How to Conduct Joint Bayesian Inference in VAR models?[*]

Andrian Yambolov[†]

29 October 2024

Abstract. Performing series of simulation experiments using widely adopted Bayesian structural vector autoregressions, this paper demonstrates that conventional pointwise quantiles substantially understate uncertainty about impulse responses when economic analysis requires simultaneous inference across multiple variables and horizons. Since recently established joint inference methods provide meaningfully distinct error bands estimates, their performance is evaluated and improved by proposing calibration routines to achieve the nominal probability content. Two practical applications illustrate the implications of these findings: (i) within a structural vector autoregression, the fiscal multiplier exhibits error bands that are 51% to 91% wider than previous estimates, and (ii) a pseudo-out-of-sample projection exercise for inflation and GDP shows that joint inference methods could effectively summarize uncertainty for forecasts as well. These results underscore the importance of using joint inference methods for more robust econometric analysis.

Keywords: vector autoregressions, impulse responses, forecasts, pointwise inference, joint inference

JEL Codes: C22, C32, C52

[†]Goethe University Frankfurt. Email: yambolov.andrian@gmail.com.

1 Introduction

The impulse response function (IRF) is the most widely used tool for analyzing economic questions within a structural vector autoregression (VAR) framework. Regardless of the estimation methodology, whether frequentist or Bayesian, practitioners typically construct pointwise (individual) bands[1] around the impulse response coefficients at each horizon separately. However, most economic questions necessitate simultaneous evaluation across multiple structural parameters.

For example, a standard monetary structural VAR model could be used to examine whether the exchange rate overshoots in response to a monetary policy shock (see Eichenbaum and Evans (1995)). Alternatively, the same model could study the conditional exchange rate pass-through by computing the ratio between the responses of the price level and the exchange rate following a monetary policy shock (see Forbes et al. (2018)). While the former question involves a consideration of cross-horizon relationship in the impulse response of a single variable, the latter requires inference method that accounts for cross-variable dependence.

This example illustrates why error bands should be viewed as context-specific objects, whose proper specification requires thorough discussion before any empirical analysis. Lütkepohl et al. (2015, 2018) have evaluated the empirical properties of frequentist methods for joint inference in structural VARs using Monte Carlo experiments. These studies are comprehensive, covering a wide range of methods that have been discussed and assessed. Overall, Lütkepohl and his coauthors demonstrate that some of the methods they examine could be used to conduct joint inference across multiple parameters. However, the practical impact of these findings is limited, as the vast majority of empirical research currently employing VARs is relying on Bayesian estimation. This raises the main question addressed in the paper: how can joint error bands be constructed to achieve a pre-specified probability content in Bayesian VAR models?

In the Bayesian literature, two recently proposed methods for conducting joint inference have emerged, with significant implications due to their adaptability to various time series models. The sup-t estimator introduced by Montiel Olea and Plagborg-Møller (2019), employs numerical optimization to determine the optimal pointwise quantile that achieves joint coverage across multiple parameters. Another method independently introduced by Akram et al. (2016) and Inoue and Kilian (2022) treats the multivariate posterior draws as individual models satisfying structural restrictions. This method uses a vector-valued loss function to retain a specified share of draws (corresponding to the credibility level) with the lowest loss values. An envelope around the impulse response functions is then constructed by taking the minimum and maximum responses across horizons and structural parameters, a step that gives this method its name — min-max. While both methods are readily applicable in practice, their empirical properties have not been thoroughly evaluated, raising questions whether the error bands achieve the desired probability content. More importantly, the sup-t and the min-max estimator produce considerably different uncertainty estimates.

By conducting series of simulation experiments using four widely adopted structural VAR models, I demonstrate that conventional pointwise quantiles understate estimation uncertainty when the analysis requires joint inference across multiple parameters. Moreover, the sup-t estimator achieves the nominal probability content jointly across multiple parameters, but the performance of the min-max estimator depends on the specific vector-valued loss function,

---

[1]In the frequentist framework, these are confidence intervals, while in the Bayesian framework, they are credible intervals.

which is employed. In particular, the min-max estimator relying on absolute, quadratic, or angular loss functions, as suggested by Inoue and Kilian (2022), tends to overstate joint uncertainty and produces more conservative estimates of the error bands. Conversely, the min-max estimator using the Chebyshev loss function, as proposed by Akram et al. (2016), shows similar coverage to the sup-t estimator. A further important contribution of this paper is that, I enhance the min-max estimator under the three loss functions proposed by Inoue and Kilian (2022) by adding an additional calibration step. The error bands computed with the enhanced min-max estimator are subsequently shown to achieve the nominal probability content, but they are tighter in comparison to the error bands computed with the sup-t estimator. Finally, I also show that the min-max estimator preserves the covariance among structural parameters, which are within the joint error bands. This supports the argument put forward by Inoue and Kilian (2022), who motivate the usage of their method by its ability to capture the shape and comovement of impulse responses, rather than concentrating only on error bands that have the desired probability content.

The remainder of the paper is organized as follows. The next section reviews the relevant literature, also discussing the frequentist methods, as advancements in this framework have significantly contributed to enhancing the Bayesian procedures proposed in this study. Section 3 describes a general reduced-form model, along with structural identification methodologies within the Bayesian framework. It also provides a detailed explanation of the joint inference methods evaluated in this paper and introduces the experimental design, which follows the simulation-based calibration procedure recently proposed by Talts et al. (2020). To the best of my knowledge, this is the first paper to apply it in macroeconometrics. Its primary practical application lies in evaluating whether any Bayesian estimation routine properly samples from the underlying posterior distribution. Section 4 outlines the data-generating processes and presents the analysis of results. Section 5 highlights the practical contributions of the joint inference methods, showing that error bands for the fiscal multiplier computed with joint inference methods could be between 51% and 91% wider, depending on the selected credibility level, compared to conventional pointwise estimates. The second example demonstrates the similarities between impulse responses and forecasts by applying the joint inference methods to a pseudo-out-of-sample projection exercise for inflation and GDP in the US and euro area using two distinct time series models. Section 6 concludes the paper.

2 Literature Review

Joint inference methods have long been a focal point of interest for econometricians. For example, the Wald test is commonly used in reduced-form regression to assess whether a set of model parameters differs significantly from a specified value. Joint inference methods for structural VAR models can be thought of in a similar way, but they have two key dimensions: they can be used either to construct error bands that achieve the nominal coverage level or to capture the joint distribution among structural parameters. The methods discussed in this section address either the first aspect or both.

Sims and Zha (1999) were among the first to recognize the limitations of pointwise inference for impulse responses. Although their solution relies on Bayesian methods, most subsequent studies have employed frequentist techniques (e.g., Staszewska (2007), Jordà (2009), and Inoue and Kilian (2016)). The simplest technique for joint inference about impulse responses is presented in the well-known textbook by Lütkepohl (2005). It builds on Bonferroni's (1936)

method, which adjusts significance levels for multiple tests by dividing the overall significance level by the number of structural parameters considered. This adjustment ensures that the probability of making one or more Type I errors across multiple comparisons remains below the desired threshold.

Lütkepohl et al. (2015) provide a review of the most commonly used frequentist approaches for conducting joint inference about impulse responses. Their objective is to compare the coverage levels of pointwise and joint confidence bands through Monte Carlo simulations. They demonstrate that conventional confidence bands exhibit very low coverage levels in SVARs when the economic analysis requires joint inference across multiple structural parameters. Additionally, they find that Bonferroni error bands are overly conservative in summarizing joint uncertainty. To address this, Lütkepohl et al. (2015) propose an iterative procedure that discards draws from the bootstrap distribution until the resulting error bands achieve the nominal joint coverage level. This procedure has inspired the calibration techniques proposed in this study to enhance Bayesian inference methods. Lütkepohl et al. (2015) also show that the adjusted Bonferroni error bands, along with another method known as the neighboring paths procedure (Staszewska (2007)), systematically achieve the desired nominal joint coverage level. Another significant contribution of their study is the finding that Scheffé bands, as proposed by Jordà (2009), do not reach the prespecified confidence level.[2] Conversely, confidence bands adapted from the method by Wolf and Wunderli (2015) for constructing multi-horizon forecasts do achieve sufficient coverage but tend to be wider compared to other joint inference techniques.[3] In a similar Monte Carlo study, Lütkepohl et al. (2018) demonstrate that joint confidence bands, constructed by evaluating the density of structural models, are also preferable when the researcher seeks a coverage level corresponding to the prespecified confidence level.

The Bayesian literature focusing on joint inference about impulse response functions in VARs remains relatively scarce. Inoue and Kilian (2013) propose a method to estimate a joint credible region by ranking the posterior density of each structural draw. Initially, draws are generated from the joint posterior distribution of the reduced-form VAR parameters, and the posterior density value is computed. Subsequently, draws that do not satisfy the theoretically motivated sign restrictions on the impulse response functions are discarded, and the remaining draws are ranked by the value of their posterior density. According to Inoue and Kilian (2013), the draw that maximises the posterior density, known as the modal model, serves as an estimator of the central tendency of the impulse response functions. Furthermore, the $(1 - \alpha)100\%$ highest posterior density credible region of admissible models is used as an estimator for the inference about the impulse response functions.

In a more recent study, Montiel Olea and Plagborg-Møller (2019) propose the so-called sup-t method for conducting joint inference about impulse responses. Their main analysis is conducted within the frequentist framework[4], but they also propose a procedure that could be applied within Bayesian estimation. This estimator requires the use of numerical optimization to find the pointwise quantiles that achieve the nominal coverage level for all structural parameters simultaneously. Montiel Olea and Plagborg-Møller (2019) also compare their method with more traditional approaches for joint inference, such as the Bonferroni (1936) and Sidak (1967) error bands. The Šidák method, similar to Bonferroni, offers a slightly more powerful adjustment by assuming the tests are independent, and adjusts the significance level using a formula that accounts for this assumption. In anycase both methods are used only as a benchmark compar-

---

[2]The Scheffé is similar to the Bonferroni approach, but adjusts the critical value instead of the significance level.
[3]Both of these methods have only frequentist interpretations and are not discussed further here for brevity.
[4]The frequentist sup-t estimator is also suggested by Inoue et al. (2023) as a reliable method for conducting joint inference about impulse responses estimated with local projections.

ison to show that the sup-t method produces tighter error bands. Although this is the only study in the Bayesian framework that provides a simulation exercise to evaluate the empirical properties of a joint error band, their empirical model does not resemble any established VAR, but rather serves as a theoretical example. Furthermore, the simulation does not provide a fully Bayesian interpretation, as it is only valid for a single point estimate. A further drawback of this method is that it only provides error bands, preventing the applied econometrician from making probability statements about the conditional shape or comovements of endogenous variables.

Another approach for modeling joint inference within a generic Bayesian framework was originally proposed by Bernardo (2011). More recently, Inoue and Kilian (2022) adapted this technique to SVARs to conduct joint inference about impulse response functions. The method addresses both shortcomings associated with pointwise impulse responses. Compared to the joint inference method proposed earlier in Inoue and Kilian (2013), it is less computationally burdensome and, crucially, more easily adaptable to different SVAR models, without a need to specify the structural likelihood function, which might not be a trivial task.

In essence, for each posterior draw of the vector of structural parameters, Inoue and Kilian (2022) calculate the average distance to all other draws under a vector-valued loss function, and then retain the first $(1 - \alpha)100\%$ draws, starting with the draw having the lowest value. The constructed lowest posterior risk region visually approximates the joint credible region as the number of draws approaches infinity. The members of this region resemble the shotgun trajectories used in ballistics rather than the conventional credible band.[5] Finally, the joint error bands are computed by constructing the upper and lower envelopes around the shotgun trajectories within the credible region, which is referred to as the min-max estimator.

In a related study, Akram et al. (2016) generate projection fan charts for a Bayesian DSGE model using a simplified version of the methodology proposed by Inoue and Kilian (2022). Their approach does not depend on the average distance across draws, but rather on the distance of each draw from a pointwise measure of central tendency—specifically, the pointwise mean. Another key difference between the two studies is that Inoue and Kilian (2022) employ a variety of vector-valued loss functions, including absolute, quadratic, angular, and Dirac, while Akram et al. (2016) use only the Chebyshev specification. The choice of loss function proves to be crucial for the performance of the min-max estimator, as demonstrated in the latter sections of this paper.

Despite the clear disadvantages of partially ignoring estimation uncertainty due to interdependence among structural parameters during inference, joint error bands have rarely been used in applied research (e.g., Herrera and Rangaraju (2020), Zhou (2020), Kilian and Zhou (2022), Arias et al. (2022), and Güntner et al. (2024)). One possible explanation is that joint error bands yield significantly wider uncertainty estimates compared to conventional pointwise methods, yet their performance has not been thoroughly assessed in simulation studies.

---

[5]The shotgun trajectories were originally proposed in the context of frequentist methods by Inoue and Kilian (2016).

3 Methodology

3.1 Model

This subsection defines the reduced-form model and introduces identification methodologies used in the simulation experiments.

Reduced-form model. The data-generating models employed in the Monte Carlo simulations follow a general reduced-form VAR setup:

$$\mathbf{y}_t = \mathbf{c} + \sum_{l=1}^{p} \mathbf{A}_l \mathbf{y}_{t-l} + \mathbf{V}\mathbf{z}_t + \mathbf{u}_t, \quad \text{for } 1 \le t \le T, \tag{1}$$

where $\mathbf{y}_t$ is $(N \times 1)$ vector of endogenous variables, $\mathbf{A}_l$ are $(N \times N)$ matrices with parameters, $\mathbf{c}$ is $(N \times 1)$ vector of intercepts, $\mathbf{z}_t$ is $(G \times 1)$ vector of exogenous variables with its corresponding $(N \times G)$ matrix of parameters denoted as $\mathbf{V}$, N is the number of endogenous variables, G is the number of exogenous variables, p is the lag length and T is the sample size. All the models employed in this study are assumed to be stable, necessitating that $det(A(z)) \ne 0$, with the lag polynomial defined as $A(L) := \mathbf{I}_{N \times N} - \mathbf{A}_1 L - \cdots - \mathbf{A}_p L^p$ and an identity matrix $\mathbf{I}_N$ of size $(N \times N)$.

The companion form of the model can be written as

$$\mathbf{Y} = \mathbf{X}\mathbf{A} + \mathbf{u}, \tag{2}$$

where $\mathbf{Y} = (\mathbf{y}_1, \mathbf{y}_2, \ldots, \mathbf{y}_T)'$, $\mathbf{X} = (\mathbf{X}_1, \mathbf{X}_2, \ldots, \mathbf{X}_T)'$, $\mathbf{X}_t = (\mathbf{i}_N', \mathbf{y}_{t-1}', \ldots, \mathbf{y}_{t-p}', \mathbf{z}_t')'$, $\mathbf{A} = (\boldsymbol{\mu}, \mathbf{A}_1, \ldots, \mathbf{A}_p)'$, $\mathbf{u} = (\mathbf{u}_1, \mathbf{u}_2, \ldots, \mathbf{u}_T)'$, and $\mathbf{i}_N$ is $(N \times 1)$ vector of all ones.

Bayesian estimation. The reduced-form VAR parameters are specified with normal-inverse Wishart prior distribution following a conventional approach (e.g., Inoue and Kilian (2013))

$$vec(\mathbf{A})|\boldsymbol{\Sigma} \sim \mathcal{N}\left(vec(\bar{\mathbf{A}}_0), \boldsymbol{\Sigma} \otimes \mathbf{N}_0^{-1}\right) \quad \text{and} \quad \boldsymbol{\Sigma} \sim IW_n(\nu_0 \mathbf{S}_0, \nu_0), \tag{3}$$

where $\mathbf{N}_0$ is $(p \cdot N \times p \cdot N)$ positive definite matrix, $\mathbf{S}_0$ is $(N \times N)$ covariance matrix and $\nu_0 > 0$. This model implies a posterior distribution of the same functional family as the prior given by

$$vec(\mathbf{A})|\boldsymbol{\Sigma} \sim \mathcal{N}\left(vec(\bar{\mathbf{A}}_T), \boldsymbol{\Sigma} \otimes \mathbf{N}_T^{-1}\right) \quad \text{and} \quad \boldsymbol{\Sigma} \sim IW_n(\nu_T \mathbf{S}_T, \nu_T), \tag{4}$$

where

$$\nu_T = T + \nu, \quad \mathbf{N}_T = \mathbf{N}_0 + \mathbf{X}'\mathbf{X}, \quad \bar{\mathbf{A}}_T = \mathbf{N}_T^{-1}(\mathbf{N}_0 \bar{\mathbf{A}}_0 + \mathbf{X}'\mathbf{X}\widehat{\mathbf{A}}),$$

$$\mathbf{S}_T = \frac{\nu_0}{\nu_T}\mathbf{S}_0 + \frac{T}{\nu_T}\widehat{\boldsymbol{\Sigma}} + \frac{1}{\nu_T}(\widehat{\mathbf{A}} - \bar{\mathbf{A}}_0)'\mathbf{N}_0 \mathbf{N}_T^{-1}\mathbf{X}'\mathbf{X}(\widehat{\mathbf{A}} - \bar{\mathbf{A}}_0),$$

$$\widehat{\mathbf{A}} = (\mathbf{X}'\mathbf{X})^{-1}\mathbf{X}'\mathbf{Y} \quad \text{and} \quad \widehat{\boldsymbol{\Sigma}} = (\mathbf{Y} - \mathbf{X}\widehat{\mathbf{A}})'(\mathbf{Y} - \mathbf{X}\widehat{\mathbf{A}})/T.$$

Throughout this paper, two distinct parametrizations of the prior distribution are considered: a "non-informative" prior, characterized by:

$$\nu_0 = 0, \quad \mathbf{S}_0 = \mathbf{O}_{N \times N}, \quad \bar{\mathbf{A}}_0 = \mathbf{O}_{p \cdot N + R \times N}, \quad \text{and} \quad \mathbf{N}_0^{-1} = \mathbf{O}_{p \cdot N + R \times p \cdot N + R},$$

where $\mathbf{O}$ denotes zero matrix with size indicated by the subindex. The second prior parametrization, referred to as a "unit root," is defined by:

$$\nu_0 = N + 2, \quad \mathbf{S}_0 = \begin{pmatrix} \widetilde{\sigma}_1^2 & & \\ & \ddots & \\ & & \widetilde{\sigma}_p^2 \end{pmatrix}, \quad \bar{\mathbf{A}}_0 = \begin{pmatrix} \mathbf{O}'_{R \times N} & \mathbf{I}'_{N \times N} & \mathbf{O}'_{p \cdot (N-1) \times N} \end{pmatrix}', \quad \text{and}$$

$$\mathbf{N}_0^{-1} = \begin{pmatrix} \omega \mathbf{I}_{R \times R} & 1^{-\alpha} \lambda \mathbf{S}_0^{-1} & 2^{-\alpha} \lambda \mathbf{S}_0^{-1} & \dots & p^{-\alpha} \lambda \mathbf{S}_0^{-1} \end{pmatrix}',$$

where $\omega = 10^5$, $\lambda = 0.2$, $\alpha = 2$, and $\widetilde{\sigma}_i^2$ is the error's variance of the regression $y_{t,\cdot} = \gamma_0 + \gamma_1 y_{t-1,\cdot} + e_t$, which is estimated using OLS.

Structural Model. The computation of the impulse response coefficients requires first defining the following companion form

$$\underbrace{\begin{pmatrix} \mathbf{y}_t \\ \mathbf{y}_{t-1} \\ \vdots \\ \mathbf{y}_{t-p} \end{pmatrix}}_{\mathbf{Y}_t} = \underbrace{\begin{pmatrix} \mathbf{A}_1 & \mathbf{A}_2 & \dots & \mathbf{A}_{p-1} & \mathbf{A}_p \\ \mathbf{I}_{N \times N} & \mathbf{O}_{N \times N} & & \mathbf{O}_{N \times N} & \mathbf{O}_{N \times N} \\ \mathbf{O}_{N \times N} & \mathbf{I}_{N \times N} & & \mathbf{O}_{N \times N} & \mathbf{O}_{N \times N} \\ \vdots & & \ddots & \vdots & \vdots \\ \mathbf{O}_{N \times N} & \mathbf{O}_{N \times N} & \dots & \mathbf{I}_{N \times N} & \mathbf{O}_{N \times N} \end{pmatrix}}_{\widetilde{\mathbf{A}}} \underbrace{\begin{pmatrix} \mathbf{y}_{t-1} \\ \mathbf{y}_{t-2} \\ \vdots \\ \mathbf{y}_{t-p-1} \end{pmatrix}}_{\mathbf{Y}_{t-1}} + \underbrace{\begin{pmatrix} \mathbf{u}_t \\ \mathbf{0}_N \\ \vdots \\ \mathbf{0}_N \end{pmatrix}}_{\mathbf{U}_t}.$$

It is well known that by successive substitution for $\mathbf{Y}_{t-l}$ the definition of the companion form reduces to

$$\mathbf{Y}_{t-l} = \mathbf{A}^{l+1} \mathbf{Y}_{t-1} + \sum_{b=0}^{l} \mathbf{A}^b \mathbf{U}_{t+l-b}$$

and left-multiplying by $\mathbf{J} = \begin{pmatrix} \mathbf{I}_{N \times N} & \mathbf{0}_{N \times (p-1)N} \end{pmatrix}$ results in:

$$\mathbf{y}_{t+l} = \mathbf{J} \mathbf{A}^{l+1} \mathbf{Y}_{t-1} + \sum_{b=0}^{l} \mathbf{J} \mathbf{A}^b \mathbf{J}' \mathbf{u}_{t+l-b}.$$

Finally, the n-th row and s-th column element $\theta_{n,s,h}$ of the impulse response coefficient matrix $\mathbf{\Theta}_h$, which quantifies the effect on variable n in response to a structural shock s of size one in period h, is defined as:

$$\mathbf{\Theta}_h = \begin{pmatrix} \theta_{n,s,h} \end{pmatrix} = \mathbf{J} \mathbf{A}^h \mathbf{J}' \mathbf{B},$$

where $\mathbf{B}$ is the structural impact matrix. A linear transformation on the reduced-form errors defines the structural shocks as $\varepsilon_t = \mathbf{B}^{-1} \mathbf{u}_t$, which are assumed to be contemporaneously uncorrelated and have unit variance. Therefore, the relationship $\mathbf{BB}' = \mathbf{\Sigma}$ must hold. However, the matrix $\mathbf{B}$ is not unique, and identification requires additional assumptions. In the simulation exercises, four different identification methodologies are considered. The first is recursive identification resulting in structural impact matrix $\mathbf{B}^{-1} = \text{chol}(\mathbf{\Sigma})$, i.e., the lower triangular matrix of the Cholesky decomposition. The second identification strategy relies on sign restrictions motivated by economic theory. To this end, for each posterior draw, i, of the error's covariance matrix, the lower triangular matrix of the Cholesky decomposition $\text{chol}(\mathbf{\Sigma}_i)$ is calculated and then an orthogonal matrix, $\mathbf{Q}_i$, is drawn such that the structural impact matrix given by $\mathbf{B}_i^{-1} = \text{chol}(\mathbf{\Sigma}_i) \mathbf{Q}_i$ implies impulse responses, which are kept only if they satisfy the

imposed sign restrictions. The described procedure is repeated L times for each posterior draw of the reduced-form model. The third methodology requires an instrumental variable, $\mathbf{z}_t$, which enters the reduced-form VAR as exogenous variable and the column vector $\mathbf{V}$ gives the first column of the structural impact matrix, $\mathbf{B}^{-1}$, while the remaining columns have zero in all entries.[6] The term $\mathbf{V}\mathbf{z}_t$ is included in the reduced-form VAR setup only for models, which rely on instrumental variables for the identification. The fourth and last identification scheme considered in the Monte Carlo simulations imposes a short-run and long-run restrictions on the impulse responses. The implementation follows the methodology of Blanchard (1989) that requires stationary vector of endogenous variables, (i.e., $\mathbf{y}_t \sim I(0)$). Then defining the matrix $\mathbf{\Omega} = \mathbf{J}(\mathbf{I}_{pN \times pN} - \widetilde{\mathbf{A}})^{-1}\mathbf{J}'$ allows to derive the structural impact matrix as $\mathbf{B} = \mathbf{\Omega}^{-1}chol(\mathbf{\Omega}\mathbf{\Sigma}\mathbf{\Omega}')$.

3.2   Conducting Inference

In this subsection, the uncertainty concepts are defined first, followed by an explanation of the construction of the error bands using all inference methods evaluated in this study.

Pointwise Inference. Pointwise inference about impulse responses is conducted by constructing credible interval that achieve predefined probability content for each structural parameter individually. For the sake of simplicity, a generalized vector of stacked structural parameters $\boldsymbol{\theta} \in \mathbb{R}^K$ is defined.

Definition 1 (Credible interval). The $(1 - \alpha)100\%$ credible interval, $\mathbf{\Theta}_j^{1-\alpha}$, for any structural parameter $\theta_j \in \mathbb{R}$ is an interval with lower and upper bounds given by $\left[\underline{\theta}_j, \bar{\theta}_j\right]$ and coverage probability $1 - \alpha$ such that

$$Pr\left(\theta_j \in \mathbf{\Theta}_j^{1-\alpha}\right) = Pr\left(\underline{\theta}_j^{1-\alpha} \leq \theta_j \leq \bar{\theta}_j^{1-\alpha}\right) = 1 - \alpha.$$

This credible interval is expected to include the actual value of θ_j with a probability $1 - \alpha$. The conventional approach of constructing the error bands is to compute the corresponding quantiles over all posterior draws.

Definition 2 (Conventional (pointwise) error band). The $(1 - \alpha)100\%$ conventional (pointwise) error band, $\widehat{\mathbf{\Theta}}_j^{1-\alpha,PW}$, for any structural parameter, $\theta_j \in \mathbb{R}$, is an interval given by

$$\widehat{\mathbf{\Theta}}_j^{1-\alpha,PW} = \left[\underline{\theta}_j^{1-\alpha,PW}, \bar{\theta}_j^{1-\alpha,PW}\right],$$

where $\underline{\theta}_j^{1-\alpha,PW}$ and $\bar{\theta}_j^{1-\alpha,PW}$ are the individual $(\alpha/2)100$-th and $(1 - \alpha/2)100$-th percentiles, respectively.

Joint Inference. Conducting joint inference results in error bands that must achieve a predefined probability content across multiple structural parameters simultaneously. Unlike previous studies, which apply joint inference across all available variables, horizons, and shocks (e.g., Inoue and Kilian (2016, 2022)) or solely across horizons for each impulse response function (e.g., Lütkepohl et al. (2015, 2018) and Montiel Olea and Plagborg-Møller (2019)), this paper argues that error bands are context-specific objects. The selection of variables included in the joint error bands should be guided by careful consideration of the specific economic question. For instance, evaluating different pricing theories in international trade (such as dominant currency,

---

[6]This instrumental variable identification methodology is discussed in details by Paul (2020) and has the usual requirements, namely relevancy with respect to the structural shock of interest and also exogeneity with respect to the remaining structural shocks.

producer currency, or local currency pricing) requires examining the simultaneous response of import and export prices to a foreign monetary policy shock across individual horizons. On the other hand, more frequently, IRF analysis focuses on the persistence of the response, like Jarociński and Karadi (2020), who argue that "real GDP and the price level both decline persistently by about 10 basis points and 5 basis points, respectively". This example implies the simultaneous evaluation of IRFs across horizons. In practice, any combination of both types of inference can appear. For example, the computation of a fiscal multiplier, often defined as the cumulative sum of the impulse response of GDP divided by that of government expenditure, requires simultaneous inference of two variables across several horizons. Furthermore, if the SVAR identifies several shocks, such as in Blanchard (1989), then simultaneous inference about a positive demand shock and a favorable supply shock on unemployment and output must be carried out, considering the responses of two variables to two shocks across several horizons. Therefore, depending on the parameters targeted by the error bands, four possible types of joint inference and their corresponding vectors of stacked structural parameters are considered. The definitions that follow are inspired by Akram et al. (2016). First, the construction of error bands that achieve predefined probability content simultaneously for all variables at a specific horizon requires a multivariate (MV) vector of structural parameters.

Definition 3 (MV vector of structural parameters). Let $\boldsymbol{\theta} \in \mathbb{R}^K$ denotes the MV vector of structural parameters, defined as

$$\boldsymbol{\theta} = (\theta_{1,s,h}, \theta_{2,s,h}, \ldots, \theta_{n,s,h})' \quad \text{for} \quad s = 1, \ldots, S \text{ and } h = 1, \ldots, H.$$

In the same vein, the time-simultaneous (TS) vector of structural parameters is used by a practitioner targeting error bands that achieve predefined probability content simultaneously for all horizons of a given variable.

Definition 4 (TS vector of structural parameters). Let $\boldsymbol{\theta} \in \mathbb{R}^K$ denotes the TS vector of structural parameters, defined as

$$\boldsymbol{\theta} = (\theta_{n,s,1}, \theta_{n,s,2}, \ldots, \theta_{n,s,H})' \quad \text{for} \quad n = 1, \ldots, N \text{ and } s = 1, \ldots, S.$$

A natural extension of the previous two cases occurs when the error bands must achieve predefined probability content simultaneously for the impulse responses at all horizons for every variable reacting to a single shock. This requires a multivariate and time-simultaneous (MV-TS) vector of structural parameters.

Definition 5 (MV-TS vector of structural parameters). Let $\boldsymbol{\theta} \in \mathbb{R}^K$ denotes the MV-TS vector of structural parameters, defined as

$$\boldsymbol{\theta} = \text{vec} \begin{pmatrix} \theta_{1,s,1} & \theta_{2,s,1} & \ldots & \theta_{N,s,1} \\ \theta_{1,s,2} & \theta_{2,s,2} & \ldots & \theta_{N,s,2} \\ \vdots & \vdots & & \vdots \\ \theta_{1,s,H} & \theta_{2,s,H} & \ldots & \theta_{N,s,H} \end{pmatrix} \quad \text{for} \quad s = 1, \ldots, S.$$

In the most comprehensive case, when the economic question requires error bands that achieve predefined probability content simultaneously at all horizons for all variables reacting to all shocks, it is necessary to construct a multivariate, time-simultaneous and cross-shock (MV-TS-CS) vector of structural parameters.

Definition 6 (MV-TS-CS vector of structural parameters). Let $\boldsymbol{\theta} \in \mathbb{R}^K$ denote the MV-TS-CS vector of structural parameters, defined as

$$
\boldsymbol{\theta} = \text{vec} \begin{pmatrix} \boldsymbol{\theta}_{1,1} & \boldsymbol{\theta}_{2,1} & \dots & \boldsymbol{\theta}_{N,1} \\ \boldsymbol{\theta}_{1,2} & \boldsymbol{\theta}_{2,2} & \dots & \boldsymbol{\theta}_{N,2} \\ \vdots & \vdots & & \vdots \\ \boldsymbol{\theta}_{1,S} & \boldsymbol{\theta}_{2,S} & \dots & \boldsymbol{\theta}_{N,S} \end{pmatrix},
$$

where $\boldsymbol{\theta}_{n,s} = (\theta_{n,s,1}, \theta_{n,s,2}, \dots, \theta_{n,s,H})'$.

Definition 7 (Joint credible region). The $(1-\alpha)100\%$ joint credible region, $\boldsymbol{\Theta}^{1-\alpha}$, for any vector of structural parameters, $\boldsymbol{\theta} \in \mathbb{R}^K$, is a Cartesian product given by

$$
\boldsymbol{\Theta}^{1-\alpha} = \bigtimes_{j=1}^{K} \left[\underline{\theta}_j^{1-\alpha}, \bar{\theta}_j^{1-\alpha} \right]
$$

and coverage probability $1-\alpha$ such that

$$
Pr\left(\boldsymbol{\theta} \in \boldsymbol{\Theta}^{1-\alpha} \right) = Pr\left(\bigcap_{j=1}^{K} \left(\underline{\theta}_j^{1-\alpha} \le \theta_j \le \bar{\theta}_j^{1-\alpha} \right) \right) = 1-\alpha.
$$

Bonferroni's (1936) inequality is the standard approach that relates joint probabilities to marginal probabilities. Its application to constructing joint confidence regions for forecasts in VARs is discussed in the seminal textbook by Lütkepohl (2005). In terms of the definitions of the credible interval and region, the inequality implies that

$$
Pr\left(\boldsymbol{\theta} \in \boldsymbol{\Theta}^{1-\alpha} \right) \ge 1 - \sum_{j=1}^{K} Pr\left(\theta_j \notin \boldsymbol{\Theta}_j^{1-\alpha} \right).
$$

This inequality can be used to construct joint error bands. For example, a 95% TS error band for a variable across 5 horizons requires the marginal probability of each structural parameter to be 0.99, since $5 \cdot (1 - 0.99) = 0.05$. Because the relationship is an inequality, the proposed error band has at least 95% probability content, but it may be more conservative. Additionally, when a large number of structural parameters are of interest, the bounds can become uninformative.

Definition 8 (Bonferroni error band). The $(1-\alpha)100\%$ Bonferroni error band, $\widehat{\boldsymbol{\Theta}}^{1-\alpha,B}$, for the vector of structural parameters, $\boldsymbol{\theta} \in \mathbb{R}^K$, is a Cartesian product given by

$$
\widehat{\boldsymbol{\Theta}}^{1-\alpha,B} = \bigtimes_{j=1}^{K} \left[\underline{\theta}_j^{1-\alpha/K,PW}, \bar{\theta}_j^{1-\alpha/K,PW} \right].
$$

Similarly to the Bonferroni's approach, Sidak (1967) proposes another correction method, which assumes that all parameters of interest are mutually independent. This method relies on the following inequality:

$$
Pr\left(\boldsymbol{\theta} \in \boldsymbol{\Theta}^{1-\alpha} \right) \ge 1 - \prod_{j=1}^{K} Pr\left(\theta_j \notin \boldsymbol{\Theta}_j^{1-\alpha} \right).
$$

This correction approach shares all the drawbacks of the previous method, but it is less conservative for a small number of structural coefficients. The respective credible region is defined similarly.

Definition 9 (Šidák error band). The $(1 - \alpha)100\%$ Šidák error band, $\widehat{\Theta}^{1-\alpha,S}$, for any vector of structural parameters, $\boldsymbol{\theta} \in \mathbb{R}^K$, is a Cartesian product given by

$$\widehat{\Theta}^{1-\alpha,S} = \underset{j=1}{\overset{K}{\times}} \left[\underline{\theta}_j^{(1-\alpha)^{1/K},PW}, \bar{\theta}_j^{(1-\alpha)^{1/K},PW} \right].$$

Montiel Olea and Plagborg-Møller (2019) propose constructing joint error bands using the sup-t method, which builds on the previous two approaches. Although their primary analysis is in the frequentist framework, relying on parametric solutions, they also propose a methodology tailored to the Bayesian framework. Specifically, Montiel Olea and Plagborg-Møller (2019) suggest using numerical optimization to solve for the pointwise quantile that achieves the joint nominal coverage level of the posterior draws.[7] The proper definition of the sup-t credible region relies on the notation of the posterior draws of the structural parameters, which, for D draws, is given by $\theta_j^{(1)}, \ldots, \theta_j^{(D)}$. Montiel Olea and Plagborg-Møller (2019) refer to the optimization procedure as calibration, while Lütkepohl et al. (2015) call a similar process adjustment. In this paper, the former terminology is adhered to. Since additional calibration methods will be introduced in the following paragraphs, each will be given a distinct name. The calibration method related to the sup-t will be referred to as pointwise quantile optimization (PQO).

Definition 10 (Sup-t (PQO) error band). The $(1 - \alpha)100\%$ sup-t (PQO) error band, $\widehat{\Theta}^{1-\alpha,Su}$, for any vector of structural parameters, $\boldsymbol{\theta} \in \mathbb{R}^K$, is is a Cartesian product given by

$$\widehat{\Theta}^{1-\alpha,Su} = \underset{j=1}{\overset{K}{\times}} \left[\underline{\theta}_j^{1-\widehat{\alpha},PW}, \bar{\theta}_j^{1-\widehat{\alpha},PW} \right],$$

where

$$\widehat{\alpha} = sup \left\{ \widehat{\alpha} \in [\alpha/(2K), \alpha/2] \,\middle|\, D^{-1} \sum_{d=1}^{D} \mathbb{1} \left(\boldsymbol{\theta}^{(d)} \in \underset{j=1}{\overset{K}{\times}} \left[\underline{\theta}_j^{1-\widehat{\alpha},PW}, \bar{\theta}_j^{1-\widehat{\alpha},PW} \right] \right) \geq 1 - \alpha \right\}$$

and $\mathbb{1}(\cdot)$ denotes indicator function.

Akram et al. (2016) and Inoue and Kilian (2022) propose two related methods based on vector-valued loss functions to construct joint error bands. In the former study, Akram et al. (2016) evaluate the distance of each posterior draw to a pointwise measure of central tendency, specifically the mean, while in the latter study, Inoue and Kilian (2022) evaluate the average distance of each posterior draw to every other draw. The main results in this paper focus on the proposal by Akram et al. (2016).

Definition 11 (Lowest posterior risk joint credible set). The $(1 - \alpha)100\%$ vector of structural parameters draws, $\boldsymbol{\theta}^{(d)} \in \mathbb{R}^K$, which are closest to a specified pointwise central tendency measure, $\widetilde{\boldsymbol{\theta}} \in \mathbb{R}^K$, under a vector-valued loss function, $L\left(\boldsymbol{\theta}^{(d)}, \widetilde{\boldsymbol{\theta}}\right)$ are members of the lowest posterior risk joint credible set denoted as $\widehat{\Theta}_L^{1-\alpha}$.

Definition 12 (Min-max error band). The $(1 - \alpha)100\%$ closest vector of structural parameters draws, $\boldsymbol{\theta}^{(d)} \in \mathbb{R}^K$, to any pointwise central tendency measure, $\widetilde{\boldsymbol{\theta}} \in \mathbb{R}^K$, evaluated under a

---

vector-valued loss function, L, is a Cartesian product given by

$$\widehat{\Theta}^{1-\alpha,M} = \underset{j=1}{\overset{K}{\times}} \left[\underline{\theta}_j^{1-\alpha,M}, \bar{\theta}_j^{1-\alpha,M} \right],$$

where

$$\underline{\theta}_j^{1-\alpha,M} = min \left\{ \theta_j^{1-\alpha,M} \middle| \boldsymbol{\theta}^{1-\alpha,M} \in \widehat{\Theta}_L^{1-\alpha} \right\} \quad and \quad \bar{\theta}_j^{1-\alpha,M} = max \left\{ \theta_j^{1-\alpha,M} \middle| \boldsymbol{\theta}^{1-\alpha,M} \in \widehat{\Theta}_L^{1-\alpha} \right\}.$$

The preliminary results indicate that this method of constructing joint error bands may significantly overestimate the nominal credibility level. Inoue and Kilian (2022) are aware of this possibility and suggest the potential use of a calibration technique similar to the proposition of Montiel Olea and Plagborg-Møller (2019). However, their study does not include a simulation exercise or a details on how to implement the calibration step. Two possible routines to optimize the coverage level of the min-max error bands are explored in this paper. The first approach is to optimize the percentage of retained draws of the lowest posterior risk set, ensuring that their effective coverage level is at least equal to the nominal credibility level. This approach is referred to as loss quantile optimization (LQO).

Definition 13 (Min-max (LQO) error band). The $(1-\alpha)100\%$ min-max (LQO) error band, $\widehat{\Theta}^{1-\alpha,M(LQO)}$, for any vector of structural parameters, $\boldsymbol{\theta} \in \mathbb{R}^K$, is a Cartesian product given by

$$\widehat{\Theta}^{1-\alpha,M(LQO)} = \underset{j=1}{\overset{K}{\times}} \left[\underline{\theta}_j^{1-\widehat{\alpha},M}, \bar{\theta}_j^{1-\widehat{\alpha},M} \right],$$

where

$$\widehat{\alpha} = sup \left\{ \widehat{\alpha} \in (\alpha, 1) \middle| D^{-1} \sum_{d=1}^{D} \mathbb{1} \left(\boldsymbol{\theta}^{(d)} \in \underset{j=1}{\overset{K}{\times}} \left[\underline{\theta}_j^{1-\widehat{\alpha},M}, \bar{\theta}_j^{1-\widehat{\alpha},M} \right] \right) \geq 1-\alpha \right\}.$$

The second optimization method applied to the min-max error bands is an iterative procedure proposed by Lütkepohl et al. (2015). After constructing the lowest posterior risk set along with the min-max error band, all draws on the boundary of the hyperrectangular box are discarded until the posterior draws within the new error band achieve the nominal credibility level. This approach is referred to as boundary draw rejection (BDR).

Definition 14 (Min-max (BDR) error band). The $(1-\alpha)100\%$ min-max (BDR) error band, $\widehat{\Theta}^{1-\alpha,M(LQO)}$, for any vector of structural parameters, $\boldsymbol{\theta} \in \mathbb{R}^K$, is a Cartesian product given by

$$\widehat{\Theta}^{1-\alpha,M(BDR)} = \underset{j=1}{\overset{K}{\times}} \left[\underline{\theta}_j^{1-\alpha,M(BDR)}, \bar{\theta}_j^{1-\alpha,M(BDR)} \right] \quad s.t.$$

$$D^{-1} \sum_{d=1}^{D} \mathbb{1} \left(\boldsymbol{\theta}^{(d)} \in \underset{j=1}{\overset{K}{\times}} \left[\underline{\theta}_j^{1-\alpha,M(BDR)}, \bar{\theta}_j^{1-\alpha,M(BDR)} \right] \right) = 1-\alpha,$$

where $\underline{\theta}^{1-\alpha,M(BDR)} \in \check{\Theta}_L^{1-\alpha}$ and $\bar{\theta}^{1-\alpha,M(BDR)} \in \check{\Theta}_L^{1-\alpha}$. The set $\check{\Theta}_L^{1-\alpha}$ is obtained as draws which belong to the boundaries of the error bands of the min-max estimator are iteratively discarded from the lowest posterior risk joint credible set, $\widehat{\Theta}_L^{1-\alpha}$, until the coverage level condition is satisfied.

The min-max estimator relies on specification of a vector-valued loss function to calculate the distance of each draw from a measure of a central tendency. The functional forms considered in this study are: absolute, quadratic, angular, and Chebyshev. The first three functions are proposed by Inoue and Kilian (2022), while the last is suggested by Akram et al. (2016). Inoue and Kilian (2022) also propose a Dirac loss function, but it is non-separable, unlike the others, making it more complex to derive the credible region for various models and it has limited practical relevance. The formal definitions of the vector-valued loss functions are provided below.

$$L(\boldsymbol{\theta}, \widetilde{\boldsymbol{\theta}}) = \begin{cases} \sum_{j=1}^{K} |\theta_j - \widetilde{\theta}_j| & \text{(absolute loss)} \\ \sum_{j=1}^{K} (\theta_j - \widetilde{\theta}_j)^2 & \text{(quadratic loss)} \\ \frac{1}{\pi NS} \sum_{i=1}^{N} \sum_{j=1}^{S} cos^{-1} \left(\frac{\boldsymbol{\theta}_{i,j} \cdot \widetilde{\boldsymbol{\theta}}_{i,j}}{\|\boldsymbol{\theta}_{i,j}\| \|\widetilde{\boldsymbol{\theta}}_{i,j}\|} \right) & \text{(angular loss)} \\ \max_{j} |\theta_j - \widetilde{\theta}_j| & \text{(Chebyshev loss)} \end{cases}$$

Since the pointwise median is the estimator that solves the real-valued absolute loss function, it is used as a measure of central tendency for the first vector-valued loss function. Similarly, the mean solves the real-valued quadratic loss function, and thus, it is used as the pointwise measure for the second function. The latter two loss functions do not have such pointwise estimators, so the pointwise median is used because it is the most frequently employed measure.

Inoue and Kilian (2022) advocate for the use of the angular loss function because it is independent of the scale of the endogenous variables, which is a highly appealing property in the context of SVARs. Similarly, the Chebyshev loss function is applied only to draws that are normalized pointwise, as the maximum distance across different structural parameters is meaningless otherwise. Therefore, this loss function is also scale-invariant, even though this property is imposed in a different way.

3.3 Illustrative Example

This section visually compares the proposed joint error bands and the conventional pointwise error bands for a bivariate normally distributed toy parameter space, which is assumed to represent a posterior distribution given by:

$$\begin{pmatrix} \theta_1 \\ \theta_2 \end{pmatrix} \sim \mathcal{N} \left(\begin{pmatrix} 0 \\ 0 \end{pmatrix}, \begin{pmatrix} 1 & 0 \\ 0 & 0.25 \end{pmatrix} \right).$$

The estimates of the error bands are computed for 20,000 draws and plotted in Figure 1. As expected, the pointwise error bands is the smallest rectangle and is contained within all joint error bands, suggesting that this method provides the least conservative estimates of uncertainty. On the other hand, the Bonferroni, Sidák, and sup-t error bands nearly coincide for both parameters, which is not necessarily the case for parameter space of interest with more parameters. For the sake of simplicity, the analysis considers only the absolute loss function for the class of min-max estimators. The non-calibrated min-max error bands encompass all other error bands, which implies that it is the most conservative estimator. In comparison, the min-max (LQO) error bands are tighter and even coincide with the pointwise error bands for θ_1. However, the area enclosed within the box of the min-max (BDR) error bands is smaller, making it a more precise estimator. These conclusions hold true for both the 68% credible

region and the 90% credible region, although the error bands for the 90% credible region are more dispersed, reflecting the higher credibility level.

Figure 1
Pointwise and joint inference about toy parameters
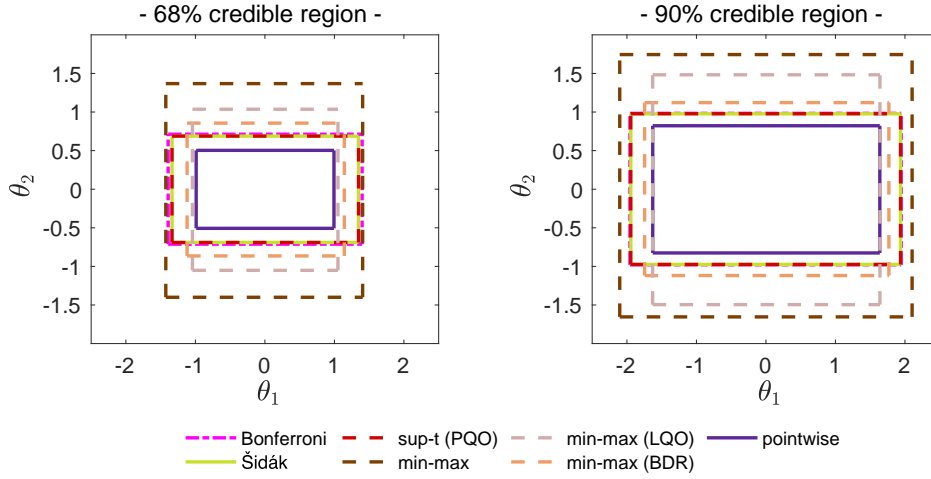
Notes: The squares depict pointwise and joint error bands for the toy parameters. The min-max estimators are computed only with absolute loss function. The figure on the left side shows results for 68% credible region and the figure on the right side shows results for 90% credible region.

The joint inference methods are formally accessed by comparing the relative difference between the width of the joint error bands and the pointwise error bands as

$$\Delta = \left(H\left((\bar{\boldsymbol{\theta}}^{1-\alpha,e} - \underline{\boldsymbol{\theta}}^{1-\alpha,e}) \oslash (\bar{\boldsymbol{\theta}}^{1-\alpha,PW} - \underline{\boldsymbol{\theta}}^{1-\alpha,PW}) \right) - 1 \right) \cdot 100 \quad \text{for } e \in \{B, S, Su, M, M(LQO), M(BDR)\},$$

where $H(\cdot)$ is the harmonic mean and $\oslash$ denotes the element-wise division. The harmonic mean is selected here as it is the most appropriate measure of the average for ratios.

First, the error bands are computed for a range of covariances in the interval $[-0.45, 0.45]$, which implies correlations in the interval $[-0.9, 0.9]$ for the toy parameter space. The relative width of the error bands against the correlation between the parameters is plotted in Figure 2. The 68% non-calibrated min-max error band is between 70% and 90% wider than the corresponding pointwise error band, but for the 90% credible interval, this difference falls to the range of 40%-70%. The min-max (LQO) method produces tighter error bands but is strongly dominated by the sup-t and min-max (BDR) methods. The error bands corresponding to the latter two methods nearly coincide being between 18% and 35% wider for the 68% credible region and between 10% and 18% wider for the 90% credible region compared to the pointwise error bands. The relative width of the min-max error bands and, to a certain extent, the min-max (LQO) error bands does not change smoothly with marginal increases in the correlation among the toy parameters. This is caused by the variability in the error bands of roughly 10% for simulations generated with different seeds of the random number generator. This issue could not be resolved by increasing the number of draws from the posterior distribution. Consequently, the min-max error bands could be less appealing for policy analysis, but the results show that the BDR method does not exhibit the same problem.

Figure 2
Relative width of the joint error bands

Notes: The lines depict relative width of the joint error bands in comparison to the pointwise error bands against correlation of the toy parameters. The min-max class of estimators are computed only with absolute loss function. The figure on the left side shows results for the 68% credible region and the figure on the right side shows results for the 90% credible region.
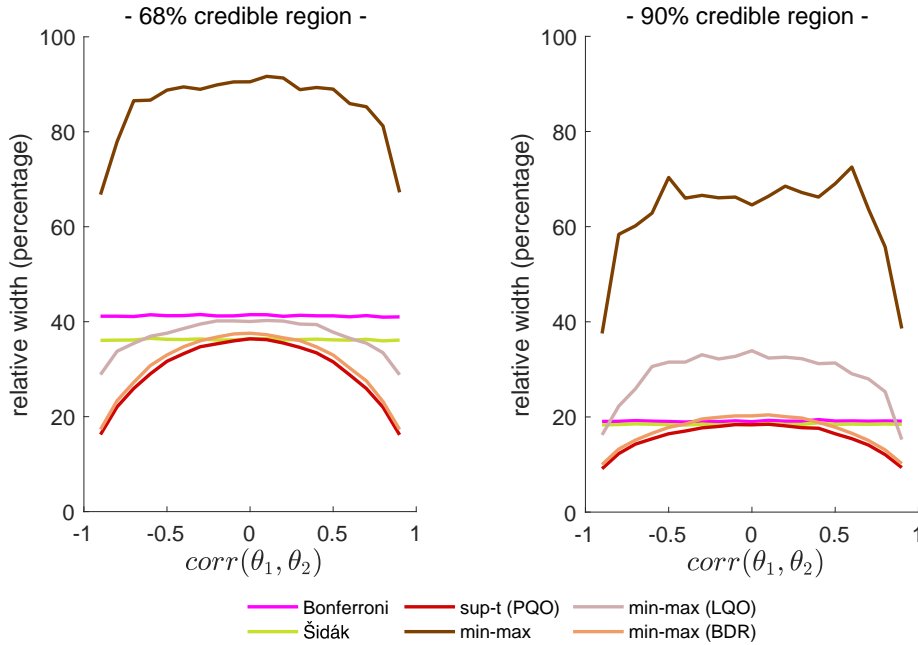
In summary, the difference between the width of the joint error bands and the pointwise error bands could be considered significant, but it is not certain that this is an argument to overlook the construction of proper error bands. Except the Bonferroni's and Sidák's error bands, all other joint inference methods produce error bands that decrease with the absolute value of the correlation between the parameters. Since VARs usually have correlated structural parameters, this strengthens the case for employing joint error bands. From this perspective the behavior of an applied econometrician who decides to ignore the join uncertainty is analogous to an 18th-century astronomer who is afraid to combine observations from different sources when making computations. Moreover, the loss of precision associated with having joint error bands is smaller for regions with higher credibility.

3.4 Simulation Design

The simulation strategy is designed with three objectives: (i) to evaluate the underestimation of uncertainty when the economic question requires joint inference but the economist reports conventional pointwise error bands, (ii) to assess the ability of different methods of conducting joint inference to properly account for different types of targeted uncertainty (i.e., MV, TS, MV-TS, and MV-TS-CS error bands), and (iii) to compare the width of the joint error bands.

Recently, Talts et al. (2020) proposed a general procedure called simulation-based calibration[8] for validating the inference of Bayesian estimation algorithms. The error bands are evaluated using the property that the credible regions for parameters retain frequentist coverage on

---

[8]Simulation-based calibration refers to the Monte Carlo procedure used in this study, which is different from the calibration procedures used to adjust the nominal coverage of the joint error bands introduced in the previous section.

average across the prior if the model is specified as in the procedure proposed by Talts et al. (2020) (see also Gelman et al. (1996), Little (2006), and Cook et al. (2006)).

The simulation exercise begins by drawing model parameters from the prior distribution. Based on these parameter values, the true impulse responses are computed, and endogenous variables are simulated. The model is then fitted to the simulated data using D posterior draws, which are saved, and the pointwise and joint error bands are computed. These steps are repeated C times, and the appropriate coverage statistics and average relative width are calculated for each type of targeted uncertainty.

For any credibility level $(1 - \alpha)100\%$, the coverage statistics is defined as

$$\Pi = C^{-1} \sum_{c=1}^{C} \mathbb{1}\left((\boldsymbol{\theta})^{(c)} \in \left[(\underline{\boldsymbol{\theta}}^{1-\alpha})^{(c)}, (\bar{\boldsymbol{\theta}}^{1-\alpha})^{(c)}\right]\right) \cdot 100,$$

where $\mathbb{1}(\cdot)$ is the indicator function defined over the true structural coefficient, $(\boldsymbol{\theta})^{(c)}$, and the estimated lower, $(\underline{\boldsymbol{\theta}}^{1-\alpha})^{(c)}$, and upper, $(\bar{\boldsymbol{\theta}}^{1-\alpha})^{(c)}$, bounds for each iteration, c. If the inference algorithm is correct the coverage statistics tends to the nominal credibility level, $(1 - \alpha)100\%$, as the total number of iterations, C, increases. Depending on the inference type, which is targeted with the error band the vector of structural parameters will include different elements and there are 5 different types of the coverage statistics: PW, MV, TS, MV-TS, and MV-TS-CS.

The simulation strategy requires a well-defined generative model. Since the data generating processes (described in the next section) resemble the real-world empirical VAR models, their parameter distributions are unknown. Therefore, using common prior parametrizations ("non-informative" or "unit root") as initial parameter distributions might cause non-defined or economically implausible models. To solve this problem, the models are first estimated using standard priors and then the the posterior distributions are used as priors in the simulation procedure.[9]

In the empirical literature authors usually report 68% and 95% credible intervals (e.g., Jarociński and Karadi (2020), Paul (2020) and Arias et al. (2019)), but there is no clear cut which one is preferable (see Kilian and Lütkepohl (2017)). The analysis in this paper is conducted for 68% and 90% credible intervals. The latter level is being selected to produce more informative statistics if the coverage is overstated.

Regardless of the large number of iterations, the experiments are still subject to uncertainty; and thus, the standard error of the coverage statistics and the relative width is calculated using bootstrap method.[10] The complete simulation procedure is formally described below.

---

[9]This strategy is proposed for a similar experiment by Gabry et al. (2019) providing an example from natural sciences where simulated air pollution data based on loose priors has extreme density.

[10]See Koehler et al. (2009) for a discussion on the uncertainty reporting for Monte Carlo experiments.

Simulation-based calibration
---
set prior distributions $\Sigma \sim IW_n(\nu_0 \mathbf{S}_0, \nu_0)$ and $vec(\mathbf{A}) \sim \mathcal{N}\left(vec(\bar{\mathbf{A}}_0), \Sigma \otimes \mathbf{N}_0^{-1}\right)$
for c in C **do**
  draw posterior $(\Sigma)^{(c)} \sim IW_n(\nu_T \mathbf{S}_T, \nu_T)$ and $(vec(\mathbf{A}))^{(c)} \sim \mathcal{N}\left(vec(\bar{\mathbf{A}}_T), (\Sigma)^{(c)} \otimes \mathbf{N}_T^{-1}\right)$
  evaluate structural coefficients, $(\theta_{n,s,h})^{(c)} \sim \pi\left(\phi_{n,s,h}|(\Sigma)^{(c)}, (vec(\mathbf{A}))^{(c)}\right)$
  simulate endogenous variables, $(y_t)^{(c)} \sim \pi\left(y_t|(\Sigma)^{(c)}, (vec(\mathbf{A}))^{(c)}\right)$
  for d in D **do**
    draw posterior, $(\Sigma)^{(d:c)} \sim IW_n(\nu_{\widetilde{T}} \mathbf{S}_{\widetilde{T}}, \nu_{\widetilde{T}})$ and $(vec(\mathbf{A}))^{(d:c)} \sim \mathcal{N}\left(vec(\bar{\mathbf{A}}_{\widetilde{T}}), (\Sigma)^{(d:c)} \otimes \mathbf{N}_{\widetilde{T}}^{-1}\right)$
    evaluate structural coefficients, $(\theta_{n,s,h})^{(d:c)} \sim \pi\left(\theta_{n,s,h}|(\Sigma)^{(d:c)}, (vec(\mathbf{A}))^{(d:c)}\right)$
  end for
  compute the pointwise and joint error bands using using $(\theta_{n,s,h})^{(d:c)}$
end for
compute the coverage statistics, Π, and average relative width, Δ.
---

4 Experiments

4.1 Data Generating Processes

This subsection introduces the specifics of the data generating processes used in the simulation exercises. In short those are: i) the quarterly fiscal model of Ramey (2019) identified with short-run exclusion restrictions; ii) the quarterly macro model of Peersman (2005) identified with sign restrictions; iii) the monthly monetary model of Jarociński and Karadi (2020) identified with an instrumental variable; and iv) the quarterly macro model of Blanchard and Quah (1989) identified with long-run exclusion restrictions.

DGP1. The first data generating process follows the fiscal model of Ramey (2019). The VAR model has three endogenous variables: government spending, gross domestic product and tax revenue with all variables in real per capita terms specified in logs. The reduced-form model is estimated initially with "non-informative" prior on quarterly data (from 1939:Q1 to 2015:Q4) with a constant and 4 lags of the endogenous variables. Only government spending shock is identified through recursive identification scheme based on the Cholesky decomposition. The proposed variables order implies that the government spending is affected contemporaneously only by its own structural shock. Figure B.1 in Appendix B shows the impulse responses of the three endogenous variables for the simulated realizations of the data generating process. A positive shock to government spending increases all endogenous variables on impact, but the effect peaks between the 5th and 12th quarter and by the end of the last observed 20th quarter is already negligible. The experiment for DGP1 is repeated 5000 times conducted on a sample size of 200 observations and impulse response horizon of 20 quarters, which is characterized by 15000 posterior draws.

DGP2. The second data generating process resembles the empirical model of Peersman (2005). The VAR model has three endogenous variables: gross domestic product, deflator and 3-month treasury bill yield with the first two variables in real per capita terms specified in logs and the third in percentage points. The reduced-form model is estimated initially with "non-informative" prior on quarterly data (from 1950:Q1 to 2006:Q4) with a constant and 4 lags of the endogenous variables. The sample period of DGP2 is shorter compared to DGP1 since it excludes periods with non-responsive monetary policy, and consequently constant yield. Unlike Peersman (2005), I identify only a supply shock due to computational considerations. The supply shock is identified with sign restrictions and is assumed to raise the interest rate on impact (holding only one period), reduce output for 4 quarters and increase prices for 4 quarters. Figure B.2 in

Appendix B shows the impulse responses of the three endogenous variables for the simulated realizations of the data generating process. In most of the simulated models, the output hits the trough 4 to 8 quarters after impact, while the deflator is experiencing prolonged effects. The interest rate reacts to the economic developments and remains elevated for significant number of periods. The experiment for DGP2 is repeated 5000 times and conducted on a sample size of 200 observations. The impulse response horizon consists of 20 quarters, which is characterized by 150 posterior draws and 150 rotations of the impact matrix.

DGP3. The third data generating process considered in this study mimics the monetary model of Jarociński and Karadi (2020). The VAR model has four endogenous variables: one year government bond yield, real gross domestic product, price deflator and excess bond premium. The yield and the excess bond premium are expressed in percentage points and the other two variables in logarithm. The reduced-form model is estimated initially with "unit root" prior on monthly data (from 1990:M2 to 2016:M12) with a constant and 6 lags of the endogenous variables. Only monetary shock is identified through the instrumental variable denoted as poor man's monetary shock by Jarociński and Karadi (2020). It is calculated as the size of the interest rate change around FOMC meetings when stock prices and interest rates move in opposite directions. Figure B.3 in Appendix B shows the impulse responses of the four endogenous variables for the simulated realizations of the data generating process. The experiment for DGP3 is repeated 5000 times and conducted on a sample size of 270 observations and impulse response horizon of 36 months, which is characterized by 15000 posterior draws.

DGP4. The last data generating process follows the macro model of Blanchard and Quah (1989). The VAR model has two endogenous variables: gross national product (specified as a growth rate) and unemployment rate; and thus the vector of endogenous variables is potentially stationary. The reduced-form model is estimated initially with "non-informative" priors on quarterly data (from 1948:Q2 to 1987:Q4) with a constant and 4 lags of the endogenous variables. In contrast to the previous data generating processes, two structural shocks, namely supply and demand, are identified imposing long-run restrictions. In particular, it is assumed that neither shock has long-run effect on unemployment and also the only supply shock has long-run effect on output. Figure B.4 in Appendix B shows the impulse responses of the two endogenous variables in response to both structural shocks for the simulated realizations of the data generating process. Structural demand shocks have a hump-shaped effect on GNP and unemployment with the effect peaking by the end of the first year and quickly vanishing by end of the fourth year. The supply shock has an even shorter lived effect on unemployment, but GNP remains elevated in the long run with maximum achieved by the end of the third year after the impact. The experiment for DGP4 is repeated 5000 times conducted on a sample size of 200 observations and impulse response horizon of 20 quarters, which is characterized by 15000 posterior draws.

4.2 Coverage Level and Relative Width

The estimated coverage levels along with their standard deviations of the PW error bands of DGP1-4 are given in Table 1. The first row shows the pointwise coverage level, i.e., the fraction of individual parameters lying within their respective credible intervals. As expected, the conventional pointwise inference achieves the nominal coverage level for the 68% and 90% credible intervals, implying that the Monte Carlo design is properly specified and the estimation procedure is correctly implemented. The coverage levels shown in the rows below answer the question: what would the effective coverage level be if the economic question requires

joint error bands, but the applied econometrician naïvely reports conventional pointwise error bands? Even in the mildest situation, when only multivariate dependence is ignored, the respective coverage level, Π^{MV}, of the PW error bands is between 26% and 57% for the 68% error band and between 70% and 85% for the 90% error band. The subsequent rows show that as the number of ignored structural parameters in the inference procedure increases, the underestimation of uncertainty when using PW error bands rises dramatically. For example, DGP1 has the worst TS coverage level, Π^{TS}, which is 21% and 64% for the 68% and 90% error bands, respectively. In the most severe case, when MV-TS-CS dependence is ignored, the estimated coverage level for DGP4 is 5% and 39% for both credibility levels.[11]

In summary, the PW error band is an appropriate choice for conducting inference about impulse responses only when the economic question does not involve a comparison of multiple structural parameters simultaneously. If this is the case, the uncertainty is significantly underestimated. The loss associated with reporting incorrectly specified error bands is larger for the 68% credible region compared to the 90% credible region.

Table 1
Coverage level of PW error bands

Coverage level	68% PW error band				90% PW error band			
	DGP1	DGP2	DGP3	DGP4	DGP1	DGP2	DGP3	DGP4
Π^{PW}	68.07	67.25	67.63	67.70	90.17	89.47	89.94	89.71
	(0.34)	(0.29)	(0.34)	(0.32)	(0.22)	(0.19)	(0.22)	(0.21)
Π^{MV}	44.64	36.18	26.05	57.06	80.00	74.92	69.75	85.31
	(0.42)	(0.48)	(0.46)	(0.37)	(0.36)	(0.41)	(0.54)	(0.27)
Π^{TS}	21.23	43.83	31.65	26.34	63.98	75.09	71.82	66.41
	(0.40)	(0.41)	(0.40)	(0.42)	(0.48)	(0.35)	(0.42)	(0.48)
Π^{MV-TS}	2.90	17.14	3.24	16.43	36.12	51.98	37.70	56.43
	(0.25)	(0.44)	(0.24)	(0.37)	(0.67)	(0.64)	(0.72)	(0.54)
$\Pi^{MV-TS-CS}$	-	-	-	4.68	-	-	-	38.86
				(0.28)				(0.65)

Notes: Coverage levels of 68% and 90% pointwise credible bands for DGP1, DGP2, DGP3 and DGP4.
Standard deviations computed with bootstrap method are in brackets.

The results of the experiments regarding the joint inference methods for all DGPs are given in Tables 2 - 5. Each table details the estimated coverage level and average relative width, along with their standard deviations, for the four different assumptions on the targeted mutual dependence between the structural parameters: MV, TS, MV-TS, and MV-TS-CS.

First, consider the results for the MV error bands in Table 2. As expected, the Bonferroni's and Sidák's error bands are the among the most conservative measures of joint uncertainty. This is particularly true for the Bonferroni method, while the Sidák method achieves the nominal coverage level of 90% for DGP2 and DGP3. Conversely, the sup-t error bands method meets the nominal rate for all DGPs and both credibility levels, indicating that it effectively summarizes multivariate estimation uncertainty. The average relative width of the sup-t error bands ranges from 19% to 65% for the 68% credible region and from 11% to 35% for the 90% credible region.

---

[11] DGP1-3 have only one identified shock; thus, the MV-TS-CS coverage level is not defined.

These results indicate that the MV error bands are not significantly wider than the conventional PW error bands, suggesting that their use does not incur substantial costs. Additionally, a comparison of the MV error bands for DGP4 in the 90% credibility region reveals that Sidák's error bands have a 92% effective coverage level and the sup-t error band has a 90% effective coverage level, with average relative widths of 21% and 11%, respectively. This demonstrates that properly constructed error bands achieving the nominal coverage level can lead to better uncertainty estimation compared to ad-hoc methods.

Next, examining the MV error bands constructed with the min-max methodology, I find that the non-calibrated error bands are conservative and tend to produce higher coverage levels than the nominal credibility level for absolute, quadratic, and angular loss functions. Furthermore, there are significant differences in the performance of error bands constructed with these loss functions. Specifically, the angular loss function performs worse than the absolute loss function, which is strongly dominated by the quadratic loss function. This finding is surprising, given that Inoue and Kilian (2022) report no significant differences among the loss functions in their empirical applications. One possible reason for this discrepancy is that their specifications consider MV-TS-CS uncertainty, which increases the dimensionality of the problem and diminishes the importance of the selected loss function. This will be evaluated later in this section when analyzing the simulation results in Table 5. Additionally, it is worth noting that, despite Inoue and Kilian (2022) proposing the angular loss function to address potential issues with evaluating vector-valued distances of structural parameters with different scales, it actually computes the most conservative error bands. Finally, among the non-calibrated min-max error bands, the Chebyshev loss function consistently achieves the nominal coverage level, with a relative width economically indistinguishable from that of the sup-t error bands.

Having examined the non-calibrated min-max error bands, I now shift the focus to other extended routines that implement an intermediate calibration step (i.e., LQO or BDR). Notably, the Chebyshev loss function achieves the nominal coverage level without implementing LQO or BDR, and the performance of the error bands in the corresponding rows aligns with the min-max error bands as the calibration steps are not initiated. For the absolute and quadratic loss functions, the LQO calibration helps achieve the nominal coverage level, but it fails for the angular loss function. However, comparing the average relative width of the error bands for the former two loss functions, the results indicate that these error bands are wider compared to the sup-t estimations, with differences reaching up to 20 percentage points for some DGPs. On the other hand, the error bands with BDR calibration not only achieve the nominal coverage level but also have widths for the absolute and angular loss functions similar to those of the sup-t error bands. The BDR calibration applied to the min-max error bands with the quadratic loss function is slightly larger than the sup-t error bands for some DGPs.

In summary, the sup-t, min-max with Chebyshev loss and min-max (BDR) with absolute and angular loss are good estimators MV uncertainty with very similar performance regarding the estimated coverage level and average relative width.

The results of the experiments regarding the TS error bands are given in Table 3. The increase of the number of elements of the structural vector of interest, changes the relative performance of the joint inference methods. The Bonferorrni's and Sidák's error bands become the most conservative methods, which greatly overestimate the uncertainty about the impulse responses. In this class of methods, the sup-t error bands is the only one, which achieves the nominal coverage level. The difference of the average relative width for the TS error bands and the MV error bands estimated with the sup-t method are large for DGP1 and DGP4, but very similar for

DGP2 and DGP3. A possible reason for that could be the different identification methodology. Similarly to the reported results in the previous table, the TS error bands with the min-max error bands with Chebyshev loss function perform the same as the sup-t error bands even without a calibration step. The same could not be said for the other loss functions within this class of joint inference methods. Without calibration step the angular loss is more conservative than the absolute loss function, which is dominated by the quadratic loss. The LQO calibration method allows for all error bands to achieve the nominal coverage level with slightly smaller average relative width for the absolute and quadratic loss functions. The results for the LQO calibration of the min-max error bands with angular loss function indicate that it is significantly wider than the sup-t error band. The BDR calibration method, however, also allows the min-max estimator to achieve the nominal coverage level, but it further decreases the width of the error bands. The min-max (BDR) with angular loss has similar or slightly wider bands than the sup-t method, but the absolute and quadratic loss functions already improve significantly on the sup-t error bands. For example, the difference of the relative width for the 68% credible region between the min-max (BDR) error bands with absolute loss and the sup-t error bands are between 7 and 14 percentage points.

Table 4 presents the result for analysis of MV-TS joint inference methods. There are no differences in the relative performance in terms of coverage levels of the different inference methods caused by the increased amount of structural parameters of interest for the MV-TS error bands in comparison to the TS error bands. It is worth noting, however, that the increase of the average relative width does not entirely depends on the number of structural parameters, but also on the specifics of the data generating processes. In particular the estimate of the 68% TS error band of the min-max (BDR) method under absolute loss function for DGP1 and DGP2 is 60% and 24%. Even though the increase of the number of structural parameters for both DGPs in the case of MV-TS inference is the same, the average relative width changes to 95% and 56%. The highest change in the average relative width from the TS error bands to the MV-TS error bands is observed for DGP3 - from 45% to 104%, which is expected since this models has the largest increase of structural parameters between both error bands. The smallest increase of the width between both types of error bands is observed for DGP4. The comparison between the results of the TS and MV-TS error bands show that an economist interested in joint inference method could observe different width among models depending on the economic question and the corresponding inference methods, which is selected.

Finally, Table 5 shows the results of the Monte Carlo simulations for MV-TS-CS joint inference, which is conducted only for DGP4 as it is the only model with more than one identified shock. The results for that inference type for the 68% credible region indicate again that Bonferroni's and Sidák's method are overestimating the uncertainty, but the sup-t error band achieves the nominal coverage level at least within 2 standard deviations of the coverage statistics. The min-max estimators without any calibration step again overestimate the coverage level. The LQO step could provide sufficiently close estimate, but as already discussed it could be very hard to achieve stable results especially for fuzzy joint distributions of some structural parameters. On the other hand, the BDR step could underestimate the uncertainty with approximately 1% depending on the applied loss functions. Even though the simulation exercise gives very clear results for the performance of the method, the economic implications are insignificant. Furthermore, the methods tend to achieve the nominal coverage level for the 90% credible region

Table 2

Multivariate coverage level (Π^{MV}) and average relative width (Δ^{MV}) of MV error bands

Method	68% MV error band								90% MV error band							
	DGP1		DGP2		DGP3		DGP4		DGP1		DGP2		DGP3		DGP4	
	Π	Δ	Π	Δ	Π	Δ	Π	Δ	Π	Δ	Π	Δ	Π	Δ	Π	Δ
Bonferroni	78.72	63.39	73.49	43.36	75.07	79.64	77.30	44.64	93.18	30.40	90.49	18.22	91.17	39.26	92.47	21.90
	(0.37)	(0.01)	(0.41)	(0.01)	(0.51)	(0.01)	(0.33)	(0.01)	(0.23)	(0.01)	(0.26)	(0.01)	(0.32)	(0.01)	(0.19)	(0.01)
Šidák	76.14	57.05	70.51	39.62	71.88	72.57	75.26	39.08	92.94	29.51	90.15	17.73	90.88	38.23	92.31	21.11
	(0.39)	(0.01)	(0.43)	(0.01)	(0.53)	(0.01)	(0.35)	(0.01)	(0.24)	(0.01)	(0.26)	(0.01)	(0.32)	(0.01)	(0.19)	(0.01)
sup-t (PQO)	67.78	38.41	67.18	35.61	68.18	65.06	67.75	18.89	90.36	21.08	89.17	16.34	89.69	34.82	89.73	11.16
	(0.44)	(0.01)	(0.44)	(0.01)	(0.55)	(0.02)	(0.37)	(0.01)	(0.27)	(0.01)	(0.28)	(0.01)	(0.34)	(0.01)	(0.22)	(0.01)
min-max: absolute loss (median)	88.25	102.33	89.82	81.62	88.62	159.88	76.03	47.25	98.00	64.33	98.17	44.07	97.38	95.38	93.08	32.88
	(0.28)	(0.03)	(0.28)	(0.04)	(0.36)	(0.06)	(0.28)	(0.02)	(0.13)	(0.03)	(0.13)	(0.04)	(0.18)	(0.04)	(0.18)	(0.01)
min-max (LQO): absolute loss (median)	68.27	47.26	67.28	41.63	67.20	85.92	67.74	27.17	90.33	27.13	89.49	24.49	89.52	57.02	89.65	19.66
	(0.40)	(0.02)	(0.49)	(0.02)	(0.55)	(0.04)	(0.31)	(0.01)	(0.25)	(0.01)	(0.28)	(0.02)	(0.34)	(0.04)	(0.21)	(0.02)
min-max (BDR): absolute loss (median)	67.85	38.97	66.86	35.05	67.80	67.63	67.66	20.63	90.39	21.14	89.24	16.95	89.56	37.19	89.62	13.38
	(0.44)	(0.01)	(0.48)	(0.02)	(0.58)	(0.02)	(0.34)	(0.01)	(0.27)	(0.01)	(0.31)	(0.01)	(0.35)	(0.01)	(0.22)	(0.01)
min-max: quadratic loss (mean)	75.69	72.05	77.64	61.13	74.64	117.61	70.20	37.70	93.66	47.52	95.04	36.19	92.33	77.47	90.75	26.64
	(0.38)	(0.02)	(0.44)	(0.02)	(0.54)	(0.03)	(0.30)	(0.02)	(0.20)	(0.02)	(0.20)	(0.02)	(0.31)	(0.03)	(0.20)	(0.02)
min-max (LQO): quadratic loss (mean)	68.30	54.17	66.91	43.32	67.13	94.14	67.66	31.49	90.24	35.22	89.46	25.88	89.41	66.68	89.63	22.79
	(0.39)	(0.02)	(0.50)	(0.02)	(0.58)	(0.04)	(0.30)	(0.02)	(0.25)	(0.02)	(0.29)	(0.02)	(0.36)	(0.04)	(0.21)	(0.02)
min-max (BDR): quadratic loss (mean)	68.19	43.66	66.78	36.41	67.29	75.89	67.65	24.95	90.24	24.74	89.24	18.25	89.54	45.65	89.60	16.96
	(0.42)	(0.02)	(0.50)	(0.02)	(0.57)	(0.03)	(0.31)	(0.01)	(0.26)	(0.01)	(0.30)	(0.01)	(0.35)	(0.03)	(0.21)	(0.02)
min-max: angular loss (median)	99.93	310.81	99.87	132.56	99.93	335.78	99.97	376.96	99.95	151.50	99.92	63.35	99.95	165.21	99.97	187.07
	(0.02)	(0.15)	(0.03)	(0.07)	(0.03)	(0.18)	(0.01)	(0.25)	(0.01)	(0.08)	(0.03)	(0.05)	(0.02)	(0.10)	(0.01)	(0.14)
min-max (LQO): angular loss (median)	97.08	171.06	86.99	73.62	95.00	179.40	97.87	195.08	97.08	62.67	91.86	26.69	95.02	66.27	97.87	72.58
	(0.14)	(0.21)	(0.30)	(0.12)	(0.23)	(0.24)	(0.09)	(0.29)	(0.14)	(0.13)	(0.24)	(0.06)	(0.23)	(0.14)	(0.09)	(0.17)
min-max (BDR): angular loss (median)	67.65	37.83	67.18	34.92	68.00	64.17	67.72	18.83	90.35	20.84	89.14	16.22	89.67	34.61	89.71	11.09
	(0.44)	(0.01)	(0.46)	(0.01)	(0.56)	(0.02)	(0.37)	(0.01)	(0.27)	(0.01)	(0.28)	(0.01)	(0.34)	(0.01)	(0.22)	(0.01)
min-max: Chebyshev loss (median)	67.70	37.91	67.43	35.23	67.96	63.93	67.64	18.23	90.50	20.81	88.99	17.17	89.67	34.83	89.66	12.02
	(0.44)	(0.01)	(0.46)	(0.02)	(0.56)	(0.02)	(0.37)	(0.01)	(0.28)	(0.01)	(0.30)	(0.01)	(0.33)	(0.01)	(0.25)	(0.01)
min-max (LQO): Chebyshev loss (median)	67.70	37.91	67.42	35.23	67.96	63.93	67.64	18.23	90.50	20.81	88.99	17.17	89.67	34.83	89.66	12.02
	(0.44)	(0.01)	(0.46)	(0.02)	(0.56)	(0.02)	(0.37)	(0.01)	(0.28)	(0.01)	(0.30)	(0.01)	(0.33)	(0.01)	(0.25)	(0.01)
min-max (BDR): Chebyshev loss (median)	67.70	37.91	67.43	35.23	67.96	63.93	67.64	18.23	90.50	20.81	88.99	17.17	89.67	34.83	89.66	12.02
	(0.44)	(0.01)	(0.46)	(0.02)	(0.56)	(0.02)	(0.37)	(0.01)	(0.28)	(0.01)	(0.30)	(0.01)	(0.33)	(0.01)	(0.25)	(0.01)

Notes: Multivariate coverage levels (Π^{MV}) and average relative width (Δ^{MV}) of 68% and 90% MV credible bands for DGP1, DGP2, DGP3 and DGP4. Standard deviations computed with bootstrap method are in brackets.

Table 3
Time-simultaneous coverage level (Π^{TS}) and average relative width (Δ^{TS}) of TS error bands

Method	68% TS error band								90% TS error band							
	DGP1		DGP2		DGP3		DGP4		DGP1		DGP2		DGP3		DGP4	
	Π	Δ	Π	Δ	Π	Δ	Π	Δ	Π	Δ	Π	Δ	Π	Δ	Π	Δ
Bonferroni	93.03	147.35	94.03	84.63	96.37	177.70	92.79	165.40	97.83	74.47	97.79	39.61	98.77	92.25	97.36	86.47
	(0.25)	(0.03)	(0.19)	(0.03)	(0.16)	(0.03)	(0.24)	(0.06)	(0.15)	(0.02)	(0.12)	(0.02)	(0.11)	(0.02)	(0.14)	(0.03)
Šidák	91.77	140.32	93.13	81.64	95.73	170.31	91.59	156.98	97.73	73.42	97.70	39.14	98.74	91.10	97.23	85.17
	(0.29)	(0.02)	(0.21)	(0.03)	(0.18)	(0.03)	(0.25)	(0.05)	(0.15)	(0.02)	(0.12)	(0.02)	(0.11)	(0.02)	(0.15)	(0.03)
sup-t (PQO)	68.07	74.60	66.60	31.71	67.81	59.47	67.74	72.10	90.58	40.35	88.80	17.58	89.67	33.83	89.56	41.20
	(0.48)	(0.01)	(0.43)	(0.02)	(0.46)	(0.03)	(0.46)	(0.07)	(0.31)	(0.01)	(0.28)	(0.01)	(0.28)	(0.01)	(0.28)	(0.04)
min-max: absolute loss (median)	78.13	94.48	77.13	46.62	77.06	79.09	81.91	118.89	94.41	60.28	93.17	25.88	93.53	50.74	94.82	72.02
	(0.45)	(0.02)	(0.31)	(0.03)	(0.44)	(0.05)	(0.38)	(0.10)	(0.26)	(0.01)	(0.20)	(0.02)	(0.24)	(0.03)	(0.25)	(0.06)
min-max (LQO): absolute loss (median)	68.23	68.00	67.12	29.94	67.47	53.72	67.70	75.90	90.18	44.28	89.56	18.00	89.72	36.28	89.44	47.75
	(0.54)	(0.01)	(0.36)	(0.02)	(0.48)	(0.03)	(0.46)	(0.09)	(0.31)	(0.01)	(0.26)	(0.02)	(0.29)	(0.02)	(0.32)	(0.05)
min-max (BDR): absolute loss (median)	67.91	60.27	66.74	24.13	67.16	45.01	67.41	64.55	90.28	35.62	89.23	13.39	89.53	27.64	89.30	37.43
	(0.52)	(0.01)	(0.37)	(0.01)	(0.47)	(0.03)	(0.48)	(0.08)	(0.31)	(0.01)	(0.25)	(0.01)	(0.29)	(0.01)	(0.32)	(0.04)
min-max: quadratic loss (mean)	73.80	79.87	74.02	39.25	73.58	66.98	76.19	107.11	92.53	50.33	92.36	22.48	92.28	44.17	92.80	65.87
	(0.48)	(0.02)	(0.33)	(0.02)	(0.46)	(0.04)	(0.40)	(0.10)	(0.28)	(0.02)	(0.22)	(0.02)	(0.27)	(0.02)	(0.25)	(0.06)
min-max (LQO): quadratic loss (mean)	68.19	66.13	66.80	27.93	67.39	51.75	67.44	83.69	90.19	42.38	89.42	16.43	89.66	35.34	89.19	52.96
	(0.53)	(0.02)	(0.38)	(0.02)	(0.50)	(0.02)	(0.48)	(0.10)	(0.32)	(0.02)	(0.27)	(0.01)	(0.29)	(0.02)	(0.31)	(0.06)
min-max (BDR): quadratic loss (mean)	67.93	60.81	66.51	24.40	67.25	46.14	67.41	72.14	90.03	36.45	89.25	13.19	89.46	28.82	89.28	42.41
	(0.53)	(0.02)	(0.37)	(0.01)	(0.48)	(0.02)	(0.49)	(0.09)	(0.32)	(0.01)	(0.26)	(0.01)	(0.29)	(0.01)	(0.31)	(0.05)
min-max: angular loss (median)	78.27	142.13	80.11	84.49	88.28	185.52	78.55	135.78	92.99	86.07	91.49	38.62	95.86	110.96	93.07	91.16
	(0.45)	(0.09)	(0.27)	(0.07)	(0.25)	(0.10)	(0.37)	(0.11)	(0.27)	(0.05)	(0.19)	(0.05)	(0.15)	(0.07)	(0.22)	(0.08)
min-max (LQO): angular loss (median)	68.93	113.02	67.19	50.90	69.32	109.68	67.58	100.42	90.14	75.31	89.29	36.00	90.00	79.37	89.39	73.33
	(0.52)	(0.05)	(0.33)	(0.05)	(0.38)	(0.12)	(0.48)	(0.13)	(0.34)	(0.04)	(0.22)	(0.05)	(0.23)	(0.10)	(0.25)	(0.07)
min-max (BDR): angular loss (median)	68.39	75.09	66.86	32.21	67.28	60.19	67.21	75.01	90.12	47.43	89.33	24.65	89.39	42.00	89.19	51.20
	(0.52)	(0.03)	(0.36)	(0.02)	(0.41)	(0.03)	(0.48)	(0.09)	(0.32)	(0.02)	(0.22)	(0.03)	(0.27)	(0.05)	(0.28)	(0.05)
min-max: Chebyshev loss (median)	67.64	73.13	66.53	31.41	67.40	58.35	67.39	70.33	90.40	39.62	88.81	15.26	89.61	33.35	89.38	39.98
	(0.47)	(0.01)	(0.43)	(0.02)	(0.45)	(0.03)	(0.47)	(0.06)	(0.34)	(0.01)	(0.27)	(0.01)	(0.27)	(0.02)	(0.30)	(0.03)
min-max (LQO): Chebyshev loss (median)	67.64	73.13	66.53	31.40	67.40	58.35	67.39	70.33	90.40	39.62	88.81	15.26	89.61	33.35	89.38	39.98
	(0.47)	(0.01)	(0.43)	(0.02)	(0.45)	(0.03)	(0.47)	(0.06)	(0.34)	(0.01)	(0.27)	(0.01)	(0.27)	(0.02)	(0.30)	(0.03)
min-max (BDR): Chebyshev loss (median)	67.64	73.13	66.53	31.41	67.40	58.35	67.39	70.33	90.40	39.62	88.81	15.26	89.61	33.35	89.38	39.98
	(0.47)	(0.01)	(0.43)	(0.02)	(0.45)	(0.03)	(0.47)	(0.06)	(0.34)	(0.01)	(0.27)	(0.01)	(0.27)	(0.02)	(0.30)	(0.03)

Notes: Time-simultaneous coverage levels (Π^{TS}) and average relative width (Δ^{TS}) of 68% and 90% TS credible bands for DGP1, DGP2, DGP3 and DGP4. Standard deviations computed with bootstrap method are in brackets.

Table 4

Multivariate, time-simultaneous coverage level (Π^{MV-TS}) and average relative width (Δ^{MV-TS}) of MV-TS error bands

Method	68% MV-TS error band								90% MV-TS error band							
	DGP1		DGP2		DGP3		DGP4		DGP1		DGP2		DGP3		DGP4	
	Π	Δ	Π	Δ	Π	Δ	Π	Δ	Π	Δ	Π	Δ	Π	Δ	Π	Δ
Bonferroni	94.56	188.45	93.82	100.97	96.80	231.49	94.32	197.52	97.90	97.03	97.44	48.97	98.68	122.56	97.94	104.43
	(0.33)	(0.03)	(0.29)	(0.04)	(0.26)	(0.04)	(0.25)	(0.07)	(0.21)	(0.02)	(0.22)	(0.03)	(0.14)	(0.03)	(0.16)	(0.04)
Šidák	93.18	181.85	93.10	98.46	96.10	224.54	93.44	189.22	97.86	96.03	97.34	48.58	98.66	121.38	97.83	103.13
	(0.36)	(0.03)	(0.30)	(0.04)	(0.29)	(0.04)	(0.26)	(0.07)	(0.21)	(0.02)	(0.22)	(0.03)	(0.15)	(0.03)	(0.16)	(0.04)
sup-t (PQO)	68.70	113.68	66.42	60.73	67.10	117.81	67.80	92.41	90.40	60.75	88.02	30.32	88.28	64.54	89.44	52.35
	(0.70)	(0.02)	(0.60)	(0.02)	(0.68)	(0.03)	(0.50)	(0.07)	(0.42)	(0.02)	(0.42)	(0.01)	(0.41)	(0.02)	(0.31)	(0.04)
min-max: absolute loss (median)	87.14	155.65	84.06	88.39	81.24	165.08	79.76	128.67	96.54	87.67	95.68	46.13	94.32	97.53	93.96	77.69
	(0.50)	(0.05)	(0.45)	(0.06)	(0.53)	(0.07)	(0.41)	(0.10)	(0.28)	(0.03)	(0.27)	(0.04)	(0.35)	(0.06)	(0.29)	(0.06)
min-max (LQO): absolute loss (median)	67.58	106.37	66.82	63.64	66.94	121.28	67.80	92.76	90.26	61.85	88.76	35.69	89.76	76.83	89.34	57.80
	(0.66)	(0.03)	(0.60)	(0.03)	(0.74)	(0.06)	(0.52)	(0.09)	(0.47)	(0.02)	(0.38)	(0.03)	(0.44)	(0.04)	(0.33)	(0.05)
min-max (BDR): absolute loss (median)	66.78	94.75	66.10	55.48	65.88	104.31	67.69	79.77	89.66	52.97	88.42	29.14	88.32	60.62	88.99	47.07
	(0.66)	(0.02)	(0.61)	(0.03)	(0.69)	(0.04)	(0.51)	(0.08)	(0.47)	(0.02)	(0.35)	(0.02)	(0.43)	(0.03)	(0.36)	(0.04)
min-max: quadratic loss (mean)	78.06	141.39	78.96	79.84	73.44	150.64	75.45	122.26	94.28	81.43	94.76	43.52	91.38	93.18	92.28	74.50
	(0.59)	(0.04)	(0.50)	(0.04)	(0.72)	(0.06)	(0.42)	(0.10)	(0.38)	(0.03)	(0.25)	(0.03)	(0.42)	(0.04)	(0.29)	(0.06)
min-max (LQO): quadratic loss (mean)	67.46	116.92	66.64	62.93	66.80	131.38	67.24	101.23	90.34	69.04	88.96	35.48	89.12	86.47	88.94	63.21
	(0.65)	(0.04)	(0.62)	(0.03)	(0.72)	(0.06)	(0.51)	(0.09)	(0.42)	(0.03)	(0.40)	(0.03)	(0.45)	(0.05)	(0.33)	(0.06)
min-max (BDR): quadratic loss (mean)	66.62	100.77	66.28	56.32	65.64	113.07	67.20	86.85	89.94	56.80	88.68	29.53	88.68	68.51	88.75	51.47
	(0.68)	(0.03)	(0.61)	(0.03)	(0.77)	(0.05)	(0.54)	(0.08)	(0.46)	(0.02)	(0.41)	(0.02)	(0.45)	(0.04)	(0.35)	(0.05)
min-max: angular loss (median)	86.30	191.84	74.94	87.93	81.98	201.12	75.39	135.15	95.84	106.46	92.62	46.84	94.00	114.76	92.00	85.82
	(0.52)	(0.09)	(0.53)	(0.06)	(0.60)	(0.14)	(0.44)	(0.11)	(0.29)	(0.06)	(0.34)	(0.05)	(0.34)	(0.11)	(0.28)	(0.08)
min-max (LQO): angular loss (median)	68.64	137.75	66.94	76.22	66.96	155.50	67.73	110.36	89.84	84.36	88.74	42.36	89.54	98.53	89.16	73.43
	(0.68)	(0.06)	(0.60)	(0.05)	(0.72)	(0.13)	(0.51)	(0.12)	(0.43)	(0.05)	(0.39)	(0.04)	(0.47)	(0.12)	(0.33)	(0.07)
min-max (BDR): angular loss (median)	67.82	104.37	66.48	64.30	65.86	118.21	67.27	89.41	89.42	59.87	88.56	35.02	89.00	71.20	88.74	56.80
	(0.63)	(0.02)	(0.59)	(0.04)	(0.70)	(0.08)	(0.51)	(0.10)	(0.44)	(0.02)	(0.40)	(0.03)	(0.48)	(0.07)	(0.33)	(0.06)
min-max: Chebyshev loss (median)	68.48	111.33	66.52	56.96	66.88	115.04	67.12	88.83	91.16	59.95	88.34	29.63	88.96	64.41	89.53	51.34
	(0.72)	(0.02)	(0.56)	(0.02)	(0.64)	(0.03)	(0.53)	(0.06)	(0.42)	(0.01)	(0.44)	(0.01)	(0.44)	(0.02)	(0.36)	(0.03)
min-max (LQO): Chebyshev loss (median)	68.48	111.33	66.52	56.96	66.88	115.04	67.12	88.83	91.16	59.95	88.34	29.63	88.96	64.41	89.53	51.34
	(0.72)	(0.02)	(0.56)	(0.02)	(0.64)	(0.03)	(0.53)	(0.06)	(0.42)	(0.01)	(0.44)	(0.01)	(0.44)	(0.02)	(0.36)	(0.03)
min-max (BDR): Chebyshev loss (median)	68.48	111.33	66.52	56.96	66.88	115.04	67.12	88.83	91.16	59.95	88.34	29.63	88.96	64.41	89.53	51.34
	(0.72)	(0.02)	(0.56)	(0.02)	(0.64)	(0.03)	(0.53)	(0.06)	(0.42)	(0.01)	(0.44)	(0.01)	(0.44)	(0.02)	(0.36)	(0.03)

Notes: Multivariate, time-simultaneous coverage levels (Π^{MV-TS}) and average relative width (Δ^{MV-TS}) of 68% and 90% MV-TS credible bands for DGP1, DGP2, DGP3 and DGP4. Standard deviations computed with bootstrap method are in brackets.

Table 5

Multivariate, time-simultaneous, cross-shock coverage level ($\Pi^{MV-TS-CS}$) and average relative width ($\Delta^{MV-TS-CS}$) of MV-TS-CS error bands

Method	68% MV-TS-CS error band		90% MV-TS-CS error band	
	DGP4		DGP4	
	Π	Δ	Π	Δ
Bonferroni	94.68 (0.28)	228.52 (0.08)	98.32 (0.17)	121.96 (0.05)
Šidák	93.42 (0.32)	220.37 (0.08)	98.28 (0.17)	120.68 (0.05)
sup-t (PQO)	67.16 (0.66)	122.92 (0.07)	88.94 (0.41)	68.65 (0.03)
min-max: absolute loss (median)	84.86 (0.54)	169.29 (0.15)	95.08 (0.32)	97.11 (0.06)
min-max (LQO): absolute loss (median)	66.98 (0.66)	115.07 (0.09)	89.06 (0.44)	68.73 (0.05)
min-max (BDR): absolute loss (median)	65.82 (0.71)	104.42 (0.09)	88.26 (0.45)	59.56 (0.05)
min-max: quadratic loss (mean)	77.58 (0.57)	156.07 (0.13)	93.32 (0.36)	91.44 (0.07)
min-max (LQO): quadratic loss (mean)	67.64 (0.68)	125.59 (0.10)	88.78 (0.45)	75.51 (0.05)
min-max (BDR): quadratic loss (mean)	66.58 (0.68)	111.80 (0.09)	88.32 (0.46)	64.29 (0.05)
min-max: angular loss (median)	78.22 (0.59)	182.89 (0.21)	92.84 (0.37)	114.39 (0.13)
min-max (LQO): angular loss (median)	67.00 (0.68)	143.98 (0.18)	88.80 (0.49)	96.16 (0.13)
min-max (BDR): angular loss (median)	66.32 (0.70)	118.35 (0.12)	87.96 (0.46)	73.46 (0.08)
min-max: Chebyshev loss (median)	66.84 (0.76)	116.46 (0.07)	89.56 (0.42)	66.04 (0.03)
min-max (LQO): Chebyshev loss (median)	66.84 (0.76)	116.46 (0.07)	89.56 (0.42)	66.04 (0.03)
min-max (BDR): Chebyshev loss (median)	66.84 (0.76)	116.46 (0.07)	89.56 (0.42)	66.04 (0.03)

Notes: Multivariate, time-simultaneous, cross-shock coverage levels ($\Pi^{MV-TS-CS}$) and average relative width ($\Delta^{MV-TS-CS}$) of 68% and 90% MV-TS-CS credible bands for DGP1, DGP2, DGP3 and DGP4. Standard deviations computed with bootstrap method are in brackets.

4.3 Shape and Comovement

This section discusses whether Bayesian estimation is capable of preserving the joint distributions of non-linear transformations, such as the impulse response functions, beyond just achieving the desired coverage level for multiple parameters simultaneously. This aligns with the argument of Inoue and Kilian (2022), who advocate for using the lowest posterior risk joint credible set to describe the shape and comovement of impulse responses rather than focusing solely on error bands.

Talts et al. (2020) show that if the prior of the Bayesian estimation procedure is specified to coincide with the distribution of the true parameters of the data-generating process, then the posterior distribution is the same as the prior distribution. The following rank statistics is used

to evaluate if both distributions coincide:

$$r(\theta_{n,s,h}) = \sum_{d=1}^{D} \mathbb{1}\left[(\theta_{n,s,h})^{(d:c)} < (\theta_{n,s,h})^{(c)}\right], \quad \text{for } n = 1,\ldots,N, s = 1,\ldots,S \text{ and } h = 1,\ldots,H,$$

where $\mathbb{1}[\cdot]$ is the indicator function defined over $(\theta_{n,s,h})^{(c)}$, $(\theta_{n,s,h})^{(d:c)}$, which denote for each Monte Carlo iteration c, the true impulse response coefficients and the d-th posterior draw, respectively. The distribution of the rank statistics is uniform across the integers $[0, D]$ if and only if the sampling procedure is correctly implemented. The rank statistics proposed by Talts et al. (2020) is applicable only to univariate random variables and allows to approach the question if each individual parameter is sampled correctly, but not if their joint distribution is approximated sufficiently well. Since evaluation of the marginal distributions involves significant amount of histograms and also this question is not the main focus of the study, then the analysis will turn towards the ability of the estimation to approximate the joint distributions of the structural parameters.

Thorarinsdottir et al. (2016) proposes several vector-valued loss statistics, which are used to evaluate the multivariate distributions of different forecasting procedures applied in the context of weather forecasting.[12] Their comparative analysis shows that the best performing measures are the band depth and average rank. The analysis in this section relies on the latter statistics. Starting with any vector of stacked impulse responses $\boldsymbol{\theta} \in \mathbb{R}^K$, the set of vectors $\mathbf{S} = \{(\boldsymbol{\theta})^{(1:c)}, (\boldsymbol{\theta})^{(2:c)}, \ldots, (\boldsymbol{\theta})^{(D:c)}, (\boldsymbol{\theta})^{(c)}\}$ is the backbone of the derivation. The first M elements of $\mathbf{S}$ denote draws from the posterior of this particular Monte Carlo iteration, c, and the last element denote the true vector of structural parameters. The prerank of the j-th coordinate of $\boldsymbol{\theta}$ in $\mathbf{S}$ is given by:

$$ru_S\left((\theta_j)^{(i:c)}\right) = \sum_{j=1}^{K} \mathbb{1}\left((\theta_j)^{(d:c)} \leq (\theta_j)^{(i:c)}\right) \quad \text{for } i = 1,\cdots,D+1.$$

Finally, the average rank statistic is simply the average over the univariate ranks

$$r_S\left((\boldsymbol{\theta})^{(i:c)}\right) = \frac{1}{K}\sum_{j=1}^{K} ru_S\left((\theta_j)^{(i:c)}\right) \quad \text{for } i = 1,\cdots,D+1,$$

where ties are resolved at random. The resulting rank of observation $(\boldsymbol{\theta})^{(D+1:c)}$ in S is uniform across the integers $[0, D+1]$ if the elements of S are independent and identically distributed (see Thorarinsdottir et al. (2016)).

For simplicity in presenting the results, average rank statistics are calculated for three types of vectors of structural parameters: TS, MV-TS, and MV-TS-CS. The MV vector of stacked impulse responses is excluded because it involves summarizing multiple quantities for each data generating process (DGP) and is already implicitly contained in two of the three specifications considered in the analysis. Following this, histograms are constructed to assess whether the statistics are uniformly distributed. This is done by checking if the bins fall within the 0.005 and 0.995 percentiles of a $Binomial(C, 1/B)$ distribution, where C represents the number of Monte Carlo iterations and B the number of bins. Under uniformity, only one bin in a hundred should fall outside these bounds (see Talts et al. (2020)).

---

[12]A similar procedure is proposed by Knüppel et al. (2022) to evaluate the joint distribution of forecasts in Bayesian VAR models.

If the histogram of the average rank statistics is ∪-shaped, the joint distribution of the impulse response function is underdispersive. On the other hand, if the histogram of the average rank statistics is ∩-shaped, the joint distribution of the impulse response functions is overdispersive. The average rank statistics under TS inference for all data generating processes is shown in Figure C.1 of Appendix C. For each DGP there is one histogram for every impulse response function - 3 for DGP1 and DGP2, 4 for DGP3 and 4 for DGP4. The bins in all histograms are within the bounds specified by the appropriate binomial distribution, therefore there is no evidence to reject the null hypothesis of uniformity. Analogously, the number of average rank statistics histograms in case of MV-TS inference for DGP1, DGP2 and DGP3 is 1 and for DGP4 is 2, all of which are presented in Figure C.2 of Appendix C. The bins of those histograms are also fully contained within the percentiles of the specified Binomial distribution, which implies that there is no evidence against the uniformity of the statistics. The same conclusion follows the analysis of the single histogram in the case of MV-TS-CS inference for DGP4 shown in Figure C.3 of Appendix C.

In summary, the uniform distribution of the average rank statistics implies that standard Bayesian inference about impulse responses is not only capable of achieving the targeted probability content jointly, but also it preserves shapes and comovements of impulse responses. This means that the lowest posterior risk joint credible set is very powerful tool to evaluate different economic theories as suggested by Inoue and Kilian (2022). On the other hand, the sup-t joint inference method and the respective draws contained within the estimated error bands do not necessarily have this property as it does not evaluate the distance across draws. This question along with the contribution of the loss function in the estimation of the lowest posterior risk set is left for future study.

5   Empirical Applications

Having established the importance of constructing context-specific error bands, this section highlights the benefits of the discussed methods for practitioners and policymakers. First, I demonstrate the implications of conducting joint inference on the fiscal multiplier in a typical SVAR model. In the second subsection, I explore the similarities between impulse response functions and forecasts, applying joint inference methods in an out-of-sample forecasting exercise for inflation and real GDP growth in the United States and the euro area.

5.1   Fiscal Multiplier

There are periods when fiscal actions by governments become the main driver of economic activity. For example, the fiscal expenditure and revenue responses to the economic implications of the Covid-19 pandemic in September 2020 amounted to 9 percent of GDP in advanced economies (see Lacey et al. (2021)). During such times, the fiscal multiplier becomes a key metric for macroeconomic forecasting and policy design, as it quantifies the impact of discretionary fiscal policy on output. Blanchard and Leigh (2013) argue that underestimation of fiscal multipliers led to persistent forecast errors of output for European economies in the beginning of the 2010s, especially in the early stages of the austerity measures implementation. Although the immediate consequences are evident, the failure to properly quantify the fiscal multiplier when implementing fiscal policy could have adverse long-term implications for the economy. For example, fiscal tightening can worsen the debt ratio in the short run, potentially leading the

government to implement additional rounds of tightening in an attempt to stabilize the debt (e.g., Eyraud and Weber (2013)). As a result, confidence in the fiscal authority decreases, and the economy may enter a vicious cycle of slow growth, deflation, and further tightening.

In the context of SVARs, the fiscal multiplier is usually computed as the ratio of the cumulative sum of the GDP response to the cumulative sum of the government expenditure response to an exogenous government expenditure shock (see Ramey (2019)). Identifying these shocks is not straightforward with sign restrictions, but DGP1 considered in the previous section provides a good example of a fiscal model identified with short-run exclusion restrictions. It consists of three endogenous variables ordered in the vector as follows: government expenditure, real GDP, and government revenue. The fiscal shock is identified through Cholesky decomposition of the variance-covariance matrix, which implies that a government expenditure shock does not affect real GDP and government revenues contemporaneously. In particular, using the previously introduced notation, the fiscal multiplier is given by

$$
m_h = \frac{\sum_{i=1}^h \theta_{2,1,i}}{\sum_{i=1}^h \theta_{1,1,i}} \cdot (Y/G),
$$

where $\theta_{2,1,i}$ denotes the impulse response of GDP in period i to an exogenous shock in fiscal expenditures, $\theta_{1,1,i}$ denotes the impulse response of government expenditure in period i with respect to the same shock, and Y/G denotes the average GDP-to-government-expenditure ratio over the entire estimation sample. Therefore, the fiscal multiplier is a non-linear function of two variables that are accumulated over time. If the error bands around the fiscal multiplier are to be constructed by considering the specifics of the economic question, they need to achieve multivariate and time-simultaneous coverage. In this case, the vector of structural parameters consists of the fiscal multiplier and the responses of government expenditure and real GDP over 20 quarters, but does not include the response of government revenue. In particular, the vector is given by:

$$
\boldsymbol{\theta} = \left(m_1, m_2, \ldots, m_H, \theta_{1,1,1}, \theta_{1,1,2}, \ldots, \theta_{1,1,H}, \theta_{2,1,1}, \theta_{2,1,2}, \ldots, \theta_{2,1,H} \right).
$$

The model is estimated with US data from 1939:Q1 to 2015:Q4 and Figure 3 illustrates the pointwise and joint error bands. The conventional pointwise error bands show that the fiscal multiplier might be higher than 1 only after 15 quarters with 68% probability and after 13 quarters with 90% probability. As expected, the Bonferroni and Sidák error bands are the most conservative. Since these methods tend to overstate the uncertainty, their implications are less important. On the other hand, the error bands computed with the sup-t method are wider even in comparison to the uncalibrated min-max error bands; therefore, the analysis will focus on the min-max (BDR) specification, as it was shown in the Monte Carlo exercise that this method achieves the desired coverage level.

A key observation is that the LQO calibration method generates error bands that are nearly identical to those produced by the BDR method. This similarity between the two methods highlights that the choice of calibration approach may not significantly affect the overall conclusions about the fiscal multiplier's behavior. Moreover, when comparing the width of the LQO and BDR error bands to the more conventional pointwise error bands, the min-max (BDR) bands are, on average, significantly wider—91% for the 68% credibility level and 51% for the 90% level. This additional width reflects the more conservative nature of the BDR method, suggesting a greater level of uncertainty about the multiplier's true value. Importantly, the joint error bands demonstrate that the multiplier is likely to exceed 1 earlier than previously expected,

containing the value of 1 within the credible region 4 quarters sooner for both credibility levels. This outcome calls for a re-assessment of fiscal policy, as it implies that government spending could have a more immediate and pronounced effect on output than standard pointwise error bands suggest. The need for careful policy design is thus underscored, as overlooking these nuances in error band width could lead to misinformed decisions regarding the timing and magnitude of fiscal interventions.

5.2   Forecasting

Structural impulse responses in VAR models are designed to trace the effect of a one-time shock to one of the variables on all other variables in the system over time. They provide detailed insights into the dynamic interactions within the system, illustrating how a shock propagates and influences each variable across different periods. In contrast, forecasts in VAR models aim to predict the future values of the variables based on their historical relationships. More formally, as shown in Section 3, any VAR model can be represented as:

$$\mathbf{y}_{t+l} = \mathbf{J}\mathbf{A}^{l+1}\mathbf{Y}_{t-1} + \sum_{b=0}^{l} \mathbf{J}\mathbf{A}^{b}\mathbf{J}'\mathbf{u}_{t+l-b}.$$

Therefore, the l-periods ahead forecast is given by:

$$\mathbb{E}\left(\mathbf{y}_{t+l}|\mathbf{y}_t, \mathbf{y}_{t-1}, \dots\right) = \mathbf{J}\mathbf{A}^{l+1}\mathbf{Y}_{t-1}$$

and the l-periods ahead IRF is defined as

$$\mathbb{E}\left(\mathbf{y}_{t+l}|\mathbf{u}_t = \mathbf{B}\varepsilon_t, \mathbf{y}_t, \mathbf{y}_{t-1}, \dots\right) = \mathbf{J}\mathbf{A}^{l+1}\mathbf{Y}_{t-1} + \mathbf{J}\mathbf{A}^{l}\mathbf{J}'\mathbf{B}\varepsilon_t.$$

This illustrates that impulse responses and forecasts in time series models are closely related, implying that methods used for joint inference on these quantities are likely to perform similarly. Given these similarities, the methods for joint inference are applied in a pseudo-out-of-sample forecasting exercise, despite not being formally evaluated in the simulation exercises of this paper.

For this exercise, I use two time series models known for their strong forecasting performance. First, I consider a time-varying bivariate model of inflation and survey-based inflation expectations proposed by Chan et al. (2018) (henceforth CCK18). In this model, the inflation gap is defined as the difference between actual inflation and the unobserved inflation trend. The inflation gap is assumed to follow an AR(1) process, while the inflation trend is modeled as a random walk. Additionally, there is a linear relationship between the survey-based inflation expectations and the inflation trend.[13]

The second model used in the forecasting exercise is a standard time-varying VAR with stochastic volatility, originally proposed by Primiceri (2005) and later refined by Del Negro and Primiceri (2015) (henceforth DP15). This VAR model includes three endogenous variables: the annualized quarter-on-quarter real GDP growth rate, the annualized quarter-on-quarter headline inflation rate, and the 1-year government bond yield. The performance of the joint inference methods in the forecasting exercise is evaluated by constructing TS error bands for inflation forecasts using CCK18 and TS and MV-TS error bands for inflation and GDP growth rates using DP15.

---

[13]Further details about the model can be found in Chan et al. (2018).

Both models are estimated using U.S. and euro area data, with samples starting in 1982:Q1 and 1999:Q1, respectively. The forecasts are evaluated through an expanding sample pseudo-out-of-sample forecasting exercise, beginning in 2010:Q1 and concluding in 2019:Q4.[14] The inflation rate is calculated using the Personal Consumption Expenditure (PCE) price index for the U.S. and the Harmonized Index of Consumer Prices (HICP) for the euro area. For the U.S., the survey-based inflation expectations are derived from the 10-year expected inflation estimated by the Federal Reserve Bank of Cleveland, while for the euro area, the 5-year expected inflation is reported by the Survey of Professional Forecasters of the European Central Bank. For DP15, in addition to the inflation rate, I use the 1-year U.S. Treasury bond yield and the 1-year German government bond yield, along with the real quarter-on-quarter GDP growth rates for both the U.S. and euro area.

Before analyzing the results of the entire sample for each joint inference method, I first focus on the graphical comparison of the error bands. In the aftermath of the Covid pandemic, inflation in both the U.S. and the euro area experienced a strong and persistent increase, driven by a combination of supply and demand factors. These effects were further exacerbated by the energy crisis triggered by Russia's invasion of Ukraine in February 2022.

A simple projection conducted at the beginning of 2022:Q2 using CCK18 for the U.S. and euro area is presented in Figure 4. The figures for the TS error bands (calculated with min-max estimator under absolute loss function) in the U.S. show that, even after such a significant disruption, the inflation recovery path would have been well-contained within the 30% credible band, except for the second quarter of 2022. This is somewhat surprising, as the model is generally expected to produce better short-term forecasts. On the other hand, the European economy was more severely affected by energy supply disruptions, causing inflation to remain elevated for an extended period. As a result, the PW error band for the headline inflation rate in the euro area is less useful as a measure of the recovery path. As the econometric model is trend-reverting by nature, both PW and TS error band indicate downside risk for the inflation, but the realized path of the series is well contained only within the TS error bands for both countries.

The results of the projection exercises conducted with the DP15 model are presented in Figure 5. It shows a comparison between the PW and MV error bands (using a calibrated min-max estimator under absolute loss function) of the joint projections of 2-quarters ahead inflation and GDP q-o-q growth rates during the period 2010:Q1-2019:Q4. The results show that the 90% MV error bands are not much wider in comparison to the corresponding PW error bands, but the difference between both starts to increase as the credibility level falls. This indicates that underestimation of uncertainty for lower credibility levels will possibly be much larger. A possible cause for that is the nature of the projected variables, namely growth rates. A further important implication of the calibration step implemented for the min-max estimators is that it stabilizes the results across different initialization of the forecasting exercises. This is not caused by not converged estimation procedure, but rather infeasible representation of the joint distribution function.

The coverage rates of the PW error bands for the CCK18 model are presented in Appendix D, and in particular Table D.1 for United States and D.2 for the euro area. The results of the out-of-sample forecasting exercise reveal that even the PW error bands for United States slightly overestimate the uncertainty especially for credible levels of 30% and 60%. On the other hand, the euro area results indicate much better performance of the model. Nonetheless, if the applied

---

[14]The sample size ends in 2019 to simplify the econometric modeling of the post-Covid heteroskedasticity.

researcher reports PW error bands, but requires TS inference the uncertainty will be significantly underestimated. The TS error bands evaluated against TS coverage levels are presented in Figure D.3 and D.4. The results indicate same conclusions regarding the performance of the models as for the error bands of the IRFs evaluated with the Monte Carlo exercises. The difference, however, is driven by the general (pointwise) forecasting performance of the model for the U.S. and euro area inflation projections. Furthermore, the TS error bands for 8-quarters ahead and smaller credibility level are significantly wider than the PW equivalent, which could make them inapplicable in practical applications.

Finally, I evaluate the results of the out-of-sample exercise of the DP15 model. The PW error bands evaluated under PW, MV and TS coverage statistics are shown in Table D.5 and D.6. Both table indicate much better PW coverage of the inflation forecasts in comparison to the GDP growth. Furthermore, there is also difference in the performance across countries - with the projections for United States performing slightly better in comparison to the euro area. The coverage levels of the MV and TS again indicate that the PW error bands are inadequate representations of the joint uncertainty.

The results for the TS error bands are presented in Table D.7 and D.8. As expected the performance of the U.S. forecasting is slightly better. The coverage rate of the min-max estimators under absolute, quadratic and angular loss functions even with BDR calibration step significantly worse in comparison to the Chebyshev loss functions. The estimates under the latter loss function are also significantly tighter. For example, U.S. results indicate that the 60% error band for $h = 4$ and GDP growth rate (i.e., y_2) estimated with min-max estimator under absolute loss function with BDR calibration covers 90% and is 188% of the PW error band. On the other hand the same measures under Chebyshev loss are 79% and 83%. The results in both tables (for United States and euro area) show similar performance in terms of magnitudes of the Chebyshev loss function. The performance of the sup-t error bands is comparable to the Chebyshev min-max estimator. The results for the MV error bands are shown on Table D.9 for United States and Table D.10 for the euro area. As the MV error bands involve less elements to conduct joint inference about, the performance of the min-max estimator under the absolute and Chebyshev loss are comparable among each other but also the sup-t estimator is good alternative.

In summary, the joint error bands about the forecasts have similar performance in comparison to the joint error bands about the IRFs. Since this represents a real-world application rather than an exercise in simulated environment, the performance of the models even in conventional pointwise sense is rather mixed. The results, however, indicate that the applied researcher could use either of the methods - sup-t, min-max under absolute loss with BDR calibration or plain min-max under Chebyshev loss and report the other error bands in robustness exercises.

Figure 3
Pointwise and joint inference about the fiscal multiplier
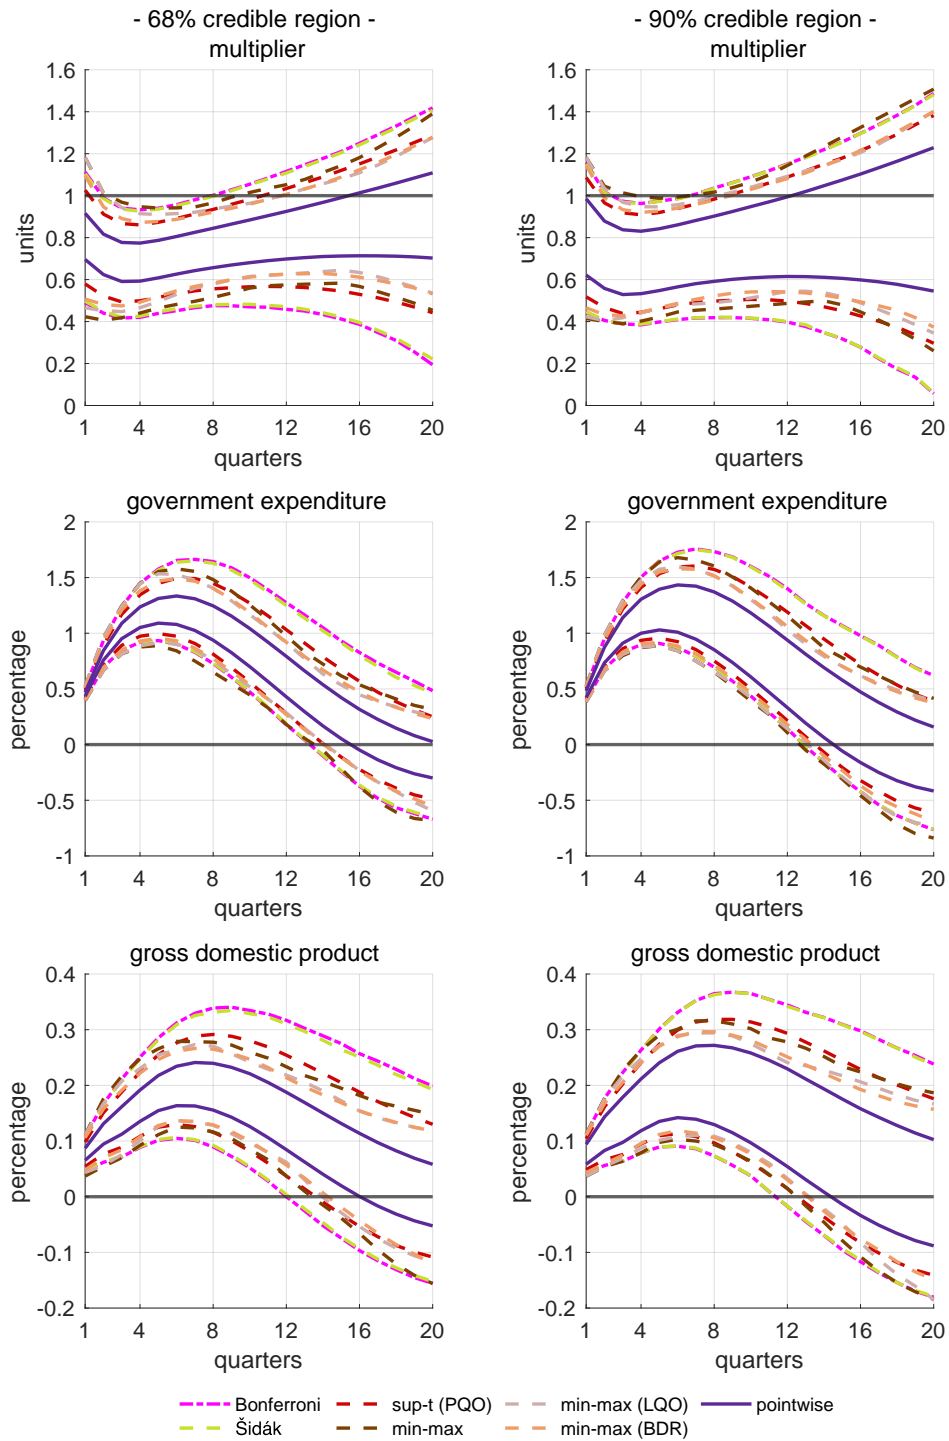
Notes: The lines depict pointwise and MV-TS joint error bands for government expenditure, gross domestic product, and the fiscal multiplier. For the sake of simplicity, the min-max estimators are computed only with absolute loss function. The figures in the left column illustrate the results for the 68% credible region, while the figures in the right column illustrate the results for the 90% credible region.

Figure 4

CCK18: inflation forecasts for United States and the euro area in 2022:Q1

- PW inference -
headline inflation

- TS inference -
headline inflation

(a) United States

- PW inference -
headline inflation

- TS inference -
headline inflation
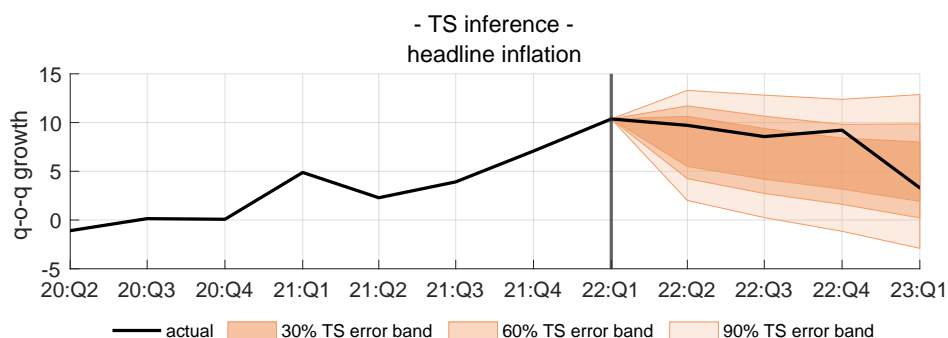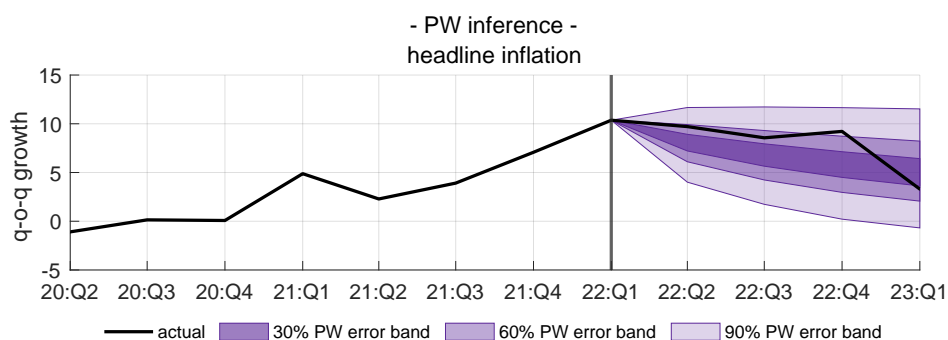
(b) euro area

Notes: The shaded areas depict pointwise and TS joint error bands for inflation forecasts of the CCK18 model. The TS joint error bands are computed with min-max (BDR) and absolute loss function. Subfigure (a) shows the results for United States and subfigure (b) shows the results for the euro area. The models are estimated with data up until 2022:Q1.

Figure 5

DP15: 2-quarters ahead inflation and gross domestic product forecasts for United States and the euro area in 2010:Q1-2019:Q4
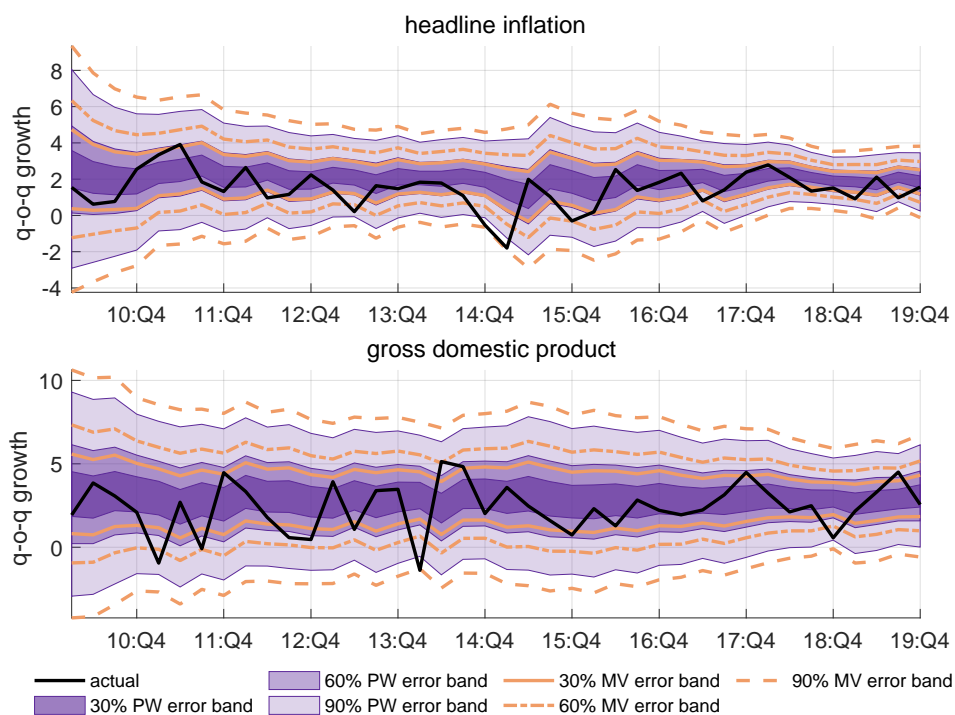
(a) United States
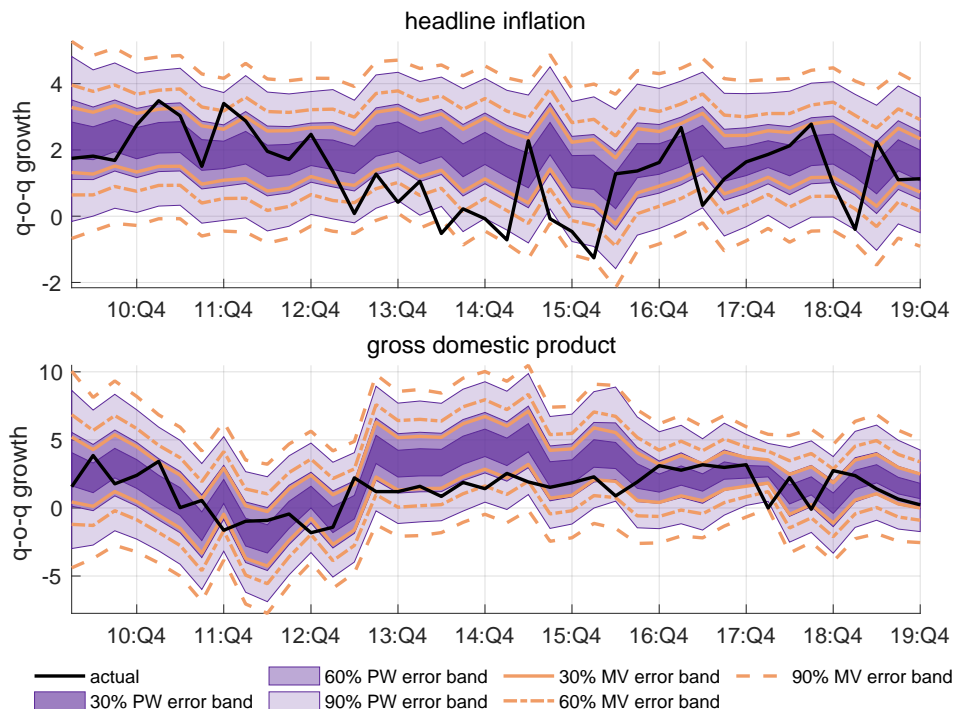
(b) euro area

Notes: The shaded areas and the dashed lines depict pointwise and MV joint error bands for 2-quarters ahead forecasts of inflation and gross domestic product of DP15. The MV joint error bands are computed with min-max (BDR) and absolute loss function. Subfigure (a) shows the results for United States and subfigure (b) shows the results for the euro area.

6 Conclusion

This paper highlights the critical need to accurately quantify uncertainty in impulse responses and forecasts when using VAR models. I discuss the construction of context-specific error bands, tailored to the economic question at hand, as different inquiries may require distinct set of structural parameters or forecasts to effectively capture the underlying uncertainty. Through series of simulation experiments, I demonstrate that conventional pointwise quantiles significantly underestimate the uncertainty in cases that require simultaneous inference about multiple parameters. The results illustrate that joint inference methods, such as the sup-t (see Montiel Olea and Plagborg-Møller (2019)) and min-max (see Inoue and Kilian (2022) and Akram et al. (2016)) estimators, offer substantial improvements in achieving the desired coverage levels across horizons and variables, compared to pointwise methods.

The joint error bands, however, are considerably wider, especially for economic questions involving multiple variables and horizons. These findings underline the potential risks associated with relying on traditional pointwise methods, which may lead to overly confident conclusions. On the other hand, the joint inference methods I evaluate provide a more reliable reflection of the underlying uncertainty, making them suitable for applied research where precise estimation of impulse responses and forecasts is crucial. This is particularly relevant for policy makers, when understanding the full scope of uncertainty is essential for informed decision-making.

Moreover, the carefully designed simulation exercise reveals that the min-max estimators, which rely on vector-valued loss functions can be overly conservative. Contrary to Inoue and Kilian (2022), I contend that the choice of the loss function plays a crucial role in determining the width of the error bands, and applied econometricians must be meticulous in selecting it. For example, the angular loss function consistently performed the worst, producing the widest error bands. In contrast, calibrated versions of the min-max estimator with quadratic or absolute loss functions performed much better, offering narrower bands while still achieving the desired coverage. Notably, the Chebyshev loss function proposed by Akram et al. (2016) stands out, as it achieves the target coverage level straight off the shelf, without requiring any additional calibration steps.

In addition, this study introduces simulation-based calibration to the field of macroeconometrics - a technique commonly used in other areas of econometrics to assess the performance of Bayesian estimators. Specifically, it allows for evaluating whether the posterior distribution is properly sampled by the proposed algorithm and the implemented software (see Gelman et al. (1996) and Talts et al. (2020)).

Unfortunately, the post-Covid period is characterized by increased volatility, resulting in excessively wide error bands. Consequently, the implications of this paper are less appealing for practitioners and policymakers. However, understating uncertainty is undesirable when investigating structural relationships or producing robust economic projections. Therefore, applied economists must carefully model heteroskedasticity (see Lenza and Primiceri (2022) and Carriero et al. (2024)) and rely on more precise estimation techniques. These methods could relate to the identification methodology—for example, using narrative information to shrink error bands with shocks identified via sign restrictions (Antolín-Díaz and Rubio-Ramírez (2018)) or improving reduced-form estimates by estimating economic relationships in data-rich environment (Kerssenfischer (2019)).

References

Akram, F., Binning, A., and Maih, J. (2016). Joint Prediction Bands for Macroeconomic Risk Management. Working Paper 7-2016, Norges Bank.

Antolín-Díaz, J. and Rubio-Ramírez, J. F. (2018). Narrative Sign Restrictions for SVARs. *American Economic Review*, 108(10):2802–2829.

Arias, J., Rubio Ramírez, J., and Waggoner, D. (2022). Uniform Priors for Impulse Responses. Working Paper 22-30, Federal Reserve Bank of Philadelphia.

Arias, J. E., Caldara, D., and Rubio-Ramírez, J. F. (2019). The Systematic Component of Monetary Policy in SVARs: An Agnostic Identification Procedure. *Journal of Monetary Economics*, 101:1–13.

Bernardo, J. M. (2011). Integrated Objective Bayesian Estimation and Hypothesis Testing. In Bernardo, J. M., Bayarri, M. J., Berger, J. O., Bernardo, J. M., Bayarri, M. J., Berger, J. O., Dawid, A. P., Heckermann, D., Smith, A. F. M., and West, M., editors, *Bayesian Statistics 9: Proceedings of the Ninth Valencia International Meeting; June 3-8, 2010*, Oxford Science Publications. Oxford University Press, Oxford.

Blanchard, O. J. (1989). A Traditional Interpretation of Macroeconomic Fluctuations. *The American Economic Review*, 79(5):1146–1164.

Blanchard, O. J. and Leigh, D. (2013). Growth Forecast Errors and Fiscal Multipliers. *American Economic Review*, 103(3):117–20.

Blanchard, O. J. and Quah, D. (1989). The Dynamic Effects of Aggregate Demand and Supply Disturbances. *The American Economic Review*, 79(4):655–673.

Bonferroni, C. (1936). *Teoria Statistica Delle Classi e Calcolo Delle Probabilità*. Pubblicazioni Del R. Istituto Superiore Di Scienze Economiche e Commerciali Di Firenze. Seeber.

Carriero, A., Clark, T. E., Marcellino, M., and Mertens, E. (2024). Addressing COVID-19 Outliers in BVARs with Stochastic Volatility. *Review of Economics and Statistics*, 106(5):1403–1417.

Chan, J. C., Clark, T. E., and Koop, G. (2018). A New Model of Inflation, Trend Inflation, and Long-Run Inflation Expectations. *Journal of Money, Credit and Banking*, 50(1):5–53.

Cook, S. R., Gelman, A., and Rubin, D. B. (2006). Validation of Software for Bayesian Models Using Posterior Quantiles. *Journal of Computational and Graphical Statistics*, 15(3):675–692.

Del Negro, M. and Primiceri, G. E. (2015). Time Varying Structural Vector Autoregressions and Monetary Policy: A Corrigendum. *The Review of Economic Studies*, 82(4):1342–1345.

Eichenbaum, M. and Evans, C. L. (1995). Some Empirical Evidence on the Effects of Shocks to Monetary Policy on Exchange Rates. *The Quarterly Journal of Economics*, 110(4):975–1009.

Eyraud, L. and Weber, A. (2013). The Challenge of Debt Reduction during Fiscal Consolidation. Working Paper 13-67, International Monetary Fund.

Forbes, K., Hjortsoe, I., and Nenova, T. (2018). The shocks matter: Improving our estimates of exchange rate pass-through. *Journal of International Economics*, 114:255–275.

Gabry, J., Simpson, D., Vehtari, A., Betancourt, M., and Gelman, A. (2019). Visualization in bayesian workflow. *Journal of the Royal Statistical Society Series A: Statistics in Society*, 182(2):389–402.

Gafarov, B., Meier, M., and Montiel Olea, J. L. (2018). Delta-method inference for a class of set-identified SVARs. *Journal of Econometrics*, 203(2):316–327.

Gelman, A., Meng, X.-L., and Stern, H. (1996). Posterior Predictive Assesment of Model Fitness via Realized Discrepancies. *Statistica Sinica*, 6(4):733–760.

Güntner, J., Reif, M., and Wolters, M. (2024). Sudden Stop: Supply and Demand Shocks in the German Natural Gas Market. *Journal of Applied Econometrics (forthcoming)*.

Herrera, A. M. and Rangaraju, S. K. (2020). The Effect of Oil Supply Shocks on US Economic Activity: What Have We Learned? *Journal of Applied Econometrics*, 35(2):141–159.

Inoue, A., Jordà, Ò., and Kuersteiner, G. M. (2023). Inference for Local Projections.

Inoue, A. and Kilian, L. (2013). Inference on Impulse Response Functions in Structural VAR Models. *Journal of Econometrics*, 177(1):1–13.

Inoue, A. and Kilian, L. (2016). Joint Confidence Sets for Structural Impulse Responses. *Innovations in Multiple Time Series Analysis*, 192(2):421–432.

Inoue, A. and Kilian, L. (2022). Joint Bayesian Inference About Impulse Responses in VAR Models. *Special Issue: The Econometrics of Macroeconomic and Financial Data*, 231(2):457–476.

Jarociński, M. and Karadi, P. (2020). Deconstructing Monetary Policy Surprises—The Role of Information Shocks. *American Economic Journal: Macroeconomics*, 12(2):1–43.

Jordà, Ò. (2009). Simultaneous Confidence Regions for Impulse Responses. *The Review of Economics and Statistics*, 91(3):629–647.

Kaido, H., Molinari, F., and Stoye, J. (2019). Confidence Intervals for Projections of Partially Identified Parameters. *Econometrica*, 87(4):1397–1432.

Kerssenfischer, M. (2019). ThePuzzling Effects of Monetary Policy in VARs: Invalid Identification or Missing Information? *Journal of Applied Econometrics*, 34(1):18–25.

Kilian, L. and Lütkepohl, H. (2017). *Structural Vector Autoregressive Analysis*. Cambridge University Press, 1 edition.

Kilian, L. and Zhou, X. (2022). The Impact of Rising Oil Prices on U.S. Inflation and Inflation expectations in 2020–23. *Energy Economics*, 113:106228.

Knüppel, M., Krüger, F., and Pohle, M.-O. (2022). Score-based Calibration Testing for Multivariate Forecast Distributions. Working Paper 50/2022, Deutsche Bundesbank.

Koehler, E., Brown, E., and Haneuse, S. J.-P. A. (2009). On the Assessment of Monte Carlo Error in Simulation-Based Statistical Analyses. *The American Statistician*, 63(2):155–162.

Lacey, E., Massad, J., and Utz, R. (2021). *A Review of Fiscal Policy Responses to COVID-19*. World Bank.

Lenza, M. and Primiceri, G. E. (2022). How to Estimate a Vector Autoregression After March 2020. *Journal of Applied Econometrics*, 37(4):688–699.

Little, R. J. (2006). Calibrated Bayes: A Bayes/Frequentist Roadmap. *The American Statistician*, 60(3):213–223.

Lütkepohl, H. (2005). *New Introduction to Multiple Time Series Analysis*. New York : Springer, Berlin.

Lütkepohl, H., Staszewska-Bystrova, A., and Winker, P. (2015). Comparison of Methods for Constructing Joint Confidence Bands for Impulse Response Functions. *International Journal of Forecasting*, 31(3):782–798.

Lütkepohl, H., Staszewska-Bystrova, A., and Winker, P. (2018). Calculating Joint Confidence Bands for Impulse Response Functions Using Highest Density Regions. *Empirical Economics*, 55(4):1389–1411.

Montiel Olea, J. L. and Plagborg-Møller, M. (2019). Simultaneous Confidence Bands: Theory, Implementation, and an Application to SVARs. *Journal of Applied Econometrics*, 34(1):1–17.

Paul, P. (2020). The Time-Varying Effect of Monetary Policy on Asset Prices. *The Review of Economics and Statistics*, 102(4):690–704.

Peersman, G. (2005). What caused the early millennium slowdown? Evidence based on vector autoregressions. *Journal of Applied Econometrics*, 20(2):185–207.

Primiceri, G. E. (2005). Time Varying Structural Vector Autoregressions and Monetary Policy. *The Review of Economic Studies*, 72(3):821–852.

Ramey, V. A. (2019). Ten Years after the Financial Crisis: What Have We Learned from the Renaissance in Fiscal Research? *Journal of Economic Perspectives*, 33(2):89–114.

Sidak, Z. (1967). Rectangular Confidence Regions for the Means of Multivariate Normal Distributions. *Journal of the American Statistical Association*, 62(318):626–633.

Sims, C. A. and Zha, T. (1999). Error Bands for Impulse Responses. *Econometrica*, 67(5):1113–1155.

Staszewska, A. (2007). Representing Uncertainty About Response Paths: The Use of Heuristic Optimisation Methods. *Computational Statistics & Data Analysis*, 52(1):121–132.

Talts, S., Betancourt, M., Simpson, D., Vehtari, A., and Gelman, A. (2020). Validating Bayesian Inference Algorithms with Simulation-Based Calibration.

Thorarinsdottir, T. L., Scheuerer, M., and Heinz, C. (2016). Assessing the Calibration of High-Dimensional Ensemble Forecasts Using Rank Histograms. *Journal of Computational and Graphical Statistics*, 25(1):105–122.

Wolf, M. and Wunderli, D. (2015). Bootstrap Joint Prediction Regions. *Journal of Time Series Analysis*, 36(3):352–376.

Zhou, X. (2020). Refining the Workhorse Oil Market Model. *Journal of Applied Econometrics*, 35(1):130–140.

Appendix A   Data

Figure A.1
DGP1 and DGP2: endogenous variables
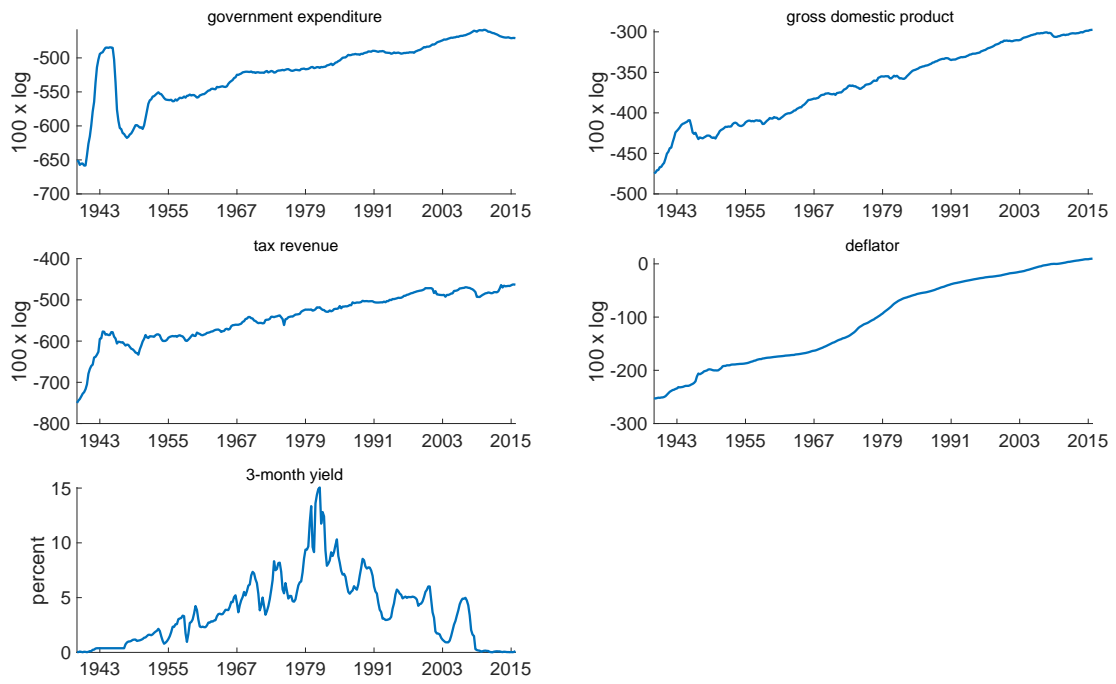
Figure A.2
DGP3: endogenous variables
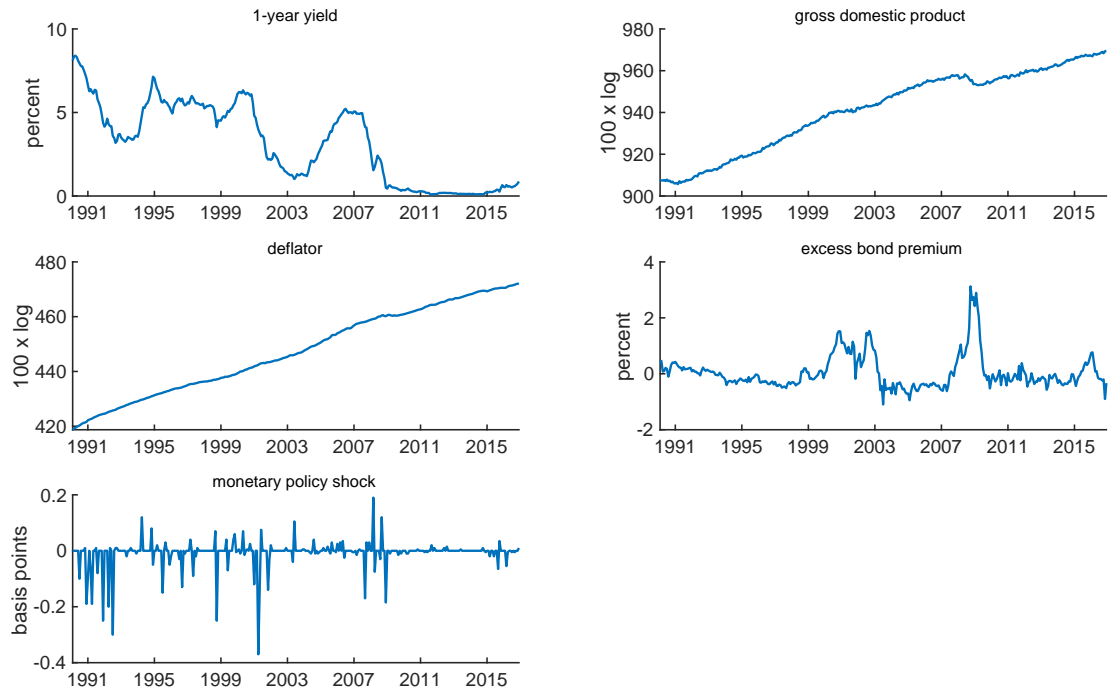
Figure A.3
DGP4: endogenous variables

gross national product (growth rate)

unemployment rate

Appendix B Model Simulations

Figure B.1
DGP1: simulated IRFs
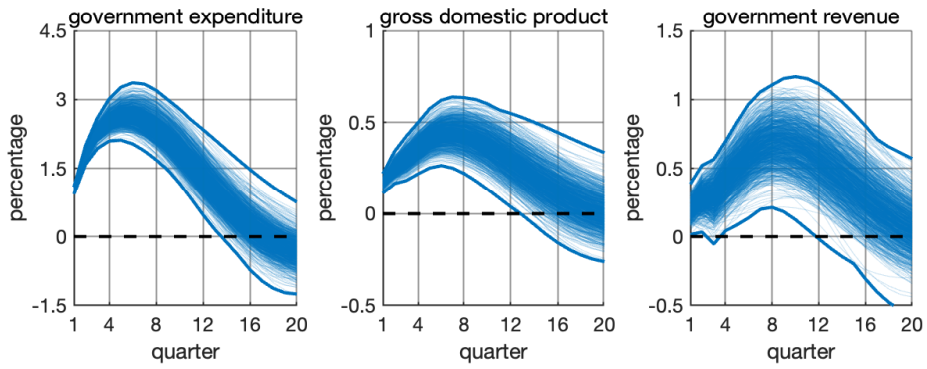
Notes: Blue solid lines depict IRFs simulated from the data generating process.

Figure B.2
DGP2: simulated IRFs
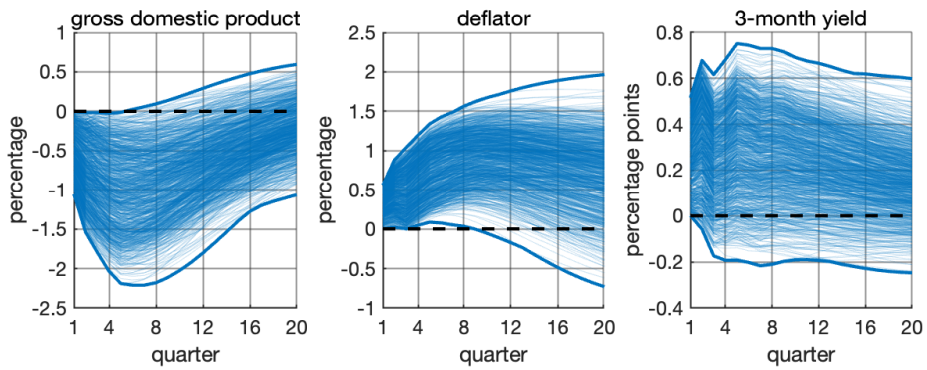
Notes: Blue solid lines depict IRFs simulated from the data generating process.

Figure B.3
DGP3: simulated IRFs
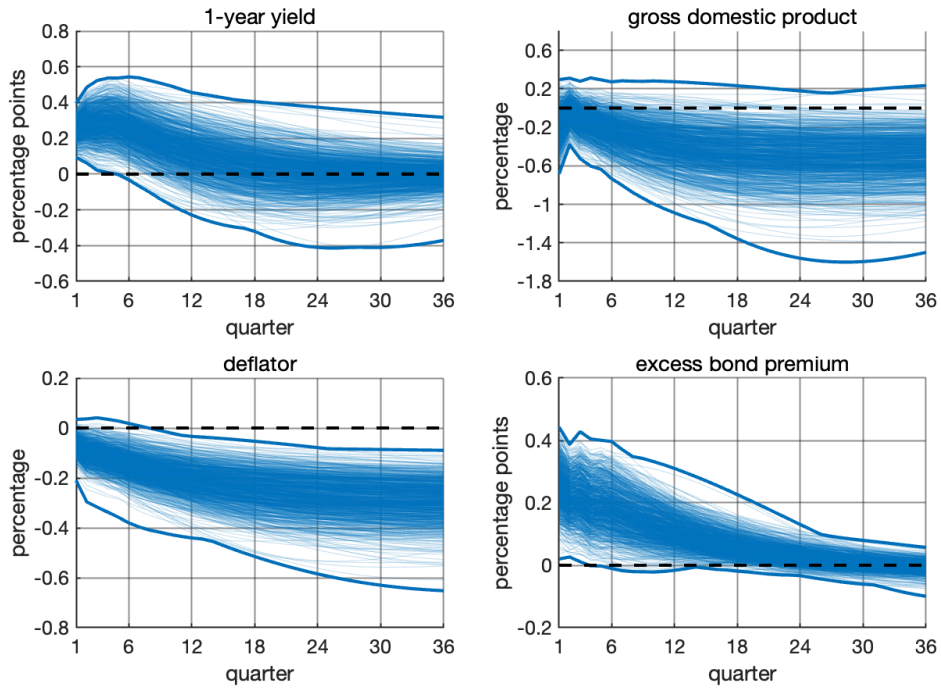
Notes: Blue solid lines depict IRFs simulated from the data generating process.

Figure B.4
DGP4: simulated IRFs
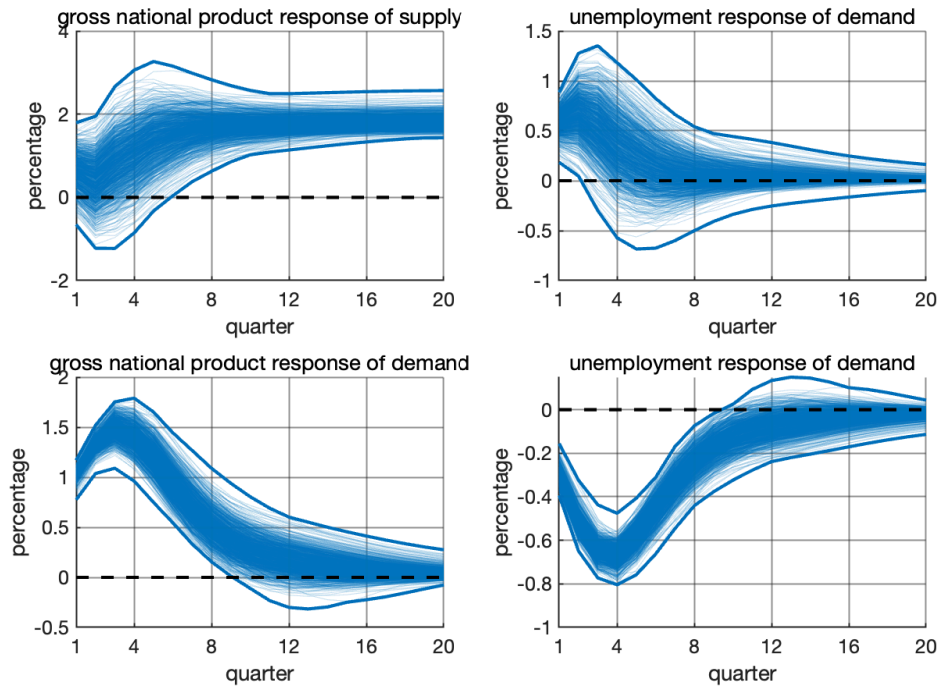
Notes: Blue solid lines depict IRFs simulated from the data generating process.

Appendix C   Shape and Comovement

Figure C.1
Average rank statistics of the TS relation
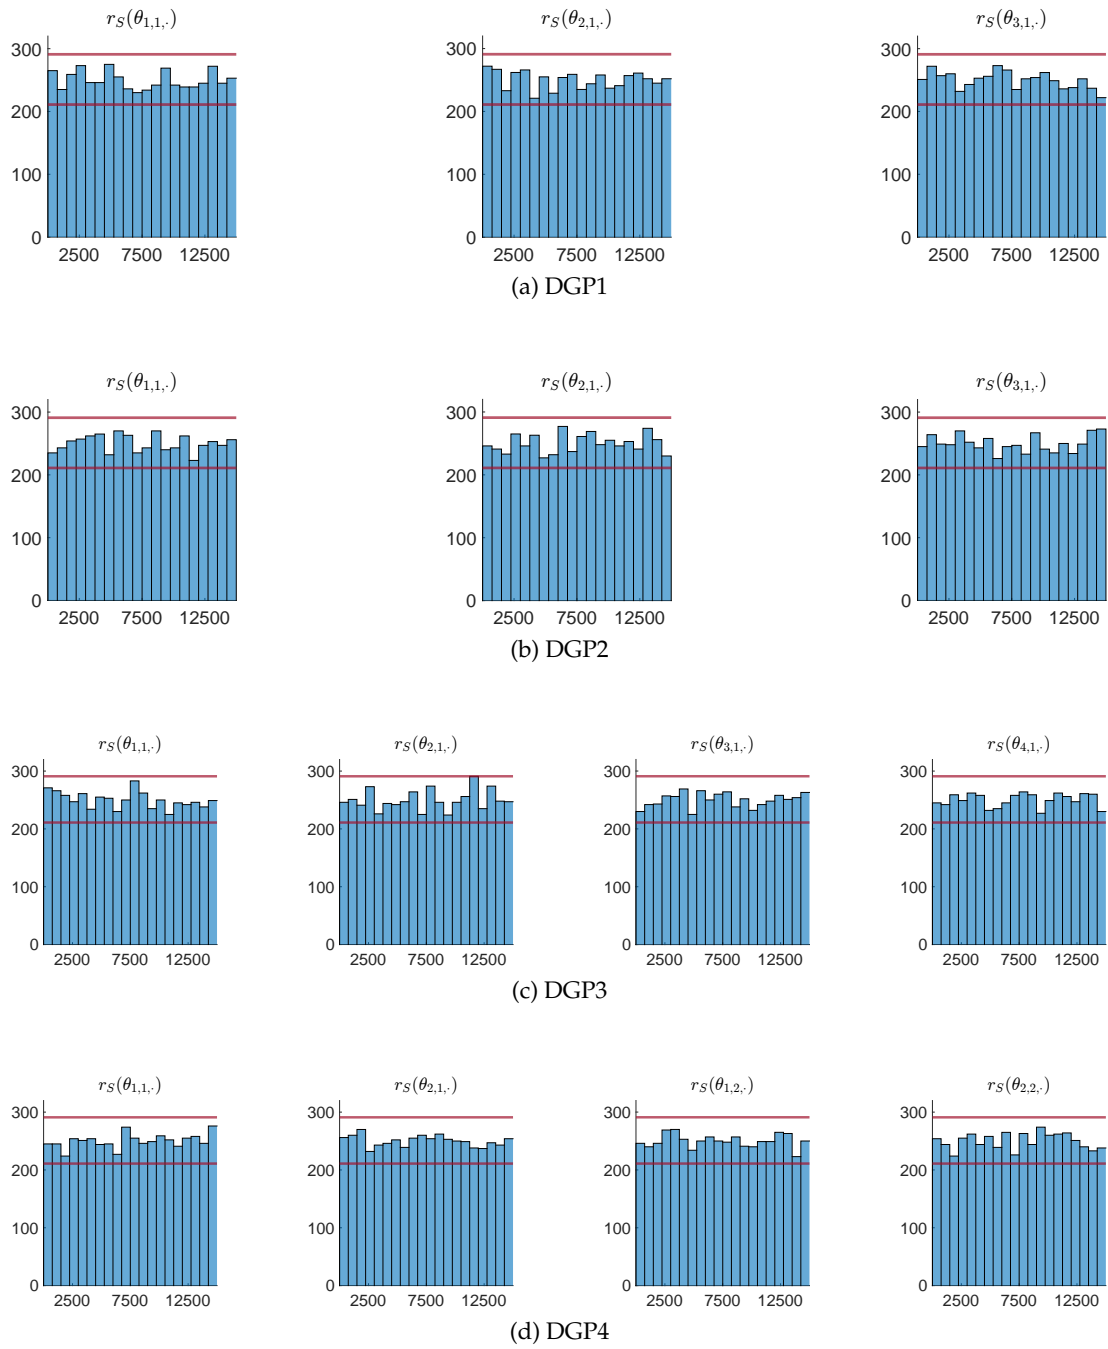
(a) DGP1

(b) DGP2

(c) DGP3

(d) DGP4

Notes: Blue bars denote the average rank. The red solid lines denote the 0.005 and 0.995 percentiles of Binomial (5000, 1/20) distribution.

Figure C.2
Average rank statistics of the MV-TS relation

(a) DGP1 (b) DGP2 (c) DGP3

(d) DGP4

Notes: Blue bars denote the average rank. The red solid lines denote the 0.005 and 0.995 percentiles of Binomial$(5000, 1/20)$ distribution.

Figure C.3
Average rank statistics of the MV-TS-CS relation

(a) DGP4

Notes: Blue bars denote the average rank. The red solid lines denote the 0.005 and 0.995 percentiles of Binomial$(5000, 1/20)$ distribution.

Appendix D   Out-of-sample Forecasting Exercise

Table D.1
CCK18: pointwise (Π^{PW}) and time-simultenous (Π^{TS}) coverage level of PW error bands for United States

Coverage level	30% PW error band			60% PW error band			90% PW error band		
	h = 2	h = 4	h = 8	h = 2	h = 4	h = 8	h = 2	h = 4	h = 8
Π^{PW}	35,90	46,15	48,72	74,36	82,05	69,23	92,31	92,31	84,62
Π^{TS}	10,26	2,56	0,00	53,85	38,46	15,38	89,74	82,05	61,54

Notes: Pointwise (Π^{MV}) and time-simultenous (Π^{TS}) coverage levels of 68% and 90% PW credible bands for 2, 4, and 8 quarters ahead out-of-sample headline inflation forecast in United States using CCK18.

Table D.2
CCK18: pointwise (Π^{PW}) and time-simultenous (Π^{TS}) coverage level of PW error bands for euro area

Coverage level	30% PW error band			60% PW error band			90% PW error band		
	h = 2	h = 4	h = 8	h = 2	h = 4	h = 8	h = 2	h = 4	h = 8
Π^{PW}	20,51	30,77	30,77	61,54	53,85	51,28	94,87	94,87	82,05
Π^{TS}	10,26	2,56	0,00	35,90	12,82	5,13	87,18	79,49	48,72

Notes: Pointwise (Π^{MV}) and time-simultenous (Π^{TS}) coverage levels of 68% and 90% PW credible bands for 2, 4, and 8 quarters ahead out-of-sample headline inflation forecast in the euro area using CCK18.

Table D.3

CCK18: time-simultaneous coverage level (Π^{TS}) of TS error bands for United States

Method	30% TS error band						60% TS error band						90% TS error band					
	h = 2		h = 4		h = 8		h = 2		h = 4		h = 8		h = 2		h = 4		h = 8	
	Π	Δ	Π	Δ	Π	Δ	Π	Δ	Π	Δ	Π	Δ	Π	Δ	Π	Δ	Π	Δ
Bonferroni	69,23	150,53	79,49	277,81	61,54	402,48	89,74	57,36	82,05	111,04	66,67	165,82	92,31	23,69	84,62	47,61	79,49	74,29
Šidák	46,15	98,75	66,67	207,18	61,54	320,69	82,05	48,15	79,49	97,35	61,54	149,90	92,31	22,87	84,62	46,39	79,49	72,65
sup-t (PQO)	43,59	87,74	61,54	176,93	61,54	262,24	82,05	42,59	79,49	84,56	61,54	125,44	92,31	20,51	84,62	41,01	76,92	62,25
min-max: absolute loss (median)	74,36	160,78	82,05	393,17	69,23	554,43	89,74	98,94	84,62	219,72	82,05	289,94	94,87	67,20	100,00	120,21	94,87	149,65
min-max (LQO): absolute loss (median)	43,59	87,37	61,54	178,51	58,97	265,51	79,49	42,71	79,49	84,24	61,54	129,01	89,74	20,57	84,62	41,80	79,49	65,26
min-max (BDR): absolute loss (mean)	43,59	87,44	58,97	176,70	61,54	261,57	82,05	42,47	79,49	84,10	61,54	125,29	92,31	20,50	84,62	40,99	74,36	62,78
min-max: quadratic loss (mean)	48,72	109,75	79,49	250,79	61,54	394,66	89,74	59,05	82,05	131,54	69,23	205,40	94,87	33,76	87,18	73,08	89,74	109,21
min-max (LQO): quadratic loss (mean)	43,59	87,56	58,97	177,11	58,97	263,23	84,62	42,48	79,49	84,49	61,54	126,34	89,74	20,63	84,62	40,88	76,92	63,80
min-max (BDR): quadratic loss (median)	43,59	87,53	58,97	176,62	61,54	261,34	82,05	42,47	79,49	84,20	61,54	125,36	89,74	20,49	84,62	40,97	74,36	62,61
min-max: angular loss (median)	84,62	745,85	71,79	754,24	61,54	856,82	84,62	333,78	79,49	361,90	69,23	439,05	92,31	189,24	87,18	200,65	100,00	226,73
min-max (LQO): angular loss (median)	76,92	456,24	53,85	408,06	46,15	620,34	79,49	152,54	74,36	300,67	61,54	321,41	92,31	128,58	84,62	142,26	69,23	149,94
min-max (BDR): angular loss (median)	38,46	116,17	46,15	250,05	51,28	335,85	79,49	75,80	74,36	124,34	61,54	151,37	89,74	29,16	84,62	47,31	71,79	67,51
min-max: Chebyshev loss (median)	43,59	87,73	58,97	177,02	58,97	262,34	82,05	42,45	79,49	84,42	61,54	125,52	92,31	20,46	84,62	40,93	76,92	62,80
min-max (LQO): Chebyshev loss (median)	43,59	87,73	58,97	177,02	58,97	262,34	82,05	42,45	79,49	84,42	61,54	125,52	92,31	20,46	84,62	40,93	76,92	62,80
min-max (BDR): Chebyshev loss (median)	43,59	87,73	58,97	177,02	58,97	262,34	82,05	42,45	79,49	84,42	61,54	125,52	92,31	20,46	84,62	40,93	76,92	62,80

Notes: Time-simultenous (Π^{TS}) coverage level of 68% and 90% TS credible bands for 2, 4, and 8 quarters ahead out-of-sample headline inflation forecast in United States using CCK18.

<div align="center">

Table D.4

CCK18: time-simultaneous coverage level (Π^{TS}) of TS error bands for euro area

</div>

Method	30% TS error band						60% TS error band						90% TS error band					
	h = 2		h = 4		h = 8		h = 2		h = 4		h = 8		h = 2		h = 4		h = 8	
	Π	Δ	Π	Δ	Π	Δ	Π	Δ	Π	Δ	Π	Δ	Π	Δ	Π	Δ	Π	Δ
Bonferroni	46,15	146,03	43,59	267,14	56,41	381,51	61,54	55,32	79,49	105,20	82,05	155,41	97,44	21,92	100,00	43,50	97,44	68,16
Šidák	30,77	96,43	30,77	200,92	38,46	307,71	56,41	46,38	76,92	92,65	79,49	140,72	97,44	21,08	100,00	42,46	97,44	66,60
sup-t (PQO)	25,64	83,24	25,64	166,72	30,77	247,34	53,85	39,51	53,85	77,26	53,85	114,74	97,44	18,30	100,00	35,93	97,44	54,12
min-max: absolute loss (median)	41,03	153,96	79,49	370,35	87,18	513,74	79,49	94,99	100,00	203,17	97,44	254,94	100,00	66,71	100,00	106,73	97,44	124,65
min-max (LQO): absolute loss (median)	25,64	82,57	23,08	165,36	30,77	254,46	53,85	39,49	53,85	77,66	51,28	116,30	100,00	19,12	100,00	37,33	94,87	56,64
min-max (BDR): absolute loss (mean)	25,64	82,99	23,08	165,86	33,33	246,09	53,85	39,35	58,97	76,52	48,72	113,85	97,44	18,31	100,00	35,96	97,44	53,62
min-max: quadratic loss (mean)	30,77	104,68	35,90	236,90	51,28	374,39	58,97	55,98	87,18	123,57	92,31	190,95	100,00	33,29	100,00	69,16	97,44	99,72
min-max (LQO): quadratic loss (mean)	25,64	82,68	25,64	166,33	28,21	247,51	53,85	39,37	56,41	76,85	53,85	114,31	100,00	19,01	100,00	36,68	97,44	54,99
min-max (BDR): quadratic loss (median)	25,64	82,84	23,08	165,85	28,21	246,32	53,85	39,36	58,97	76,59	51,28	113,85	100,00	18,50	100,00	36,14	97,44	53,51
min-max: angular loss (median)	61,54	652,78	41,03	677,53	28,21	787,46	66,67	278,04	53,85	303,64	79,49	375,43	97,44	134,80	100,00	159,17	100,00	183,21
min-max (LQO): angular loss (median)	56,41	405,78	28,21	370,68	12,82	609,20	58,97	142,08	41,03	239,31	30,77	306,45	84,62	100,00	94,87	116,02	87,18	135,11
min-max (BDR): angular loss (median)	25,64	105,01	25,64	226,90	20,51	323,87	53,85	62,40	41,03	113,42	43,59	144,94	84,62	31,72	94,87	44,57	97,44	60,26
min-max: Chebyshev loss (median)	25,64	83,35	25,64	166,70	28,21	247,84	53,85	39,44	53,85	76,88	51,28	114,52	97,44	18,29	100,00	35,90	97,44	54,03
min-max (LQO): Chebyshev loss (median)	25,64	83,35	25,64	166,70	28,21	247,84	53,85	39,44	53,85	76,88	51,28	114,52	97,44	18,29	100,00	35,90	97,44	54,03
min-max (BDR): Chebyshev loss (median)	25,64	83,35	25,64	166,70	28,21	247,84	53,85	39,44	53,85	76,88	51,28	114,52	97,44	18,29	100,00	35,90	97,44	54,03

Notes: Time-simultenous (Π^{TS}) coverage level of 68% and 90% TS credible bands for 2, 4, and 8 quarters ahead out-of-sample headline inflation forecast in the euro area using CCK18.

Table D.5
DP15: various coverage levels of PW error bands for United States

Coverage level	Variable	30% PW error band			60% PW error band			90% PW error band		
		h = 2	h = 4	h = 8	h = 2	h = 4	h = 8	h = 2	h = 4	h = 8
Π^{PW}	y_1	33,33	30,77	30,77	69,23	74,36	66,67	94,87	92,31	84,62
	y_2	41,03	48,72	46,15	76,92	74,36	74,36	97,44	92,31	89,74
Π^{MV}	-	10,26	17,95	15,38	51,28	56,41	56,41	92,31	87,18	79,49
Π^{TS}	y_1	7,69	2,56	0,00	48,72	28,21	17,95	92,31	84,62	64,10
	y_2	15,38	5,13	0,00	58,97	35,90	15,38	92,31	84,62	71,79

Notes: Pointwise (Π^{PW}), multivariate (Π^{MV}), and time-simultenous (Π^{TS}) coverage level of 68% and 90% PW credible bands for 2, 4, and 8 quarters ahead out-of-sample headline inflation and real GDP forecast in United States using DP15.

Table D.6
DP15: various coverage levels of PW error bands for euro area

Coverage level	Variable	30% PW error band			60% PW error band			90% PW error band		
		h = 2	h = 4	h = 8	h = 2	h = 4	h = 8	h = 2	h = 4	h = 8
Π^{PW}	y_1	28,21	25,64	23,08	53,85	58,97	56,41	84,62	79,49	82,05
	y_2	33,33	48,72	43,59	66,67	89,74	87,18	100,00	94,87	89,74
Π^{MV}	-	15,38	17,95	7,69	41,03	53,85	51,28	84,62	76,92	74,36
Π^{TS}	y_1	10,26	0,00	0,00	35,90	17,95	2,56	74,36	56,41	38,46
	y_2	10,26	0,00	0,00	48,72	43,59	35,90	94,87	89,74	79,49

Notes: Pointwise (Π^{PW}), multivariate (Π^{MV}), and time-simultenous (Π^{TS}) coverage level of 68% and 90% PW credible bands for 2, 4, and 8 quarters ahead out-of-sample headline inflation and real GDP forecast in the euro area using DP15.

Table D.7

DP15: time-simultaneous coverage level (Π^{TS}) of TS error bands for United States

Method	Variable	30% TS error band						60% TS error band						90% TS error band					
		h = 2		h = 4		h = 8		h = 2		h = 4		h = 8		h = 2		h = 4		h = 8	
		Π	Δ	Π	Δ	Π	Δ	Π	Δ	Π	Δ	Π	Δ	Π	Δ	Π	Δ	Π	Δ
Bonferroni	y_1	56,41	155,54	74,36	300,36	64,10	478,17	74,36	61,76	84,62	125,42	69,23	202,71	97,44	26,52	89,74	57,01	87,18	95,93
	y_2	66,67	145,34	74,36	264,79	74,36	377,86	84,62	54,59	84,62	103,72	82,05	154,25	97,44	21,35	94,87	42,43	84,62	66,35
Šidák	y_1	66,67	145,34	74,36	264,79	74,36	377,86	84,62	54,59	84,62	103,72	82,05	154,25	97,44	21,35	94,87	42,43	84,62	66,35
	y_2	35,90	100,98	51,28	221,37	61,54	369,93	71,79	51,20	82,05	109,33	69,23	181,32	97,44	25,46	89,74	55,35	87,18	94,09
sup-t (PQO)	y_1	35,90	100,98	51,28	221,37	61,54	369,93	71,79	51,20	82,05	109,33	69,23	181,32	97,44	25,46	89,74	55,35	87,18	94,09
	y_2	48,72	95,86	58,97	198,18	61,54	305,32	82,05	45,94	79,49	91,67	79,49	139,51	97,44	20,60	94,87	41,29	84,62	64,51
min-max: absolute loss (median)	y_1	48,72	95,86	58,97	198,18	61,54	305,32	82,05	45,94	79,49	91,67	79,49	139,51	97,44	20,60	94,87	41,29	84,62	64,51
	y_2	33,33	83,62	43,59	166,87	38,46	255,61	66,67	41,30	79,49	83,77	64,10	131,12	97,44	21,19	89,74	44,31	82,05	72,13
min-max (LQO): absolute loss (median)	y_1	33,33	83,62	43,59	166,87	38,46	255,61	66,67	41,30	79,49	83,77	64,10	131,12	97,44	21,19	89,74	44,31	82,05	72,13
	y_2	46,15	88,31	56,41	178,73	56,41	269,55	79,49	42,12	79,49	83,34	79,49	123,84	94,87	19,05	94,87	37,89	84,62	57,56
min-max (BDR): absolute loss (mean)	y_1	46,15	88,31	56,41	178,73	56,41	269,55	79,49	42,12	79,49	83,34	79,49	123,84	94,87	19,05	94,87	37,89	84,62	57,56
	y_2	48,72	149,42	79,49	332,39	66,67	458,90	89,74	94,10	89,74	188,08	76,92	252,02	97,44	64,55	97,44	106,94	89,74	138,85
min-max: quadratic loss (mean)	y_1	48,72	149,42	79,49	332,39	66,67	458,90	89,74	94,10	89,74	188,08	76,92	252,02	97,44	64,55	97,44	106,94	89,74	138,85
	y_2	66,67	154,87	84,62	360,98	84,62	495,28	92,31	93,42	94,87	192,81	84,62	249,33	100,00	60,75	94,87	97,00	84,62	114,48
min-max (LQO): quadratic loss (mean)	y_1	66,67	154,87	84,62	360,98	84,62	495,28	92,31	93,42	94,87	192,81	84,62	249,33	100,00	60,75	94,87	97,00	84,62	114,48
	y_2	33,33	82,57	48,72	166,32	43,59	258,15	71,79	41,54	71,79	84,48	66,67	134,61	94,87	21,85	89,74	45,61	79,49	75,63
min-max (BDR): quadratic loss (median)	y_1	33,33	82,57	48,72	166,32	43,59	258,15	71,79	41,54	71,79	84,48	66,67	134,61	94,87	21,85	89,74	45,61	79,49	75,63
	y_2	48,72	87,62	51,28	179,32	56,41	275,40	76,92	41,95	79,49	83,72	76,92	128,25	97,44	19,57	94,87	39,27	84,62	61,64
min-max: angular loss (median)	y_1	48,72	87,62	51,28	179,32	56,41	275,40	76,92	41,95	79,49	83,72	76,92	128,25	97,44	19,57	94,87	39,27	84,62	61,64
	y_2	33,33	82,53	48,72	164,73	41,03	251,72	69,23	40,91	79,49	83,24	66,67	129,39	94,87	20,82	89,74	43,97	84,62	70,67
min-max (LQO): angular loss (median)	y_1	33,33	82,53	48,72	164,73	41,03	251,72	69,23	40,91	79,49	83,24	66,67	129,39	94,87	20,82	89,74	43,97	84,62	70,67
	y_2	48,72	87,54	53,85	177,19	56,41	266,36	79,49	41,65	79,49	82,39	76,92	122,65	97,44	18,92	94,87	37,40	84,62	56,82
min-max (BDR): angular loss (median)	y_1	48,72	87,54	53,85	177,19	56,41	266,36	79,49	41,65	79,49	82,39	76,92	122,65	97,44	18,92	94,87	37,40	84,62	56,82
	y_2	35,90	103,76	56,41	228,38	64,10	351,47	74,36	56,67	87,18	125,31	69,23	191,30	97,44	34,96	92,31	71,93	87,18	105,89
min-max: Chebyshev loss (median)	y_1	35,90	103,76	56,41	228,38	64,10	351,47	74,36	56,67	87,18	125,31	69,23	191,30	97,44	34,96	92,31	71,93	87,18	105,89
	y_2	53,85	108,99	69,23	246,32	71,79	377,35	84,62	57,45	89,74	124,68	84,62	187,87	100,00	31,45	94,87	64,29	84,62	92,01
min-max (LQO): Chebyshev loss (median)	y_1	53,85	108,99	69,23	246,32	71,79	377,35	84,62	57,45	89,74	124,68	84,62	187,87	100,00	31,45	94,87	64,29	84,62	92,01
	y_2	30,77	82,80	46,15	166,15	43,59	253,66	71,79	41,15	79,49	83,32	66,67	130,48	94,87	21,39	89,74	44,57	82,05	73,92
min-max (BDR): Chebyshev loss (median)	y_1	30,77	82,80	46,15	166,15	43,59	253,66	71,79	41,15	79,49	83,32	66,67	130,48	94,87	21,39	89,74	44,57	82,05	73,92
	y_2	51,28	87,66	53,85	177,88	58,97	268,78	79,49	42,07	79,49	83,07	74,36	124,10	97,44	19,27	94,87	37,94	84,62	58,05

Notes: Time-simultenous (Π^{TS}) coverage level of 68% and 90% TS credible bands for 2, 4, and 8 quarters ahead out-of-sample headline inflation and real GDP forecast in United States using DP15.

Table D.8

DP15: time-simultaneous coverage level (Π^{TS}) of TS error bands for euro area

Method	Variable	30% TS error band						60% TS error band						90% TS error band					
		h = 2		h = 4		h = 8		h = 2		h = 4		h = 8		h = 2		h = 4		h = 8	
		Π	Δ	Π	Δ	Π	Δ	Π	Δ	Π	Δ	Π	Δ	Π	Δ	Π	Δ	Π	Δ
Bonferroni	y_1	41,03	144,90	43,59	260,79	48,72	363,98	53,85	53,80	56,41	100,12	56,41	144,38	84,62	20,75	87,18	40,37	97,44	59,84
	y_2	51,28	147,72	82,05	267,46	79,49	375,11	87,18	55,60	89,74	104,40	82,05	153,08	97,44	22,16	92,31	43,29	84,62	67,25
Šidák	y_1	51,28	147,72	82,05	267,46	79,49	375,11	87,18	55,60	89,74	104,40	82,05	153,08	97,44	22,16	92,31	43,29	84,62	67,25
	y_2	33,33	96,43	30,77	197,48	33,33	297,31	51,28	45,54	51,28	88,78	53,85	131,44	84,62	20,04	87,18	39,09	94,87	58,92
sup-t (PQO)	y_1	33,33	96,43	30,77	197,48	33,33	297,31	51,28	45,54	51,28	88,78	53,85	131,44	84,62	20,04	87,18	39,09	94,87	58,92
	y_2	41,03	98,51	64,10	200,65	76,92	304,13	82,05	46,58	89,74	91,94	82,05	139,18	97,44	21,19	92,31	42,09	84,62	66,03
min-max: absolute loss (median)	y_1	41,03	98,51	64,10	200,65	76,92	304,13	82,05	46,58	89,74	91,94	82,05	139,18	97,44	21,19	92,31	42,09	84,62	66,03
	y_2	33,33	90,31	30,77	178,52	25,64	263,82	51,28	42,53	46,15	80,65	46,15	117,51	84,62	18,92	84,62	35,88	92,31	53,41
min-max (LQO): absolute loss (median)	y_1	33,33	90,31	30,77	178,52	25,64	263,82	51,28	42,53	46,15	80,65	46,15	117,51	84,62	18,92	84,62	35,88	92,31	53,41
	y_2	33,33	82,25	48,72	163,07	64,10	240,99	71,79	39,09	87,18	75,51	79,49	110,99	97,44	18,29	92,31	34,54	82,05	52,29
min-max (BDR): absolute loss (mean)	y_1	33,33	82,25	48,72	163,07	64,10	240,99	71,79	39,09	87,18	75,51	79,49	110,99	97,44	18,29	92,31	34,54	82,05	52,29
	y_2	43,59	160,65	58,97	366,07	66,67	482,56	71,79	93,98	84,62	185,42	92,31	234,51	100,00	59,18	100,00	91,82	100,00	106,94
min-max: quadratic loss (mean)	y_1	43,59	160,65	58,97	366,07	66,67	482,56	71,79	93,98	84,62	185,42	92,31	234,51	100,00	59,18	100,00	91,82	100,00	106,94
	y_2	51,28	146,70	89,74	321,75	82,05	419,08	94,87	91,62	92,31	172,42	84,62	209,86	100,00	60,17	94,87	87,13	84,62	98,04
min-max (LQO): quadratic loss (mean)	y_1	51,28	146,70	89,74	321,75	82,05	419,08	94,87	91,62	92,31	172,42	84,62	209,86	100,00	60,17	94,87	87,13	84,62	98,04
	y_2	33,33	90,88	28,21	178,60	28,21	270,80	51,28	42,44	43,59	81,73	43,59	122,06	84,62	19,41	84,62	37,67	87,18	56,77
min-max (BDR): quadratic loss (median)	y_1	33,33	90,88	28,21	178,60	28,21	270,80	51,28	42,44	43,59	81,73	43,59	122,06	84,62	19,41	84,62	37,67	87,18	56,77
	y_2	38,46	81,58	56,41	162,52	74,36	240,22	71,79	39,09	92,31	78,92	82,05	114,18	97,44	18,52	94,87	38,42	84,62	56,04
min-max: angular loss (median)	y_1	38,46	81,58	56,41	162,52	74,36	240,22	71,79	39,09	92,31	78,92	82,05	114,18	97,44	18,52	94,87	38,42	84,62	56,04
	y_2	30,77	89,79	28,21	176,80	25,64	261,49	51,28	42,19	48,72	79,70	46,15	116,91	84,62	18,59	87,18	35,63	89,74	52,66
min-max (LQO): angular loss (median)	y_1	30,77	89,79	28,21	176,80	25,64	261,49	51,28	42,19	48,72	79,70	46,15	116,91	84,62	18,59	87,18	35,63	89,74	52,66
	y_2	38,46	81,57	51,28	160,23	74,36	233,94	71,79	38,41	87,18	74,40	82,05	109,01	97,44	17,84	94,87	35,00	84,62	50,07
min-max (BDR): angular loss (median)	y_1	38,46	81,57	51,28	160,23	74,36	233,94	71,79	38,41	87,18	74,40	82,05	109,01	97,44	17,84	94,87	35,00	84,62	50,07
	y_2	33,33	111,99	38,46	245,10	46,15	371,88	53,85	57,96	64,10	121,47	79,49	180,02	89,74	30,89	97,44	62,10	100,00	86,12
min-max: Chebyshev loss (median)	y_1	33,33	111,99	38,46	245,10	46,15	371,88	53,85	57,96	64,10	121,47	79,49	180,02	89,74	30,89	97,44	62,10	100,00	86,12
	y_2	43,59	102,16	82,05	222,97	79,49	334,61	79,49	54,97	92,31	115,77	84,62	168,24	97,44	33,34	94,87	62,22	84,62	81,39
min-max (LQO): Chebyshev loss (median)	y_1	43,59	102,16	82,05	222,97	79,49	334,61	79,49	54,97	92,31	115,77	84,62	168,24	97,44	33,34	94,87	62,22	84,62	81,39
	y_2	30,77	89,43	28,21	178,71	25,64	263,86	51,28	41,91	46,15	80,09	46,15	118,46	82,05	18,77	84,62	36,99	84,62	54,57
min-max (BDR): Chebyshev loss (median)	y_1	30,77	89,43	28,21	178,71	25,64	263,86	51,28	41,91	46,15	80,09	46,15	118,46	82,05	18,77	84,62	36,99	84,62	54,57
	y_2	38,46	81,79	58,97	161,10	74,36	238,14	74,36	39,02	89,74	76,42	82,05	112,03	97,44	18,87	94,87	37,72	84,62	54,66

Notes: Time-simultenous (Π^{TS}) coverage level of 68% and 90% TS credible bands for 2, 4, and 8 quarters ahead out-of-sample headline inflation and real GDP forecast in the euro area using DP15.

Table D.9
DP15: multivariate coverage level (Π^{MV}) of MV error bands for United States

Method	30% TS error band						60% TS error band						90% TS error band					
	h=1:2		h=1:4		h=1:8		h=1:2		h=1:4		h=1:8		h=1:2		h=1:4		h=1:8	
	Π	Δ	Π	Δ	Π	Δ	Π	Δ	Π	Δ	Π	Δ	Π	Δ	Π	Δ	Π	Δ
Bonferroni	61,54	151,23	64,10	152,88	64,10	155,49	82,05	58,31	82,05	59,44	74,36	61,30	97,44	24,01	92,31	24,84	84,62	26,00
Šidák	43,59	99,01	46,15	99,98	43,59	101,60	79,49	48,79	74,36	49,69	74,36	51,16	97,44	23,12	92,31	23,99	84,62	25,13
sup-t (PQO)	43,59	93,83	46,15	93,49	43,59	93,75	79,49	45,72	74,36	46,27	71,79	46,93	97,44	21,93	92,31	22,45	84,62	23,52
min-max: absolute loss (median)	61,54	157,21	64,10	157,21	53,85	157,74	89,74	96,70	87,18	97,39	79,49	98,17	100,00	61,93	94,87	61,77	89,74	62,74
min-max (LQO): absolute loss (median)	51,28	89,52	51,28	90,40	46,15	91,39	76,92	48,67	74,36	48,74	76,92	49,83	94,87	28,12	92,31	27,89	87,18	28,27
min-max (BDR): absolute loss (mean)	48,72	91,26	51,28	91,21	46,15	92,14	79,49	46,95	76,92	47,24	74,36	47,94	94,87	23,33	92,31	23,61	84,62	24,20
min-max: quadratic loss (mean)	58,97	110,48	53,85	110,42	48,72	110,99	84,62	62,42	79,49	62,63	76,92	63,73	97,44	37,42	94,87	37,31	89,74	37,74
min-max (LQO): quadratic loss (mean)	51,28	89,85	51,28	90,29	46,15	90,76	76,92	48,53	74,36	48,61	76,92	49,52	97,44	29,38	92,31	28,62	89,74	28,61
min-max (BDR): quadratic loss (median)	51,28	90,23	51,28	90,18	46,15	91,24	76,92	47,79	74,36	48,11	76,92	48,74	94,87	26,35	92,31	26,10	87,18	26,28
min-max: angular loss (median)	100,00	1142,00	94,87	1241,78	92,31	1364,39	100,00	521,43	94,87	558,19	94,87	621,17	100,00	218,71	97,44	223,34	94,87	249,88
min-max (LQO): angular loss (median)	94,87	599,16	92,31	597,71	84,62	611,16	94,87	210,42	92,31	207,47	84,62	213,48	94,87	50,81	92,31	48,61	84,62	49,52
min-max (BDR): angular loss (median)	43,59	93,16	43,59	92,60	43,59	93,13	82,05	45,58	74,36	46,12	71,79	46,61	94,87	21,75	92,31	22,11	84,62	23,31
min-max: Chebyshev loss (median)	46,15	93,64	46,15	93,25	46,15	93,74	79,49	45,89	76,92	46,39	74,36	47,36	97,44	21,49	92,31	22,04	84,62	23,12
min-max (LQO): Chebyshev loss (median)	46,15	93,64	46,15	93,25	46,15	93,74	79,49	45,89	76,92	46,39	74,36	47,36	97,44	21,49	92,31	22,04	84,62	23,12
min-max (BDR): Chebyshev loss (median)	46,15	93,64	46,15	93,25	46,15	93,74	79,49	45,89	76,92	46,39	74,36	47,36	97,44	21,49	92,31	22,04	84,62	23,12

Notes: Multivariate (Π^{MV}) coverage level of 68% and 90% MV credible bands for 2, 4, and 8 quarters ahead out-of-sample headline inflation and real GDP forecast in United States using DP15.

Table D.10
DP15: multivariate coverage level (Π^{MV}) of MV error bands for euro area

Method	30% TS error band						60% TS error band						90% TS error band					
	h=1:2		h=1:4		h=1:8		h=1:2		h=1:4		h=1:8		h=1:2		h=1:4		h=1:8	
	Π	Δ	Π	Δ	Π	Δ	Π	Δ	Π	Δ	Π	Δ	Π	Δ	Π	Δ	Π	Δ
Bonferroni	48,72	147,32	58,97	146,10	53,85	147,08	66,67	54,91	71,79	54,88	64,10	55,44	89,74	21,30	87,18	21,57	82,05	22,20
Šidák	33,33	97,34	51,28	97,29	41,03	97,48	61,54	46,34	69,23	46,22	56,41	46,54	89,74	20,58	87,18	20,85	82,05	21,34
sup-t (PQO)	33,33	92,49	46,15	92,63	38,46	91,40	58,97	44,06	69,23	43,88	56,41	43,63	89,74	19,79	84,62	19,85	82,05	20,01
min-max: absolute loss (median)	53,85	151,73	66,67	150,88	56,41	149,08	84,62	93,02	92,31	92,15	89,74	92,32	100,00	58,15	94,87	57,21	89,74	56,29
min-max (LQO): absolute loss (median)	33,33	86,97	38,46	85,45	28,21	83,95	64,10	45,08	79,49	45,92	82,05	46,55	100,00	28,61	94,87	30,64	89,74	32,44
min-max (BDR): absolute loss (mean)	35,90	87,33	38,46	87,19	33,33	85,85	64,10	43,89	69,23	44,09	71,79	44,15	97,44	20,81	92,31	21,82	89,74	22,10
min-max: quadratic loss (mean)	43,59	106,96	53,85	105,63	38,46	103,69	69,23	58,92	84,62	59,41	82,05	59,70	100,00	37,25	94,87	38,28	89,74	39,85
min-max (LQO): quadratic loss (mean)	30,77	87,03	38,46	85,61	28,21	83,57	64,10	45,16	79,49	45,95	82,05	46,79	100,00	30,28	94,87	32,49	89,74	34,65
min-max (BDR): quadratic loss (median)	33,33	86,61	38,46	86,14	28,21	84,15	64,10	44,56	74,36	44,94	76,92	45,54	100,00	24,98	94,87	25,85	89,74	27,42
min-max: angular loss (median)	100,00	917,05	97,44	913,63	92,31	999,70	100,00	434,76	97,44	481,09	92,31	547,61	100,00	177,54	97,44	201,37	92,31	248,71
min-max (LQO): angular loss (median)	94,87	537,70	82,05	521,81	87,18	535,36	94,87	188,71	82,05	182,92	87,18	187,21	94,87	44,77	82,05	43,73	87,18	44,66
min-max (BDR): angular loss (median)	28,21	91,71	43,59	91,43	41,03	90,53	61,54	43,76	69,23	43,25	56,41	43,22	89,74	19,54	84,62	19,62	82,05	19,87
min-max: Chebyshev loss (median)	35,90	91,82	51,28	91,73	38,46	91,00	64,10	43,88	69,23	43,61	56,41	43,18	89,74	19,98	84,62	19,75	82,05	19,92
min-max (LQO): Chebyshev loss (median)	35,90	91,82	51,28	91,73	38,46	91,00	64,10	43,88	69,23	43,61	56,41	43,18	89,74	19,98	84,62	19,75	82,05	19,92
min-max (BDR): Chebyshev loss (median)	35,90	91,82	51,28	91,73	38,46	91,00	64,10	43,88	69,23	43,61	56,41	43,18	89,74	19,98	84,62	19,75	82,05	19,92

Notes: Multivariate (Π^{MV}) coverage level of 68% and 90% MV credible bands for 2, 4, and 8 quarters ahead out-of-sample headline inflation and real GDP forecast in the euro area using DP15.