), Benes et al. (2008) and Maih (2010) for extensions to DSGE models.

[23] An alternative approach, which involves extending the state space representation of the DSGE model with auxiliary information, is suggested by Monti (2008).

[24] Strictly speaking there are no exogenous variables in NEMO. In what follows, we will refer to exogenous variables as observable variables that are modelled as autoregressive (AR) processes and hence are not affected by the evolution of other variables in the model.

For a central bank that publishes its interest rate forecast, forecasting and monetary policy analysis are closely interrelated. The interest rate path has both a descriptive and normative element; it represents the central bank's best forecast of the future key policy rate and at the same time, reflects the policymakers' view on what is the appropriate interest rate path given the objectives of monetary policy. Below we describe how NEMO is used as a tool to help policymakers decide on an appropriate interest rate path.

3.1 Conditional forecasting and monetary policy analysis at Norges Bank

The discussion of the interest rate path is centred around two questions. First, what are the implications of the new data and the conditioning information for the interest rate path? Second, which interest rate path best meets the objectives of the central bank? This question relates to what is the appropriate *level* of the interest rate.

In analysing the implications of new information we take as given the view of the monetary policy transmission mechanism and the preferences of the policymaker implicit in the most recent interest rate path. This involves computing forecasts based on the same model specification and the same loss function for monetary policy as in the previous forecast round. Keeping the model and the loss function constant serves as a disciplining device for the discussion of the interest rate path and helps ensure consistency over time. However, although the interest rate path that comes out of this exercise is a useful benchmark, the policymakers' choices are obviously not restricted to this path. First, the model equations and the parameters could be adapted to reflect new insight about the functioning of the economy. Second, the loss function is only a crude approximation to the preferences of the Executive Board and could be misspecified in ways that could not be inferred from past interest rate decisions.

The first step in every forecast round is to assess how new and revised historical data affect the interpretation of recent economic developments. Technically, this involves running the Kalman-filter on the state-space representation of the model up to the start of the forecast horizon. The Kalman-filter will produce new estimates of the historical disturbances affecting the economy (e.g., technology shocks, demand shocks, mark-up shocks) and unobservable variables such as the output gap. This estimate of the output gap from the model is cross-checked against estimates from statistical models such as the Hodrick-Prescott filter, unobserved component models and the production function method.

The second step is to analyse the implications of the new conditioning information. In NEMO the conditioning information includes some of the exogenous variables (e.g., foreign variables, government spending, oil investments) over the entire forecast horizon and short-term forecasts for observable endogenous variables. We use an informal approach to choose a sequence of shocks to make the forecasts from NEMO consistent with the conditioning information. The choice is based on our assessment of the underlying driving forces in the economy. Our baseline assumption is that the conditioning information is anticipated.[25]

[25] We do not, however, allow the conditioning information to affect the estimate of the state of the economy at the beginning of the forecast period.

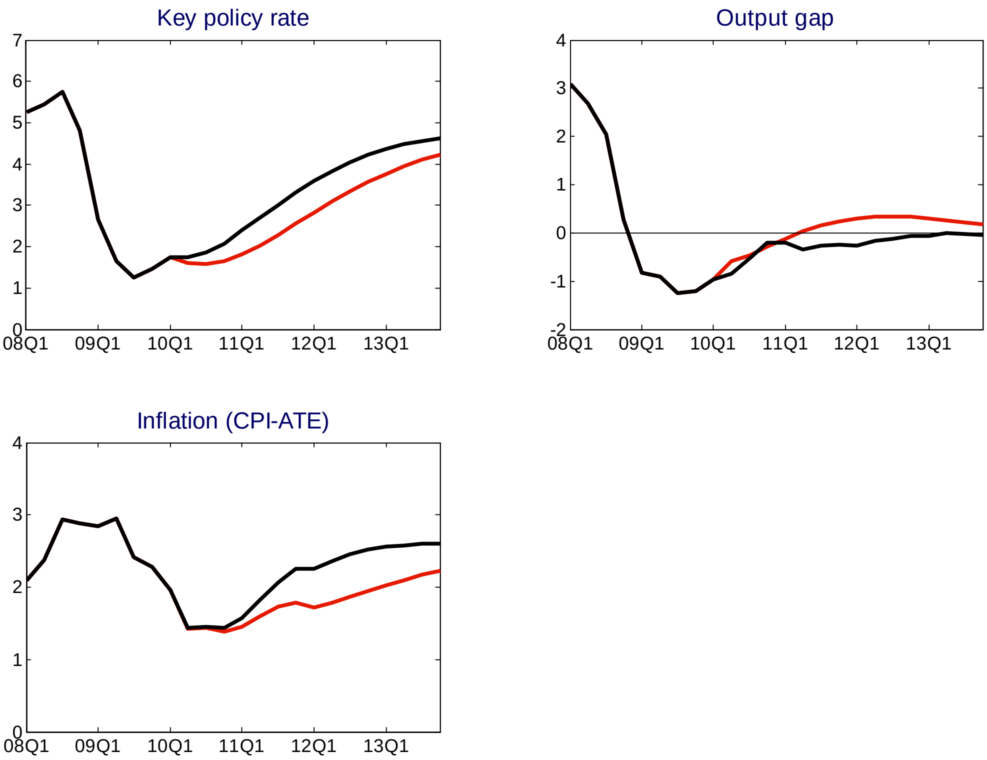

Figure 8: Unconditional (red lines) and conditional (black lines) forecasts from NEMO March 2010.

This ensures that the central bank will not be surprised by, and monetary policy will not react to, outcomes that turn out as projected.

In practice, the forecasting process is iterative. The first step involves computing forecasts from NEMO given the initial short-term forecasts provided by the sector experts. Then, based on the implications of the short-term forecasts for the structural shocks and the endogenous variables, the sector experts may revise their short-term forecasts. Subsequently, the revised short-term forecasts are used as new conditioning information in NEMO. The iteration continues until convergence is reached. For some variables, the sector experts also produce forecasts beyond the short-term horizon that serve as cross-checks for the medium-term NEMO forecasts.

As an additional exercise we also produce unconditional forecasts from NEMO in each forecast round (see figure 8). These provide valuable insight into the mechanisms in the model and serve as a cross-check on the short-term forecasts. Moreover, they allow us to assess the amount of judgement added to the forecasts and the implications of that judgement for the interest rate path.

In the MPR the uncertainty associated with the point forecasts is illustrated using fan charts. Forecast densities taking account of parameter and/or shock uncertainty are easily available from NEMO. So far, however, the fan charts published in the reports have been based on estimated historical disturbances to the supply and demand side in the

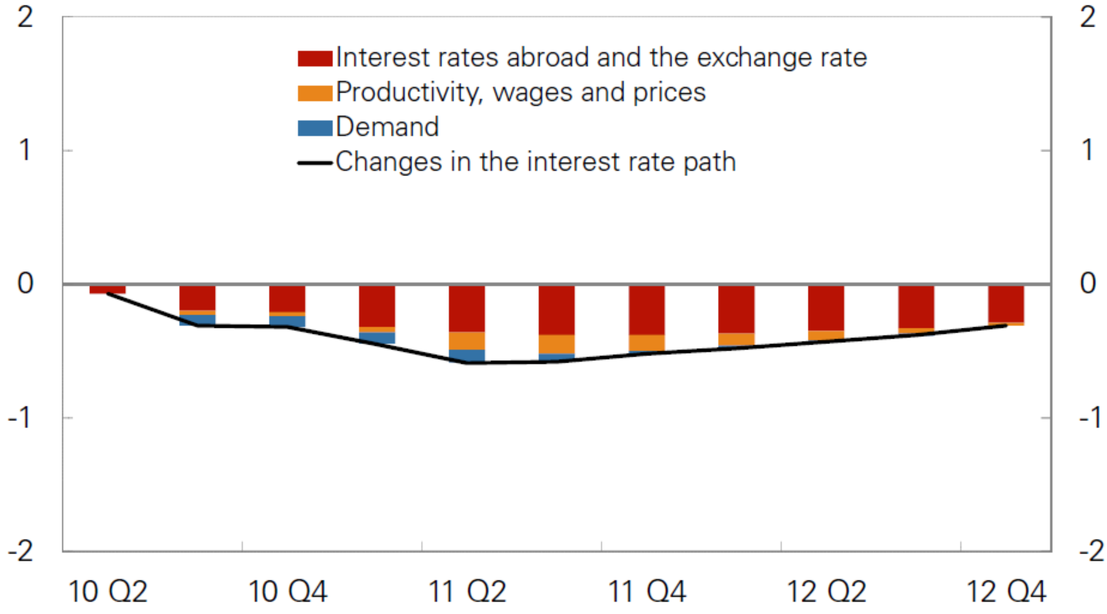

Figure 9: Interest rate account MPR 1/10

Norwegian economy identified from a small macroeconomic model.[26] In the MPR we also present scenarios based on alternative conditioning assumptions. The scenarios serve to highlight assumptions that have received particular attention in the course of the forecast process. The exact specification of the scenarios differ from one Report to the next, but the shifts in the interest rate, and the corresponding scenarios for inflation and the output gap give an indication of how the Bank responds. The shifts are specified such that, should these outcomes materialise, the alternative interest rate path is the Bank's best estimate of how monetary policy would respond. The shifts are consistent with the main scenario in the sense that they are based on the same loss function guiding the response of the central bank.

In addition to analysing the implications of new conditioning information for a given model specification, we also produce alternative interest rate paths based on different assumptions about monetary policy. The discussion is structured around the set of criteria for an appropriate interest rate path published in the MPR. The alternative interest rate paths are based on alternative loss functions for monetary policy, e.g. different weights on output gap variability, interest rate changes or deviations from simple instrument rules.[27]

A key ingredient in Norges Bank's communication approach is the so-called interest rate account (see figure 9). The interest rate account is a technical model-based illustration of how the change in the interest rate forecast from the previous report can be decomposed into the contributions from different exogenous shocks to the model. In the Report the

[26]See Inflation Report 3/05 for details. In normal circumstances, the fan charts are symmetrical and there is no distinction between the mean, mode and the median forecasts. During the recent financial crisis, the key policy rate was reduced to a historically low level. Since the key policy rate in principle has a lower bound close to zero, we set all outcomes implying a negative interest rate, to zero. Technically, the mean value for the interest rate was then marginally higher than the interest rate forecast, which could be interpreted as the median forecast.

[27]See box in Monetary Policy Report 2/10.

disturbances are grouped together in a few main categories; demand shocks, shocks to prices, costs and productivity, shocks to the exchange rate risk premium and foreign interest rates and shocks to money market spreads. If parameters in the model are changed from one forecast round to the next, the contribution from that change will be reflected in the account.[28] In special circumstances, the policymaker may wish to deviate from a normal reaction pattern. This was the case in October 2008 when the reduction in the key policy rate was moved forward because of an unusually high level of uncertainty and a desire to stave off particularly adverse outcomes. The contribution from this change in policy preferences was made explicit in the interest rate account in MPR 3/08. Since the interest rate account follows from a specific model, the exact decomposition is model-dependent and should thus be interpreted as a model-based illustration rather than a precise description of the Executive Board's reaction pattern.

3.2 A new integrated approach to density-conditional forecasting

The next step in conditional forecasting will allow us to produce model consistent density forecasts (fan charts) under a wide range of conditioning assumptions based on our core DSGE model. A main ingredient is the procedure for density-conditional forecasting for DSGE models developed in Maih (2010). This procedure allows us to condition on short-term density forecasts from SAM to produce density-conditional forecasts in a formal and consistent fashion.[29]

The forecasts from a DSGE model are computed from the policy and transition function which has a state space representation

$$y_t = A(\theta)y_{t-1} + B(\theta)\varepsilon_t. \tag{7}$$

where y_t is an $m_A \times 1$ vector of endogenous variables and ε_t is an $m_\varepsilon \times 1$ vector of structural shocks. The matrices A and B represent the reduced form parameters of the model, which again are functions of the vector of structural parameters, θ.

To allow for the possibility of anticipated future disturbances, we can generalize equation (7) to[30]

$$y_t = A(\theta)y_{t-1} + \sum_{j=1}^{n} B_j(\theta)\varepsilon_{t+j-1}. \tag{8}$$

where n denotes the anticipating horizon of agents in the economy. When $n = 1$, we have that all shocks are unanticipated. The k-step forecast at time T for agents can now be

[28] E.g., effects of changes in the parameters in the Euler equation for consumption would be attributed to the category 'demand shocks'.

[29] Since NEMO is estimated using Bayesian techniques, the resulting forecast densities can be used to construct the exact finite-sample distribution of conditional forecasts. Hence, it makes sense to ask questions like "given the historical data, what is the probability that inflation will be below target in one year". This is much harder in models using the classical approach to statistical inference. There are ways, however, from a classical perspective to create hybrids of confidence intervals and probability intervals that in some sense incorporate parameter uncertainty into measures of forecast uncertainty—but nonetheless do not result in probability intervals conditioned on the data. See Sims (2003) for a discussion of this point.

[30] See Maih (2010) for details on the derivation of the A and B matrices.

written as

$$y_{T+k} = A^k y_T + \sum_{j=1}^{n} \sum_{s=1}^{k} A^{k-s} B_j \varepsilon_{T+j+s-1} \tag{9}$$

which can be written more compactly as

$$Y = \overline{Y} + \Phi\varepsilon \tag{10}$$

where $Y \sim N\left(\overline{Y}, \Phi\Phi'\right)$.

Conditioning information on the Y vector translates into restrictions on ε. Suppose we are given the following general restriction

$$DY \sim TN\left(\mu, \Omega, [L, U]\right)$$

where D is a $q \times mk$ restriction matrix, assumed to be of full rank, μ denotes the mean of the truncated normal distribution and Ω is the corresponding covariance matrix. From (10) this implies that

$$D\overline{Y} + R\varepsilon \sim TN\left(\mu, \Omega, [L, U]\right)$$

where $R \equiv D\Phi$ is of rank $q < h \equiv (n + k - 1)\, m_\varepsilon$. D could simply be a selection matrix or a more complicated matrix if we allow for cross-variable restrictions in the endogenous variables. Using the model properties to translate the restrictions on Y into restrictions on the shocks, the above expression implies that

$$R\varepsilon \sim TN\left(\mu - D\overline{Y}, \Omega, [\underline{r}, \overline{r}]\right)$$

where $\underline{r} \equiv L - D\overline{Y}$ and $\overline{r} \equiv U - D\overline{Y}$. In other words, we have mapped the conditioning information into restrictions on the structural shocks.

Since $q < h$, there is no unique way of representing the leading information in terms of the shocks. Maih (2010) shows how, by exploiting the singular value decomposition of R, one can characterize the distribution of shocks that is required to fulfil the restrictions. In particular, under hard conditions, the procedure will find a unique combination of shocks with the smallest variance.[31]

The use of a truncated normal distribution in which the bounds could be finite or infinite permits a characterization of the conditional forecasts in a continuous fashion where both hard and soft conditions are special cases. It also allows us to avoid a wasteful rejection sampling scheme when constructing the distribution of soft-conditional forecasts. In the case where the moments of the conditioning information (μ and Ω) are not known, one can use the theoretical counterparts given by the model. This framework is flexible as it allows for the conditioning information to come in the form of an interval or a truncated density. At any rate, incorporating short-term conditional information from SAM or sectoral experts will be straightforward.

The algorithm to generate conditional forecasts can be described as follows

[31] For details, see Maih (2010)

- Draw a parameter vector $\theta^{(i)}$ from the posterior distribution
- Calculate all relevant matrices that are functions of θ and recover the starting values for the unobservable variables using the Kalman filter
- Draw $\varepsilon^{(i)}$ and generate forecasts $\left\{y_{T+1}^{(i)}, y_{T+2}^{(i)}, ..., y_{T+k}^{(i)}\right\}$

As shown in Maih (2010), conditioning does not necessarily improve the forecasting performance of the model and can in some cases even lead to a deterioration in forecast accuracy. This happens when the dynamics of the model is at odds with the data or when the correlation between the conditioning information and the other variables in the model is insignificant. On the other hand, in the presence of good conditioning information, even a misspecified model can have its forecasting performance improved if it adequately nails the dynamics of the data or the correlation between the conditioning information and variables of interest. In the presence of model misspecification, hard conditioning is not necessarily the best way to go, no matter how accurate the conditioning information is. Tight cross-equation restrictions implied by the model that are forced upon the forecasts in hard conditioning can be relaxed with soft conditioning.

4 Concluding remarks

In this paper, we have given a broad outline of the analytical framework and the forecasting process underlying the policy projections at Norges Bank, focusing on the way we conduct conditional forecasting. We have also outlined a new integrated system for conditional forecasting and policy analysis currently under development. The new system will allow us to produce model consistent density forecasts (fan charts) based on our core DSGE model under a wide range of conditioning assumptions.

A recurring challenge for monetary policy analysis is how to deal with model uncertainty. We continually strive to improve NEMO and reduce the degree of misspecification. A key issue, not least in light of the recent financial turmoil, concerns the interplay between financial variables and the real economy and how these interrelations can be included in the modelling framework. Designing a monetary policy that is robust to model misspecification is clearly a challenging task, and the literature on robustness does not provide any clear guidance. Ideally, one should address the issue of robustness by considering a set of different models. Optimal interest rate paths in models based on alternative assumptions about the functioning of the economy could serve as useful cross-checks. A related exercise would be to condition on the interest rate path derived from NEMO and deriving the implications for inflation and the output gap in the competing models of the economy. Going forward, we plan to go further in this direction, keeping in mind, however, that for a small central bank it may be too costly to develop and maintain a large set of structural models.

References

Adolfson, M., Laséen, S., Lindé, J., & Villani, M. (2005). Are constant interest rate forecasts modest policy interventions? Evidence from a dynamic open-economy model. *International Finance*, 8, 509–544.

Alstadheim, R., Bache, I. W., Holmsen, A., & Røisland, Ø. (2010). Monetary policy analysis in practice. Unpublished mimeo, Norges Bank.

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

Andersson, M. & Löf, M. (2007). The Riksbank's new indicator procedures. *Sveriges Riksbank Economic Review*, 2007:1, 76–95.

Bache, I. W., Brubakk, L., & Maih, J. (2010). *Simple Rules versus Optimal Policy: What Fits?* Norges Bank Working Paper 2010/3.

Bache, I. W., Jore, A. S., Mitchell, J., & Vahey, S. P. (2009a). *Combining VAR and DSGE Forecast Densities.* Norges Bank Working Paper 2009/23.

Bache, I. W., Mitchell, J., Ravazzolo, F., & Vahey, S. P. (2009b). *Macro Modeling with Many Models.* Norges Bank Working Paper 2009/15.

Benes, J., Binning, A., & Lees, K. (2008). *Incorporating Judgement with DSGE Models.* Reserve Bank of New Zealand Discussion Paper 2008/10.

Bjørnland, H., Gerdrup, K., Smith, C., Jore, A. S., & Thorsrud, L. A. (2010). *Weights and pools for a Norwegian density combination.* Norges Bank Working Paper 2010/06.

Bjørnland, H. C., Gerdrup, K., Jore, A. S., Smith, C., & Thorsrud, L. A. (2009). There is more than one weight to skin a cat: Combining densities at Norges Bank. Unpublished manuscript, Norges Bank.

Bjørnland, H. C., Jore, A. S., Smith, C., & Thorsrud, L. A. (2008). *Improving and Evaluating Short Term Forecasts at the Norges Bank.* Norges Bank Staff Memo 4/2008.

Bloor, C. (2009). The use of statistical forecasting models at the Reserve Bank of New Zealand. *Reserve Bank of New Zealand Bulletin*, 72 (2), 21–26.

Brubakk, L., Husebø, T. A., Maih, J., Olsen, K., & Østnor, M. (2006). *Finding NEMO: Documentation of the Norwegian Economy Model.* Norges Bank Staff Memo 2006/6.

Christoffel, K., Coenen, G., & Warne, A. (2007). Conditional versus unconditional forecasting with the New Area Wide Model of the euro area. Unpublished manuscript, European Central Bank.

Coletti, D. & Murchison, S. (2002). Models in policy-making. *Bank of Canada Review*, Summer, 19–26.

Doan, T., Litterman, R., & Sims, C. (1984). Forecasting and conditional projection using realistic prior distributions. *Econometrics Reviews*, 3, 1–100.

Garratt, A., Mitchell, J., & Vahey, S. P. (2009). *Measuring Output Gap Uncertainty.* Reserve Bank of New Zealand Discussion Paper DP2009/15.

Gerdrup, K. R., Jore, A. S., Smith, C., & Thorsrud, L. A. (2009). *Evaluating Ensemble Density Combination - Forecasting GDP and Inflation.* Norges Bank Working Paper 2009/19.

Hall, S. G. & Mitchell, J. (2007). Combining density forecasts. *International Journal of Forecasting*, 23(1), 1–13.

Hall, S. G. & Mitchell, J. (2009). Recent developments in density forecasting. In T. C. Mills & K. Patterson (Eds.), *Palgrave Handbook of Econometrics*, volume 2.

Hampton, T. (2002). The role of the Reserve Bank's macro model in the formation of interest rate projections. *Reserve Bank of New Zealand Bulletin*, 65 (2).

Holmsen, A., Qvigstad, J. F., & Røisland, Ø. (2007). *Implementing and Communicating Optimal Monetary Policy.* Staff Memo 2007/3, Norges Bank.

Holmsen, A., Qvigstad, J. F., Røisland, Ø., & Solberg-Johansen, K. (2008). *Communicating Monetary Policy Intentions: The Case of Norges Bank.* Norges Bank Working Paper 2008/20.

Jore, A. S., Mitchell, J., & Vahey, S. P. (2010). Combining forecast densities from VARs with uncertain instabilities. *Journal of Applied Econometrics (forthcoming).* Available as NIESR Discussion Paper No. 303.

Kapetanios, G., Labhard, V., & Price, S. (2008). Forecast combination and the Bank of England's suite of statistical forecasting models. *Economic Modeling*, 25, 772–792.

Laxton, D. & Pesenti, P. (2003). Monetary rules for small, open, emerging economies. *Journal of Monetary Economics*, 50, 1109–1146.

Maih, J. (2010). *Conditional Forecasts in DSGE Models.* Norges Bank Working Paper 2010/07.

Monti, F. (2008). *Forecast with Judgment and Models.* National Bank of Belgium Working Paper Research 153.

Robertson, J. C., Tallman, E. W., & Whiteman, C. H. (2005). Forecasting using relative entropy. *Journal of Money, Credit, and Banking*, 37, 383–401.

Rotemberg, J. J. (1982). Monopolistic price adjustment and aggregate output. *Review of Economic Studies*, 49, 517–531.

Sims, C. A. (2003). Probability models for monetary policy decisions. Unpublished manuscript available Sims' home page at Princeton University.

Timmermann, A. (2006). Forecast combinations. In G. Elliott, C. W. J. Granger, & A. Timmermann (Eds.), *Handbook of Economic Forecasting* (pp. 136–96). Amsterdam: Elsevier.

Waggoner, D. F. & Zha, T. (1999). Conditional forecasts in dynamic multivariate models. *Review of Economics and Statistics*, 81, 639–651.

Winkler, R. (1981). Combining probability distributions from dependent information sources. *Management Science*, 27, 479–488.

A NEMO

This appendix provides more details on the main structure of NEMO. The model is continually under development and so the exact specification of the model used to produce the forecasts in the MPRs may differ from the model below.

Final goods sector The perfectly competitive final goods sector consists of a continuum of final good producers indexed by $x \in [0,1]$ that aggregates composite domestic intermediate goods, Q, and imports, M, using a constant elasticity of substitution (CES) technology:

$$A_t(x) = \left[\eta^{\frac{1}{\mu}} Q_t(x)^{1-\frac{1}{\mu}} + (1-\eta)^{\frac{1}{\mu}} M_t(x)^{1-\frac{1}{\mu}}\right]^{\frac{\mu}{\mu-1}}, \tag{11}$$

The degree of substitutability between the composite domestic and imported goods is determined by the parameter $\mu > 0$, whereas η $(0 \leq \eta \leq 1)$ measures the steady-state share of domestic intermediates in the final good for the case where relative prices are equal to 1.

The composite good $Q(x)$ is an index of differentiated domestic intermediate goods, produced by a continuum of firms $h \in [0,1]$:

$$Q_t(x) = \left[\int_0^1 Q_t(h,x)^{1-\frac{1}{\theta_t}} dh\right]^{\frac{\theta_t}{\theta_t-1}}, \tag{12}$$

where the time-varying elasticity of substitution between domestic intermediates is captured by θ_t and evolves according to an AR(1) process.

Similarly, the composite imported good is a CES aggregate of differentiated import goods indexed by $f \in [0,1]$:

$$M_t(x) = \left[\int_0^1 M_t(f,x)^{1-\frac{1}{\theta^f}} df\right]^{\frac{\theta^f}{\theta^f-1}}, \tag{13}$$

where $\theta^f > 1$ is the steady-state elasticity of substitution between imported goods.

Intermediate goods sector Each intermediate goods firm h is assumed to produce a differentiated good $T_t(h)$ for sale in domestic and foreign markets using the following CES production function:

$$T_t(h) = \left[(1-\alpha)^{\frac{1}{\xi}} \left(Z_t z_t^L l_t(h)\right)^{1-\frac{1}{\xi}} + \alpha^{\frac{1}{\xi}} \overline{K}_t(h)^{1-\frac{1}{\xi}}\right]^{\frac{\xi}{\xi-1}}, \tag{14}$$

where $\alpha \in [0,1]$ is the capital share and ξ denotes the elasticity of substitution between labour and capital. The variables $l_t(h)$ and $\overline{K}_t(h)$ denote, respectively, hours used and effective capital of firm h in period t. There are two exogenous shocks to productivity in the model: Z_t refers to an exogenous permanent (level) technology process, which grows at the gross rate π_t^z, whereas z_t^L denotes a temporary (stationary) shock to productivity (or labour utilization). The technology processes are modelled as

$$\ln(Z_t) = \ln(Z_{t-1}) + \ln(\pi^z) + \ln\left(\frac{\pi_t^z}{\pi^z}\right), \tag{15}$$

where

$$\ln\left(\frac{\pi_t^z}{\pi^z}\right) = \lambda_z \ln\left(\frac{\pi_{t-1}^z}{\pi^z}\right) + \varepsilon_t^z, \qquad 0 \leq \lambda_z < 1, \quad \varepsilon_t^z \sim iid\left(0, \sigma_z^2\right), \tag{16}$$

and

$$\ln\left(\frac{z_t^L}{z^L}\right) = \lambda_L \ln\left(\frac{z_{t-1}^L}{z^L}\right) + \varepsilon_t^L, \qquad 0 \leq \lambda_L < 1, \quad \varepsilon_t^L \sim iid\left(0, \sigma_L^2\right). \tag{17}$$

The variable $K_t(h)$ is defined as firm h's capital stock that is chosen in period t and becomes productive in period $t+1$. Firm h's *effective* capital in period t is related to the capital stock that was chosen in period $t-1$ by

$$\overline{K}_t(h) = u_t(h) K_{t-1}(h), \tag{18}$$

where $u_t(h)$ is the endogenous rate of capital utilization. When adjusting the utilization rate the firm incurs a cost of $\gamma_t^u(h)$ units of final goods per unit of capital. The cost function is

$$\gamma_t^u(h) = \phi_1^u \left(e^{\phi_2^u(u_t(h)-1)} - 1\right), \tag{19}$$

where ϕ_1^u and ϕ_2^u are parameters determining the cost of deviating from the steady state utilization rate. The steady state utilization rate is normalized to one.[32]

Firm h's law of motion for physical capital reads:

$$K_t(h) = (1-\delta) K_{t-1}(h) + \kappa_t(h) K_{t-1}(h), \tag{20}$$

where $\delta \in [0,1]$ is the rate of depreciation and $\kappa_t(h)$ denotes capital adjustment costs. The adjustment costs take the following form:

$$\begin{aligned} \kappa_t(h) \quad = \quad & \frac{I_t(h)}{K_{t-1}(h)} - \frac{\phi_1^I}{2}\left[\left(\frac{I_t(h)}{K_{t-1}(h)} - \left(\frac{I}{K} + z_t^I\right)\right)\right]^2 \\ & - \frac{\phi_2^I}{2}\left(\frac{I_t(h)}{K_{t-1}(h)} - \frac{I_{t-1}}{K_{t-2}}\right)^2, \end{aligned} \tag{21}$$

where I_t denotes investment and z_t^I is an investment shock[33] that evolves according to an AR(1) process

The labour input is a CES aggregate of hours supplied by a continuum of infinitely-lived

[32]Note that ϕ_1^u is not a free parameter. It is set to ensure that the marginal cost of utilisation is equal to the rental rate of capital in steady-state.

[33]This shock could e.g., represent changes in the relative price of consumption and investment.

households indexed by $j \in [0,1]$:

$$l_t(h) = \left[\int_0^1 l_t(h,j)^{1-\frac{1}{\psi_t}} dj \right]^{\frac{\psi_t}{\psi_t - 1}}, \tag{22}$$

where ψ_t denotes the elasticity of substitution between different types of labour that evolves according to an AR(1) process.

Firms sell their goods in markets characterised by monopolistic competition. International goods markets are segmented and firms set prices in the local currency of the buyer. An individual firm h charges $P_t^Q(h)$ in the home market and $P_t^{Mf}(h)$ abroad, where the latter is denoted in foreign currency. Nominal price stickiness is modelled by assuming that firms face quadratic costs of adjusting prices,

$$\gamma_t^{PQ}(h) \equiv \frac{\phi_1^Q}{2} \left[\frac{P_t^Q(h)}{\pi P_{t-1}^Q(h)} - 1 \right] + \frac{\phi_2^Q}{2} \left[\frac{P_t^Q(h) / P_{t-1}^Q(h)}{P_{t-1}^Q / P_{t-2}^Q} - 1 \right]^2, \tag{23}$$

$$\gamma_t^{PMf}(h) \equiv \frac{\phi_1^{Mf}}{2} \left[\frac{P_t^{Mf}(h)}{\pi P_{t-1}^{Mf}(h)} - 1 \right] + \frac{\phi_2^{Mf}}{2} \left[\frac{P_t^{Mf}(h) / P_{t-1}^{Mf}(h)}{P_{t-1}^{Mf} / P_{t-2}^{Mf}} - 1 \right]^2, \tag{24}$$

in the domestic and foreign market, respectively and where π denotes the steady-state inflation rate in the domestic economy. In every period cash-flows are paid out to the households as dividends.

Firms choose hours, capital[34], investment, the utilization rate and prices to maximize the present discounted value of cash-flows, adjusted for the cost of changing prices, taking into account the law of motion for capital, and demand both at home and abroad, $T_t^D(h)$. The latter is given by:

$$T_t^D(h) = \int_0^1 Q_t(h,x)dx + \int_0^1 M_t^f(h,x^f)dx^f \tag{25}$$

Households There are two types of households in the economy: 'spenders' (or liquidity constrained households) and 'savers'. The spenders simply consume their disposable income. Total consumption is a weighted average of the consumption levels of the two types of households.[35]

The savers' utility function is additively separable in consumption and leisure. The lifetime expected utility of household j is:

$$U_t(j) = E_t \sum_{i=0}^{\infty} \beta^i \left[z_{t+i}^u u\left(C_{t+i}(j)\right) - v\left(l_{t+i}(j)\right) \right], \tag{26}$$

[34] Capital is firm-specific, but since all firms are identical and there is no price dispersion this assumption does not affect the linearised dynamics of the model.

[35] We assume that the spenders' wage rate is equal to the savers' (average) wage and that they supply whatever is demanded of their type of labour.

where C denotes consumption, l is hours worked and β is the discount factor $0 < \beta < 1$. The consumption preference shock, z_t^u, evolves according to an AR(1) process.

The current period utility functions for consumption and labour choices, $u(C_t(j))$ and $v(l_t(j))$, are

$$u\left(C_t\left(j\right)\right) = \left(1 - b^c/\pi^z\right) \ln \left[\frac{\left(C_t\left(j\right) - b^c C_{t-1}\right)}{1 - b^c/\pi^z}\right], \tag{27}$$

and

$$v\left(l_t\left(j\right)\right) = \frac{1}{1+\zeta} l_t\left(j\right)^{1+\zeta}. \tag{28}$$

where the degree of external habit persistence in consumption is governed by the parameter b^c $(0 < b^c < 1)$ and the disutility of supplying labour is governed by the parameter $\zeta > 0$.

Each household is the monopolistic supplier of a differentiated labour input and sets the nominal wage subject to the labour demand of intermediate goods firms and subject to quadratic costs of adjustment, γ^W:

$$\gamma_t^W\left(j\right) \equiv \frac{\phi^W}{2}\left[\frac{W_t\left(j\right)/W_{t-1}\left(j\right)}{W_{t-1}/W_{t-2}} - 1\right]^2 \tag{29}$$

where W_t is the nominal wage rate.

The flow budget constraint for household j is:

$$\begin{aligned} P_t C_t\left(j\right) + S_t B_{H,t}^f\left(j\right) + B_t\left(j\right) \leq W_t\left(j\right) l_t\left(j\right)\left[1 - \gamma_t^W\left(j\right)\right] \\ + \left[1 - \gamma_{t-1}^{B^f}\right]\left(1 + r_{t-1}^f\right) S_t B_{H,t-1}^f\left(j\right) \\ + \left(1 + r_{t-1}\right) B_{t-1}\left(j\right) + DIV_t\left(j\right) - TAX_t\left(j\right), \end{aligned} \tag{30}$$

where S_t is the nominal exchange rate, $B_t\left(j\right)$ and $B_{H,t}^f\left(j\right)$ are household j's end of period t holdings of domestic and foreign bonds, respectively. Only the latter are traded internationally. The domestic short-term nominal interest rate is denoted by r_t, and the nominal return on foreign bonds is r_t^f. The variable DIV includes all profits from intermediate goods firms and nominal adjustment costs, which are rebated in a lump-sum fashion. Finally, home agents pay lump-sum (non-distortionary) net taxes, TAX_t, denominated in home currency.

A financial intermediation cost, γ^{B^f}, is introduced to guarantee that aggregate net foreign assets follow a stationary process. This cost depends on the aggregate net foreign asset position of the domestic economy. Specifically, the intermediation cost takes the following form[36]

$$\gamma_t^{B^f} = \phi^{B1} \frac{\exp\left(\phi^{B2}\left(\frac{S_t B_{H,t}^f}{P_t Z_t}\right)\right) - 1}{\exp\left(\phi^{B2}\left(\frac{S_t B_{H,t}^f}{P_t Z_t}\right)\right) + 1} + z_t^B, \tag{31}$$

where $0 \leq \phi^{B1} \leq 1$ and $\phi^{B2} > 0$. The exogenous 'risk premium', z_t^B, evolves according to an AR(1) process.

[36] See e.g., Laxton & Pesenti (2003) for a discussion of this specification of the intermediation cost.

Government The government purchases final goods financed through a lump-sum tax. Real government spending (adjusted for productivity), $g_t \equiv G_t/Z_t$, is modelled as an AR(1) process. The central bank sets a short-term nominal interest rate, r_t^*. We consider two alternative specifications of monetary policy. First, we assume that the behaviour of the central bank can be represented by a simple instrument rule. Specifically, the central bank sets the interest rate according to a rule which in its log-linearised version takes the form

$$r_t^* = \lambda_r r_{t-1}^* + (1-\lambda_r)\left[\lambda_\pi \pi_t + \lambda_{\Delta\pi}\Delta\pi_t + \lambda_y y_t + \lambda_{\Delta y}\Delta y_t\right], \tag{32}$$

where r_t^* refers to the key policy rate, π_t is the inflation rate and y_t denotes the output gap. The parameter $\lambda_r \in [0,1\rangle$ determines the degree of interest rate smoothing. Output (y_t) is measured in deviation from the stochastic productivity trend, the remaining variables are in deviation from their steady-state levels.

The alternative assumption about monetary policy is that the central bank sets the interest rate to minimise the intertemporal loss function.

$$L_t = E_t \sum_{i=0}^{\infty} \beta^i \left[\pi_{t+i}^2 + \omega_y y_{t+i}^2 + \omega_{\Delta r}\left(\Delta r_{t+i}^*\right)^2\right], \tag{33}$$

where β is the discount factor of the central bank. The loss function is minimized subject to the log-linearised first-order conditions of the private sector and the exogenous shock processes.

The interest rate that enters into the decisions of households and firms, r_t, equals the interest rate set by the monetary policy authority, r_t^*, plus a shock, z_t^r, that is

$$r_t = r_t^* + z_t^r \tag{34}$$

where z_t^r is modelled as an AR(1) process. This shock could be interpreted e.g., as variations in the banks interest rate margins or in the spread between the key policy rate and the short-term interest rate in the money market.

Foreign variables The foreign variables that enter the model are the real marginal cost of foreign firms, mc_t^f, the output gap, y_t^f, the interest rate r_t^f and the inflation rate π_t^f. The foreign variables are assumed to follow AR(1) processes.