WORKING paper

BANQUE DE FRANCE
EUROSYSTÈME

Quantifying Uncertainty in France's Debt Trajectory: A VAR-Based Analysis

Kéa Baret, Frédérique Bec & Marion Cochard[1]

November 2025, WP #1019

Abstract

We propose a simple, simulation-based framework for stochastic debt-sustainability analysis. Estimating a parsimonious vector autoregression (frequentist and Bayesian) on quarterly French data (1990:Q1–2023:Q4) for the debt's key drivers, we generate predictive fan charts and probability statements for debt-to-GDP outcomes. Median VAR projections are close to a hypothetical deterministic baseline derived from the deterministic debt sustainability analysis framework. Assuming this illustrative central scenario, historical relationships estimated by our VAR models imply a corresponding confidence band around the debt trajectory. The BVAR yields slightly wider cones and lower tail probabilities than the frequentist VAR, with cone widths between those reported by the European Commission and the ECB. Our analysis, which does not reflect the most recent developments in public finance, suggests that an ambitious fiscal consolidation effort would be required to materially enhance the prospects of stabilizing the debt-to-GDP ratio over the medium term.

Keywords: Debt Sustainability, Stochastic Analysis, VAR Model, Bayesian forecasting, Density Forecasts.

JEL classification: C3, E6, H6

[1] Baret: Banque de France. Bec : CY Cergy Paris University, THEMA (UMR 8184), CNRS. Cochard: Banque de France. This paper reflects the opinions of the authors and does not necessarily express the views of the Banque de France. We thank Marie Bessec, Vladimir Borgy, Othman Bouabdallah, Javier Ferri, Olivier Garnier, Christophe Hurlin, Caroline Jardet, Yannick Kalantzis, Hervé Le Bihan and Heino Bohn Nielsen, as well as participants to the 73rd Congress of the French Economic Association (AFSE), the 7th Quantitative Finance and Financial Econometrics conference and Banque de France seminar for fruitful comments and discussions.. All remaining errors are ours.

NON-TECHNICAL SUMMARY

This paper develops a simple, transparent method to quantify the uncertainty around forecasts of public debt. Rather than producing a single "best-guess" debt path, the approach generates a range of plausible futures and associates probabilities with different outcomes. This probabilistic view is intended to give policymakers and analysts a clearer sense of fiscal risk and the likelihood that particular policies will succeed in stabilizing public debt. Based on 2023 vintage data, this analysis does not reflect recent fiscal developments and is therefore not intended to inform the current public debate on public debt. Its purpose is to improve SDSA methodologies.

Our approach models the joint behavior of the key drivers of the debt-to-GDP ratio — namely, the primary balance (revenues minus expenditures, excluding interest), nominal GDP growth, and short- and long-term nominal interest rates — using a standard Vector Auto-Regression (VAR) estimated on quarterly French data from 1990 through 2023. We implement two variants of this model: a conventional (frequentist) specification and a Bayesian version that incorporates mild prior information. They are used to simulate a large number (10,000) of future paths for these drivers, drawing shocks that reflect their historical volatility and interdependence. Combining these simulated paths with the familiar debt accounting identity produces a "fan chart" for the debt-to-GDP ratio: a visual and quantitative representation of the range of outcomes and their probabilities.

The median debt trajectories produced by both model variants closely track the deterministic baseline projection for 2024–2028, which follows the deterministic debt sustainability analysis framework of Bouabdallah et al. (2017). Crucially, that baseline lies comfortably within the middle of the distribution produced by our simulations, suggesting the baseline is a plausible central scenario given known historical dynamics. The Bayesian model yields slightly wider uncertainty bands than the frequentist model for the horizon we study; for example, the Bayesian 10–90 percent fan cone for 2028 is modestly larger than its frequentist counterpart. The two model variants also assign high probabilities that the 2028 debt ratio will remain above its 2023 level, although the Bayesian model produces a somewhat lower probability than the frequentist model. Overall, the magnitude of our uncertainty bands falls between the measures reported by the European Commission and the European Central Bank, lending further credibility to the quantitative scale of the results.

Figure 1. Primary balance ratio fan charts and debt-stabilizing primary balance ratio (red line)

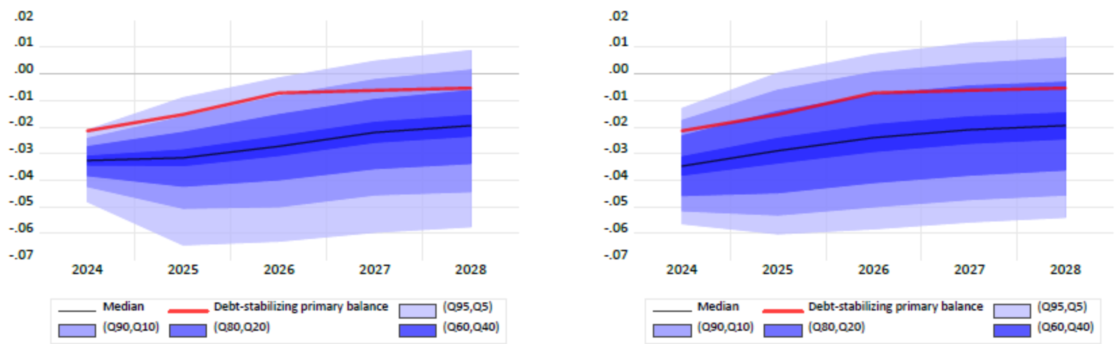

(a) VAR(4)

(b) BVAR(1)

The probabilistic framework makes it straightforward to ask policy-relevant questions such as: what is the probability that a given fiscal path will stabilize debt within a target horizon? Applied to France (see Figure 1) our simulations indicate that an earlier and larger consolidation than in the baseline would meaningfully improve the odds of bringing debt dynamics onto a stable path.

The proposed method deliberately favors parsimony and transparency over structural complexity: it quantifies risk conditional on historical relationships among the main drivers. An important extension is to embed the analysis within the European Union's new fiscal framework (the Economic Governance Review adopted in February 2024), which will alter both baseline trajectories and the policy metrics used to judge sustainability.

By converting deterministic projections into probabilistic assessments, the approach offers a practical, replicable complement to standard institutional debt projections and helps clarify the degree of policy effort needed to change the odds of stabilizing public debt.

Quantifier l'incertitude sur la trajectoire de la dette française : une analyse VAR

RÉSUMÉ

Nous proposons un cadre simple, fondé sur la simulation, pour l'analyse stochastique de la soutenabilité de la dette. En estimant un VAR parcimonieux (fréquentiste et bayésien) sur des données françaises trimestrielles (1990:T1–2023:T4) pour les principaux déterminants de la dette, nous produisons des fan-charts prédictifs et des mesures de probabilité pour les trajectoires du ratio dette/PIB. Les trajectoires médianes issues des VAR sont proches d'un scénario déterministe hypothétique, dérivé du cadre d'analyse déterministe de la soutenabilité de la dette. Le BVAR génère des cônes légèrement plus larges et des probabilités de dépasser la dernière valeur observée plus faibles que le VAR fréquentiste, les largeurs de cône se situant entre celles publiées par la Commission Européenne et la BCE. Notre analyse, qui ne prend pas en compte les développements budgétaires les plus récents, suggère qu'un effort ambitieux de consolidation budgétaire serait nécessaire pour améliorer de manière significative les perspectives de stabilisation du ratio dette/PIB à moyen terme.

Mots-clés : soutenabilité de la dette, analyse stochastique, modèles VAR, prévisions bayesiennes, distribution des prévisions

---

1 Introduction

The trajectory of public debt is central to macroeconomic policy and political debate. In the euro area—where debt-to-GDP ratios are tightly monitored and constrained—the sovereign-debt crisis of the early 2010s underscored the consequences of adverse debt dynamics. More recently, the COVID-19 pandemic and the war in Ukraine have pushed public indebtedness well above historical norms (see Figure 1), renewing questions about how and when debt paths can be realigned. In a September 12, 2024 address, the President of the European Central Bank urged

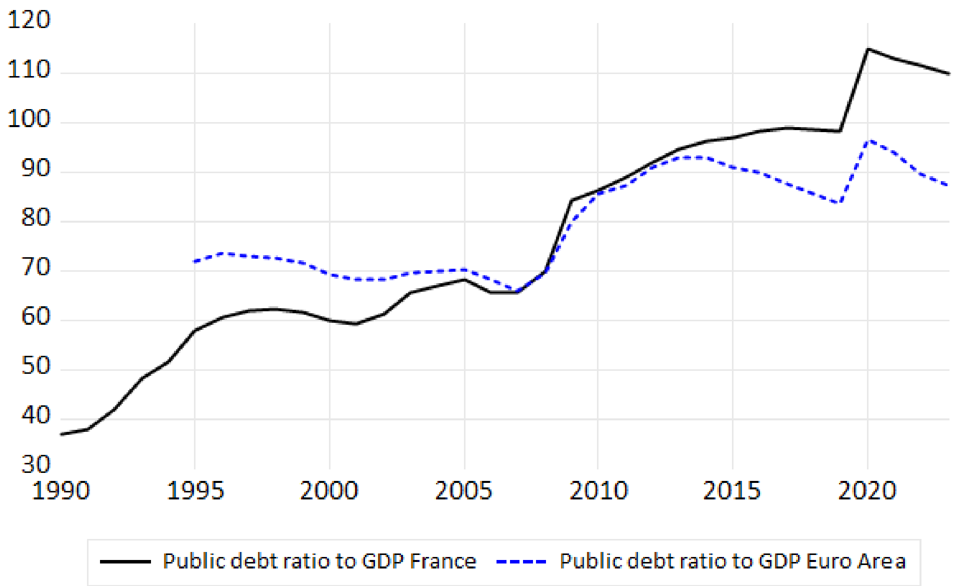

FIGURE 1: *Public debt ratio to GDP in France and Euro Area from 1990 to 2023*

national authorities to adopt prompt budgetary and fiscal measures, highlighting the practical urgency of reliable debt forecasts and transparent measures of their uncertainty.

Forecasting future debt trajectories inevitably requires an assessment of uncertainty. As Blanchard et al. [2021] note, "That debt forecasts and thus debt sustainability assessments are made under substantial uncertainty is obvious." (p.10 therein). This uncertainty reflects both the future evolution of key drivers—interest rates, growth, and primary balances—and the occurrence of shocks and policy responses that alter those drivers. Capturing this uncertainty is essential for informed policy design and credible communication.

This paper addresses stochastic debt sustainability analysis (SDSA), distinguishing it from the conventional deterministic debt sustainability analysis (DDSA) that produces a single baseline path from fixed assumptions. SDSA complements DDSA by generating probabilistic projections that reflect the joint uncertainty in macro-fiscal determinants. Building on recent simulation-based work (e.g., Cuerpo and Ramos [2015]; Cherif and Hasanov [2018]; Benjamin Carton and Fouejieu [2020]; Steel [2021]; Papaoikonomou [2025]), our contribution is twofold. First, we survey existing SDSA approaches and clarify their relative strengths and limitations. Second, we introduce a parsimonious, statistically disciplined procedure—implemented using both frequentist and Bayesian VAR models—to quantify uncertainty around debt-to-GDP projections and to compute policy-relevant probabilities, such as the likelihood that a given primary balance would stabilize the debt ratio. Note that, based on 2023-vintage data, this analysis does not reflect recent fiscal developments and is therefore not intended to inform the current public debate on public debt. Its purpose is to improve SDSA methodology.

The paper is organized as follows. Section 2 presents a compact accounting framework for debt dynamics. Section 3 reviews the state of the art in SDSA. Section 4 describes our empirical strategy and data. Sections 5 and 6 report results from the frequentist and Bayesian VAR implementations, respectively. Section 7 compares our projections with those of major European institutions and illustrates how the proposed approach can assess the probability of attaining a primary balance consistent with debt stabilization. Section 8 concludes.

2 A simple accounting framework

Before turning to the existing and proposed approaches of Stochastic Debt Sustainability Analysis (SDSA), let us first introduce notation through a brief reminder of basic concepts of public debt dynamics from a simple accounting framework. The typical debt accumulation equation, expressed in nominal terms, is given by

$$B_t = (1 + i_t)B_{t-1} - PB_t + DDA_t, \tag{1}$$

where B_t denotes the amount of public debt at time t, i_t the "effective" or average nominal interest rate charged on government debt, PB_t is the primary balance including the charge of interest and DDA_t is the deficit-debt, or stock-flow adjustment which gathers all other factors which affect the debt but are not included in the primary balance (e.g. acquisitions or sales of financial assets, valuation effects, etc.). For simplicity sake, it is assumed that all the debt is issued in domestic currency. When expressed in terms of nominal GDP percentage, Eq. (1)

becomes:

$$\frac{B_t}{Y_t} = (1+i_t)\frac{B_{t-1}}{Y_t} - \frac{PB_t}{Y_t} + \frac{DDA_t}{Y_t},$$

with Y_t the nominal GDP at time t. Substituting Y_t by $(1+g_t)Y_{t-1}$ on the right-hand side, where g_t is the growth rate of the nominal GDP, and using lower cases to denote the variables expressed as a percentage of GDP, the accumulation equation of the debt-to-GDP ratio is obtained as:

$$b_t = \frac{(1+i_t)}{(1+g_t)}b_{t-1} - pb_t + dda_t, \tag{2}$$

which in turn can be rewritten as:

$$\Delta b_t = \frac{(i_t - g_t)}{(1+g_t)}b_{t-1} - pb_t + dda_t. \tag{3}$$

As dda_t basically stems from discretionary decisions or exogenous shocks, the change in the debt-to-GDP ratio (Δb_t) is mainly driven by changes in the primary balance and the interest rate-growth differential. Assuming that $dda_t \simeq 0$, it can be seen from Eq. (3) that the debt can remain stable (i.e. $\Delta b_t = 0$) despite a primary fiscal deficit ($pb_t < 0$), as long as g_t is large enough compared to i_t, so that the first term on the right-hand side of Eq. (3) compensates for the second. In contrast, when $i_t > g_t$, a primary fiscal surplus is required to achieve the stability of the debt ratio.

Traditional assessments of public-debt sustainability have predominantly relied on deterministic projections of the debt–dynamics identity given by Eq. (3) over a finite horizon. Under this approach, assumed paths for i_t, g_t and pb_t determine a unique forecast for the debt-to-GDP ratio: the baseline projection. That baseline is constructed from what are judged to be "reasonable" or "standard" assumptions about the right-hand-side variables; its credibility is then evaluated by considering alternative assumptions, less optimistic forecasts, and plausible shock scenarios.

This Deterministic Debt Sustainability Analysis (DDSA) thus produces a central debt trajectory that serves as the reference case for policy discussion. Institutions typically augment the simple accounting identity with institution-specific inputs — estimates of potential output, country-specific interest-payment dynamics, and calibrated fiscal multipliers—to operationalize the baseline and its variants.[1]. The realism of the DDSA baseline is therefore tested *ex*

[1] The Banque de France's DDSA model, based on Bouabdallah et al. [2017], implements the macro-accounting

post through sensitivity analysis and scenario exercises rather than by an explicit probabilistic assessment of future outcomes.

3 The SDSA

Stochastic debt-sustainability analysis (SDSA) was developed to quantify the uncertainty that surrounds deterministic debt projections. In practice, SDSA proceeds by simulating a large ensemble of future paths for the key drivers of debt dynamics, drawing shocks for those drivers from distributions informed by their historical volatility, persistence, and cross-correlations. This large set of simulated outcomes yields a full predictive distribution for the debt-to-GDP ratio rather than a single-point trajectory. Two broad classes of approaches are common. The first — which we label "model-free" — directly exploits the historical variance–covariance properties of the drivers to generate shocks. The second employs explicit econometric models, typically vector autoregressions, to capture the joint dynamics of the drivers and to recover the historical innovations that drive simulation.

3.1 The model-free approach of SDSA

3.1.1 The European Commission

The approach of SDSA by the European Commission (EC hereafter) described in its 2022 *Debt Sustainability Monitor*[2] consists of simulating a large number (2,000) of annual nominal debt to GDP trajectories based on Gaussian random draws of the drivers of the debt dynamic equation. As the EC considers a set of 27 countries, some of them do not belong to the Euro Area. Hence, for the latter, Equation (2) is augmented to account for the debt denominated in foreign currency which involves the introduction of the nominal exchange rate, e_t, as follows:

$$b_t = \alpha^n \frac{(1+i_t)}{(1+g_t)} b_{t-1} + \alpha^f \frac{(1+i_t)}{(1+g_t)} \frac{e_t}{e_{t-1}} b_{t-1} - pb_t + dda_t, \tag{4}$$

where α^n and α^f denote the share of total debt denominated in national and foreign currencies, respectively.

To measure the uncertainty surrounding the future debt trajectory, the EC uses past quar-

identity described above using its own assumptions—an internal estimate of potential output, a France-specific treatment of interest payments on public debt, interest-rate projections consistent with the ECB outlook, and calibrated fiscal multipliers. Starting from an unchanged macro-fiscal framework, the model produces a hypothetical scenario that serves as the reference debt-to-GDP trajectory for subsequent sensitivity analysis.

[2] European Commission Institutional Paper 199, April 2023. See also Berti [2013].

terly data of the primary balance (pb_t), the nominal short- and long-term interest rates from which the effective interest rate (i_t) is calculated, and the nominal GDP growth rate (g_t).[3] Stochastic shocks are then defined as the first difference of these series. The covariance matrix of these quarterly shocks, Σ, is then calculated using available observations and used to draw 2,000 random shocks from a multivariate normal distribution $\mathcal{N}(0, \Sigma)$. For each draw, the annual shocks are then obtained by cumulating the four quarterly shocks of the same year. These annual shocks are added to the baseline value of the corresponding variable, where the baseline is determined from the standard deterministic projections, namely, the central scenario. The empirical distribution of the future debt-to-GDP ratio path is finally obtained from these 2,000 simulated annual trajectories.[4]

Although this method has the advantage of not involving any other estimation process than the empirical variance-covariance matrix of these series, it suffers from two important limitations. The first is the assumption that the vector $\Delta X_t = (\Delta pb_t, \Delta i_t^{LT}, \Delta i_t^{ST}, \Delta g_t)'$ has a Gaussian distribution, which probably does not hold in general. This choice has two countervailing effects on the distribution: first, assuming Gaussian shocks reduces the frequency of extreme draws; second, it spreads the variance associated with large shocks across the entire distribution, thereby diluting their impact. The second limit is the implicit assumption that these quarterly series display no persistence, which again does not hold in general.[5]

3.1.2 The International Monetary Fund

The current approach of the IMF, as described in its 2022 *Staff Guidance Note on the Sovereign Risk and Debt Sustainability Framework for Market Access Countries*, departs from the EC one mainly in two ways. First, it relies completely on annual data, and second, bootstrap is used instead of Gaussian simulations. Due to the large number of heterogeneous countries monitored by the IMF, the debt dynamic equation considered is also given by Equation (4). It is then

[3] Data for the nominal exchange rate is also used for non-EA countries.

[4] In its 2023 *Debt Sustainability Monitor* — European Commission Institutional Paper 271, March 2024 — the EC has introduced a few changes regarding the dataset. First, the period covered by the data sample was harmonized across all Member States and now begins in 2000Q1. Second, the treatment of outliers has been changed as follows: observations of a series outside its 5th and 95th percentile thresholds are considered outliers and are replaced by the nearest percentile value. Finally, the number of replications for the simulations has been increased to 10,000.

[5] In a forthcoming Discussion Paper, Bec et al. [2025] propose solutions to these two limitations of the EC approach to the SDSA. First, to account for the persistence of shocks, they introduce a pre-filtering method for the shocks so as to work with series of non-autocorrelated shocks. Second, the assumption of a Gaussian distribution is abandoned and replaced by a bootstrap method.

expressed in real terms, which gives, for Δb_t:

$$\Delta b_t = \frac{(r_t - \rho_t)}{(1+\rho_t)} b_{t-1} + \frac{z_t \alpha_t^f b_{t-1}}{(1+\rho_t)(1+\pi_t^f)} + \frac{(\pi_t^d - \pi_t^f)\alpha_t^f b_{t-1}}{(1+\pi_t^f)\rho_t} - pb_t + dda_t, \tag{5}$$

where r_t is the real effective interest rate, ρ_t is the real GDP growth rate, z_t is the real effective exchange rate, π_t^f and π_t^d are the foreign and domestic inflation rates.[6] The first term accounts for the real growth-interest differential, the second for the real exchange rate, and the third for the relative inflation component. The SDSA then consists of the block bootstrapping of all the variables entering Equation (5), namely r_t, ρ_t, pb_t, z_t, π_t^d, and π_t^f, the latter being measured by the US inflation rate. Let us denote by X_t the column vector that includes these six variables at time t. The block bootstrap approach retained by the IMF consists, for one simulated path, in randomly choosing three blocks of two consecutive observations of X, say (X_{t-1}, X_t), to build a debt ratio trajectory of six years (the current year and the five subsequent years). This process is repeated 10,000 times to obtain the empirical distribution of the future paths of the debt ratio. The latter is presented as a so-called "historical fan chart", which is used to evaluate the realism of the baseline scenario. When there is no realism concern, which means that the deterministically projected debt ratio lies between the 20^{th} and the 80^{th} percentiles of the fan chart, the historical fan chart is used to construct a final baseline-centered fan chart: the demeaned distribution of the historical fan chart is simply added to the baseline.[7]

The main advantage claimed by this approach compared to the one implemented by the EC, is that it captures the persistence of the series on top of their cross-correlation. Nevertheless, in our opinion, it also suffers from two limits. The first one is that for the bootstrap to provide a good approximation of the empirical distribution, the object to bootstrap has to be *i.i.d.* and stationary (see e.g. Bühlmann [1997] or Politis [2003]). However, these conditions are likely violated for the X_t vector defined above. This issue is even more salient when using the block-bootstrap approach. Note that with such short blocks of two years, the bootstrapped series are likely to display jumps between the blocks and more generally nonlinearities which are not present in the original series, hence resulting in less persistence than the observed one. The second limit of the IMF approach is the small size of the sample used. Annual data from 2000 to 2022, *i.e.* twenty-three observations, leave only twenty-two overlapping blocks of two years for

[6] See details of the calculations in Box 3, page 30, of the SRDSF 2022 guidance note.

[7] If the baseline scenario is close to the top of the fan chart, hence revealing some pessimism, it has to be justified by e.g. an expected loosening in future policies. If this pessimism cannot be justified, a correction of the baseline is requested. Similarly, if the baseline scenario is close to the bottom of the fan chart, revealing optimism, the baseline has to be revised.

the random drawing. Without replacement, this gives $22\times21\times20 = 9240$ different bootstrapped trajectories. Consequently, some of the 10,000 bootstrapped X_t must be identical. Of course, this small sample issue would be even more pronounced if a local block bootstrap method[8] was used: it would leave even less different bootstrapped trajectories to draw.

3.2 The VAR-based approach of SDSA

3.2.1 Gaussian approach

To our knowledge, the econometric, VAR-based approach of the SDSA dates back to Garcia and Rigobon [2004] paper, focusing on the Brazilian case. Starting from the equation of accumulation of debt to GDP, expressed in real terms:

$$b_t = (1 + r_t - \rho_t)b_{t-1} - pb_t + \varepsilon_t,$$

where ε_t denotes debt shocks, defined as $b_t - (1 + r_t - \rho_t)b_{t-1} - pb_t$, their goal is to project the future paths of the drivers from the following VAR model:

$$\begin{aligned} X_t &= c + B(L)X_{t-1} + \nu_t \\ X_t &\equiv (r_t, \rho_t, pb_t, \varepsilon_t, z_t, \pi_t^d)' \\ \nu_t &\sim \mathcal{N}(0, \Sigma). \end{aligned}$$

Even though they are not direct drivers of the debt, the real exchange rate and the domestic inflation rate are added to the VAR model since they "could generate comovement in the variables entering the debt accumulation equation" (Garcia and Rigobon [2004], p.9). Monthly data from January 1994 to October 2003 were used to estimate the parameters of this VAR model and the variance-covariance matrix of the shocks Σ. These estimates are then used to generate 500 Monte Carlo replications of 120 months of the VAR residuals, which are in turn used to generate 500 simulated paths of 120 future months of the elements of X_t. From the debt-to-GDP accumulation equation, 500 simulated paths of b_t are obtained. The latter are used to compute the probabilities of reaching a debt larger than 66, 75, 85, 95 and 100 percent of GDP in the following 10 years. Although this study does not show the standard indicators retained in current SDSA, such as fan charts or the probability that the debt will be larger at

[8] As stressed in e.g. Paparoditis and Politis [2002], a local block-resampling procedure should be preferred if the nonstationarity can be captured by a slowly changing stochastic structure. This local block bootstrap resamples only blocks that are adjacent (or nearby) to one another, thereby avoiding large jumps between blocks.

some point in the future than at the last observation, it has paved the way for the subsequent papers on this topic.

The first paper promoting the use of fan charts in the VAR-based approach is the one by Celasun et al. [2006].[9] The approach retained by these authors departs slightly from the previous one in that the VAR model only includes non-fiscal determinants of the public debt dynamics, typically real interest rates, real GDP growth rate, and effective real exchange rate. The latter is introduced because this study focuses on 34 emerging market economies. This VAR model is estimated using quarterly data from 1990 until 2004. It is used to get an estimate of the residuals variance-covariance matrix and to simulate 1,000 future paths of the variables it contains. Following Garcia and Rigobon [2004], the shocks used to generate the simulated paths are drawn from a Gaussian distribution with the same variance-covariance matrix as the residuals. Celasun et al. [2006] introduce a second block in their analysis, namely a yearly fiscal reaction function (FRF), to simulate future paths of the primary balance ratio. The FRF relates the annualized primary balance ratio to the last observed debt to GDP ratio, the current output gap, and other control variables. The simulated annual output gap is calculated from the growth differential between the predicted GDP growth and the steady-state (annualized) growth rate produced by the VAR. Finally, the simulated VAR paths are annualized and used together with the simulated outcome of the FRF to obtain the future simulated path of the debt ratio, using an accumulation equation similar to Equation (5) above. From a large number of simulated paths, the empirical distribution of the debt ratio future path can be calculated and summarized using fan charts.

3.2.2 Bootstrap approach

The same lines as above have been retained by Jooste et al. [2011] and Medeiros [2012], with the main difference that a bootstrap approach of the VAR simulation was substituted for the Gaussian approach described above. Jooste et al. [2011] analysis focuses on the South African economy, and only three variables enter the VAR model, namely the real GDP, the real interest rate, and the GDP deflator. This model is estimated from quarterly data from 1995Q1 to 2010Q1, and with two lags only. Then it is bootstrapped 1,000 times to capture the empirical joint distribution of the future path of these three variables. The fiscal reaction function and the debt accumulation equation are then used to obtain the empirical distribution of future debt-

[9] See also Celasun and Keim [2010].

to-GDP ratio paths.[10] As emphasized in Berti [2013], the issue with the estimation of the FRF is that the primary balance depends, among other drivers, on the level of debt in the previous year and therefore must be estimated using annual data. This poses a problem as a result of the small sample size. Medeiros [2012] circumvents that issue by estimating the FRF from a panel of 15 member states of the European Union[11]. Then, he used the longest quarterly data sample available for each country to estimate a VAR model that included the real effective exchange rate, the German GDP growth rate, and the real interest rate in addition to the three variables retained by Jooste et al. [2011]. Finally, 2,000 bootstrapped drawings of the VAR residuals are used to approach the empirical distribution of the debt ratio. A similar approach is employed, for example, by Cuerpo and Ramos [2015] for Spanish data and by Cherif and Hasanov [2018] for U.S. data; however, rather than supplementing the VAR with a fiscal reaction function, these authors include the primary balance — or its components — directly in the estimated VAR.

More recently, Bouabdallah et al. [2017] described the methodological framework retained by the European Central Bank for debt sustainability analysis, which includes both the deterministic and stochastic blocks. Unfortunately, the stochastic part of the DSA is very briefly described (see p.25-26 therein). It relies on a VAR that includes four drivers, namely the real short-term and long-term interest rates, the real growth of the GDP, and the growth of the GDP deflator. It is estimated using quarterly data from 2001q1 until 2015q4. The estimated residuals are then bootstrapped 5,000 times to generate as many future paths of the drivers for the next five years using the VAR estimates. These bootstrapped paths are annualized, and "future debt paths are consequently calculated using the same debt aggregation model as in the deterministic benchmark scenario" (p.25). Hence, the ECB approach avoids the challenge of estimating a fiscal reaction function. Instead, "the change in the cyclical developments implied by each simulated GDP path influences the path of the structural balance according to the fiscal effort matrix [...]" (p.25). Although the lag order retained for the estimated VAR is not mentioned, since the estimation sample size is 60, this does not leave too many degrees of freedom. Moreover, the serial correlation of residuals jeopardizes the validity of this bootstrap approach. Indeed, as stressed earlier, the object to be bootstrapped has to be *i.i.d.* and stationary. Furthermore, the validity of the stationarity condition can be questioned due to the presence of the two real interest rates series.[12]

[10] More recently, Benjamin Carton and Fouejieu [2020] explore the Dutch SDSA using an identical VAR specification. However, the 5,000 bootstrapped shock series are simply added to the baseline forecasts of the VAR's three variables.

[11] BE, DK, DE, ES, FR, IT, NL, AT, PL,PT, SI, SK, FI, SE, and the UK.

[12] Since that publication, the European Central Bank has revised its procedure, replacing the frequentist VAR

4 The proposed VAR-based framework

In what follows, we will delve into the VAR approach under both frequentist and Bayesian perspectives. We begin by providing a general overview of our approach and dataset before presenting detailed findings in the subsequent sections.

4.1 Departures from existing approaches

Our methodology differs from those previously described in at least three ways.

First, it relies fully on the VAR model estimates: *i)* no fiscal reaction function is used *ad hoc* and *ii)* the simulated debt ratio trajectories are not filtered through the same debt aggregation model as in the deterministic benchmark scenario. Indeed, the estimation of a fiscal reaction function is delicate as one of the main explanatory variables is the lagged debt which is measured at the annual frequency only. Moreover, it is most likely non-stationary. Another important explanatory variable is the output gap, which can only be measured approximately as it is not observable. Then, feeding the debt aggregation model of the deterministic benchmark scenario with a large number of simulated debt drivers trajectories would significantly increase the computational burden of the SDSA.

Second, the variables included in the VAR model are those involved in Eq. (3), namely the nominal short- and long-term interest rates, the nominal GDP growth rate, and the nominal primary balance expressed as a percentage of nominal GDP[13]. The use of nominal variables in the VAR specification is motivated by considerations of data reliability and parsimony. Nominal fiscal aggregates — such as the primary deficit or public debt — are directly observed and tend to be more accurately measured than their real counterparts, which require the use of deflators or price indices that are often subject to revision and may introduce additional measurement error (Orphanides [2001], Croushore and Stark [2001]). Moreover, working in nominal terms avoids adding a price deflator to the VAR, thereby preserving degrees of freedom and improving estimation efficiency. Because the model uses nominal series, price dynamics are implicitly accounted for and explicit deflation of each variable is unnecessary. This choice then allows us to build the debt ratio trajectory very straightforwardly, using Eq. (2) and its last observed value

with a Bayesian VAR; however, the details of the new methodology have not yet been released, but are described in Bouabdallah and Cozmanca [2025]'s forthcoming working paper.

[13] We have tested the inclusion of various additional exogenous variables in the VAR model, such as the price of oil, the U.S. 3-month interest rate (to capture delayed effects of U.S. monetary policy on European rates), the Banque de France's energy price index and the real effective exchange rate. None of these variables significantly improved the predictive performance of the VAR. Since the model uses nominal variables, inflation effects are already embedded in the system's dynamics.

as the initial condition. As the primary balance is included in the VAR model, the corresponding equation can be interpreted as a mild version of the fiscal reaction function: the primary balance is indeed explained by its own past values, by the lagged nominal GDP growth rate and the lagged interest rates.

Third, the use of data for a longer period of time allows us to include enough autoregressive lags in the frequentist VAR to eliminate serial correlation of residuals, so that the bootstrap of residuals is a suitable method.

4.2 The basic steps of our methodology

Regardless of whether a frequentist or Bayesian approach is adopted, simulated future paths of the debt ratio are calculated as follows[14]:

1. First, the estimated VAR residuals are either bootstrapped or drawn from Bayesian sampling say S times.

2. This allows us to calculate S simulated future paths of the VAR variables from $T+1$ to $T+H$, where T is the end date of the sample used for the estimation and H is the number of quarters to forecast. These simulated paths are obtained using the VAR estimated parameters, the last observations of the variables used for the estimation step, and the simulated series of residuals.

3. The S simulated quarterly trajectories of the debt-ratio drivers are annualized and fed into Eq. (2) to produce S annual trajectories of the debt ratio, from which the predictive distribution of the future debt path is constructed.

4.3 The Data

The quarterly data used in the subsequent study cover the period 1990Q1-2023Q4. All data used in this study — historical and forecast — are based on information available in spring 2024. Regarding the short-term interest rate, the data correspond to the 3-month PIBOR until 1995 and to the 3-month Euribor after 1995. It is denoted TX_3M. The long term rate is the 10-year French government bonds rate, and is denoted TX_10Y. The nominal primary balance and GDP series, expressed in billions of euros, come from the French national quarterly accounts released by the INSEE (National Institute of Statistics and Economic Studies). The variable denoted

[14]All our exercises use Eviews as software.

SOLDEP_P is the nominal primary balance to the nominal GDP ratio. The quarter-on-quarter growth rate of the nominal GDP is denoted G_V_QOQ. These series are plotted in Fig. 2.

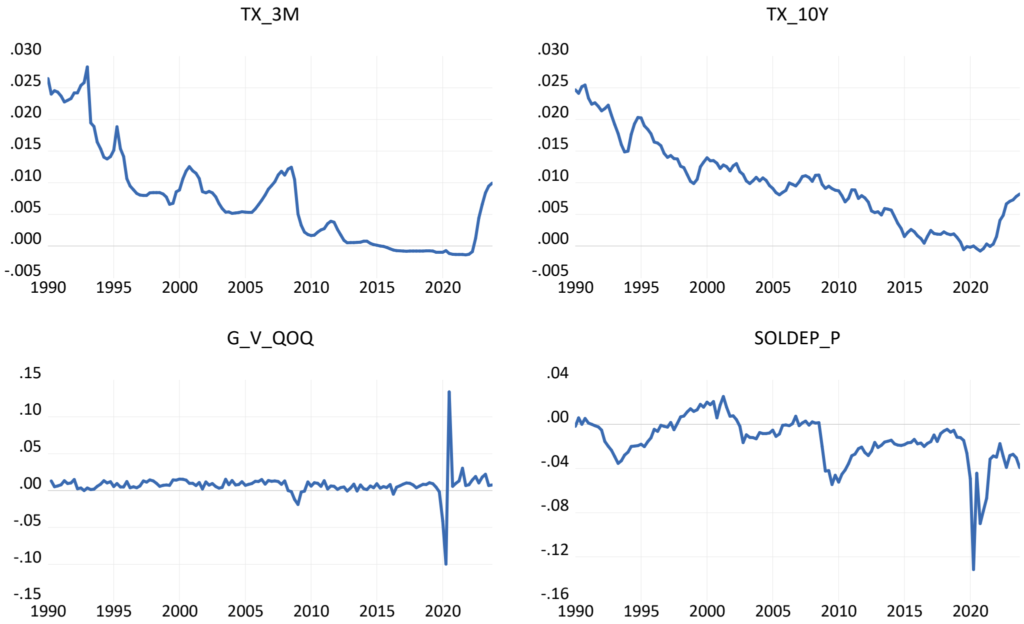

FIGURE 2: *Data (1990Q1-2023Q4)*

Although quarterly fiscal series can exhibit greater measurement noise and short-term revisions than annual aggregates, using quarterly data increases the sample size and improves the precision of VAR estimation while better capturing dynamic propagation and persistence. This choice aligns with the fiscal series used by the ECB or European Commission. Moreover, results are ultimately presented at an annual frequency, which helps smooth potential short-term volatility.

5 The frequentist VAR model

5.1 Estimation results

First, a VAR model including the 3-month and 10-year rates, SOLDEP_P and G_V_QOQ is estimated including a constant term. Four lags are necessary to remove any serial correlation up to order 4. However, this VAR model turned out to be nonstationary with one cointegration relation, according to Johansen's trace test (Johansen [1991]). Henceforth, the VAR used for the SDSA analysis includes the 3-month rate in first difference (DTX_3M) and the interest rate spread (denoted SPREAD) defined as the difference between the ten-year and the three-month

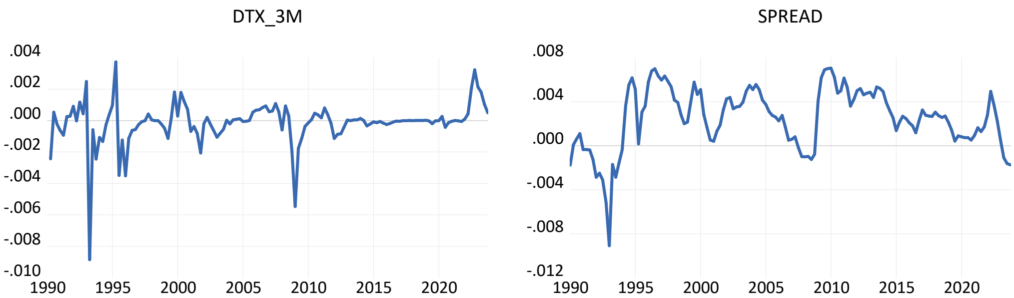

FIGURE 3: *Data (1990Q1-2023Q4)*

rates, together with G_V_QOQ and SOLDEP_P. Notice that this VAR representation is equivalent to a VECM that would include both interest rates in first differences, the growth rate of nominal GDP, the primary balance-to-GDP ratio, and the spread of interest rates as the cointegration vector. It has the advantage of requiring fewer parameters to estimate, which is important in our study.[15] The series DTX_3M and SPREAD are plotted in Fig. 3. The column vector (DTX_3M,SPREAD,G_V_QOQ,SOLDEP_P)' is found to be stationary at the 5%-level. This VAR estimation results are reported in Table A1, see the appendix. Table 1 reports the p-values of the Granger causality tests. The corresponding LR statistics are distributed as $\chi^2(4)$. Expectedly, the three-month interest rate changes and the spread cause each other at the 1%-level. They are also caused by the nominal GDP growth rate, at the 6%-level for DTX_3M and the 3%-level for SPREAD. The joint test that all lags of the primary balance ratio are zero in the spread equation is strongly rejected, with a p-value less than 1%. Finally, the primary balance ratio and the nominal GDP growth rate influence each other as SOLDEP_P Granger causes G_V_QOQ at the 2%-level while SOLDEP_P Granger causes G_V_QOQ at the 3%-level.

TABLE 1: GRANGER CAUSALITY TESTS

	Dependent variable			
	DTX_3M	SPREAD	G_V_QOQ	SOLDEP_P
DTX_3M	0.00	0.00	0.43	0.49
SPREAD	0.00	0.00	0.25	0.16
G_V_QOQ	0.06	0.03	0.13	0.03
SOLDEP_P	0.18	0.00	0.02	0.00

NOTE: *p-values of the LR statistics of the test of the null that the variables in line do not Granger cause the variables in column.*

[15] As stressed in Campbell and Shiller [1987], p. 1066, a VAR model including the two interest rates in first differences would also be stationary, but using these variables instead of one of the interest rate in first difference together with the interest rate spread, one would lose information on the relative levels of the 3-month and 10-year rates.

5.2 Bootstrap results

If the VAR residuals were jointly Gaussian, one could construct fan charts by drawing shocks from a multivariate normal distribution. However, the Jarque and Bera [1980] test reported in Table A1 strongly rejects joint normality ($p \leq 0.0001$). Accordingly, we rely on a nonparametric residual bootstrap to simulate shocks — an approach that preserves the empirical distribution of the residuals (including skewness, excess kurtosis, and cross-correlations) without imposing Gaussianity.

Figure 4 shows for each year from 2024 until 2028 the fan chart obtained for 10,000 bootstrapped trajectories generated from the VAR model described above, together with the baseline (*i.e.* no policy change) scenario of the deterministic DSA (red line). The latter is a hypothetical scenario computed according to the deterministic debt-sustainability framework of Bouabdallah et al. [2017], using 2023-vintage data available in spring 2024. To avoid unrealistic trajectories, the bootstrapped values of the interest rates have been restricted to be greater than -0.01 for the 3-month rate and strictly positive for the 10-year rate[16]. The resulting fan chart is the VAR-based analogue of the so-called historical fan chart presented by the IMF. It can be seen that our deterministic baseline scenario lies between the 20^{th} and the 80^{th} percentiles of the 10,000 simulated debt ratio paths: the baseline seems slightly optimistic until 2026, but becomes very close to the simulations' median afterward. When looking at the fan charts of the debt drivers, Fig. 5, it can be seen that the nominal GDP growth rate (top left panel), short- and long-term interest rates (top right and bottom left panels) and the primary balance ratio (bottom right panel) of the hypothetical scenario from DDSA all lie between the 20^{th} and the 80^{th} percentiles for all forecasting horizons. Their DDSA paths are slightly lower than their respective bootstrapped medians. Fig. A1 in the Appendix shows the implicit interest rate calculated from the DDSA together with the bootstrapped distribution of its analogue, calculated as the weighted average of the 3-month and 10-year rates. The weights used correspond to the historical average shares of short- and long-term debt emitted by the French government during the period considered, namely 42% and 58%, respectively. The calculation of the implicit interest rate being much more sophisticated in the DDSA, which involves more variables and parameters, no wonder its path strays so far from the median of our simulations, particularly in the first two years. Indeed, the first observation in 2023 illustrates the discrepancy between the two kinds of measurement of the implicit interest rate. This might be the reason why the bootstrapped median of the primary balance ratio is slightly above the DDSA trajectory, as its evaluation of

[16]These restrictions have also been imposed by the ECB.

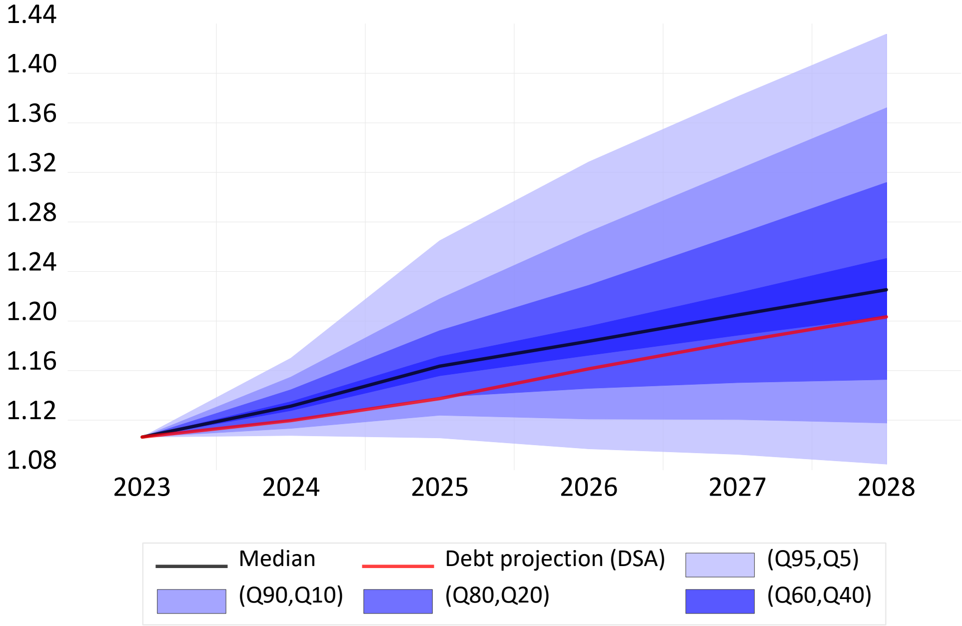

FIGURE 4: *Historical fan chart of the debt ratio*

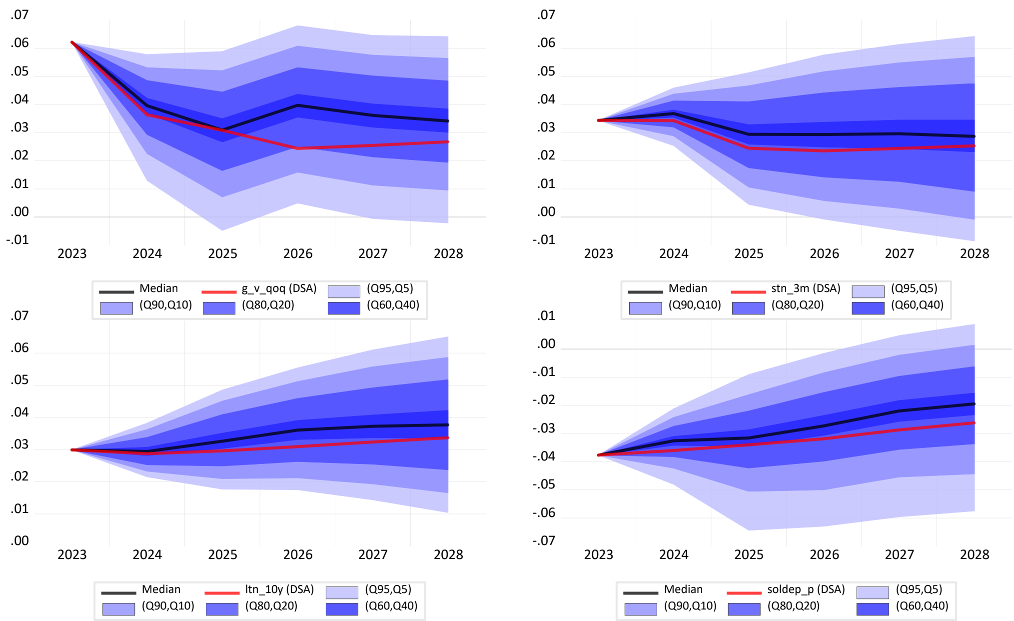

FIGURE 5: *Historical fan chart of the drivers*

the charge of interest is greater. All in all, the above-mentioned imperfections in the simulated drivers trajectories balance each other out in the end, as the DDSA debt ratio never departs too much from the bootstrapped median; see Fig. 4.

Finally, the baseline-centered fan chart can be calculated. To this end, the median corrected distributions of the historical bootstrapped drivers are simply added to their DDSA baseline. The resulting distributions are then used to calculate the 10,000 baseline-centered debt ratio paths. The corresponding fan chart is presented in Figure 6. The uncertainty of the debt forecast

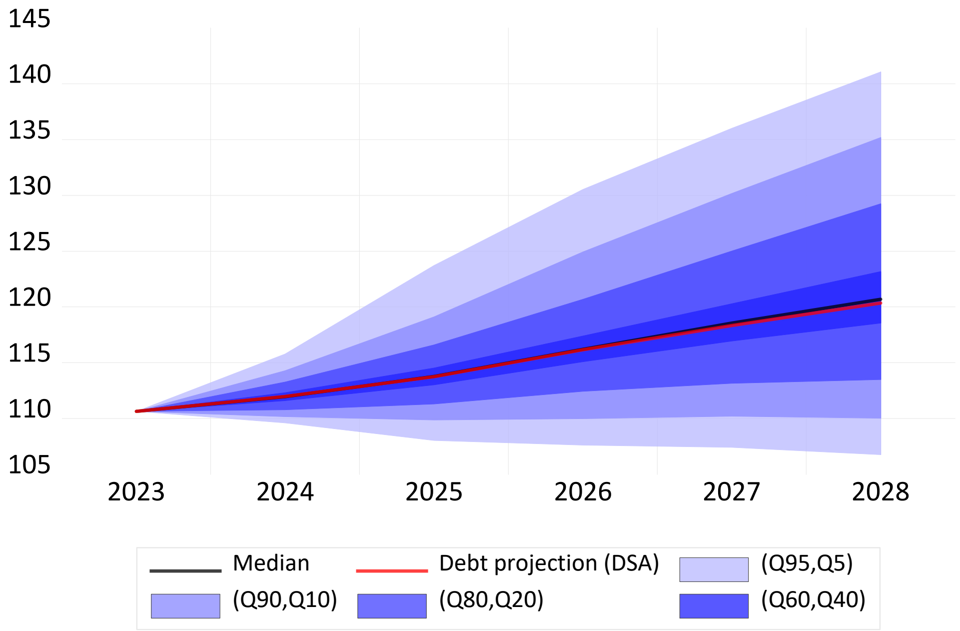

FIGURE 6: *Baseline-centered fan chart of the debt ratio*

increases as the horizon lengthens, but not too much according to the size of the confidence intervals in the fan chart depicted in Figure 6. To measure the uncertainty surrounding the projected trajectory of the French public debt, it is possible to examine the width of the 10%-90% cone, that is, the difference between the 90th and 10th percentiles of the debt ratio. In 2028, this cone width is equal to 25.2 percentage points of GDP.

6 A Bayesian VAR model

Although pioneered in the 1980s by the works of Doan et al. [1984] and Litterman [1986], the use of Bayesian VAR models (BVAR hereafter) for macroeconomic forecasting has really taken off in the past two decades[17]. In this framework, the VAR parameters are considered random variables with respect to which the econometrician has priors. The Bayesian approach involves estimating the posterior probability distribution of the model's parameters based on a sample of

[17]See Carriero et al. [2009], Banbura et al. [2010], Karlsson [2013], Koop [2013], Giannone et al. [2014], Carriero et al. [2015], Giannone et al. [2019], Del Negro et al. [2020], Lenza and Primiceri [2020], Crump et al. [2021] or Cimadomo et al. [2022], among others.

observations and the chosen priors. This approach has proven particularly fruitful for forecasting from a VAR model, as the latter is very demanding in terms of the number of parameters to estimate. In fact, if we denote n as the number of variables introduced into the model and p as the number of autoregressive lags, the VAR implies the estimation of $n \times n \times p$ parameters. This so-called "curse of dimensionality" can cause a severe overfitting issue. In this case, while the in-sample fit might be excellent — or even perfect — the out-of-sample forecasts will in general be very poor. In the Bayesian approach, the information given by the prior beliefs on the parameters might address this problem.

6.1 Definition of the priors

A very common prior retained in the BVAR literature applied to macroeconomics is the one proposed by Kadiyala and Karlsson [1997] and Sims and Zha [1998], often referred to as the normal-inverted Wishart (N-IW hereafter) prior. This is an extension of the Minnesota prior promoted by Litterman [1986] in that it relaxes the assumption that the covariance matrix of the VAR residuals is diagonal, fixed, and known. This assumption excludes any correlation among the n residual series of the VAR model. Instead, Kadiyala and Karlsson [1997] and Sims and Zha [1998] assume that this covariance matrix is distributed as an inverted Wishart.

To fix ideas, let us consider the following VAR model[18]:

$$x_t = b_0 + b_1 x_{t-1} + b_2 x_{t-2} + \cdots + b_p x_{t-p} + \nu_t \tag{6}$$

where $x_t = (x_{1t}\ x_{2t} \cdots x_{nt})'$ is a vector that includes the n variables at time t and $\nu_t \sim i.i.d.\ \mathcal{N}(0, \Sigma)$. To write Eq.(6) in a more compact form, we define $X = [x_1\ x_2 \cdots x_T]'$, $B = [b_0\ b_1 \cdots b_p]$, $Z = [z_1\ z_2 \cdots z_T]'$ with $z_t = (1\ y'_{t-1} \cdots y'_{t-p})'$, and finally $V = [\nu_1\ \nu_2 \cdots \nu_T]'$ to obtain the following:

$$X = ZB + V \tag{7}$$

The N-IW conjugate prior corresponds to:

$$B|\Sigma \sim \mathcal{N}(B_0, \Sigma \otimes \Omega_0) \text{ and } \Sigma \sim IW(S_0, v_0) \tag{8}$$

As the variables included in our VAR model are mean reverting, the prior for the expectation

[18] For more details about the setting of the priors, see Banbura et al. [2010] or Carriero et al. [2015] among others.

of the coefficient matrices is $E[B_k^{(ij)}] = 0$, where $B_k^{(ij)}$ is the element $(i,j)^{th}$ of matrix B_k, $k = 1, \cdots, p$. This amounts to the prior belief that all equations are centered around a white noise. The prior for the standard deviation is:

$$SD[B_k^{(ij)}] = \begin{cases} \frac{\lambda_1 \lambda_2}{k} \frac{\sigma_i}{\sigma_j}, & k = 1, \cdots p \\ \lambda_0 \sigma_i, & k = 0 \end{cases} \tag{9}$$

The overall tightness of the prior is governed by λ_1. This parameter is crucial because it determines the informativeness of the prior. The closer it is to zero, the less the data will influence the estimates. On the other hand, as its value goes to infinity, the estimation is not influenced by the prior and the posterior expectations coincide with the frequentist OLS approach. λ_2 is the parameter that governs the shrinkage between variables, that is, the shrinkage of the parameters of the lags of other variables but the ones of the dependent variable. Regarding λ_0 and σ_i, the scale parameter, they are commonly set respectively to one and to either the standard deviation of the residuals from a univariate autoregression or the diagonal of the frequentist VAR residual covariance. Finally, S_0 and v_0 are chosen so that the prior expectation of Σ corresponds to the fixed residual covariance matrix of the Minnesota prior: $E[\Sigma] = diag(\sigma_1^2, \cdots, \sigma_n^2)$.

6.2 Calibration of the hyperparameters and selection of the lag length

As shown in Carriero et al. [2015], the payoffs of optimizing the lag length are larger than those of optimizing the shrinkage parameters. According to them, the reason why "is probably that the tried-and-true values [of the shrinkage parameters] have been established largely on the basis of forecasting performance"(p.58 therein). Consequently, the BVAR hyperparameters are first set to their common values[19]: $\lambda_0 = 1$, $\lambda_1 = 0.2$ and $\lambda_2 = 1$. Later in this work, the robustness of the results to this choice of the overall tightness parameter — the most influential one — will be checked.

In order to choose the lag length of the BVAR model, we focus on the pseudo-out-of-sample forecasting performance, hence following Banbura et al. [2010] or Carriero et al. [2015]. Since the main goal of this paper is to evaluate the uncertainty surrounding the future trajectory of the debt ratio, both the point forecast and the density forecast are considered. To this end, a recursive estimation window is used. The first sample stops in 2013Q4 while the last one stops in 2023Q4. For each sample, we compute one (respectively four) quarter ahead forecasts,

[19] See for example Sims and Zha [1998].

which results in a total of 40 one-step (resp. 36 four-step) ahead forecasts per variable. The point forecast is calculated as the median of the distribution. More precisely, to evaluate the performance of the point forecasts, we calculate the root mean squared forecast error (RMSFE) over the 40 (or 36 for $h = 4$) forecasts for each variable, for BVARs with lag duration $p = 1, \cdots, 4$, and the average of the RMSFE obtained for the four variables.[20] Regarding the evaluation of the density forecast, we use the average score after Carriero et al. [2015]. To this end, the log score in predicting variable i at horizon h is first calculated as follows:

$$s_t(x^i_{t+h}) = -0.5[\ln(2\pi) + \ln(V^i_{t+h|t}) + (x^i_{t+h} - \bar{x}^i_{t+h|t})^2 / V^i_{t+h|t}],$$

with $\bar{x}^i_{t+h|t}$ and $V^i_{t+h|t}$ respectively the posterior mean and variance of the simulated forecast distribution for x^i_{t+h}. Then, the average score is obtained by averaging the log scores over the 40 one-step ahead forecasts and 36 four-step ahead forecasts. The results obtained for $h = 1$ and $h = 4$ are reported in tables 2 and 3. For comparison's sake, the statistics obtained from the frequentist VAR(4) are also presented. Interestingly, BVARs outperform frequentist VARs in terms of density forecast performance. Indeed, unlike the BVARs, the standard VAR(4) might suffer from the curse of dimensionality as in the first estimation sample, 1991Q1-2013Q4, only

Table 2: RMSFE and score as a function of p ($h = 1$)

	DTX_3M	SPREAD	G_V_QOQ	SOLDEP_P	average
			RMFSE		
VAR(4)	0.995	0.477	0.514	0.438	0.606
BVAR(1)	1.191	0.539	0.687	0.560	0.744
BVAR(2)	0.995	0.483	0.663	0.548	0.672
BVAR(3)	0.991	0.480	0.656	0.539	0.667
BVAR(4)	*0.989*	*0.463*	*0.655*	*0.532*	*0.660*
			Score		
VAR(4)	5.234	4.969	2.436	2.198	3.709
BVAR(1)	*6.042*	*5.662*	*2.973*	2.392	*4.267*
BVAR(2)	5.969	5.576	2.937	2.404	4.221
BVAR(3)	5.997	5.566	2.920	2.404	4.222
BVAR(4)	6.001	5.540	2.926	*2.408*	4.219

[20]In order to prevent disproportionately greater emphasis on variables of larger magnitude, the RMSFE's are normalized by the variables standard deviations.

91 observations are left for estimation. Then, from the top panel of Table 2, it can be seen that the RMSFE criterion points to $p = 4$ as this lag length provides the smallest values. In contrast, from the bottom panel of Table 2, it can be seen that the score results point to one lag: this lag length maximizes the score obtained for all variables but the primary balance ratio — whose score is not very sensitive to the lag length — and for the average of the four scores. Because the analysis focuses on fan charts, we place greater weight on density forecasts than on point forecasts when evaluating predictive performance. For the 4 steps ahead forecasts, Table 3, the BVAR with one lag minimizes all RMSFEs and maximizes all scores statistics. Hence, we proceed with the study by setting p = 1.

Table 3: RMSFE and score as a function of p $(h = 4)$

	DTX_3M	SPREAD	G_V_QOQ	SOLDEP_P	average
			RMFSE		
VAR(4)	3.201	2.585	2.040	3.699	2.881
BVAR(1)	*1.365*	*1.420*	*1.078*	*1.256*	*1.280*
BVAR(2)	1.413	1.494	1.080	1.298	1.321
BVAR(3)	1.432	1.537	1.081	1.309	1.340
BVAR(4)	1.445	1.540	1.080	1.311	1.344
			Score		
VAR(4)	4.358	3.939	1.219	0.491	2.502
BVAR(1)	*6.162*	*5.169*	*3.220*	*1.939*	*4.123*
BVAR(2)	5.937	4.865	3.149	1.853	3.951
BVAR(3)	5.854	4.814	3.116	1.822	3.902
BVAR(4)	5.819	4.883	3.147	1.830	3.920

6.3 Bayesian sampling forecasts

We present here the out-of-sample forecasting results from the BVAR(1) selected above. The forecasts are obtained by Bayesian sampling using 10,000 draws, after excluding 10 percent burn-in draws. Similarly to the frequentist VAR bootstrapped fan charts, we begin with the historical fan chart of the debt to GDP ratio presented in Fig. 7. As can be seen, the median of the debt ratio forecast slightly overestimates the baseline in 2024 and 2025, then becomes very close to the baseline over the next two years — well in the 40-60% confidence interval — before it falls under it in 2028. However, the forecast median still lies in the 20-80% confidence interval in 2028. All in

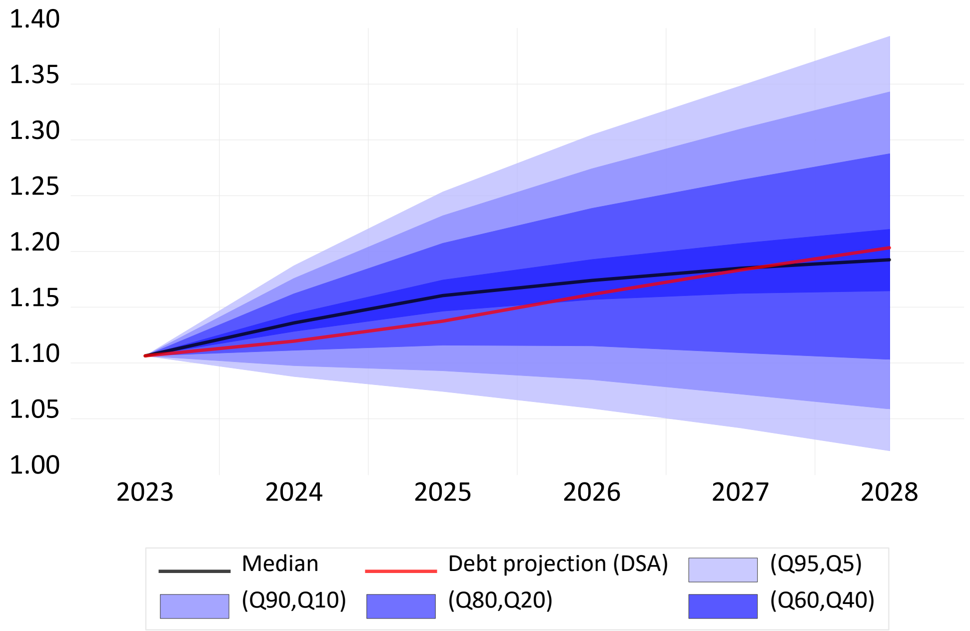

FIGURE 7: *Historical fan chart of the debt ratio (BVAR)*

all, although slightly less precise than its frequentist analogue (Fig. 4), this debt ratio trajectory is satisfying in that it never departs too much from the deterministic no-policy-change scenario. The debt drivers historical fan charts are reported in Fig. A2, see Appendix. Overall, they look rather similar to the ones obtained from the standard VAR bootstrap approach. Finally, the baseline-centered debt ratio fan chart is plotted in Fig. 8. Here, the 10%-90% cone width is equal to 28.5 p.p. of GDP in 2028, which is moderately larger than the 25.2 p.p. obtained from the frequentist VAR. This result probably stems from the fact that with 131 observations in our estimation sample, the standard VAR does not suffer from the overfitting problem. To check the robustness of this finding, Table 4 reports the 10 and 90 percentiles, the corresponding cone size, and the RMSFE between the deterministic baseline and the median forecast of BVAR (1) for the five projection years for various values of the overall tightness parameter λ_1. For this exercise, Bayesian sampling is limited to 5000 draws, after excluding 10 percent burn-in draws.

TABLE 4: $q_{90\%} - q_{10\%}$ CONE SIZE AS A FUNCTION OF λ_1

λ_1	0.01	0.05	0.1	0.15	0.2	0.3	0.5	0.75	1	5
$q_{90\%}$	138.7	137.9	137.4	137.4	137.9	138	139	139	139	139
$q_{10\%}$	109.2	109.6	109.6	109.5	109.5	109	109	109	109	109
$q_{90\%} - q_{10\%}$	29.5	28.3	27.8	27.9	28.4	29	30	30	30	30
RMSFE	0.025	0.017	0.011	0.012	0.013	0.015	0.015	0.016	0.015	0.016

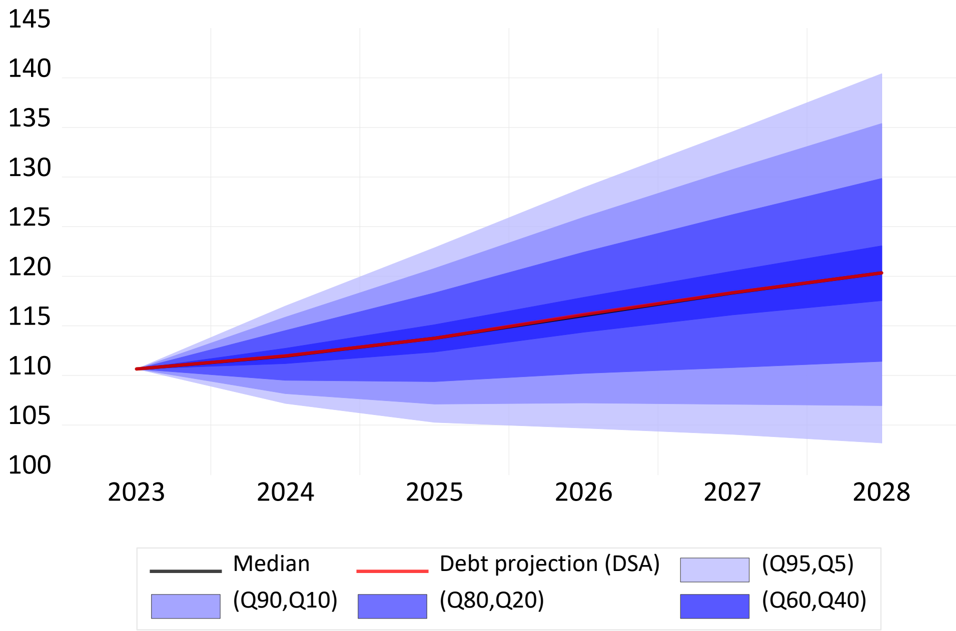

FIGURE 8: *Baseline-centered fan chart of the debt ratio (BVAR)*

As can be seen from this table, the cone size is not too much affected by the overall tightness: It ranges from 27.8 p.p. ($\lambda_1 = 0.1$) to 30 p.p. ($\lambda_1 \geq 0.5$). Similarly, except for the very tight prior $\lambda_1 = 0.01$, the RMSFE is between 0.11 ($\lambda_1 = 0.1$) and 0.17 ($\lambda_1 = 0.05$). In fact, considering these two criteria, the λ_1 values between 0.1 and 0.2 give comparable results. It is worth mentioning that even for $\lambda_1 = 0.1$, the value that minimizes both the cone width and the RMSFE, the standard VAR approach cone width remains slightly smaller. Henceforth, it seems that our estimation sample size does not prevent the latter from achieving good precision of the estimated coefficients.

7 Comparison of results and debt stabilization condition

7.1 Comparison of results

In this section, we compare our results to those obtained by the European Commission and the European Central Bank in terms of the width of the cone $q_{90\%} - q_{10\%}$ and, when available, the probability that the last projected year of the debt ratio is less than the last observed value. We will focus on the last two publicly available 5-year ahead projections, namely years 2027 and 2028. For the European Commission, the corresponding figures can be found in the "Debt Sustainability Monitor", EC Institutional Paper 199, April 2023 and EC Institutional Paper

271, March 2024. For the European Central Bank, they come from the ECB SDSA exercises. As explained in Bouabdallah and Cozmanca [2025]'s paper, the ECB's SDSA method now relies on the estimation of a Bayesian vector autoregression (BVAR). In order to compare the results obtained for the year 2027, we repeated the VAR(4) and BVAR(1) analysis restricting the estimation sample to end in 2022Q4. These results are summarized in Table 5.

Table 5: Results comparison

		EC[a]	ECB[b]	VAR(4)	BVAR(1)
2028	cone width	19.5	37.8	25.2	28.5
	prob.	81%	na	88.6%	82.1%
2027	cone width	21.7	34	24.8	27.1
	prob.	51%	na	80.9%	74.4%

The $q_{90\%} - q_{10\%}$ cone width is measured in percentage points of GDP. 'Prob.' refers to the probability that the last projected year debt ratio is greater than its last observed value. [a] : source EC [b] : source ECB

As can be seen from this table, the cone widths obtained from our VAR(4) and BVAR(1) approaches lie in between the values obtained by the EC and ECB respectively. Since the EC approach does not rely on model estimation but solely on the calculation of the empirical covariance of the first differences of the drivers, this may explain why the resulting cone widths are the smallest ones in this table: unlike the other approaches, there is no uncertainty stemming from modeling step. With the exception of the EC, all approaches exhibit a widening of the cone between 2027 and 2028, thereby suggesting a heightened uncertainty about the future path of the debt-to-GDP ratio. With respect to the EC, it is important to note that the data processing methodology has evolved between these two projection exercises, which might explain the slight decrease obtained between 2027 and 2028.[21] When comparing our VAR and BVAR approaches, relying exactly on the same data set, it is noteworthy that the cone width is slightly smaller according to the VAR (25.2 pp versus 28.5 pp for the BVAR in 2028). This finding is somewhat unexpected, but, as mentioned earlier, our sample size might be large enough for the VAR to achieve good precision of the estimated coefficients. Another possible explanation is that our sample contains essentially two outliers (2020Q2 and 2020Q3) among 131 quarterly

[21] Specifically, outliers, which were formerly discarded if they crossed a predetermined threshold, are now winsorized (they are capped at the 95th or 5th percentile, whichever they are closest to). Moreover, while the longest available sample was used for each driver previously, the sample period is now shorter as it begins in 2000 for all countries. Finally, the primary balance is now seasonally adjusted, which removes a significant part of the volatility.

observations. Consequently, when using the bootstrap method with replacement, the probability of not drawing any of these outliers in a sequence of 20 quarters is high — 73.5% — whereas the probability of drawing both is only 3.4%.[22] In contrast, in the Bayesian approach, the volatility induced by these outliers propagates throughout the entire distribution, thereby affecting the Bayesian sampling.

If we now examine the evolution of the probability that the projected debt-to-GDP ratio in five years will be higher than its initial value, we can observe that the results obtained from the VAR(4) and BVAR(1) models, while increasing, do so to a much lesser extent than those of the EC approach. In fact, both VAR (4) and BVAR (1) indicate an increase in the probability that the debt ratio for the last projected year exceeds its last observed value of approximately 7.5% between the two exercises. The sharp increase in the probability of higher debt levels around 2027–2028 is not driven by a structural change in policy expectations, but reflects the delayed impact of increasing interest payments. As debt matures and is refinanced at higher rates, fiscal pressure builds up, increasing the likelihood of unfavorable debt outcomes. In addition, the economic consequences of recent shocks - such as the war in Ukraine - remain embedded in the data and continue to affect fiscal dynamics. These factors, combined with the stochastic nature of the model, contribute to wider uncertainty and a growing upper tail in the debt distribution over the medium term. In the EC approach, an increase of 30% is found. Again, this might be due to the change in their data-processing methodology.

It is important to emphasize that comparing our results with those of these two European institutions has limitations. Indeed, in addition to the fact that these rely on different methodologies (model-free and VAR-based), the variables used in their SDSA approaches are not identical (nominal for the EC versus real for the ECB) and do not cover the same period as our data. In other words, the metrics used are different, which limits the relevance of these comparisons.

[22] The GDP growth rate series and the primary surplus to GDP ratio series were particularly affected by the emergence of the COVID-19 crisis and the subsequent confinement period. For this reason, we also considered introducing dummy variables for the second and third quarters of 2020. These dummies are highly significantly different from zero in the equations for these two variables, leading to an increase in the cone width by 6 pp and a decrease in the probability that the debt ratio in 2028 exceeds its 2023 level to 82%. The decision not to include these dummies in our VAR model is based on two arguments. The first is that these shocks were indeed observed and it seems important to account for them in evaluating future uncertainty. The second stems from the surprising result of an increase in the cone width despite the removal of these outliers from the bootstrap residuals, which could suggest that the loss of eight degrees of freedom due to added parameters to estimate might affect the overfitting issue.

7.2 Debt-stabilizing primary balance

As a by-product, the analytical framework developed in Sections 5 and 6 above allows us to compute the probability of reaching the debt-stabilizing primary balance. This balance is defined as the budget surplus or deficit that stabilizes the debt at its last observed level. Using the benchmark deterministic projections calculated from the Debt Dynamic Sustainability Analysis (DDSA) for the implicit interest rate, nominal GDP growth rate, and the debt ratio in Equation (3), taken at the steady state, it follows that the debt-stabilizing primary balance ratio is given by:

$$pb_t = \frac{(i_t - g_t)}{(1 + g_t)} b_{t-1}$$

The benchmark trajectory from the DDSA, as well as the debt-stabilizing level of the primary balance ratio, are depicted in Figure 9 below. We observe that for the data set available in spring

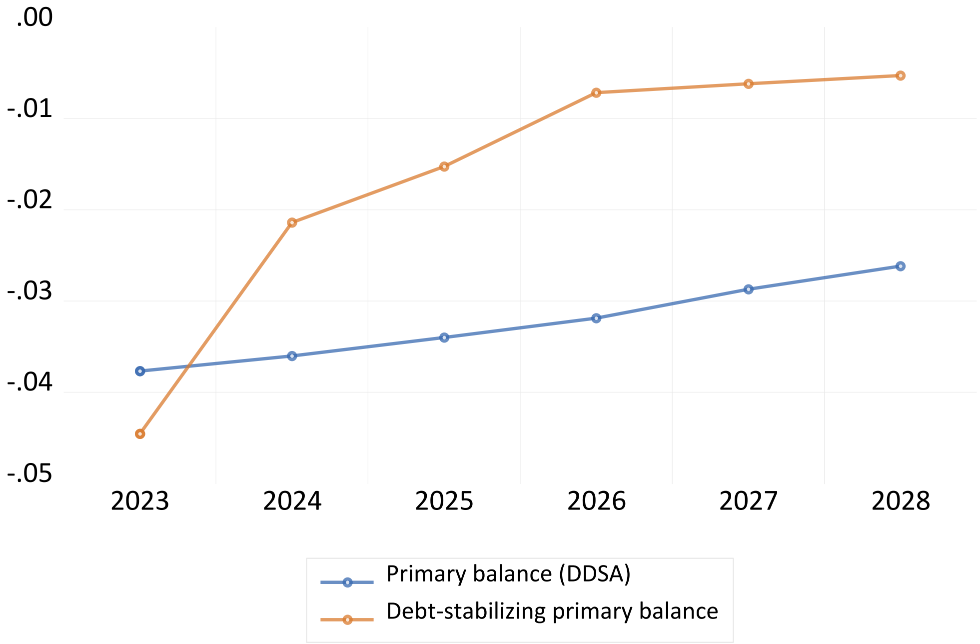

FIGURE 9: *Baseline and debt-stabilizing primary balance ratios*

2024, the primary deficit that would have stabilized the debt from one year to the next is smaller than the projection made by the DDSA from 2024 on. Based on the probability distributions of the primary balance obtained by simulation of VAR (4) and BVAR (1), we can calculate the probability of reaching the necessary deficit to stabilize the debt. More precisely, we calculate the percentage of the 10,000 simulated primary balance trajectories that lie on or above the stabilizing level. These results are presented in Table 6. It can be seen that the percentage of

Table 6: Probabilities to stabilize or decrease the debt-ratio

	2024	2025	2026	2027	2028
VAR(4)	5.23%	8.95%	8.94%	14.49%	18.43%
BVAR(1)	15.90%	22.21%	18.81%	22.44%	23.27%

These probabilities are calculated from 10,000 simulated primary balance trajectories.

simulated primary balance trajectories that lie on or above the stabilizing level is higher under the Bayesian VAR than under the frequentist VAR, especially in the first four projection years. For example, the probability of reaching the primary balance ratio that would have stabilized the debt between 2023 and 2024 is just over 5% according to the VAR(4) and nearly 16% according to the BVAR(1). To interpret this result, Figure 10 shows fan charts of the primary balance generated from simulations of the frequentist VAR (left) and the BVAR (right). The red line denotes the debt-stabilizing primary balance. Although the VAR's cone is narrower than that of the BVAR, the probability of achieving a debt-stabilizing primary balance is lower under the VAR because the simulated distribution is asymmetric. As the figure illustrates, the 10th and 5th percentiles lie much farther below the median than the 90th and 95th percentiles lie above it. In contrast, the BVAR fan chart is nearly symmetric by construction, which mechanically places the red line closer to the median. However, the results of these two models tend to converge later

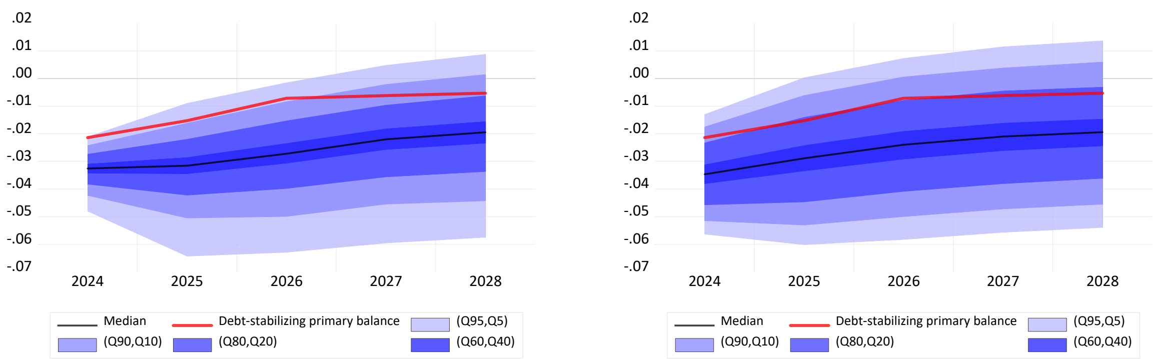

(a) VAR(4) (b) BVAR(1)

Figure 10: Primary balance ratio fan charts and debt-stabilizing primary balance ratio

in the period, where this probability is approximately 18.5% for the VAR(4) and 23.3% for the BVAR(1) for the year 2028. In any case, these results from the stochastic analysis suggest that a more substantial adjustment effort would be necessary to increase the likelihood of stabilizing the French public debt.

8 Concluding remarks

We propose a simple, transparent procedure to quantify uncertainty around projected debt-to-GDP trajectories. The method builds on the estimation — either frequentist or Bayesian — of a parsimonious vector autoregression that includes the principal drivers of the debt ratio, estimated on quarterly data from 1990:Q1 to 2023:Q4, and uses future trajectories simulation (10,000 draws) to construct predictive fan charts.

Two empirical findings stand out. First, the median trajectories produced by both VAR implementations closely match our deterministic baseline for 2024–2028; moreover, that baseline lies inside the 20–80 percent predictive interval generated by our simulations. Second, the BVAR yields a marginally wider 10–90 percent cone for the debt ratio in 2028 (28.5 percentage points) than the frequentist VAR (25.2 percentage points), while assigning a slightly lower probability that the 2028 debt ratio exceeds its 2023 level (82.1 percent versus 88.6 percent). The cone widths we obtain are between those reported by the European Commission and the European Central Bank, suggesting that our estimates are quantitatively plausible.

Future work will extend the framework to incorporate the European Union's new fiscal architecture — the Economic Governance Review adopted in February 2024 — which is likely to affect both baseline projections and the metrics used to evaluate their uncertainty.

References

Banbura, M., D. Giannone, and L. Reichlin (2010). Large Bayesian vector auto regressions. *Journal of Applied Econometrics 25*, 71–92.

Bec, F., F. Courtoy, P. Mohl, and F. Opitz (2025). The stochastic simulations of the commission's debt sustainability analysis: An improved approach. Economic papers, European Commission. Forthcoming.

Benjamin Carton, B. and A. Fouejieu (2020). Assessing dutch fiscal and debt sustainability. IMF Working Papers 2020/269, International Monetary Fund.

Berti, K. (2013). Stochastic public debt projections using the historical variance-covariance matrix approach for eu countries. European Economy - Economic Papers 2008 - 2015 480, Directorate General Economic and Financial Affairs (DG ECFIN), European Commission.

Blanchard, O., A. Leandro, and J. Zettelmeyer (2021). Redesigning EU Fiscal Rules: From Rules to Standards. Working Paper 21-1, Peterson Institute for International Economics.

Bouabdallah, O., C. Checherita-Westphal, F. Drudi, R. Setzer, R. De Stefani, T. Warmedinger, and A. Westphal (2017). Debt sustainability analysis for euro area sovereigns: a methodological framework. Occasional Paper Series 185, European Central Bank.

Bouabdallah, O. and B. Cozmanca (2025). Fiscal sustainability challenges: a stochastic approach for the euro area sovereigns. Occasional paper series, European Central Bank. Forthcoming.

Bühlmann, P. (1997). Sieve bootstrap for time series. *Bernoulli 3*, 123–148.

Campbell, J. and R. Shiller (1987). Cointegration and tests of present value models. *Journal of Political Economy 95*, 1062–1088.

Carriero, A., T. Clark, and M. Marcellino (2015). Bayesian VARs: Specification Choices and Forecast Accuracy. *Journal of Applied Econometrics 30*, 46–73.

Carriero, A., G. Kapetanios, and M. Marcellino (2009). Forecasting exchange rates with a large bayesian var. *International Journal of Forecasting 25*, 400–417.

Celasun, O., X. Debrun, and J. Ostry (2006). Primary surplus behavior and risks to fiscal sustainability in emerging market countries: A "fan-chart" approach. *IMF Staff Papers 53*, 401–425.

Celasun, O. and G. Keim (2010). The Us Federal Debt Outlook: Reading the Tea Leaves. IMF Working Papers 2010/062, International Monetary Fund.

Cherif, R. and F. Hasanov (2018). Public debt dynamics: the effects of austerity, inflation, and growth shocks. *Empirical Economics 54*, 1087–1105.

Cimadomo, J., D. Giannone, M. Lenza, F. Monti, and A. Sokol (2022). Nowcasting with large Bayesian vector autoregressions. *Journal of Econometrics 231*, 500–519.

Croushore, D. and T. Stark (2001). A real-time data set for macroeconomists. *Journal of Econometrics 105*(1), 111–130.

Crump, R., S. Eusepi, D. Giannone, E. Qian, and A. Sbordone (2021). A Large Bayesian VAR of the United States Economy. Staff Reports 976, Federal Reserve Bank of New York.

Cuerpo, C. and J. Ramos (2015). Spanish public debt sustainability analysis. *Hacienda Publica Espanola / Review of Public Economics 215*, 95–118.

Del Negro, M., M. Lenza, G. Primiceri, and A. Tambalotti (2020). What's Up with the Phillips Curve? *Brookings Papers on Economic Activity 51*, 301–373.

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

Garcia, M. and R. Rigobon (2004). A Risk Management Approach to Emerging Market's Sovereign Debt Sustainability with an Application to Brazilian Data. NBER Working Papers 10336, National Bureau of Economic Research, Inc.

Giannone, D., M. Lenza, D. Momferatou, and L. Onorante (2014). Short-term inflation projections: A bayesian vector autoregressive approach. *International Journal of Forecasting 30*, 635–644.

Giannone, D., M. Lenza, and L. Reichlin (2019). Money, Credit, Monetary Policy, and the Business Cycle in the Euro Area: What Has Changed Since the Crisis? *International Journal of Central Banking 15*, 137–173.

Jarque, C. and A. Bera (1980). Efficient tests for normality homoscedasticity and serial independence of regression residuals. *Econometrics Letters 6*, 255–259.

Johansen, S. (1991). Estimation and hypothesis testing of cointegration vectors in gaussian vector autoregressive models. *Econometrica 59*, 1551–1580.

Jooste, C., A. Cuevas, I. Stuart, and P. Burger (2011). Fiscal sustainability and the fiscal reaction function for South Africa. IMF Working Papers 2011/069, International Monetary Fund.

Kadiyala, K. and S. Karlsson (1997). Numerical Methods for Estimation and Inference in Bayesian VAR-Models. *Journal of Applied Econometrics 12*, 99–132.

Karlsson, S. (2013). Chapter 15 - forecasting with bayesian vector autoregression. In G. Elliott and A. Timmermann (Eds.), *Handbook of Economic Forecasting*, Volume 2 of *Handbook of Economic Forecasting*, pp. 791–897. Elsevier.

Koop, G. (2013). Forecasting with Medium and Large Bayesian VARS. *Journal of Applied Econometrics 28*, 177–203.

Lenza, M. and G. Primiceri (2020). How to estimate a VAR after March 2020. Working Paper Series 2461, European Central Bank.

Litterman, B. (1986). Forecasting with bayesian vector autoregressions: Five years of experience. *Journal of Business & Economic Statistics 4*, 25–38.

Medeiros, J. (2012). Stochastic debt simulation using VAR models and a panel fiscal reaction function – results for a selected number of countries. European Economy - Economic Papers 2008 - 2015 459, Directorate General Economic and Financial Affairs (DG ECFIN), European Commission.

Orphanides, A. (2001). Monetary policy rules based on real-time data. *American Economic Review 91*(4), 964–985.

Papaoikonomou, D. (2025). Stochastic debt sustainability analysis: A methodological note. Working Paper 338, Bank of Greece.

Paparoditis, E. and D. Politis (2002). Local block bootstrap. *Comptes Rendus Mathematique 335*, 959–962.

Politis, D. N. (2003). The impact of bootstrap methods on time series analysis. *Statistical Science 18*, 219–230.

Sims, C. and T. Zha (1998). Bayesian methods for dynamic multivariate models. *International Economic Review 39*, 949–68.

Steel, D. (2021). Evaluating forecast uncertainty with stochastic simulations. Working Paper 17, Office for Budget Responsibility.

Appendix

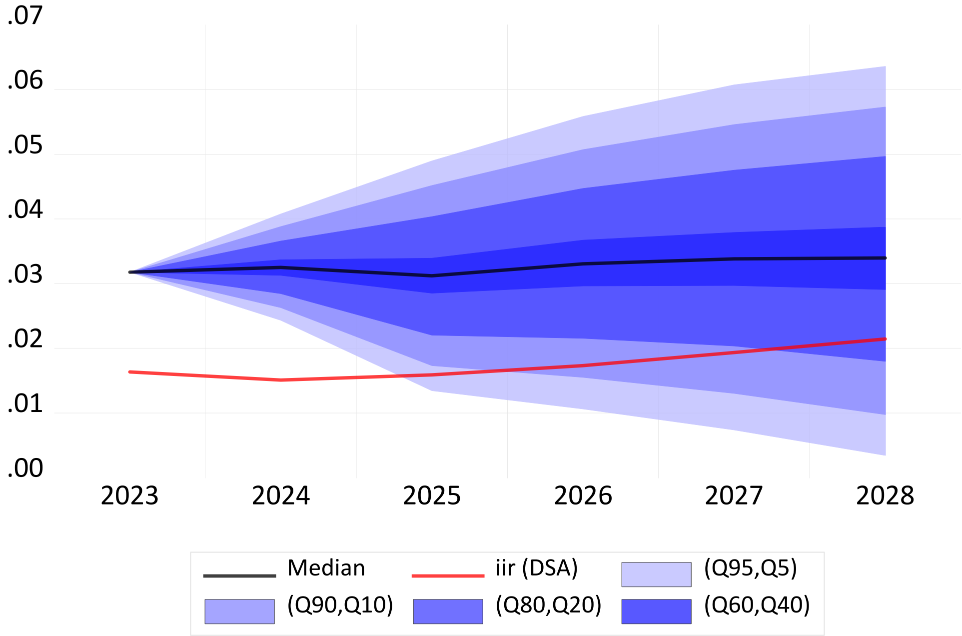

Figure A1: *Historical fan chart of the implicit interest rate*

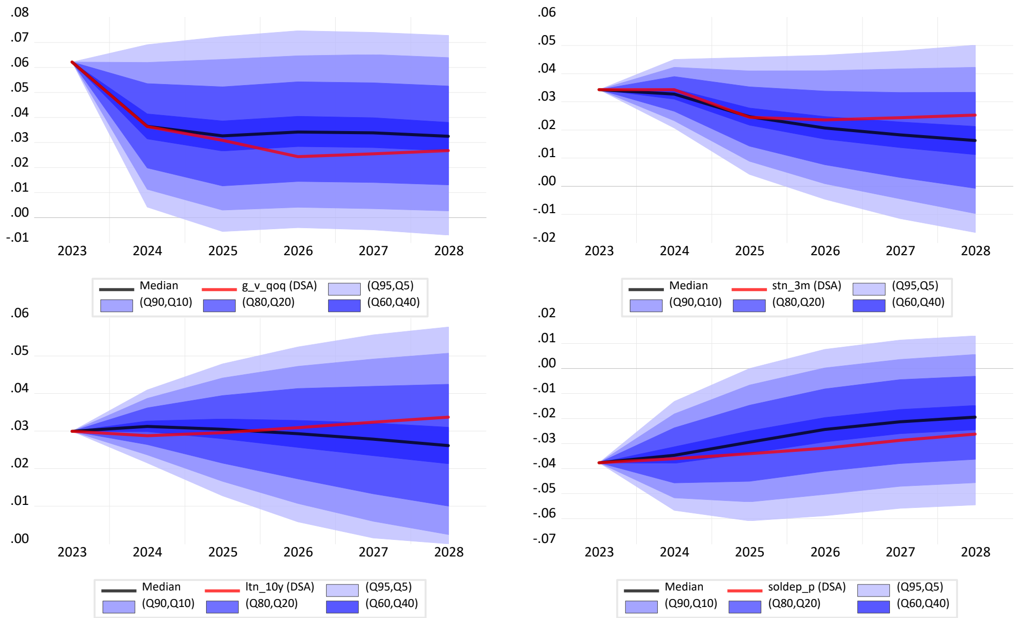

Figure A2: *Historical fan chart of the drivers (BVAR)*

Table A1: VAR estimation results (1991Q2 - 2023Q4)

	DTX_3M	SPREAD	G_V_QOQ	SOLDEP_P
DTX_3M(-1)	0.402^{**}	0.000	3.332	0.306
DTX_3M(-2)	0.508^{**}	-0.632^{**}	-1.274	-2.974^{+}
DTX_3M(-3)	-0.087	0.073	2.338	0.589
DTX_3M(-4)	0.031	-0.125	-0.112	0.074
SPREAD(-1)	0.485^{**}	0.904^{**}	3.351^{+}	1.046
SPREAD(-2)	-0.042	-0.502^{*}	-4.103	-2.596
SPREAD(-3)	-0.324	0.466^{*}	3.899	3.064
SPREAD(-4)	0.011	-0.046	-2.388	-0.522
G_V_QOQ(-1)	0.005	0.010	0.231	0.180
G_V_QOQ(-2)	0.034^{*}	-0.040^{**}	-0.342^{+}	-0.020
G_V_QOQ(-3)	0.010	-0.012	0.282	0.263^{*}
G_V_QOQ(-4)	0.007	-0.006	0.028	0.172^{+}
SOLDEP_P(-1)	0.001	-0.013	-0.692^{**}	0.321^{+}
SOLDEP_P(-2)	-0.039	0.068^{*}	0.759^{*}	0.644^{*}
SOLDEP_P(-3)	0.023	-0.024	-0.453	-0.177
SOLDEP_P(-4)	0.015	-0.035^{+}	0.270	0.017
C	0.001^{*}	0.001^{*}	0.003	-0.010^{**}
R-squared	0.389	0.837	0.188	0.751
Adj. R-squared	0.303	0.814	0.074	0.716
Sum sq. resids	0.000	0.000	0.028	0.015
S.E. equation	0.001	0.001	0.016	0.012
Log likelihood	714.121	705.824	368.090	407.185
Akaike AIC	−10.643	−10.516	−5.360	−5.957
Schwarz BIC	−10.270	−10.143	−4.987	−5.584
Jarque-Bera p-val.	0.000	0.463	0.000	0.107
Log likelihood		2361.711		
Akaike AIC		−35.018		
Schwarz BIC		−33.526		
LM(1) (p-val.)		8.63 (0.93)		
LM(4) (p-val.)		18.64 (0.29)		
Jarque-Bera (p-val.)		5761.14 (0.00)		
Number of coefficients		68		

Note: *Included observations: 131. Superscripts **, * and $^{+}$ denote significance at the 1%, 5% and 10% levels respectively. LM(k) is the LM test statistic of the null hypothesis of no serial correlation at order k.*