WORKING paper

BANQUE DE FRANCE
EUROSYSTÈME

Stochastic Simulation of the FR-BDF Model and an Assessment of Uncertainty around Conditional Forecasts

Harry Turunen[1], Anastasia Zhutova[2],

Matthieu Lemoine[3]

July 2023, WP #920

Abstract

This paper presents a framework to introduce uncertainty into the FR-BDF model, used for macroeconomic projections and policy analysis at the Banque de France. Belonging to the semi-structural class of large-scale macroeconomic models, it is only fair to assume that FR-BDF may suffer from various types of misspecification. We do not seek to correct the latter, but instead we study its systematic nature using unobserved component models for the residuals of FR-BDF. Stochastic simulations based on random draws of innovations of these models allow us to work with applications that describe probabilities of events and risk in general. Applying this framework to the December 2022 forecast exercise of Banque de France, based on the available information at that time, the highest probability of observing a technical recession occurs in 2023Q2 and reaches 42%.

Keywords: Semi-Structural Modelling, Stochastic Simulation, Unobserved Component Model

JEL classification: C54, E37

[1] Banque de France, Harri.TURUNEN@banque-france.fr
[2] Banque de France, Anastasia.ZHUTOVA@banque-france.fr
[3] Banque de France, Matthieu.LEMOINE@banque-france.fr.
We would like to thank Jean-Francois Ouvrard, Florian Pelgrin and Benjamin Favetto for their comments and fruitful discussions.

NON-TECHNICAL SUMMARY

This paper focuses on developing a framework for the stochastic simulation of the FR-BDF model, which is used for medium-term economic forecasting and policy analysis. By incorporating randomness into the model, various new applications become possible, such as constructing confidence bands and fancharts around point forecasts, as well as analyzing probabilities of events and overall risk.

There are different ways to implement a stochastic framework, but the key requirement is the ability to generate random numbers that can be used as innovations in the model equations. One simple approach is bootstrapping, where random numbers are drawn from the unconditional distribution of historical residuals obtained from estimating the model. This method is easy to implement and ensures that the generated innovations align with the properties of the model and the available data.

However, bootstrapping has a limitation. It ignores potential intertemporal dependence within the residuals, which should also be present in the random innovations applied to the model. The FR-BDF model, like many others in its class, is susceptible to suffer from such issues because of possible misspecifications, such as omitted variables or incorrect functional forms. This potential misspecification can lead to residuals that have predictable patterns due to unmodeled features of the data.

To address this challenge, we propose a solution that leverages the systematic nature of the misspecification. We assume that this systematic behavior can be captured using an Unobserved Components Model (UCM), which splits the residuals into two components: one persistent and one transitory. These components are both unobservable but follow a Gaussian distribution. The task is to estimate the parameters of the UCM to distinguish between the two components.

Similar frameworks have been implemented by other central banks with semi-structural models. The ECB, for example, uses a similar approach on the residuals of the ECB-BASE model, while the Federal Reserve's FRB/US model employs different sampling frameworks, including bootstrapping and state-contingent methods.

The estimation results for the FR-BDF model indicate some degree of misspecification. The UCM parameter estimates reveal significant persistence in the residuals from one period to the next, although the degree of persistence varies across equations. Some equations have economically insignificant persistence, while others, such as those related to exchange rates, public sector employment, and non-energy imports, exhibit larger persistence.

A practical application of the framework involves analyzing the uncertainty bands based on the December 2022 macroeconomic projections of the Banque de France. The simulations are designed to retain all conditioning information by inverting the full forecast and obtaining baseline paths for exogenous stochastic shocks. This allows for a close replication of the forecast, with the mean of the stochastic simulations matching the forecast.

The usefulness of the proposed framework is demonstrated by estimating, on a quarterly basis, the probability of a technical recession, defined as negative GDP growth in two consecutive quarters. Based on the available information at that time, we calculate the probability of a technical recession in France from 2023Q1 to 2025Q4. We find a peak of 42% in 2023Q2 and a decrease to 16% by the end of the period, see Figure 1. The probability increases notably in 2025 due to convergence to the trend. Note that, based on available GDP data, technical recession did not realize in early 2023.

Figure 1. Quarterly probability of technical recession based on December 2022 macroeconomic projections of the Banque de France.

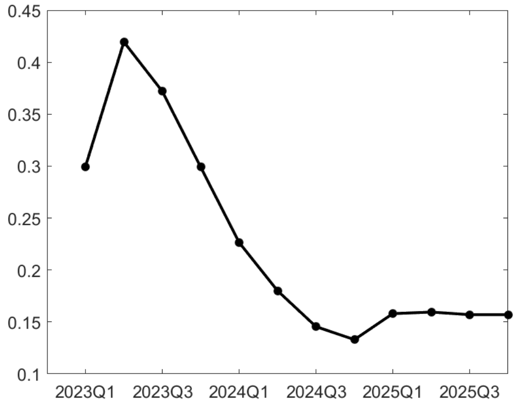

Simulation stochastique du modèle FR-BDF et évaluation de l'incertitude entourant les prévisions conditionnelles

RÉSUMÉ

Cet article présente un cadre permettant d'introduire de l'incertitude dans le modèle FR-BDF, un modèle macroéconomique de grande taille utilisé pour les projections macroéconomiques et l'analyse des politiques économiques à la Banque de France. Appartenant à la famille des modèles semi-structurels, il est légitime de supposer que le modèle FR-BDF peut souffrir de divers problèmes de mauvaise spécification. Nous ne cherchons pas ici à corriger ces derniers, mais nous étudions plutôt leur nature systématique en utilisant des modèles à composantes inobservables pour les résidus du modèle FR-BDF. Les simulations stochastiques fondées sur des tirages aléatoires des innovations de ces modèles nous permettent de mener des applications évaluant la probabilité d'événements et le risque de manière générale. En appliquant ce cadre à l'exercice de prévision de la Banque de France de décembre 2022, nous avons estimé, sur la base des informations disponibles à ce moment-là, que la probabilité la plus élevée d'observer une récession technique survient au 2023Q2 et atteint 42%.

Mots-clés : modélisation semi-structurelle, simulation stochastique, modèle à composantes inobservables

1 Introduction

The aim of this paper is to construct and apply a framework for the stochastic simulation of the FR-BDF model. This model, as described in Lemoine et al. (2019) is a tool intended for medium-term forecasting and counterfactual policy analysis; being able to conduct experiments with randomness enables a variety of new applications. These include the construction confidence bands and fancharts around point forecasts. On the analysis side, it enables applications that describe probabilities of events and risk in general.

Such a stochastic framework can be implemented in a variety of parametric and nonparametric ways – the only component absolutely necessary is to be able to produce random numbers that can be applied as innovations to the equations of the model. A simple example of the nonparametric approach is to draw from the unconditional distribution of historical, observed residuals obtained from estimating the model, i.e. bootstrapping. This approach has the convenient features of being both very easy to implement and producing innovations that by construction respect the properties of the model, e.g. by having moments that conform with the data and the equations.

The method of bootstrapping, however, also suffers from a notable limitation: by drawing from the unconditional distribution it ignores any potential intertemporal dependence within the residuals that should also be present in the random innovations applied to the model. The class of models within which FR-BDF resides is particularly likely to be subject to such issues.

This is due to the fact that the equations of FR-BDF – even though informed by economic theory – are likely to suffer from various types of misspecification, such as omitted variables or incorrect functional form. For example, the set of FR-BDF equations for business investment relate its long run target to output and the user cost of capital, but hypothetically this target could also depend on other things or the dependence on the factors present in the equation could take a different form. A key consequence of this potential misspecification is that the residuals of such equations could be predictable – they may have a systematic component that may be due to the unmodeled features of the data.

While such misspecification is almost by definition very difficult to correct by e.g. re-estimation, alternative techniques can be developed to account for it in applied work. In this paper we describe a solution that attempts to work around the problem: we do not attempt to correct the misspecification, but instead exploit its systematic nature to construct our framework for stochastic simulations.

More specifically, we assume that this systematic behavior can be modeled using an Unobserved Components Model (UCM), which is fundamentally a type of state space model. This class of models assumes that the observations – in our case residuals – can be linearly split into two stochastic components, one persistent, the other transitory, both Gaussian and unobservable directly. The econometric challenge is then to estimate the parameters describing these two processes in order to distinguish the two components from each other.

There are two leading techniques for solving this task. The first is to apply maximum likelihood-based estimation with a likelihood constructed using the Kalman filter. Our preferred method is however the Gibbs sampler, which we describe in further detail below. This choice is motivated by practical experience in applying the frequentist method: estimates based on the classical framework exhibit considerable instability depending on implementation details. That is, even minor changes in e.g. the initialization of the numerical algorithm can lead to drastically different outcomes.

Other central banks with semi-structural models have also implemented such frameworks. Angelini et al. (2019a) describe the methodology applied at the ECB on the residuals of the ECB-BASE model; these ideas have strongly inspired our approach and the use of the Gibbs sampler. In their approach the goal is also to exploit misspecification to obtain a temporally consistent model for the stochastic component. The FRB/US model of the Federal Reserve can be used with several different sampling frameworks, including bootstrapping, sampling from the unconditional distribution of residuals, and a state-contingent method, where the distribution is dependent on macroeconomic conditions – see González-Astudillo & Vilan (2019) for details on these methods. A recent example of an application of these ideas relevant to France can be found in Bourgeois & Favetto (2022), who use bootstrap methods to construct confidence intervals for the impulse responses of the Mesange model.

Our estimation results indicate that there is some misspecification present in FR-BDF. Several parameter estimates from the UCM imply significant persistence within the residuals from one period to the next, although there is also considerable variation in its degree from one equation to the next - for some equations the estimated persistence is different from zero, but small enough to be economically insignificant. Examples of the case where the persistence is large include the equations for exchange rates, public sector employment and the volume of non-energy imports.

To explore this framework for stochastic exercise, we start by analyzing the uncertainty bands obtained with it based on the scenario of December 2022 Banque de France macroeconomic projections. We construct our simulations such that all of the conditioning information – assumptions for certain variables and expert judgment – is retained by inverting the full forecast to obtain baseline paths for our exogenous stochastic shocks. This allows us to replicate almost exactly the forecast in the sense that the mean of the stochastic simulations corresponds to the forecast. We then conduct a practical application: we estimate the probability of technical recession, quarter by quarter, within the forecast period.

The rest of the paper is organized as follows. The next section discusses the concept of misspecification in FR-BDF with a concrete example. Section 3 describes the UCM for the residuals of FR-BDF and its estimation using Bayesian methods, while section 4 presents our practical exercises. Section 5 concludes and describes avenues for further work.

2 Potential misspecification in FR-BDF

Consider, as an example, the equations describing the behavior of business investment.[1] In FR-BDF this quantity is determined by two principal equations, one describing its long run behavior

$$\log I^*_{B,t} = \alpha_0 + q_t - \sigma \log r_{KB,t-1} + \log \frac{I^*}{K^*} \tag{1}$$

and the other short run dynamics

$$\begin{aligned} \Delta \log I_{B,t} \quad = \quad & \beta_0 \log \left(\frac{I^*_{B,t-1}}{I_{B,t-1}} \right) + \beta_1 \Delta \log I_{B,t-1} + \beta_2 \Delta \log I_{B,t-2} \\ & + \mathrm{PV} \left(\Delta \hat{q} \right)_{t|t-1} - \sigma \mathrm{PV} \left(\Delta \log \tilde{r}_{KB} \right)_{t|t-1} \\ & + \left(1 - \beta_1 - \beta_2 \right) \left(\Delta \hat{q}_{t-1} - \sigma \Delta \log \bar{r}_{KB,t-1} \right) \\ & + \beta_3 \left(\Delta q_{t-1} - \Delta \bar{q}_{t-1} \right) \end{aligned} \tag{2}$$

The first equation relates the behavior of target or desired investment $I^*_{B,t}$ to real corporate value added q_t and the user cost of capital $r_{KB,t-1}$. α_0 and σ are estimated coefficients and $\frac{I^*}{K^*}$ optimal investment-capital ratio and is assumed to be equal to the historical mean.

The second equation describes the dynamics of actual investment $I_{B,t}$. The change in log investment is assumed to depend on the ratio of target investment to actual investment in the previous period, two lags of the log change of investment, two terms representing expectations regarding future business value added $\mathrm{PV} \left(\Delta \hat{q} \right)_{t|t-1}$ and future user cost of capital $\mathrm{PV} \left(\Delta \log \tilde{r}_{KB} \right)_{t|t-1}$ and a term representing demand for business output $(\Delta q_{t-1} - \Delta \bar{q}_{t-1})$. Finally, β_0, β_1, β_2 and β_3 are estimated coefficients.

It is easy to see that there are many details in these equations that are potential sources of misspecification. While Lemoine et al. (2019) present clear, well-justified arguments for the various choices and assumptions within these equations, these decisions may be erroneous, in some cases simply by random chance that appears in the specificities of the data used to estimate them.

An example strongly linked to economic theory is presented in (1).[2] This equation rests on the assumption

[1] For an in-depth discussion of these equations and their role in FR-BDF, see Lemoine et al. (2019). We omit many details for brevity, such as the definition of the expectation terms, which are themselves sources of potential misspecification.

[2] While this equation is estimated, we do not include its residual in our subsequent empirical analysis. That said, this equation does play a role in our study: the variable $I^*_{B,t}$ is constructed using (1) and thus the residual of (2) is affected by the choices made in the specification of (1).

that the aggregate production function within the model has the Constant Elasticity of Substitution (CES) form, which together with some additional assumptions can be manipulated to obtain (1). An implication is that the coefficients in front of the terms q_t and $\log r_{KB,t-1}$ are 1 and $-\sigma$, respectively. In contrast, assuming that the aggregate production function has the Cobb-Douglas form would imply that the two coefficients are both equal to unity in absolute value. More exotic production functions would, in turn, imply that an alternative specification of (1) could include some additional terms: a more appropriate model for the production function might need to take into account e.g. public capital, the cost of energy or some sort of a split of either labor or capital into different types, such as high skilled and low skilled labor.

The great difficulty then in resolving these potential problems is that neither data nor theory present a clear path forward: the data imply that a particular specification offers better performance – at the very least in-sample – while discriminating between various theoretical approaches is always to some extent a subjective choice, even if informed by empirical observation. Thus, in the absence of convincing arguments against this subjective choice or against the features of the data, the researcher is forced to be satisfied with his model and all of its imperfections.

As we have no such suggestions for improvement, we proceed by attempting to circumvent these problems by constructing an auxiliary model that can be used in various applications to improve the practical performance of FR-BDF.

3 Construction and estimation of the UCM model

In this section we describe, in turn, the application of the UCM to the residuals of FR-BDF, how this model can be interpreted within a Bayesian framework and the estimation of the model using the Gibbs sampler.

3.1 The unobserved components model for the residuals of FR-BDF

The starting point for this model is the idea that the residuals of FR-BDF, suffering from misspecification, are predictable. Furthermore, in an UCM, this predictability is assumed to be such that the residuals can be split into two components, transitory and persistent, and that the persistent component can be modelled as an autoregressive process.

In more detail, we assume that the residuals e_t of FR-BDF follow the vector process given by the state equation

$$e_t = c_t + \epsilon_t \tag{3}$$

where $\epsilon_t \sim N(0, \Sigma)$ is the transitory component and c_t is the persistent component. Its movement is described by

$$c_t = Ac_{t-1} + \eta_t \tag{4}$$

with $\eta_t \sim N(0, \Omega)$. We assume that A is a diagonal matrix, with elements given by the vector $\rho = \text{diag}(\rho_1, \rho_2, ..., \rho_k)$ so that the components of c_t are individually AR(1). Furthermore, the initial state is given by $c_1 \sim N(c_0, \Omega_0)$.

We make the further assumptions that both Σ and Ω are positive definite and diagonal and that the elements of Ω are given by the vector $\omega = \text{diag}(\omega_1^2, \omega_2^2, ..., \omega_K^2)$.

3.2 The Bayesian framework of the unobserved components model

Notice first that the state equations can be stacked as

$$e = c + \epsilon \tag{5}$$

with $e = (e_1, e_2, ...e_T)'$, $c = (c_1, c_2, ...c_T)'$, $\epsilon = (\epsilon_1, \epsilon_2, ...\epsilon_T)'$ and $\epsilon \sim N(0, I_T \otimes \Sigma)$. The implied log likelihood of the data e is

$$\ln\, p(e|c, \Sigma) = -\frac{T}{2}\ln|\Sigma| - \frac{1}{2}(e - c)'(I_T \otimes \Sigma)^{-1}(e - c) \tag{6}$$

Similarly, we have a stacking for the unobserved persistent component c

$$Hc = u \tag{7}$$

with $u \sim N(0, S)$ where

$$H = \begin{bmatrix} I & 0 & 0 & \dots & 0 \\ -A & I & 0 & \dots & 0 \\ 0 & -A & I & \dots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & -A & I \end{bmatrix}$$

and

$$S = \begin{bmatrix} \Omega_0 & 0 & 0 & \dots & 0 \\ 0 & \Omega & 0 & \dots & 0 \\ 0 & 0 & \Omega & \dots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & 0 & \dots & \Omega \end{bmatrix}$$

As H is invertible, from (7) we obtain the prior distribution of c:

$$c = H^{-1}u \sim N\left[0, (H'SH)^{-1}\right] \tag{8}$$

implying that the joint log density of c is

$$\ln\, p\left(c|\rho,\Omega\right) \propto -\frac{T-1}{2}\ln|\Omega| - \frac{1}{2}c'H'S^{-1}Hc \tag{9}$$

Finally, we assume that the prior distributions of ρ and Σ are unconditionally independent:

$$p(\rho,\Sigma) = p(\rho)p(\Sigma) \tag{10}$$

and that the prior distributions of ρ_j and Σ are

$$\rho_j \sim N(\bar{\rho}, \sigma_\rho^2) \tag{11}$$

$$\Sigma \sim IW(L,\nu) \propto |\Sigma|^{-(\nu+k+1)/2}\text{exp}\left[-\frac{1}{2}tr\left(L\Sigma^{-1}\right)\right] \tag{12}$$

where IW refers to the inverse Wishart distribution.

Furthermore, notice that so far we have omitted the parameter vector ω from the discussion within the Bayesian framework. This is due to the fact that our estimation methodology differs from common procedure on this part. An implication of our assumptions – detailed in section 3.4.3 – is that the distribution of ω, conditional on the other parameters, collapses to a single point. That is, given the other parameters, ω is a constant.

Under these assumptions the joint posterior can be characterized as proportional to the product of the conditional likelihoods of the data e and the unobservable persistent component c (given by (6) and (9)) with the prior distributions of the parameters ρ, Σ:

$$p\left(c,\rho,\Sigma|e,\Omega\right) \propto p\left(e|c,\Sigma\right)p\left(c|\rho,\Omega\right)p\left(\rho\right)p\left(\Sigma\right) \tag{13}$$

This distribution is intractable, implying that classical methods are not feasible for the estimation of ρ and Σ. However, as the conditional distributions of the data, the persistent component and the parameters can be derived from the prior distributions, we apply the Gibbs sampler to this estimation task.

3.3 Gibbs sampling for the UCM

The Gibbs sampling procedure for the UCM is contingent on characterizing the conditional distributions of the parameters $p(c|e,\rho,\Sigma,\Omega)$, $p(\rho|e,c,\Sigma,\Omega)$ and $p(\Sigma|e,c,\rho,\Omega)$. Given these distributions, the procedure simply cycles through drawing from them until some iteration limit has been reached, after which the posterior distribution can be characterized using Monte Carlo methods.

The distribution of the persistent component conditional on the parameters can be obtained from the sum of (6) and (9):

$$\begin{aligned} \ln\, p(c|e,\rho,\Sigma,\Omega) &\propto -\frac{1}{2}\left[(e-c)'\left(I_T\otimes\Sigma\right)^{-1}(e-c)+c'\left(H'S^{-1}H\right)c\right]-\frac{T}{2}\ln\,|\Sigma|-\frac{T-1}{2}\ln\,|\Omega| \quad (14)\\ &\propto -\frac{1}{2}\left[(c-\hat{c})'P(c-\hat{c})\right]-\frac{T}{2}\ln\,|\Sigma|-\frac{T-1}{2}\ln\,|\Omega| \end{aligned}$$

where $P=H'S^{-1}H+(I_T\otimes\Sigma)^{-1}$ and $\hat{c}=P^{-1}(I_T\otimes\Sigma)^{-1}e$. Furthermore, we know that $p(c|e,\rho,\Sigma,\Omega)=N\left(\hat{c},P^{-1}\right)$.

By adding together (9) and the logarithm of (11), we can derive the conditional log distribution of each persistence parameter ρ_j, i.e. the elements of $p(\rho|e,c,\Sigma,\Omega)$ as

$$\begin{aligned} \ln\, p(\rho_j|e,c,\Sigma,\Omega) &\propto -\frac{1}{2}\left[\omega_j^{-2}\sum_{t=2}^{T}\left(c_{j,t}-\rho_j c_{j,t-1}\right)^2+\sigma_\rho^{-2}\left(\rho_j-\bar{\rho}\right)^2\right]-\frac{T}{2}\ln\,|\sigma_\rho^2|-\frac{T-1}{2}\ln\,|\Omega| \quad (15)\\ &\propto -\frac{1}{2}\left[\frac{\left(\rho_j-\hat{\rho}_j\right)^2}{\hat{\sigma}_\rho^2}\right]-\frac{T}{2}\ln\,|\sigma_\rho^2|-\frac{T-1}{2}\ln\,|\Omega|=N\left(\hat{\rho}_j,\hat{\sigma}_\rho^2\right) \end{aligned}$$

where

$$\hat{\sigma}_\rho^2=\left(\omega_j^{-2}\sum_{t=2}^{T}c_{j,t-1}^2+\sigma_\rho^{-2}\right)^{-1}$$

and

$$\hat{\rho}_j=\hat{\sigma}_\rho^2\left(\omega_j^{-2}\sum_{t=2}^{T}c_{j,t}c_{j,t-1}+\sigma_\rho^{-2}\bar{\rho}\right)$$

Note that to obtain (15) we exploit the fact that each individual $c_{j,t}$ is independent of other $c_{i,t}$, allowing us to characterize each distribution individually. Furthermore, the result hinges on the point that the product of two Normal distributions is also Normal.

The conditional distribution of the variance of the innovation to the transitory component, i.e. $p(\Sigma|e,c,\rho,\Omega)$, can be derived from the sum of (6) and the logarithm of (12):

$$\ln\, p(\Sigma|e,c,\rho) \;\propto\; -\frac{\nu+T+k+1}{2}\ln|\Sigma|-\frac{1}{2}\left[tr\left(\sum_{t=2}^{T}\left(e_t-c_t\right)\left(e_t-c_t\right)'\Sigma^{-1}\right)+tr\left(L\Sigma^{-1}\right)\right]$$

which implies

$$p(\Sigma|e, c, \rho) = IW\left[\left(L + \sum_{t=2}^{T}(e_t - c_t)(e_t - c_t)'\right), \nu + T\right] \tag{16}$$

as the inverse Wishart distribution is a conjugate distribution with the Normal distribution.

A standard Gibbs sampling procedure iterates over drawing sequentially from (14), (15) and (16) and for some set amount of draws M given some initial draw of the full parameter vector. Further details on our implementation of this procedure and the special treatment of ω are provided in section 3.4.3.

3.4 Data, prior parameters and implementation

3.4.1 Data

Our estimation of the UCM is based on using the same dataset as was used originally to estimate FR-BDF. More specifically, we use the detailed results of quarterly national accounts of 2018Q1, which were published in June 2018, together with the values of the observables as they were measured at that time for the variables that do not appear in the national accounts, e.g. some financial variables, such as interest rates on corporate bonds.

Our data set includes a total of $T = 52$ observations for $K = 26$ residuals corresponding to the FR-BDF equations that have estimated coefficients. Although the full set of estimated equations in FR-BDF is much larger, we focus on this subset for practical reasons related to the construction of the primary forecast, which follows a standard procedure set out by the ECB. This procedure entails that certain macroeconomic quantities, such as e.g. interest rates, are presumed to follow a pre-set path. As one of the purposes of this note is to conduct a comparison between this primary forecast and a stochastic forecast, the stochastic forecast has to be conducted using methods that produce a fair comparison, which implies that this procedure set out by the ECB is followed as closely as possible. Thus we exclude any variables affected by this procedure from our analysis. Further details on these assumptions are described in section 4.

The start date of our sample is the first quarter of 2005. This is the first date for which the residuals are available for the whole set we analyze. The end date of the sample on which the equations were estimated is 2017Q4, which is also the end date of the sample we analyze.

Table 3.4.1 presents some descriptive statistics for our data set. First, notice that for some of the series the mean of the residuals deviates clearly from zero – particularly strong examples are the two residuals corresponding to the equations describing housing credit held by households. Second, we present the variances of the residuals, within which we can see considerable variation. For the most part they range from 10^{-5} to 10^{-7}. However, there are also exceptions, which tend to occur together with deviations of the estimated mean from zero.

Finally, table 3.4.1 also contains estimated persistence parameters for simple AR(1)-type equations for each residual, which indicate that they are somewhat predictable, hinting at misspecification. Many of the estimates differ substantially from zero although there are also plenty of examples where the parameter does not deviate from zero in a economically meaningful way. However, it should be kept in mind that the observed residuals are mixtures of the persistent and transitive components of the UCM, and as such these estimates to some extent understate the stickiness of the persistent component.

Table 3.4.1: Descriptive statistics of the observed residuals

Residual	**Mean**	**Variance**	**Persistence**
Demand volumes			
Household consumption	$6.9 * 10^{-4}$	$1 * 10^{-5}$	0.16
Business investment	0.003	$1.1 * 10^{-4}$	0.14
Household investment	$5.4 * 10^{-4}$	$2.4 * 10^{-5}$	0.21
Exports	-0.002	$8.8 * 10^{-5}$	-0.11
Non-energy imports	0.006	$1.3 * 10^{-4}$	0.2
Energy imports	$7.1 * 10^{-4}$	0.002	-0.13
Deflators			
Value added	$2.5 * 10^{-4}$	$7.3 * 10^{-6}$	-0.03
Public consumption	$-2.1 * 10^{-4}$	$2.5 * 10^{-6}$	0.63
Household consumption	$2 * 10^{-4}$	$1.7 * 10^{-6}$	-0.03
Business investment	$-8.3 * 10^{-4}$	$4.8 * 10^{-6}$	-0.13
Household investment	0.003	$3 * 10^{-5}$	0.45
Exports	$-2.5 * 10^{-5}$	$7 * 10^{-6}$	-0.21
Non-energy imports	$2.7 * 10^{-5}$	$7.3 * 10^{-6}$	-0.05
Energy imports	$2.9 * 10^{-4}$	$5.5 * 10^{-4}$	0.05
Labour market			
Public sector wage	$-9 * 10^{-4}$	$1.4 * 10^{-5}$	0.31
Private sector wage	$3.5 * 10^{-4}$	$9.5 * 10^{-6}$	-0.11
Minimum wage	$6.2 * 10^{-4}$	$5.7 * 10^{-5}$	-0.1
Public sector employment	$-9 * 10^{-4}$	$6.3 * 10^{-6}$	0.44
Private sector employment	$-1.8 * 10^{-4}$	$9 * 10^{-7}$	0.07
Financial market			
CAC40 index	$7.7 * 10^{-4}$	0.002	-0.1
New household loans for housing	-0.02	0.007	-0.02
Price of housing	$8.5 * 10^{-4}$	$2.5 * 10^{-5}$	-0.06
Repayments of housing credit	-0.01	0.02	-0.08
Spread, BBB bonds	$2.3 * 10^{-5}$	$3 * 10^{-8}$	0.16
Spread, cost of equity	$2.4 * 10^{-5}$	$5.8 * 10^{-8}$	0.42

3.4.2 Prior parameters

Our chosen parameters for the prior distribution of the persistence ρ_j reflect our strong prior belief of misspecification potentially being present in the equations. This is represented especially by the rather high value for $\bar{\rho}_j = 0.9$ and the low value of $\sigma_\rho^2 = 0.01$, implying a particularly tight prior distribution for the ρ_j. We assume that this same parametrization is valid for each series.

The parametrization of the prior distribution of Σ, on the other hand, is based on the desire to let the data speak as much as possible as our prior beliefs on the innovations are rather weak. That is, we seek to avoid having an overtly informative prior. To do so, we follow a common strategy (e.g. Kadiyala & Karlsson (1997), Bańbura et al. (2010)) found in the BVAR literature and parametrize the diagonal elements of Σ as the sample variance of the corresponding series, while the off-diagonal elements are set to zero. Our chosen value for ν is close to the theoretical minimum for a well-specified inverse Wishart distribution, $K + 2$.

An argument for why these are a reasonable choice is presented by Schuurman et al. (2016). The mean of a random variable Σ following the inverse Wishart distribution with parameters L and ν is given by

$$E\left[\Sigma\right] = \frac{L}{\nu - K - 1} \tag{17}$$

while the variances of the diagonal elements Σ_{ii} of Σ are

$$\text{Var}\left(\Sigma_{ii}\right) = \frac{2L_{ii}^2}{\left(\nu - K - 1\right)^2\left(\nu - K - 3\right)} \tag{18}$$

where the L_{ii} are the diagonal elements of L. Thus the informativeness of the prior is increasing in the degrees of freedom ν and decreasing in the L_{ii}. On the other hand, the size of the elements of L also affects the mean as can be seen from (17). If L is set to e.g. I_K – a seemingly uninformative prior – the prior may come to dominate the data if the variance of the data is very small in comparison, which is exactly the case here, as seen above in Table 3.4.1. Schuurman et al. (2016) suggest that in such a situation L is specified as the maximum likelihood estimate of the variance of the data, and that ν is chosen to be as small as possible, both to minimize informativeness and to reduce the effect of using the same data twice.[3]

3.4.3 Calibration and implementation

We have so far ignored the estimation of the vector ω. This is because we choose to instead calibrate this term. Our calibration of the ω_j^2 is based on assuming that there is a constant relationship between signal and noise in the model. In technical terms, we assume that there is a constant relationship between the two

[3]This dual use of data generates a false semblance of certainty in the estimates.

unconditional variance terms:

$$\frac{\omega_j^2/(1-\rho_j^2)}{\sigma_j^2} = R \tag{19}$$

where R is the calibrated signal-to-noise ratio and σ_j^2 is the jth diagonal element of Σ.

Thus our modified Gibbs sampling procedure cycles for M iterations, on each iteration first drawing sequentially from (14), (15) and (16), and then computing each ω_j^2 with the aid of (19) given the most recent draws for the other parameters. This process is initialized with a "zeroth" draw of the parameters ρ and Σ; the values of this draw are all set equal to their prior means, while the ω_j^2 are constructed using (19) given the other initial values. Our chosen calibration for R is 1. Furthermore, we estimate the UCM for each residual separately in sequence, as there is no cross-equation dependence, given the diagonality of Σ and Ω and the structure of (4). Finally, our chosen M is 5000, and we discard the first 500 draws as burn-in to ensure that our estimates are unaffected by the initializing draw.

3.5 Estimation results

Tables 3.5.1, 3.5.2 and 3.5.3 present the results of the estimation exercise for the estimates of the persistence parameter ρ_j, the variances of the transitory component σ_j^2 and the variances of the persistent component ω_j^2, respectively.

As can be seen from table 3.5.1, the estimated persistences ρ_j are quite often clearly different from zero, and occasionally as large as 0.8, indicating strong predictability in these cases. Comparing the full set of estimates to table 3.4.1, it can also be seen that there is a connection between the degree of misspecification hinted at 3.4.1 and the size of the estimates in 3.5.1. Notably the estimated persistences of the residuals themselves in 3.4.1 are somewhat smaller than those seen here in 3.5.1, which is due to two facts: the effect of the prior and disentangling the persistent and transitory components.

Based on table 3.5.1, it can be concluded though that misspecification is not a universal feature of FR-BDF: many equations have very low estimated persistences, but it should be noted though that there are also many cases where the opposite is true. Notable cases where strong misspecification appears to be present are for example the equations for the volume of non-energy imports, the price of public sector consumption and the minimum wage.

Looking at the two right-most columns of table 3.5.1, it can be seen that while there is considerable variation in the precision of the posterior mean – e.g. the 5th and 95th percentiles of the draws of the persistences corresponding to the residuals of the equation for the price of public consumption are closer to the estimated mean than e.g. those of the residuals of the volume of exports. It is also typical, particularly in cases where the absolute value of the posterior mean is large, that zero is not within the range of the 5th to 95th percentile, providing further evidence for the posterior belief of misspecification in these equations.

In tables 3.5.2 and 3.5.3 we present the estimated variances of the transitory and persistent components. As is to be expected, the size of the transitory variances σ_j^2 is linked to the variance estimates seen for the residuals themselves in table 3.4.1. Furthermore, notice that the persistent variances ω_j^2 are estimated to be small – compared to the corresponding σ_j^2 – when the estimated persistence is high. This, however, is a simple consequence of the assumption made for the signal-noise relationship between the two variance terms.

Table 3.5.1: Posterior parameters: persistence

Residual	**Posterior mean**	**5%**	**95%**
Demand volumes			
Household consumption	0.35	0.04	0.61
Business investment	0.33	0.05	0.58
Household investment	0.42	0.16	0.62
Exports	-0.16	-0.4	0.11
Non-energy imports	0.81	0.61	0.92
Energy imports	0.17	-0.05	0.38
Deflators			
Value added	-0.06	-0.39	0.29
Public consumption	0.81	0.71	0.87
Household consumption	-0.06	-0.36	0.54
Business investment	-0.43	-0.7	-0.08
Household investment	0.69	0.53	0.81
Exports	-0.52	-0.76	-0.2
Non-energy imports	-0.09	-0.37	0.19
Energy imports	0.23	-0.03	0.48
Labour market			
Public sector wage	0.53	0.3	0.7
Private sector wage	-0.17	-0.41	0.1
Minimum wage	-0.15	-0.38	0.1
Public sector employment	0.83	0.71	0.91
Private sector employment	0.14	-0.16	0.42
Financial market			
CAC40 index	0.36	0.11	0.56
New household loans for housing	0.64	0.51	0.75
Price of housing	-0.11	-0.44	0.21
Repayments of housing credit	0.75	0.67	0.81
Spread, BBB bonds	0.32	0.08	0.52
Spread, cost of equity	0.28	0.05	0.48

Table 3.5.2: Posterior parameters: transitory variance

Residual	Posterior mean	5%	95%
Demand volumes			
Household consumption	$5.6 * 10^{-6}$	$3.2 * 10^{-6}$	$8.4 * 10^{-6}$
Business investment	$6.3 * 10^{-5}$	$3.8 * 10^{-5}$	$9.6 * 10^{-5}$
Household investment	$1.2 * 10^{-5}$	$7.8 * 10^{-6}$	$1.9 * 10^{-5}$
Exports	$4.7 * 10^{-5}$	$2.8 * 10^{-5}$	$7.4 * 10^{-5}$
Non-energy imports	$9.3 * 10^{-5}$	$6.2 * 10^{-5}$	0.0001
Energy imports	0.0009	0.0006	0.0015
Deflators			
Value added	$3.9 * 10^{-6}$	$2.2 * 10^{-6}$	$5.8 * 10^{-6}$
Public consumption	$9.7 * 10^{-7}$	$6.4 * 10^{-7}$	$1.4 * 10^{-6}$
Household consumption	$9 * 10^{-7}$	$4.9 * 10^{-7}$	$1.4 * 10^{-6}$
Business investment	$3 * 10^{-6}$	$1.9 * 10^{-6}$	$4.4 * 10^{-6}$
Household investment	$1.6 * 10^{-5}$	$1.0 * 10^{-5}$	$2.4 * 10^{-5}$
Exports	$3.8 * 10^{-6}$	$2.3 * 10^{-6}$	$5.8 * 10^{-6}$
Non-energy imports	$3.8 * 10^{-6}$	$2.3 * 10^{-6}$	$5.8 * 10^{-6}$
Energy imports	0.0003	0.0002	0.0004
Labour market			
Public sector wage	$7.4 * 10^{-6}$	$4.8 * 10^{-6}$	$1.2 * 10^{-5}$
Private sector wage	$5 * 10^{-6}$	$2.9 * 10^{-6}$	$7.8 * 10^{-6}$
Minimum wage	$3 * 10^{-5}$	$1.7 * 10^{-5}$	$4.6 * 10^{-5}$
Public sector employment	$3.4 * 10^{-6}$	$2.3 * 10^{-6}$	$4.8 * 10^{-6}$
Private sector employment	$4.9 * 10^{-7}$	$2.5 * 10^{-7}$	$8.1 * 10^{-7}$
Financial market			
CAC40 index	0.001	0.0007	0.002
New household loans for housing	0.005	0.003	0.007
Price of housing	$1.4 * 10^{-5}$	$8.4 * 10^{-6}$	$2.1 * 10^{-5}$
Repayments of housing credit	0.02	0.01	0.026
Spread, BBB bonds	$4.9 * 10^{-7}$	$3.6 * 10^{-7}$	$8.1 * 10^{-7}$
Spread, cost of equity	$1.4 * 10^{-6}$	$8.1 * 10^{-7}$	$1.9 * 10^{-6}$

Table 3.5.3: Posterior parameters: persistent variance

Residual	Posterior mean	5%	95%
Demand volumes			
Household consumption	$4.7 * 10^{-6}$	$2.6 * 10^{-6}$	$7.3 * 10^{-6}$
Business investment	$5.4 * 10^{-5}$	$3.3 * 10^{-5}$	$8.7 * 10^{-5}$
Household investment	$9.6 * 10^{-6}$	$5.8 * 10^{-6}$	$1.6 * 10^{-5}$
Exports	$4.4 * 10^{-5}$	$2.6 * 10^{-5}$	$7.1 * 10^{-5}$
Non-energy imports	$3.1 * 10^{-5}$	$1.2 * 10^{-5}$	$6.2 * 10^{-5}$
Energy imports	0.0009	0.0005	0.0013
Deflators			
Value added	$3.7 * 10^{-6}$	$2.5 * 10^{-6}$	$5.8 * 10^{-6}$
Public consumption	$3.6 * 10^{-7}$	$1.6 * 10^{-7}$	$6.4 * 10^{-7}$
Household consumption	$8.1 * 10^{-7}$	$4.9 * 10^{-7}$	$1.4 * 10^{-6}$
Business investment	$2.3 * 10^{-6}$	$1.2 * 10^{-6}$	$3.6 * 10^{-6}$
Household investment	$8.5 * 10^{-6}$	$4.4 * 10^{-6}$	$1.4 * 10^{-5}$
Exports	$2.6 * 10^{-6}$	$1.4 * 10^{-6}$	$4.4 * 10^{-6}$
Non-energy imports	$3.7 * 10^{-6}$	$2.3 * 10^{-6}$	$5.8 * 10^{-6}$
Energy imports	0.00026	0.00016	0.0004
Labour market			
Public sector wage	$5.2 * 10^{-6}$	$2.8 * 10^{-6}$	$1 * 10^{-5}$
Private sector wage	$4.8 * 10^{-6}$	$2.8 * 10^{-6}$	$7.3 * 10^{-6}$
Minimum wage	$2.8 * 10^{-5}$	$1.7 * 10^{-5}$	$4.4 * 10^{-5}$
Public sector employment	$1.1 * 10^{-6}$	$4.9 * 10^{-7}$	$1.9 * 10^{-6}$
Private sector employment	$4.9 * 10^{-7}$	$2.5 * 10^{-7}$	$8.1 * 10^{-7}$
Financial market			
CAC40 index	0.001	0.0007	0.0016
New household loans for housing	0.003	0.0017	0.004
Price of housing	$1.3 * 10^{-5}$	$7.8 * 10^{-6}$	$2 * 10^{-5}$
Repayments of housing credit	0.008	0.005	0.011
Spread, BBB bonds	$4.9 * 10^{-7}$	$2.5 * 10^{-7}$	$8.1 * 10^{-7}$
Spread, cost of equity	$1.2 * 10^{-6}$	$8.1 * 10^{-7}$	$1.9 * 10^{-6}$

4 Stochastic simulations around the December 2022 projections of Banque de France

In this section, we present a sequence of applications centered on the December 2022 projections of Banque de France.[4]. Hence, all results presented here are conditional on the information available at that time. In particular, it does not take into account the further ease of tensions on energy, which occurred afterward. In what follows, first, we present the implementation of the exercise and make necessary adjustments to FR-BDF in order to take into account uncertainty coming from recent energy crisis. We then explore the confidence bands of stochastic simulations. And finally, we compute the probability of a technical recession over the forecast horizon.

4.1 Implementation

Forecasting exercises, implemented quarterly at Banque de France in coordination with the Eurosystem, are fundamentally conditional simulations of macroeconomic models such as FR-BDF over the forecasting horizon of 3 to 4 years. The conditioning information set consists of two main components: common assumptions of the paths of certain quantities[5] provided by the ECB, and expert judgment as determined by the staff of the institutions implementing the country-level forecasts.

In our analysis we take for the most part this conditioning information set as a given, and implement stochastic simulations of FR-BDF around it. More concretely, we run stochastic simulations of FR-BDF with an unobserved component representation of the residuals around the December 2022 projection exercise. To do so, first, we invert the Banque de France projection exercise, i.e. we compute the paths of residuals needed to replicate the forecast. Second, using the estimated UCM coefficients, we filter the unobserved (persistent) components of the UCM model, – c_t from equation 4, – using recovered above values of corresponding transitory components, – e_t from equation 3, – considered as being a source of uncertainty and listed in Table 3.5.2. Third, we recover what we call "inverted" values for the innovation terms of UCM stochastic components – ϵ_t and η_t from equations 3 and 4, – needed to replicate the forecast using the model augmented with unobserved component. Finally, we draw 10 000 values for UCM innovation terms centered around those "inverted" values obtained in step 3 from the standard normal distribution and simulate the model.

We deviate from the full set of forecasting assumptions by extending our baseline stochastic framework with stochastic energy prices. That is, we relax the exogeneity of the price of oil and natural gas, further assuming that they are subject to stochastic shocks. This extension is described in subsection 4.2.

[4]More details about the December 2022 Banque de France forecast is available here: https://publications.banque-france.fr/sites/default/files/macroeconomic-projections_december-2022.pdf

[5]E.g. the short- and long-run interest rates, foreign variables, the price of oil and natural gas, etc.

Note that we do not use exactly the same model as in the projection exercise as we cannot condition the stochastic simulation exercise on public finance and HICP variables generated using two other Banque de France models – MAPU[6] and MAPI[7] – as is done during the standard forecasting process. This is due to practical concerns related to the operation of these external models: for each stochastic draw, a corresponding simulation of the two other models would be needed, which is not feasible due to lack of automation in the use of these models. Instead, we build a pseudo-conditional model: a version of the FR-BDF model in which equations corresponding to the ECB assumptions are exogenized in order to replicate the forecasting environment as closely as possible, but where we retain the endogeneity of public finance and HICP variables. In the sequel we will refer to this version of the model as pseudo-conditional UCM-augmented FR-BDF, – PC-UCM FR-BDF.

4.2 Accounting for energy price uncertainty

Motivated by the recent economic upheaval in energy markets following the Russian invasion of Ukraine, for this exercise an ad hoc extension was added to capture the new, additional uncertainty in the price of oil and natural gas. In the case of the price of natural gas this extension also involves adding an equation to the model describing its dynamics, as in the baseline version of FR-BDF it is assumed fully exogenous. All estimates in this section are based on the sample 1996Q1–2019Q4 in order to avoid issues related to the COVID-19 pandemic and the large increase in the volatility of energy prices after the invasion.[8]

For the price of oil in euros the baseline equation – which relates the price in euros to the exogenous price in dollars and the exchange rate – is replaced with

$$brenteuro = \alpha_{brent} + z_{\pi} * t + \nu_{brent} \tag{20}$$

where $\alpha_{brent} = 3.51$ is estimated and $z_{\pi} = 0.005$ is the common trend for prices in FR-BDF. $brenteuro$ denotes the log price of a barrel of Brent crude oil in euros and ν_{brent} is a shock.

The treatment of the log price of natural gas is symmetric, relating the unit price (gas), denominated in euros, of natural gas to a trend and a shock ν_{gas}:

$$gas = \alpha_{gas} + z_{\pi} * t + \nu_{gas} \tag{21}$$

[6]An aggregated public finance model.

[7]Reference model for the analysis and projection of inflation.

[8]This choice might underestimate the uncertainty related to energy prices. However, given the lack of hindsight, this conservative choice has the advantage of avoiding the potential structural change of the variance of energy prices, which would be very difficult to assess precisely. We leave this issue for further research once enough data will be available for carrying out this assessment.

where the estimate of α_{gas} is 3.36.

Both ν_{brent} and ν_{gas}, the stochastic components of these processes, are assumed to follow

$$\nu_{brent,t} = \rho_{brent}\nu_{brent,t-1} + \eta_{brent,t} \tag{22}$$

$$\nu_{gas,t} = \rho_{gas}\nu_{gas,t-1} + \eta_{gas,t} \tag{23}$$

with $\eta_{brent,t} \sim N\left(0, \sigma^2_{brent}\right)$ and $\eta_{gas,t} \sim N\left(0, \sigma^2_{gas}\right)$. We estimate ρ_{brent}, σ_{brent}, ρ_{gas} and σ_{gas} as 0.95, 0.14, 0.89 and 0.15, respectively.

4.3 Results

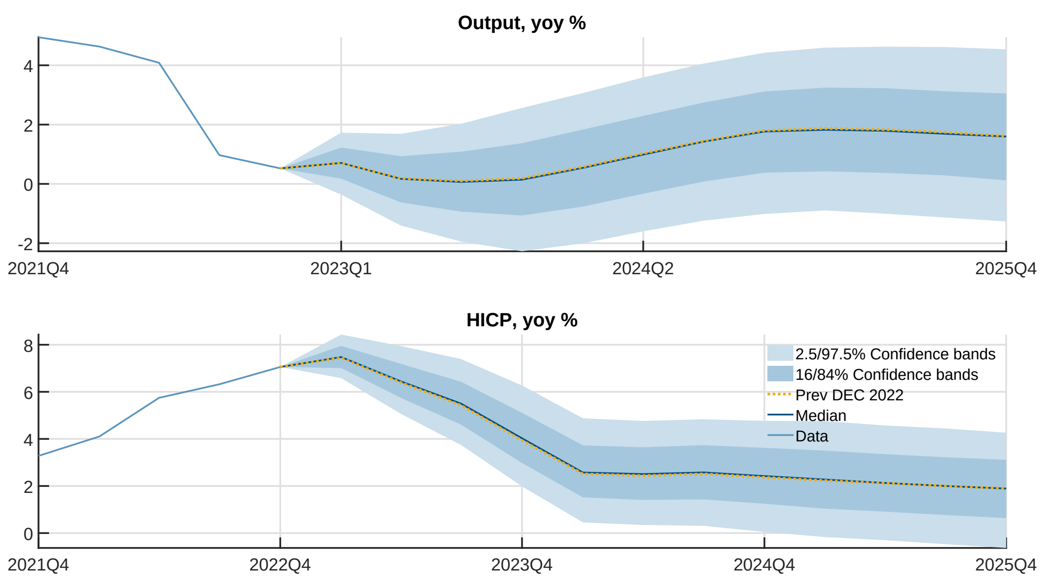

Figure 4.3.1: Stochastic simulation of FR-BDF with an unobserved component representation of the residuals around December 2022 projection exercise.

The result of this stochastic exercise for pseudo-conditional UCM-augmented FR-BDF of December 2022 forecast is presented in Figure 4.3.1. To get a flavor on how big the obtained confidence intervals for the FR-BDF model are, we compare them with those of an alternative model, Structural BVAR, see A.1. To ease illustration of further discussion, table 4.3.1 presents the point forecasts for annual output growth and inflation together with the 50%, 68% and 95% confidence intervals computed in this stochastic exercise.[9] The width of the ranges presented demonstrates that there is a large uncertainty associated with the forecast,

[9]Notice that there is a slight difference in the methodologies applied to produce confidence intervals for the associated forecast note and those applied here. More specifically the volatilities σ^2_{brent} and σ^2_{gas} were in the note estimated to be somewhat larger using a sample that included the post-COVID period. The ensuing bands are however equal up to a rounding.

as even the 50% range for output growth has a width of more than 1 percentage point in 2023. As is to be expected, the width of the intervals grows over time, as additional stochastic innovations accumulate in the simulations. For example, in the case of inflation the width of the 95% confidence band is 3 pp in 2023, while in 2025 it is 4.2 pp.

Table 4.3.1: The FR-BDF point forecasts and associated confidence intervals for annual inflation and output growth in 2023-2025

	Point forecast	**50% conf. interval**	**68% conf. interval**	**95% conf. interval**
2023, inflation	6.0%	5.6% ; 6.6%	5.3% ; 6.8%	4.7% ; 7.7%
2024, inflation	2.5%	1.9% ; 3.2%	1.6% ; 3.5%	0.6% ; 4.6%
2025, inflation	2.1%	1.3% ; 2.8%	1% ; 3.2%	0% ; 4.2%
2023, output growth	0.3%	-0.3% ; 0.8%	-0.5% ; 1.1%	-1.3% ; 1.8%
2024, output growth	1.2%	0.4% ; 1.9%	0% ; 2.3%	-1.2% ; 3.4%
2025, output growth	1.8%	0.9% ; 2.7%	0.5% ; 3.1%	-0.7% ; 4.4%

In Figure 4.3.2 we look at this phenomenon from a different perspective, presenting histograms of annual output growth (left panel) and inflation (right panel) in 2023, 2024 and 2025. In 2023 the simulated output growth rate has a mean very close to zero, while the distributions in 2024 and 2025 indicate an economic recovery with mean growth approaching 2%. In the case of inflation it can be seen that its distribution is quite narrow in 2023 with a mean of 6%, while the distributions in 2024 and 2025 are much wider, with means closer to 2%, again demonstrating the phenomenon of additional uncertainty accumulating.

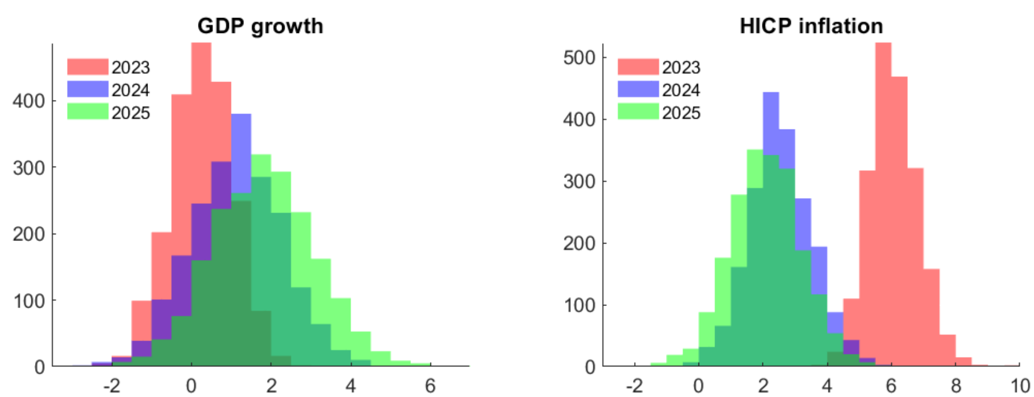

Figure 4.3.2: Histograms of annual output growth (left panel) and inflation (right panel), 2023-2025.

4.4 What is the probability of a technical recession over the forecast horizon?

Figure 4.4.1 presents a plot of the probability of technical recession in France over the period 2023Q1 – 2025Q4. By technical recession we mean the event that GDP growth is negative in two consecutive quarters.

The probability is at its peak at 42% in 2023Q2.[10] It then bottoms at the end of the period at 16%. In 2025, there is a notable increase in the probability in comparison to 2024. This is related to the decrease in forecasted quarterly growth in 2025, which can be seen in Table 4.3.1. Finally, the probability of a technical recession occurring over the period, i.e. that at a technical recession occurs at least once, is 75%.

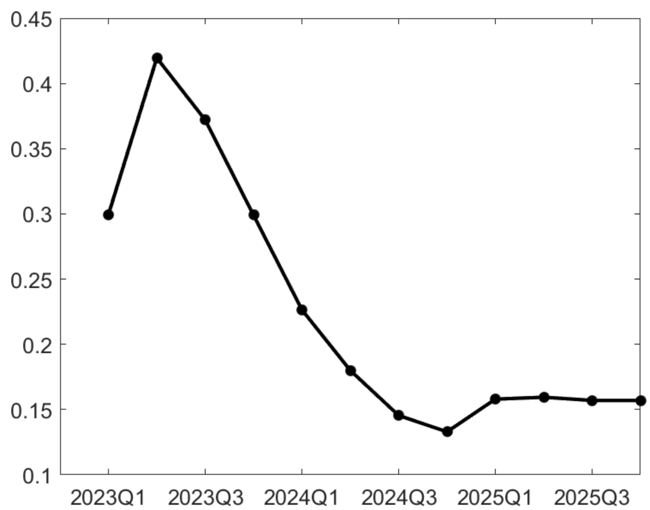

Figure 4.4.1: Quarterly probability of technical recession based on December 2022 macroeconomic projections of the Banque de France.

5 Conclusion

In order to enable applications that describe probabilities of events and risk in general for large semi-structural FR-BDF model used for forecast and policy work, we have chosen to explore its misspecification intrinsic to this type of model by using an unobserved component model. This tool allows at the same time studying uncertainty due to model misspecification as well as setting environment for structural shocks. We find that the source of stochastic volatility – residuals with significant persistence from one period to the next – are mainly present in the equations for exchange rates, the public sector, employment and the volume of non-energy imports.

We demonstrate the usefulness of this tool by providing confidence bands around a conditional forecast and estimating the probability of technical recession in the context of the Banque de France December 2022 projection exercise. As further research, we will explore the forecasting power of the nonjudgmental forecasts, i.e. of point forecasts of FR-BDF augmented with an UCM for specific residuals. We will also

[10] Note again, that this assessment is based on the information available at the time of the December 2022 forecast and that energy tensions have eased afterward, which would modify our assessment of the risk of recession, if we used more recent data.

explore uncertainty due to cost-push shocks by focusing only on residuals with unobserved components from price and wage equations. Finally, bootstrapping remains a valid alternative to our approach, and merits investigation as an alternative or even complementary method of simulating FR-BDF.

References

Angelini, E., Bokan, N., Christoffel, K. P., Ciccarelli, M., & Zimic, S. (2019a). *"Introducing ECB-BASE: The blueprint of the new ECB semi-structural model for the euro area.* Working Paper 2315, European Central Bank.

Angelini, E., Lalik, M., Lenza, M., & Paredes, J. (2019b). Mind the gap: A multi-country bvar benchmark for the eurosystem projections. *International Journal of Forecasting*, 35(4), 1658–1668.

Bańbura, M., Giannone, D., & Reichlin, L. (2010). Large bayesian vector auto regressions. *Journal of applied Econometrics*, 25(1), 71–92.

Bourgeois, A. & Favetto, B. (2022). *Construction d'intervalles de confiance et relecture du passe avec le modele Mesange.* Technical report.

Carriero, A., Clark, T. E., & Marcellino, M. (2016). *Large Vector Autoregressions with Stochastic Volatility and Flexible Priors.* Working Papers (Old Series) 1617, Federal Reserve Bank of Cleveland.

González-Astudillo, M. & Vilan, D. (2019). *A New Procedure for Generating the Stochastic Simulations in FRB/US.* Technical report, Board of Governors of the Federal Reserve System (US).

Kadiyala, K. R. & Karlsson, S. (1997). Numerical methods for estimation and inference in bayesian var-models. *Journal of Applied Econometrics*, 12(2), 99–132.

Lemoine, M., Turunen, H., Chahad, M., Lepetit, A., Zhutova, A., Aldama, P., Clerc, P., & Laffargue, J.-P. (2019). *The FR-BDF Model and an Assessment of Monetary Policy Transmission in France.* Working papers 736, Banque de France.

Schuurman, N. K., Grasman, R. P. P. P., & Hamaker, E. L. (2016). A comparison of inverse-wishart prior specifications for covariance matrices in multilevel autoregressive models. *Multivariate Behavioral Research*, 51(2-3), 185–206.

A Appendix

A.1 A reference BVAR model

In what follows, we first describe our choice of a reference model to compare our results with, and then study its forecast characteristics by looking at its out-of-sample performance. Finally, we compare its confidence intervals to those of unconditional UCM augmented FR-BDF.

We consider a Bayesian VAR consisting of two blocks: France and the rest of the euro area (REA). It includes 20 variables – (i) French and REA variables: real GDP, real total investment, consumer prices, GDP deflator, long-run and short-run interest rates, wages, foreign demand; (ii) US variables: short-term rate and real GDP, (iii) oil price and dollar/euro exchange rate. The order of the VAR is 5. We use annual data from 1999Q1 to 2017Q4 which corresponds to the estimation period for most of FR-BDF equations.

Priors. We closely follow Angelini et al. (2019b). The model is estimated using a normal-Wishart prior that implies an inverse-Wishart prior on the covariance matrix of the residuals and Minnesota prior on the beta-coefficients. The degree of freedom in the covariance is set automatically to N+2 as in Angelini et al. (2019b), a scaling matrix in the covariance is estimated with Inverse-Gamma specified for annual data.

Cross-country linkages. To check the importance of the cross-country linkages, we re-estimate the model with different settings for priors: the coefficients that link the variables of one area to another are forced to tend to zero. The difficulty of this procedure is the slow-down of the computation process and the inexistence of the marginal likelihood in a closed-form. The first issue is addressed by estimating the model equation by equation. We implement this by using the algorithm of Carriero et al. (2016), which is based on a simple triangularization of the VAR model. We address the second problem by setting the hyperparameters to their modes.

To judge which model fits better the data – with or without linkages – we compute a ratio of average (over the sample) continuous ranked probability scores based on 10 000 draws. The evaluation is based on the one-year-ahead forecast for the sample 2012Q1-2018Q1, i.e. 21 forecasts for 4 quarters each. A ratio of less than one indicates that the forecast of the baseline scenario fits better the data for a certain variable. From Figure A.1 we can see that only the forecast of total investment of the rest of the euro area would have been better without considering linkages between the European countries.[A.1] Hence, in what follows we rely only on the BVAR with cross-country linkages.

Out-of-sample performance: conditional BVAR vs. Banque de France forecast. Figure A.2 compares the Banque de France forecast (black line with stars) with the one of BVAR for two year-on-year growth rates: French GDP and consumer price deflator in order to evaluate the relative accuracy of the

[A.1] Note that we account for the interactions with the US variables in both specifications.

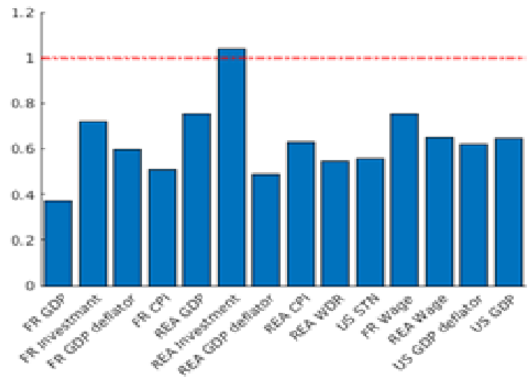

Figure A.1: Ratio of continuous ranked probability scores.

BVAR in conditional forecasts of headline variables. We present two types of forecast for BVAR model: distribution of conditional exercise (shaded area) which uses the real-time assumptions provided by the ECB and a median response of unconditional exercise (blue dashed line). The assumptions concern the following list of variables: world demand addressed to France, French long-run interest rate, oil price, dollar-euro exchange rate and ECB short-run interest rate. Note that the Banque de France forecast is computed using the Mascotte model – the previous forecasting model applied at Banque de France – which was estimated (1) on a different sample, (2) using revised data.

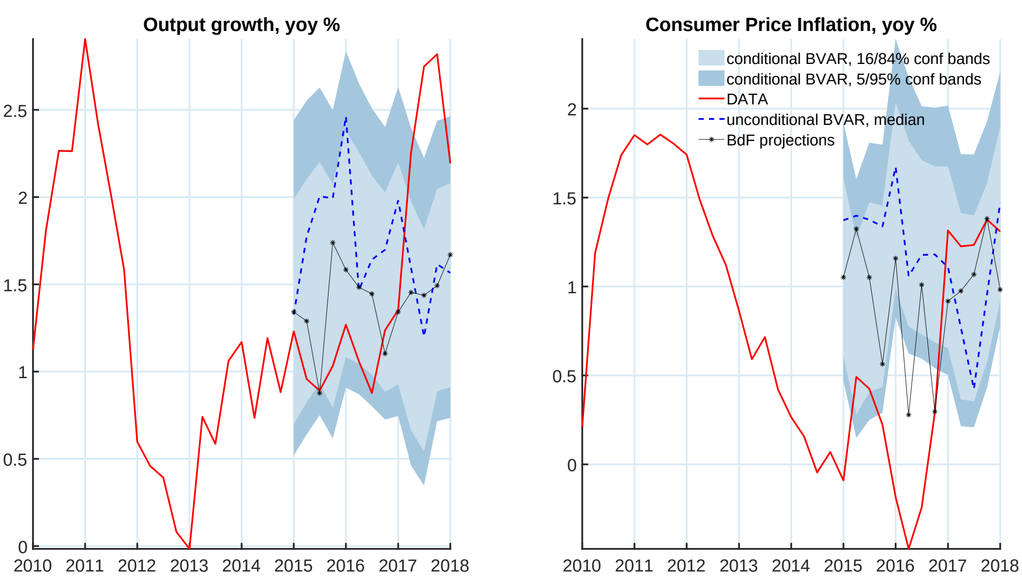

Figure A.2: BVAR conditional forecast with real-time assumptions and Banque de France projections. Note: black line with stars – Banque de France forecast; shaded area – distribution of conditional exercise; blue dashed line – median response of unconditional exercise.

The results presented in this section are based on the rolling windows of one-year-ahead forecast for the sample 2014Q1-2017Q4 with the number of increments between successive rolling windows equal to 1 quarter. This means that at the end we have 16 forecasts for 4 quarters each. To be more specific, the procedure we apply, step by step: (i) we estimate the model on a sample from 1999Q1 to 2014Q1, (ii) we simulate the model for 4 quarters (2014Q2-2015Q1), (iii) we add a quarter to estimation sample (1999Q1-2014Q2) and re-estimate the model, (iv) we simulate from 2014Q3 to 2015Q2, and so on.[A.2] In each projection exercise, we replace the historical sample of the assumptions by revised data in order to be consistent with the real-time assumptions we use for conditional forecast.[A.3] Revisions concern mainly world demand addressed to France, while the changes in other variables are not very notable.

The one-year ahead conditional BVAR density forecasts for year-on-year French GDP growth rate (shaded area) correspond rather well to the forecasted values, as the Banque de France projections lie close to the median response of conditional BVAR. It is important to note that the actual observed data is within the 90% confidence bands except in 2017Q2 and 2017Q3. It is more difficult to draw conclusions with respect to the density forecasts for the growth of French consumer prices. The conditional BVAR produces predictions that are closer to the Banque de France projections than the unconditional BVAR, but all of them are very far from the observations, especially in 2015-2017, where realized inflation was much lower than all of the forecasts. A similar figure but with the BVAR conditional distribution computed using actual observations (see Figure A.3 in the appendix) suggests that errors come for a large part from the set of ECB hypotheses. Here, the conditional median response shows good forecasting accuracy for output growth, and as for inflation, we see that the data lies within the plotted confidence bands. In the end, we conclude that the BVAR has the ability to mimic the path of the official projections and appears to be a valid alternative forecast model.

In Figure A.4 we compare the width of confidence intervals between the unconditional UCM FR-BDF model (UCM FR-BDF) and the alternative BVAR model. We present the results for year-on-year growth rates of French output and CPI; the plots represent the difference (or gap) between the upper and lower bounds of the 95% and 68% confidence intervals.

First, we notice that our uncertainty bands obtained with this framework are comparable to (and even outperform) those implied by the alternative BVAR model. To be more precise, in the case of the year-on-year output growth, we obtain a 5 percentage point gap with UCM FR-BDF (evaluated as a difference between the 2.5th and 97.5th quantiles of the distribution) against a 5.7 percentage point gap obtained with the unconditional BVAR model. To put these numbers into context, based on the sample spanning 1999Q1-

[A.2] The historical data for estimation is corresponds to a vintage from 2018, which means that the interest rate from this database does not necessarily correspond to the historical data from any other vintage of real-time assumptions. In order to disentangle the out-of-sample performance with the questions related to the data revision, we replace the data used for estimation with the one obtained from the early databases without re-estimation.

[A.3] Note that the Banque de France forecast is based on the same real-time assumption we used to construct the BVAR conditional forecast.

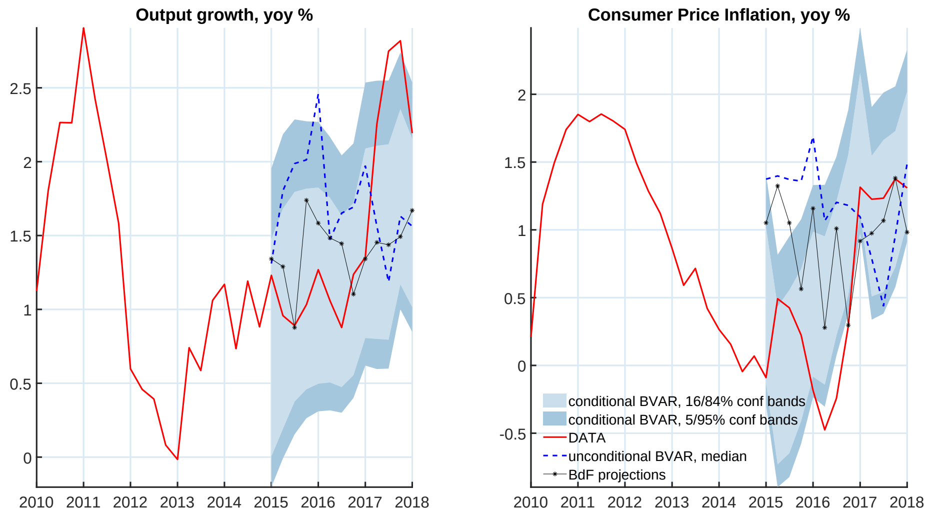

Figure A.3: BVAR conditional forecast with actual assumptions and Banque de France projections. Note: black line with stars – Banque de France forecast; shaded area – distribution of conditional exercise; blue dashed line – median response of unconditional exercise.

2019Q4, two standard deviations of output growth roughly correspond to 5%. Regarding the year-on-year CPI growth rate, the corresponding differences are 4.6 and 4.8 percentage points for UCM FR-BDF and BVAR, respectively. It's worth noting that two standard deviations of this variable equate to approximately 3.6% in the sample from 1999Q1-2019Q4.

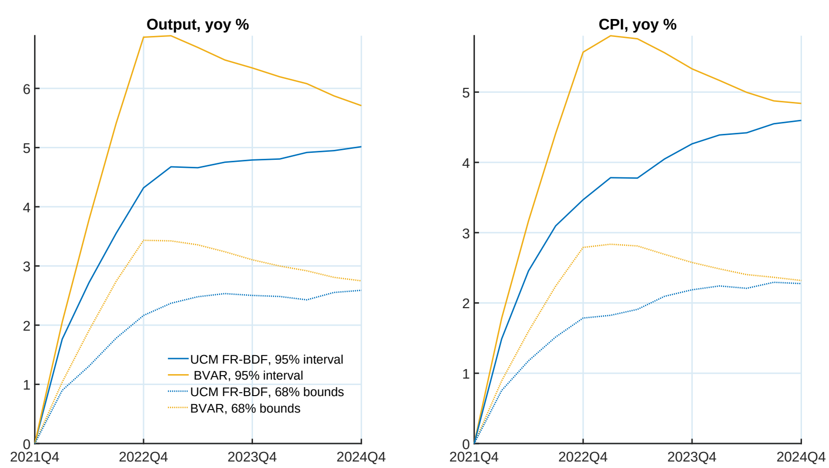

Figure A.4: Confidence intervals: solid lines correspond to the 95% confidence bands, while dashed lines correspond to 68% confidence bands.