

# Scenario Analysis with Multivariate Bayesian Machine Learning Models

Michael PFARRHOFER  
WU Vienna

Anna STELZER  
Oesterreichische Nationalbank

We present an econometric framework that adapts tools for scenario analysis, such as variants of conditional forecasts and generalized impulse responses, for use with dynamic nonparametric models. The proposed algorithms are based on predictive simulation and sequential Monte Carlo methods. Their utility is demonstrated with three applications: (1) conditional forecasts based on stress test scenarios, measuring (2) macroeconomic risk under varying financial stress, and estimating the (3) asymmetric effects of financial shocks in the US and their international spillovers. Our empirical results indicate the importance of nonlinearities and asymmetries in relationships between macroeconomic and financial variables.

**JEL:** C32, C53, E44

**Keywords:** conditional forecast, generalized impulse response function, Bayesian additive regression trees, nonlinearities, structural inference

*Contact:* [michael.pfarrhofer@wu.ac.at](mailto:michael.pfarrhofer@wu.ac.at), Department of Economics, WU Vienna University of Economics and Business. We would like to thank Christiane Baumeister, Niko Hauzenberger, Philip Lane, Massimiliano Marcellino, and seminar participants at the OeNB, Swiss National Bank (SNB), and IMIM online seminar for useful comments. *Disclaimer:* The views expressed in this paper do not necessarily reflect those of the Oesterreichische Nationalbank or the Eurosystem.

# 1. INTRODUCTION

In this paper, we discuss how to conduct scenario analysis with dynamic multivariate models in macroeconomics when the functional form of the conditional mean is nonlinear and/or unknown. While related tools exist in linear and traditional nonlinear frameworks (e.g., regime switching or time-varying parameter models, see Fischer *et al.* (2023) for a recent example), they rely on potentially restrictive parametric assumptions and are not directly applicable in nonparametric models. The latter feature in many recent macroeconometric applications that employ Bayesian machine learning (ML) methods (see, e.g., Huber and Rossini, 2022; Clark *et al.*, 2023; Huber *et al.*, 2023; Hauzenberger *et al.*, 2025b; Chernis *et al.*, 2025; Lima *et al.*, 2025).

We use the term scenario analysis broadly to refer to different counterfactual experiments. These include versions of conditional forecasts (CFs) and structural scenario analysis (as in Antolin-Diaz *et al.*, 2021), and several variants of nonlinear impulse response functions (IRFs). While some aspects and issues in the Bayesian ML context have been discussed in isolation in the aforementioned papers, a unified or comprehensive treatment is not yet available. We bridge this gap by developing a framework and estimation algorithms for scenario analysis that are widely applicable when the conditional mean is flexibly learned from the data.

Conditional forecasts simulate future paths of variables under scenarios encoded as constraints on other observables, structural shocks, or both. Related techniques have been used by academics and practitioners since the 1980s in linear models (see, e.g., Doan *et al.*, 1984), with subsequent refinements (see Waggoner and Zha, 1999; Andersson *et al.*, 2010; Jarociński, 2010; Bańbura *et al.*, 2015; Antolin-Diaz *et al.*, 2021; Chernis *et al.*, 2024; Chan *et al.*, 2025; Crump *et al.*, 2025).<sup>1</sup> Breaking the assumption of linearity, however, complicates matters: there is no general representation of nonlinear time series as functions

---

<sup>1</sup> For a Bayesian decision analysis perspective and discussion of additional related literature, see West (2024).

of shocks (see, e.g., [Potter, 2000](#)), and unknown nonlinearities make it difficult to derive multi-step ahead predictive distributions. Obtaining closed form solutions is often impossible, but in this case one can resort to Monte Carlo methods. We show how common approaches can be adapted to nonparametric settings. Akin to [Bańbura \*et al.\* \(2015\)](#), our approach casts conditional forecasting as a nonlinear state space problem. We use Particle Gibbs with Ancestor Sampling (PGAS, [Lindsten \*et al.\*, 2014](#)), which combines generality with computational efficiency, making it suitable for a wide range of nonlinear multivariate models.

An IRF can be defined as the difference between conditional forecasts ([Gallant \*et al.\*, 1993](#); [Koop \*et al.\*, 1996](#); [Rambachan and Shephard, 2021](#); [Jordà and Taylor, 2025](#)). Our framework thus lends itself to estimating nonlinear dynamic effects (see [Gonçalves \*et al.\*, 2021, 2024](#); [Kolesár and Plagborg-Møller, 2025](#), for related discussions) in the form of generalized IRFs (GIRFs), and thereby contributes to improving interpretability of ML methods. Compared with IRFs in linear settings (which are symmetric, shape invariant and history independent) the GIRFs do not feature such potentially restrictive properties. Specifically, we consider (1) unorthogonalized GIRFs (the difference between conditional and unconditional forecasts subject to certain restrictions on observables), we explore (2) structural GIRFs to shocks identified with approaches typically used in the structural VAR (SVAR) literature, and (3) restricted GIRFs, which can be used to quantify the contributions of specific sets of transmission channels in shock propagation.<sup>2</sup>

The proposed approach is generally applicable in multivariate dynamic models. This means that there are many potential candidates for estimating conditional mean functions. These choices include (but are not limited to) regression trees or Gaussian process priors (see

---

<sup>2</sup> A related approach derived from [Antolin-Diaz \*et al.\* \(2021\)](#) is discussed in [Breitenlechner \*et al.\* \(2022\)](#). Counterfactual experiments in this spirit relate to alternative policy rules and are subject to the [Lucas \(1976\)](#) critique (for a recent discussion, see [McKay and Wolf, 2023](#)). While our approach is also susceptible to these concerns, nonparametric modeling offers flexibility and a safeguard; see also [Section 2.3](#).

Marcellino and Pfarrhofer, 2024, for a recent review). Such nonparametric methods are more robust to misspecification that may arise from assuming tightly parameterized likelihoods, and can be used to offer the possibility of capturing modeling genuinely nonlinear patterns in the data. For our empirical work, we use Bayesian Additive Regression Trees (BART, Chipman *et al.*, 2010) as a specific nonparametric implementation. We pick this sum-of-trees model because tree-based approaches have proven particularly capable of producing accurate forecasts when used with time series for the US economy, with datasets structured similarly to the one we use in this paper (see, e.g., Medeiros *et al.*, 2021; Goulet Coulombe *et al.*, 2022; Clark *et al.*, 2023; Goulet Coulombe, 2024). Our approach to predictive inference is developed under the assumption of an additive multivariate Gaussian error term with a time-varying covariance matrix.

We illustrate our framework through several empirical applications. Our dataset comprises about 25 quarterly macroeconomic and financial variables for the US economy ranging from the mid-1970s to the last quarter of 2024. In one of our explorations, we also add some international variables for the euro area (EA) and the United Kingdom (UK). The applied work assesses and illustrates the role of nonlinearities when interest centers on CFs, and we explore asymmetries in the propagation of shocks of different types, signs and magnitudes. Specifically, we provide three empirical applications. First, inspired by Chan *et al.* (2025), we use a subset of the macroeconomic assumptions underlying the annual stress test conducted by the Federal Reserve System and compute CFs using constraints on observables for different scenarios, comparing predictive densities from linear and nonlinear models. Second, reflecting the growth-at-risk literature (Adrian *et al.*, 2019), we study the unorthogonalized counterfactual implications of varying financial conditions on tail risks of output growth, inflation, and employment. Third, we identify a US-based financial shock recursively (as in, e.g., Barnichon *et al.*, 2022), and compute GIRFs to shocks of different signs and magnitudes that are allowed to propagate internationally. We then use a restricted

GIRF approach to gauge the role of spillovers and spillbacks, inspired by the empirical work of Breitenlechner *et al.* (2022).

The rest of this paper is structured as follows. Section 2 addresses the challenges and solutions to obtain predictive inference in the presence of nonlinearities of unknown form. We discuss how to impose constraints on forecasts, and how these constraints may be used to construct scenarios through GIRFs. Section 3 presents an econometric implementation using BART. Section 4 provides three empirical illustrations. The last section concludes.

## 2. NONLINEARITY AND PREDICTIVE DISTRIBUTIONS

Let  $\mathbf{y}_t = (y_{1t}, \dots, y_{nt})'$  collect  $n$  variables for  $t = 1, \dots, T$ , and  $\mathbf{x}_t = (\mathbf{y}'_{t-1}, \dots, \mathbf{y}'_{t-p})'$  is a  $k = np$  vector of lags. Interest centers on dynamic multivariate models of the form:

$$\mathbf{y}_t = \mathbf{F}(\mathbf{x}_t) + \boldsymbol{\epsilon}_t, \quad \boldsymbol{\epsilon}_t \sim \mathcal{N}(\mathbf{0}_n, \boldsymbol{\Sigma}_t), \quad (1)$$

where  $\mathbf{F}(\mathbf{x}_t) = (f_1(\mathbf{x}_t), \dots, f_n(\mathbf{x}_t))'$  is an  $n$ -vector of conditional mean functions  $f_i(\mathbf{x}_t) : \mathbb{R}^k \rightarrow \mathbb{R}$  for  $i = 1, \dots, n$ , such that  $\mathbf{F}(\mathbf{x}_t) : \mathbb{R}^k \rightarrow \mathbb{R}^n$ . We assume iid reduced form Gaussian errors  $\boldsymbol{\epsilon}_t$ , with  $n \times n$  time-varying covariance matrix  $\boldsymbol{\Sigma}_t$ . One may assume functional forms for the  $f_i(\mathbf{x}_t)$ 's or treat them as unknown and estimate them nonparametrically. The methods we propose are designed specifically for the latter case.

All scenario analyses we discuss in this paper rely on computing functions of the (conditional) moments of  $\mathbf{y}_{t+h}$  based on (1). Consider the single-lag case  $p = 1$ , without loss of generality. We may write:

$$\mathbf{y}_{t+h} = \tilde{\mathbf{F}}_{(h)}(\mathbf{y}_t, \boldsymbol{\epsilon}_{t+1}, \boldsymbol{\epsilon}_{t+2}, \dots, \boldsymbol{\epsilon}_{t+h}) = \mathbf{F}(\tilde{\mathbf{F}}_{(h-1)}(\mathbf{y}_t, \boldsymbol{\epsilon}_{t+1}, \dots, \boldsymbol{\epsilon}_{t+(h-1)})) + \boldsymbol{\epsilon}_{t+h}, \quad (2)$$

with  $\mathbf{y}_{t+1} = \tilde{\mathbf{F}}_{(1)}(\mathbf{y}_t, \boldsymbol{\epsilon}_{t+1}) = \mathbf{F}(\mathbf{y}_t) + \boldsymbol{\epsilon}_{t+1}$ , and  $\tilde{\mathbf{F}}_{(h)}(\bullet)$  denotes the  $h$ -step composition of  $\mathbf{F}(\bullet)$  which is defined recursively. This expression is related to the Wold decomposition,

and expresses  $\mathbf{y}_{t+h}$  as a function of the initial condition  $\mathbf{y}_t$  and a sequence of white noise shocks,  $\{\boldsymbol{\epsilon}_{t+1}, \dots, \boldsymbol{\epsilon}_{t+h}\}$ ; see also [Gourieroux and Lee \(2023\)](#) and [Appendix A](#).<sup>3</sup>

Our framework assumes that the functional form of  $\mathbf{F}(\bullet)$  is unknown and estimated, so it is not necessarily additively separable. Even if it were, applying a nonlinear function to a Gaussian shock does not necessarily yield a random variable that follows a well-known distribution. Thus, it is not generally possible to derive closed form higher-order moments. In order to explore the predictive distributions nevertheless, one can resort to simulation-based methods.

## 2.1. Predictive Simulation

Define a vector  $\boldsymbol{\Xi}$  that contains all coefficients and latent variables necessary to parameterize (1). At time  $\tau$ , the one-step ahead predictive distribution is:

$$p_\tau(\mathbf{y}_{\tau+1} | \mathcal{I}) = \int p_\tau(\mathbf{y}_{\tau+1} | \mathcal{I}, \boldsymbol{\Xi}) p(\boldsymbol{\Xi} | \mathcal{I}) d\boldsymbol{\Xi}, \quad (3)$$

where  $\mathcal{I}$  denotes the information set used to infer  $\boldsymbol{\Xi}$ , and the subscript on  $p_\tau(\bullet)$  marks the forecast origin. For typical out-of-sample exercises,  $\mathcal{I} = \{\mathbf{y}_t\}_{t=1}^\tau$ , and we are interested in  $p_\tau(\mathbf{y}_{\tau+h} | \{\mathbf{y}_t\}_{t=1}^\tau)$  for  $h = 1, 2, \dots$ , steps ahead (see [Geweke and Amisano, 2010](#)). In other cases we condition on  $\mathcal{I} = \{\mathbf{y}_t\}_{t=1}^T$ , to compute scenarios in-sample for  $\tau \in \{1, 2, \dots, T\}$  using  $p_\tau(\mathbf{y}_{\tau+h} | \{\mathbf{y}_t\}_{t=1}^T)$  with parameters updated based on the full sample, but for initial conditions at time  $\tau$ . In either case (unless necessary for clarity, we omit the  $\tau$  index for the forecast origin below),  $p_\tau(\mathbf{y}_{\tau+1} | \mathcal{I})$  generally does not take a well-known form, and neither does the distribution of higher-order forecasts for  $h \geq 2$ .

---

<sup>3</sup> Under some assumptions about the behavior of  $\tilde{\mathbf{F}}_{(h)}(\bullet)$ , it is noteworthy that expansions such as the Volterra series could be used as an approximation and a close equivalent to the Wold representation for nonlinear time series, see, e.g., [Potter \(2000, Section 4\)](#) and [Jordà \(2005, Section II\)](#).

However, we may obtain random samples from them via simulation. This involves exploiting the fact that even though  $p(\mathbf{y}_{\tau+1} | \mathcal{I})$  is unknown,  $p(\mathbf{y}_{\tau+1} | \mathcal{I}, \Xi)$  takes a known form under the model in (1). The one-step ahead predictive distribution is:

$$p(\mathbf{y}_{\tau+1} | \mathcal{I}, \Xi^{(m)}) = \mathcal{N}(\mathbf{F}^{(m)}(\mathbf{x}_{\tau+1}), \Sigma_{\tau+1}^{(m)}), \quad (4)$$

where  $\mathbf{x}_{\tau+1} = (\mathbf{y}'_{\tau}, \dots, \mathbf{y}'_{\tau-p+1})'$ . Here,  $x^{(m)}$  refers to the  $m$ th draw of a random variable  $x$  — in most cases throughout this paper,  $m$  indexes a draw from some posterior or predictive distribution when running a Markov chain Monte Carlo (MCMC) algorithm. In what follows we omit the  $m$ -indexing for the parameters, but stress that these computations are carried out in each sweep of the algorithm and thus account for the posterior uncertainty of parameters. For  $h \geq 2$  we iterate forward, conditioning recursively on the draws for preceding horizons, by setting the predictors to  $\mathbf{x}_{\tau+h}^{(m)} = (\mathbf{y}_{\tau+h-1}^{(m)'}, \mathbf{y}_{\tau+h-2}^{(m)'}, \dots)'$ , and obtain:

$$p(\mathbf{y}_{\tau+h} | \mathcal{I}, \mathbf{y}_{\tau+1:\tau+h-1}^{(m)}, \Xi) = \mathcal{N}(\mathbf{F}(\mathbf{x}_{\tau+h}^{(m)}), \Sigma_{\tau+h}), \quad (5)$$

where  $\mathbf{y}_{\tau+1:\tau+h-1}^{(m)}$  denotes the path of the variables from  $\tau + 1$  to  $\tau + h - 1$  and  $\mathbf{y}_{\tau+1:\tau+h} = (\mathbf{y}'_{\tau+1}, \dots, \mathbf{y}'_{\tau+h})'$ . This relates to the recursive composition discussed in the context of (2), and exploits the fact that:

$$p(\mathbf{y}_{\tau+1:\tau+h} | \mathcal{I}) = \int p(\mathbf{y}_{\tau+1} | \mathcal{I}, \Xi) \prod_{j=2}^h p(\mathbf{y}_{\tau+j} | \mathbf{y}_{\tau+1:\tau+j-1}, \mathcal{I}, \Xi) p(\Xi | \mathcal{I}) d\Xi, \quad (6)$$

that is, the joint predictive distribution factors into the product of the conditional one-step ahead densities. Simulating the process forward, sampling from the distribution in (5) across horizons  $h = 1, 2, \dots$ , in each sweep of an MCMC algorithm, delivers draws from  $p(\mathbf{y}_{\tau+1:\tau+h} | \mathcal{I})$  via a Monte Carlo approach.

## 2.2. Conditional Forecasts

As their name suggests, CFs imply an additional conditioning argument in (6). In this context, we denote by  $\mathcal{C}_h$  a set that defines desired restrictions at horizon  $h = 1, 2, \dots, H$ , where  $H$  is the maximum forecast horizon of interest,  $\mathcal{C}_{1:H} = \{\mathcal{C}_1, \dots, \mathcal{C}_H\}$  collects all restrictions over the full forecast path, and the unconditional forecast results when  $\mathcal{C}_h = \emptyset$  for all  $h$ . The joint distribution of interest is  $p(\mathbf{y}_{\tau+1:\tau+H} | \mathcal{I}, \mathcal{C}_{1:H})$ , and can be obtained by marginalizing  $p(\mathbf{y}_{\tau+1:\tau+H} | \mathcal{I}, \mathcal{C}_{1:H}, \Xi)$  over the parameters via MCMC sampling, as discussed previously. When there are no restrictions (i.e., for unconditional forecasts), the model can be simulated forward according to the law of motion in (1), so that any jointly realized random path over  $(\tau + 1) : (\tau + H)$  is consistent with the model’s dynamics. In the presence of restrictions, however, the key challenge in the nonlinear context is that we cannot work directly with the joint distribution to impose them across all horizons simultaneously, as is possible in linear settings; and a decomposition like in (6) is generally not applicable.

We propose to use a particle MCMC algorithm (see [Andrieu \*et al.\*, 2010](#)), adapted from the literature on nonlinear state space models, as a solution. Specifically, our approach is based on PGAS as proposed by [Lindsten \*et al.\* \(2014\)](#), akin to the filtering-based conditional forecast implementations of [Bańbura \*et al.\* \(2015\)](#), and capable of enforcing the restrictions discussed in more detail below on an  $h$ -by- $h$  basis. In the spirit of [Andersson \*et al.\* \(2010\)](#), we define  $\mathcal{C}_h$  as a stochastic restriction:<sup>4</sup>

$$\mathbf{R}_h \mathbf{y}_{\tau+h} \sim \mathcal{N}(\mathbf{r}_h, \mathbf{\Omega}_h), \tag{7}$$

---

<sup>4</sup> The restrictions can also be written as a regression,  $\mathbf{r}_h = \mathbf{R}_h \mathbf{y}_{\tau+h} + \boldsymbol{\eta}_h$  with  $\boldsymbol{\eta}_h \sim \mathcal{N}(\mathbf{0}, \mathbf{\Omega}_h)$ , which is equivalent in distribution. The regression illustrates the interpretation as a type of measurement equation. Exact (hard) restrictions can be approximated by setting the respective elements of  $\mathbf{\Omega}_h$  to, e.g.,  $10^{-8}$ , which has computational and numerical advantages over using exact zeroes.

where  $\mathbf{R}_h$  is an  $r_h \times n$  horizon-specific matrix selecting or encoding  $r_h$  restrictions on linear combinations of endogenous variables, and  $\mathbf{r}_h$  and  $\mathbf{\Omega}_h$  are the  $r_h \times 1$  mean and  $r_h \times r_h$  covariance matrix of the restrictions.

**Structural scenarios.** In many cases it is desirable to impose a mix of restrictions on future observables as well as structural shocks, for structural scenario analysis in the spirit of Antolin-Diaz *et al.* (2021). This allows for imposing that specific structural shocks (rather than an unrestricted combination of them) are responsible for the respective scenario in terms of the observables. Econometrically, this can be achieved by forcing a subset of the structural shocks to their unconditional distribution (non-driving shocks), and leave the remaining ones (driving shocks) free. The latter thus may deviate from their unconditional distribution, thereby acting as the offsetting forces that yield the restrictions imposed on the observables. The same intuition can be used for computing dynamic responses to specific structural shocks, see also Breitenlechner *et al.* (2022).

As discussed above, the one-step ahead prediction of our model conditional on any preceding information up to  $\tau + h - 1$  can be written as  $\mathbf{y}_{\tau+h} = \boldsymbol{\mu}_{\tau+h}^{(m)} + \boldsymbol{\epsilon}_{\tau+h}$ , with  $\boldsymbol{\mu}_{\tau+h}^{(m)} = \mathbf{F}(\mathbf{y}_{\tau+h-1}^{(m)}, \dots, \mathbf{y}_{\tau+h-p}^{(m)}) = \mathbf{F}(\mathbf{x}_{\tau+h}^{(m)})$  and  $\boldsymbol{\epsilon}_{\tau+h} \sim \mathcal{N}(\mathbf{0}, \boldsymbol{\Sigma}_{\tau+h})$ . We abstract below from time-variation in the variances to simplify notation, but note that this framework is applicable also in the presence of heteroskedastic shocks. A structural form of the model can be obtained as  $\mathbf{H}\mathbf{y}_{\tau+h} = \mathbf{H}\boldsymbol{\mu}_{\tau+h} + \mathbf{u}_{\tau+h}$ , with iid shocks  $\mathbf{u}_\tau \sim \mathcal{N}(\mathbf{0}_n, \mathbf{I}_n)$ , where  $\mathbf{I}_n$  is an  $n$ -dimensional identity matrix. Then,  $\boldsymbol{\epsilon}_\tau = \mathbf{H}^{-1}\mathbf{u}_\tau$ , and  $\boldsymbol{\Sigma} = \mathbf{H}^{-1}\mathbf{H}^{-1'}$ .<sup>5</sup> Both types of restrictions can be implemented with distributional assumptions about the restricted vector of endogenous variables, as  $r_h = r_h^{(y)} + r_h^{(u)}$  restrictions on:

- (i) observables:  $\mathbf{R}_h^{(y)}\mathbf{y}_{\tau+h} \sim \mathcal{N}(\mathbf{r}_h^{(y)}, \boldsymbol{\Omega}_h^{(y)})$ , imposing  $r_h^{(y)}$  restrictions;

<sup>5</sup> An alternative, see also Section 3, is to parameterize (1) directly as  $\mathbf{H}\mathbf{y}_t = \mathbf{F}(\mathbf{x}_t) + \mathbf{u}_t$ ,  $\mathbf{u}_t \sim \mathcal{N}(\mathbf{0}_n, \mathbf{I}_n)$ , where  $\mathbf{H}$  is nonsingular, see also Arias *et al.* (2023); Chan *et al.* (2024). Time-varying variances of structural shocks, or more general forms of heteroskedasticity, can be addressed by using  $\mathbf{u}_\tau \sim \mathcal{N}(\mathbf{0}_n, \mathbf{S}_\tau)$ , where  $\mathbf{S}_\tau = \text{diag}(s_{1\tau}^2, \dots, s_{n\tau}^2)$ , such that  $\boldsymbol{\Sigma}_\tau = \mathbf{H}^{-1}\mathbf{S}_\tau\mathbf{H}^{-1'}$ .

(ii) structural shocks:  $\mathbf{R}^{(u)}\mathbf{u}_{\tau+h} \sim \mathcal{N}(\mathbf{r}_h^{(u)}, \mathbf{\Omega}_h^{(u)})$ , or equivalently,  $\mathbf{R}^{(u)}\mathbf{H}\mathbf{y}_{\tau+h} \sim \mathcal{N}(\mathbf{r}_h^{(u)} + \mathbf{R}^{(u)}\mathbf{H}\boldsymbol{\mu}_{\tau+h}, \mathbf{\Omega}_h^{(u)})$ , imposing  $r_h^{(u)}$  restrictions.

We may write these in terms of (7) by stacking both types in terms of their implications about the observables. The operator  $\text{bdiag}(\bullet)$  outputs a block diagonal matrix, with the inputs arranged along the main block diagonal; we obtain  $\mathbf{r}_h = [\mathbf{r}_h^{(y)}; \mathbf{r}_h^{(u)} + \mathbf{R}^{(u)}\mathbf{H}\boldsymbol{\mu}_{\tau+h}]$ ,  $\mathbf{R}_h = [\mathbf{R}_h^{(y)}; \mathbf{R}_h^{(u)}\mathbf{H}]$  and  $\mathbf{\Omega}_h = \text{bdiag}(\mathbf{\Omega}_h^{(y)}, \mathbf{\Omega}_h^{(u)})$ .

**Particle Gibbs with Ancestor Sampling (PGAS).** Under these assumptions we obtain the joint distribution of the forecast and the restrictions at horizon  $h$ :

$$\begin{bmatrix} \mathbf{y}_{\tau+h} \\ \mathbf{r}_h \end{bmatrix} \sim \mathcal{N} \left( \begin{bmatrix} \mathbf{I}_n \\ \mathbf{R}_h \end{bmatrix} \boldsymbol{\mu}_{\tau+h}^{(m)}, \begin{bmatrix} \boldsymbol{\Sigma} & \boldsymbol{\Sigma}\mathbf{R}'_h \\ \mathbf{R}_h\boldsymbol{\Sigma} & \mathbf{R}_h\boldsymbol{\Sigma}\mathbf{R}'_h + \mathbf{\Omega}_h \end{bmatrix} \right). \quad (8)$$

which encodes a nonlinear state combined with a linear measurement equation. This casts conditional forecasting and structural scenario analysis as a state space problem, similar to Bańbura *et al.* (2015), and allows to apply the PGAS algorithm of Lindsten *et al.* (2014) with a few adjustments.<sup>6</sup> The idea in such algorithms is to use hypothetical realizations, a set of particles  $v = 1, \dots, V$ , combined with a scheme to score how plausible these particles are in light of the measurements/restrictions. A key ingredient of particle Gibbs or PGAS is that we condition on a fixed trajectory for one of the particles. This reference particle is chosen as the previous draw for the respective states in the encompassing MCMC sampler, ensuring that the Markov kernel preserves the correct invariant distribution (see Andrieu *et al.*, 2010). The “ancestor” part of PGAS adds a resampling step for the parents of the reference trajectory, which alleviates path degeneracy by maintaining variability in the ancestral lineages, see Lindsten *et al.* (2014) for details. This improves mixing, especially with few particles, making the algorithm computationally attractive.

---

<sup>6</sup> Our framework is also directly applicable for nowcasting problems (see, e.g., Cimadomo *et al.*, 2022).

The  $h$ -specific distribution of the forecast conditional on the restrictions is available in closed form (see, e.g., West and Harrison, 1997, chapter 16) when assuming additive Gaussian errors and restrictions as in (7):

$$\begin{aligned} \mathbf{y}_{\tau+h} \mid \mathbf{r}_h, \mathbf{R}_h, \boldsymbol{\Omega}_h, \bullet &\sim \mathcal{N}(\mathbf{r}_h^*, \boldsymbol{\Omega}_h^*) \\ \mathbf{r}_h^* &= \boldsymbol{\mu}_{\tau+h}^{(m)} + \boldsymbol{\Sigma} \mathbf{R}_h' (\mathbf{R}_h \boldsymbol{\Sigma} \mathbf{R}_h' + \boldsymbol{\Omega}_h)^{-1} (\mathbf{r}_h - \mathbf{R}_h \boldsymbol{\mu}_{\tau+h}^{(m)}), \\ \boldsymbol{\Omega}_h^* &= \boldsymbol{\Sigma} - \boldsymbol{\Sigma} \mathbf{R}_h' (\mathbf{R}_h \boldsymbol{\Sigma} \mathbf{R}_h' + \boldsymbol{\Omega}_h)^{-1} \mathbf{R}_h \boldsymbol{\Sigma}. \end{aligned} \tag{9}$$

While (9) places the respective restriction only at a single horizon, this conditional Gaussian distribution can be used as an “optimal proposal” within our particle algorithm, in the spirit of a forward-filtering update. This allows to generate particles that satisfy the restrictions locally without being exposed to excessive weight degeneracy (the situation in which only very few, if any, particles are compatible with the restrictions). Used within an encompassing MCMC algorithm, PGAS propagates these particles in such a way that the restrictions are enforced jointly over the entire forecast path, and we obtain draws from  $p(\mathbf{y}_{\tau+1:\tau+H} \mid \mathcal{I}, \mathcal{C}_{1:H})$ .

Details about PGAS and explicit computations appear in Appendix A. We provide a brief summary of the main steps below. The initial conditions at each forecast origin  $\tau$  are known, and we can generate  $V - 1$  candidate particles to be considered alongside the reference particle from the previous MCMC draw. We loop through:

- (i) *Resampling and ancestor sampling*: To only retain particles with a high probability of having generated the measurements, i.e., the restrictions in (7), we sample ancestors for the non-reference particles using weights from the previous horizon, see step (iii) below. For the reference particle, we sample an ancestor index from a categorical distribution with support  $\{1, \dots, V\}$ , where the weights are proportional to the likelihood of particle  $v = 1, \dots, V$ , having generated the reference. Put simply, we obtain a random ancestry that is likely to have generated the reference path.

- (ii) *Propagation*: We then propagate the “surviving” non-reference particles from the resampling step one period forward. In case there are no restrictions ( $\mathcal{C}_h = \emptyset$ ) this is done using (5). In case there are restrictions, we use (9) to draw particles already conditioned on them at  $h$ . This avoids generating unsuitable candidate realizations and reduces the required number of particles  $V$ .
- (iii) *Weighting*: We update the weights used in step (i) for the next horizon based on (7). Specifically, we compute the likelihood of the lineage of the respective particles in light of the restrictions. In case there are no restrictions, we trivially obtain equal weights.

Using the weights at the maximum forecast horizon  $H$ , we draw a particle index and generate a fully smoothed trajectory by tracing back through the stored ancestor indices. This trajectory becomes the new reference path in the next iteration. These steps are run in each sweep of the main MCMC algorithm. In some cases we require the expectation of the predictive distribution in each MCMC sweep. We obtain this moment via exploiting a backward recursion (see, e.g., Godsill *et al.*, 2004) to obtain smoothing weights, and use the full set of particle trajectories for related computations, see also Appendix A. We provide a comparison of the precision sampler of Chan *et al.* (2025) with our PGAS approach in a linear VAR (with closed form solutions) using artificial data in Appendix B.

### 2.3. Generalized Impulse Response Functions

Various types of IRFs are widely used for both academic and policy analysis. These types of dynamic effects can generally be defined as the difference between two forecasts, see Koop *et al.* (1996). This provides a natural link to our previous discussions on predictive distributions. Indeed, due to the recursive nonlinear structure of our model, we need to

resort to a simulation-based variants of GIRFs.<sup>7</sup> We consider three main GIRF variants, which are nested in the expression:

$$\Delta_{\tau} = \mathbb{E}(\mathbf{y}_{\tau+1:\tau+H} \mid \mathcal{C}_{1:H}^{(s)}, \mathbf{x}_{\tau+1}, \mathcal{I}) - \mathbb{E}(\mathbf{y}_{\tau+1:\tau+H} \mid \mathcal{C}_{1:H}^{(b)}, \mathbf{x}_{\tau+1}, \mathcal{I}), \quad (10)$$

and differentiated by distinct conditioning assumptions. In line with the terminology in Crump *et al.* (2025), the first expectation refers to the “scenario” (s) forecast imposed with the restrictions  $\mathcal{C}_{1:H}^{(s)}$ , and the latter is the “baseline” (b) forecast subject to  $\mathcal{C}_{1:H}^{(b)}$ . Note that the conditioning on  $\mathcal{I}$  here refers to the fact that these quantities are computed with draws for parameter conditioning on all available information; the expectation is taken at time  $\tau$ , as indicated by conditioning on  $\mathbf{x}_{\tau+1}$ . For later reference, define horizon specific GIRFs  $\delta_{\tau,h}$  based on the structure of  $\Delta_{\tau} = (\delta'_{\tau,1}, \dots, \delta'_{\tau,H})'$ . That is, the one-step ahead forecast horizon is associated with the *contemporaneous* impact of the restrictions (see Breitenlechner *et al.*, 2022, for a similar timing convention).

Equation (10) nests our main variants:<sup>8</sup> the (1) unorthogonalized GIRF (UGIRF) which can be obtained from a fully reduced form model, but which may be extended to a structural scenario analysis when additionally identifying and restricting structural shocks; the (2) structural GIRF (SGIRF) in response to a structurally identified shock imposed exclusively with restrictions on the contemporaneous ( $h = 1$ ) shocks; and, the (3) restricted GIRF (RGIRF), using an SGIRF as a baseline but systematically manipulating transmission channels in the spirit of assessing alternate policy rules, by additionally imposing restrictions on observables in a certain way.

---

<sup>7</sup> Simulation-based methods are often used to compute dynamic causal effects in parametric nonlinear (e.g., Baumeister and Peersman, 2013; Alessandri and Mumtaz, 2019), or nonparametric (e.g., Huber and Rossini, 2022; Clark *et al.*, 2025; Hauzenberger *et al.*, 2025b) models.

<sup>8</sup> Other related work includes Bernanke *et al.* (1997); Hamilton and Herrera (2004); Sims and Zha (2006); Baumeister and Benati (2013); Baumeister and Kilian (2014); Adrian *et al.* (2025).

**Unorthogonalized GIRF.** This version is closely related to the original GIRF of [Koop et al. \(1996\)](#). It is obtained by restricting only observables as the scenario, which can be compared to the unconditional forecast as a baseline. Without any further restrictions, UGIRFs simply reflect a likely combination of structural shocks that drive the change in the respective observable(s), see also [Crump et al. \(2025\)](#). In case one is willing to impose additional structure in the model via identifying  $\mathbf{H}$ , restrictions on observables can be complemented with restrictions on shocks, see also [Antolin-Diaz et al. \(2021\)](#). Specifically, one may restrict a subset of the shocks to their unconditional distribution (non-driving shocks), while the unrestricted subset (driving shocks) is allowed to deviate to deliver the structural scenario in terms of the observables. In light of the assumptions about the structural form above, this implies  $\mathbf{r}_h^{(u)} = \mathbf{0}$  and  $\mathbf{\Omega}_h^{(u)} = \mathbf{I}$  for the *non-driving* shocks.

**Structural GIRF.** This variant relies on orthogonalized structural economic shocks. We follow the literature (see, e.g., [Jordà and Taylor, 2025](#)), and define the SGIRF in response to the  $j$ th structural shock of magnitude  $d$  as:

$$\Delta_{j\tau}^{(d)} = \mathbb{E}(\mathbf{y}_{\tau+1:\tau+H} \mid u_{j\tau+1} = d_0 + d, \mathbf{x}_{\tau+1}, \mathcal{I}) - \mathbb{E}(\mathbf{y}_{\tau+1:\tau+H} \mid u_{j\tau+1} = d_0, \mathbf{x}_{\tau+1}, \mathcal{I}). \quad (11)$$

The parameter  $d_0$  is the baseline level of the shock. Since the structural errors enter linearly in our model, choices about the baseline level do not matter and we implicitly use  $d_0 = 0$  in most of our subsequent discussions. Equation (11) is a special case of (10), which can be obtained by defining  $\mathcal{C}_1^{(s)}$  with  $\mathbf{r}_1^{(u)} = (d_0 + d) \cdot \mathbf{e}'_j$ , and  $\mathcal{C}_1^{(b)}$  with  $\mathbf{r}_1^{(u)} = d_0 \cdot \mathbf{e}'_j$ ;  $\mathbf{e}'_j$  is the  $j$ th column of  $\mathbf{I}_n$ . For both forecasts we use  $\mathbf{\Omega}_1^{(u)} = \text{diag}(1, \dots, 1, 10^{-8}, 1, \dots, 1)$  with an approximately binding restriction in the  $j$ th position,  $\mathbf{R}_1^{(u)} = \mathbf{e}_j$ , and we have  $\mathcal{C}_h = \emptyset$  for  $h > 1$ . The GIRF at time  $\tau$  and horizon  $h$  in response to the  $j$ th structural shock of size  $d$  is referred to as  $\delta_{j\tau,h}^{(d)}$ . In [Appendix A](#) we discuss an alternative computation method which can be used without running PGAS, and [Appendix B](#) again provides an illustration and comparison in a linear context using artificial data.

**Restricted GIRF.** The final variant we consider combines the SGIRF with restrictions on observables, so that we may impose  $\mathbb{E}(\mathbf{R}_h \boldsymbol{\delta}_{\tau,h}^{(d)}) = \mathbf{0}_{r_h}$  along the desired dimensions. This approach switches off specific transmission channels of structural shocks, by partially matching the observables of the scenario forecast conditional on the shock with those of the forecasted observables in the baseline predictive distribution. Breitenlechner *et al.* (2022) provide a discussion of this approach in a linear context, which allows for directly restricting the IRF. In our nonlinear setting, we need to account for varying initial conditions and work with the two conditional forecasts, as we cannot directly manipulate the IRF.

From an implementation perspective, we may obtain a draw  $\mathbf{y}_{\tau+1:\tau+H}^{(b,m)}$  from the baseline distribution  $p(\mathbf{y}_{\tau+1:\tau+H} | \mathcal{C}_1^{(b)}, \mathbf{x}_{\tau+1}, \mathcal{I})$ , where  $\mathcal{C}_1^{(b)}$  features the impact restrictions as in (11) plus any desired restrictions on the future sequence of non-driving shocks (this is however not a necessary requirement, as otherwise all shocks may be allowed to deviate from their unconditional distribution). For computing the scenario forecast, we then augment the shock impact restriction in (11) with a restriction on the respective observable, so that  $\mathbf{r}_h^{(y)} = \mathbf{R}_h^{(y)} \mathbf{y}_{\tau+h}^{(b,m)}$ , to define  $\mathcal{C}_{1:H}^{(s)}$ . Individual *realized paths* from the scenario and baseline forecast distributions are identical along the restricted dimensions. The restricted GIRF then arises as the difference in *expected values* as in (10), which has an expected value of zero, but this cannot strictly be enforced for each individual draw. This is because otherwise the reference trajectory in each sweep of the PGAS algorithm may not necessarily respect the restrictions.

**Conditional and unconditional GIRFs.** One may obtain all three of these GIRFs conditional on each point in time  $\tau$ . In principle, one may thus consider “time-varying” dynamic effects of shocks for each period individually, or averages for different known subsets of observations (e.g., recessions vs. expansions). This time variation, however, is exclusively due to variation across initial conditions, since the functional form is nonlinear but time-invariant. Unconditional (G)IRFs can be computed in various ways (e.g., by randomizing

over initial conditions and averaging, in bootstrap-type approaches); see also Kilian and Lütkepohl (2017, chapter 18), and Gonçalves *et al.* (2021, 2024) for definitions. In case we require estimates of unconditional versions, our preferred approach is to compute the GIRF at each point in time, and then average out this source of randomness,  $\bar{\Delta} = \sum_{\tau=1}^T \Delta_{\tau}/T$ .

**Lucas critique.** Conditional forecasts and the GIRF variants relate to work on evaluating counterfactual policy rules, and may be exposed to the Lucas critique. This is the case particularly when imposing highly unusual conditioning scenarios, as opposed to mere “modest policy interventions,” which may leave behavioral patterns of economic agents unchanged (see Leeper and Zha, 2003). A recent discussion of these aspects is provided in McKay and Wolf (2023): the issue is that any imposed counterfactual paths on observables in this spirit require specific offsetting *future* shocks, and the restrictions hold only ex post and not necessarily in expectation. Since our approach is based on the same ideas as Antolin-Diaz *et al.* (2021), similar concerns as in linear models apply. It is worth noting that by virtue of a nonparametric approach, structural shocks may locally affect relationships among variables, even if the underlying functional form of the respective relationship remains fixed. Thus, while our approach is susceptible to the Lucas critique, it is likely more robust to related concerns than linear models.

### 3. MODEL SPECIFICATION AND ESTIMATION ALGORITHM

#### 3.1. Multivariate System Estimation

To estimate the model in (1) we rely on a conditional representation of its  $n$  equations. Let  $\mathbf{e}_i$  of size  $1 \times n$  denote the  $i$ th row of  $\mathbf{I}_n$ , and  $\mathbf{E}_i$  of size  $(n-1) \times n$  results from deleting the  $i$ th row of  $\mathbf{I}_n$ . Using  $\mathbf{y}_{-it} = \mathbf{E}_i \mathbf{y}_t$  we may write  $\mathbf{y}_t = \mathbf{e}'_i y_{it} + \mathbf{E}'_i \mathbf{y}_{-it}$ . Under the assumptions of (1),  $p(y_{it} | \mathbf{y}_{-it}, \bullet) \propto \exp \left\{ -(\mathbf{e}_i \boldsymbol{\Sigma}_t^{-1} \mathbf{e}'_i y_{it}^2 - 2y_{it} \mathbf{e}_i \boldsymbol{\Sigma}_t^{-1} (\mathbf{F}(\mathbf{x}_t) - \mathbf{E}'_i \mathbf{y}_{-it}))/2 \right\}$ , which is a Gaussian with variance  $\zeta_{it}^2 = (\mathbf{e}_i \boldsymbol{\Sigma}_t^{-1} \mathbf{e}'_i)^{-1}$  and mean  $\mu_{it} = \zeta_{it}^2 (\mathbf{e}_i \boldsymbol{\Sigma}_t^{-1} (\mathbf{F}(\mathbf{x}_t) - \mathbf{E}'_i \mathbf{y}_{-it}))$ . Indeed, this

distribution is equivalent to a common representation of the conditional multivariate Gaussian (see, e.g., Cong *et al.*, 2017, section 2). The mean can alternatively be written as  $\mu_{it} = f_i(\mathbf{x}_t) + \varsigma_{it}^2(\mathbf{e}_i \boldsymbol{\Sigma}_t^{-1} \mathbf{E}_i'(\mathbf{y}_{-it} - \mathbf{E}_i \mathbf{F}(\mathbf{x}_t)))$ , and we define  $\tilde{\mu}_{it} = \varsigma_{it}^2(\mathbf{e}_i \boldsymbol{\Sigma}_t^{-1} \mathbf{E}_i'(\mathbf{y}_{-it} - \mathbf{E}_i \mathbf{F}(\mathbf{x}_t)))$ , i.e.,  $\mu_{it} = f_i(\mathbf{x}_t) + \tilde{\mu}_{it}$ .<sup>9</sup> The  $i$ th equation of the multivariate model in regression form, conditional on all other equations, is then given by:

$$(\mathbf{y}_{it} - \tilde{\mu}_{it}) = f_i(\mathbf{x}_t) + u_{it}, \quad u_{it} \sim \mathcal{N}(0, \varsigma_{it}^2), \quad (12)$$

which can be used in a Gibbs sampler to update the conditional mean relationships by looping through equations  $i = 1, \dots, n$ . This approach is similar to the one of Esser *et al.* (2024) and allows to treat each equation of the multivariate system individually, conditional on all other equations.

## 3.2. Bayesian Additive Regression Trees

The approach we discuss in Section 2 works with any implementation of multivariate models with additive and jointly Gaussian errors. That is, assuming a linear functional form for  $\mathbf{F}(\mathbf{x}_t)$  combined with suitable priors results in a standard Bayesian VAR (BVAR).<sup>10</sup> In case

---

<sup>9</sup> Related papers often either use a mapping between the structural and reduced form of the VAR to enable equation-by-equation estimation (see, e.g., Hauzenberger *et al.*, 2025b), or rely on factor models for the reduced form errors (see, e.g., Clark *et al.*, 2023). These approaches come with computational and inferential advantages and disadvantages. The former is simple to implement but requires parameterizing a structural form, which may cause issues such as inadvertently (instead of purposefully to achieve structural identification) breaking order-invariance of the equations. The latter allows for order-invariant inference but gives rise to the usual identification challenges of factor models. Our approach uses an order-invariant reduced form model for estimating the conditional mean relationships; see also Arias *et al.* (2023); Chan *et al.* (2024) for discussions.

<sup>10</sup>When we consider linear versions of our model for comparisons, we implement this setting with  $\mathbf{F}(\mathbf{x}_t) = \mathbf{A}\mathbf{x}_t$  where  $\mathbf{A}$  is an  $n \times k$  matrix of reduced form VAR coefficients. We assume a horseshoe prior with a single global shrinkage component on these parameters, see also Hauzenberger *et al.* (2024); (12) can be used to update the VAR coefficients equation-by-equation from Gaussian posteriors.

we treat  $\mathbf{F}(\mathbf{x}_t)$  nonparametrically, several options are available. Due to its versatility and established favorable empirical properties we mentioned earlier, we use BART to approximate the equation-specific functions in our applied work. That is, we consider a sum of  $s = 1, \dots, S$ , tree functions  $\ell_{is}(\mathbf{x}_t | \mathcal{T}_{is}, \mathbf{m}_{is})$  such that  $f_i(\mathbf{x}_t) \approx \sum_{s=1}^S \ell_{is}(\mathbf{x}_t | \mathcal{T}_{is}, \mathbf{m}_{is})$  where  $\mathcal{T}_{is}$  are regression trees and  $\mathbf{m}_{is}$  is a vector of terminal node parameters (which serve as fitted values). Instead of having a single but complex tree, BART is akin to ensemble methods, and uses a sum of many simple trees (“weak learners”), which has been shown to work well.

Using BART requires an algorithm that estimates splitting variables and thresholds for which we specify suitable priors that together yield  $p(\mathcal{T}_{is})$ ; we further need a prior on the terminal node parameters  $p(\mathbf{m}_{is} | \mathcal{T}_{is})$ . Our setup follows [Chipman \*et al.\* \(2010\)](#) and we first define the probability that a tree ends at a specific node at depth  $d = 0, 1, 2, \dots$ , as  $\alpha/(1+d)^\beta$ , with  $\alpha \in (0, 1)$  and  $\beta \in \mathbb{R}^+$ . This prevents trees from getting overly complex and provides regularization (here, we rely on the default values  $\alpha = 0.95$ , and  $\beta = 2$ , which perform well across many datasets). For the splitting variables, we choose a uniform prior. This implies that each predictor is equally likely to be selected as a splitting variable. We further assign a uniform prior to all thresholds within the splitting rules, based on the range of the respective splitting variable.

Next we specify the prior for the terminal node parameters. On these parameters  $\mathbf{m}_{is,l}$ , for  $l = 1, \dots, \#\text{TN}_{is}$ , where  $\#\text{TN}_{is}$  denotes the number of terminal node parameters of tree  $s$  in equation  $i$ , we impose independent conjugate Gaussian priors that are symmetric across trees and identical for all terminal nodes. As suggested by [Chipman \*et al.\* \(2010\)](#) the moments of these priors are chosen in a data-driven manner, such that 95% of the prior probability lies in the interval  $(\min(\mathbf{y}_i), \max(\mathbf{y}_i))$ , where  $\mathbf{y}_i = (y_{i1}, \dots, y_{iT})'$ , and such that shrinkage increases the more trees  $S$  are chosen for estimation. We choose  $S = 250$  trees which has been shown to work well for typical macroeconomic time series applications (see, e.g., [Huber \*et al.\*, 2023](#)). Additional details are provided in [Appendix A](#).

### 3.3. Other Priors and Sampling Algorithm

Inspired by [Carriero \*et al.\* \(2024\)](#), we assume that  $\Sigma_t = s_t \Sigma$  — the covariance structure only varies proportionally over time and  $s_t$  is used to capture outliers. Our prior setup for the constant part of the covariance matrix follows [Esser \*et al.\* \(2024\)](#). Specifically, we use a hierarchical inverse Wishart prior  $\Sigma \mid \{a_i\}_{i=1}^n \sim \mathcal{W}^{-1}(s_0, \mathbf{S}_0)$ , where  $s_0 = \nu + n - 1$ ,  $\mathbf{S}_0 = 2\nu \cdot \text{diag}(1/a_1, \dots, 1/a_n)$  and  $a_i \sim \mathcal{G}^{-1}(1/2, 1/A_j^2)$  for  $i = 1, \dots, n$ , and a fixed scale parameter  $A_j > 0$ . Setting  $\nu = 2$  implies a comparatively uninformative prior about the implied correlation structure, different from fixed-hyperparameter versions of this prior which has a tendency to overshrink. If applicable, we assume that  $\Pr(s_t^{1/2} = 1) = 1 - \mathbf{p}$  and  $\Pr(s_t^{1/2} \sim \mathcal{U}(2, \bar{s})) = \mathbf{p}$ , where  $\mathcal{U}(2, \bar{s})$  is a discrete uniform distribution with (integer) support between 2 and  $\bar{s} = 6$  and  $\mathbf{p} \sim \mathcal{B}(a_{\mathbf{p}}, b_{\mathbf{p}})$  is the probability of observing an outlier. We set  $a_{\mathbf{p}} = 1$ ,  $b_{\mathbf{p}} = 50$  which a priori implies about 2% of the observations are outliers. Alternative and more flexible stochastic volatility specifications are available in this context (e.g., [Chan, 2023](#)). We obtain draws from the joint posterior of our model using a fairly straightforward MCMC algorithm. Details are provided in [Appendix A](#).

## 4. EMPIRICAL APPLICATIONS

We employ the proposed framework in three related yet distinct applications. First, we use the annual stress test scenarios conducted by the Federal Reserve System and compute forecasts conditional on multiple observables. Second, we study the implications of varying financial conditions on tail risks of output growth, inflation, and employment. Third, we identify a US-based financial shock and gauge the role of spillovers and spillbacks.<sup>11</sup>

---

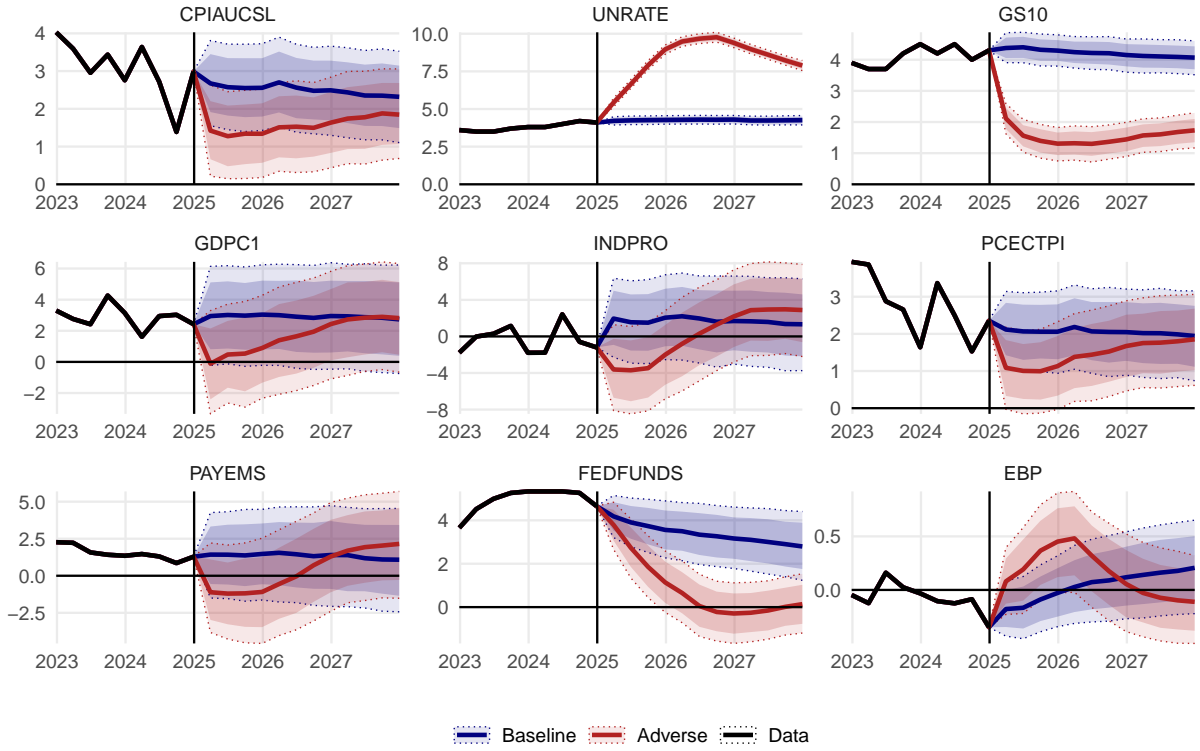
<sup>11</sup>We sometimes compare a *homoskedastic* (BART-hom,  $s_t = 1$  for all  $t$ ) and *heteroskedastic* BART (BART-het) with a *heteroskedastic* BVAR (BVAR-het), featuring the outlier specification. The linear BVAR serves for comparisons as it is a popular workhorse model in related contexts (e.g., [Crump \*et al.\*, 2025](#)); due to our sample featuring the Covid-19 pandemic, we disregard its homoskedastic version.

## 4.1. Stress Testing Scenarios for the US Economy

In our first application, we conduct a scenario analysis for the US economy inspired by the 2025 version of the *Dodd-Frank Act* (DFA) stress test assumptions. This annual exercise is conducted and published by the Board of Governors of the Federal Reserve System. Details about the underlying dataset are provided in Appendix B. The information set features about 25 broad variables (capturing economic activity, labor market, prices, housing and the financial sector). We estimate our models using quarterly data ranging from 1976Q1 to 2024Q4, and subsequently consider a baseline and adverse scenario for the period from 2025Q1 to 2027Q4. These scenarios are imposed via constraints on the path of the unemployment rate (UNRATE), CPI inflation (CPIAUCSL), and 10-year government bond yields (GS10), inspired by *Chan et al. (2025)*. We set the tightness of the restrictions in  $\Omega_{1:H}$  using the marginal variances on the diagonal of  $\Sigma$  in each MCMC draw.

The posterior median forecasts and 50 and 68 percent posterior credible sets, obtained with **BART-het** (conditional forecast distributions for other specifications are in Appendix B), for the restricted and selected unrestricted variables, are shown in Figure 1: real GDP (GDPC1), industrial production (INDPRO), personal consumption expenditure (PCE) inflation (PCECTPI), payroll employment (PAYEMS), federal funds rate (FEDFUNDS) and the *Gilchrist and Zakrajšek (2012)* excess bond premium (EBP). The baseline scenario draws from the consensus projections from 2025 *Blue Chip Financial Forecasts* and *Blue Chip Economic Indicators*; the adverse scenario is characterized by a recession. Figure 2 displays corresponding UGIRFs across model specifications, which are computed as the indicated scenario minus the unconditional forecast.

The unconditional forecasts from **BART-het** approximately coincide with the baseline scenario (the UGIRF credible sets cover 0 in most cases). For the adverse scenario, a different picture emerges — the scenario forecasts differ significantly from the unconditional forecasts. There is an immediate downturn of economic activity and a reduction in payroll employment.




**Figure 1:** Conditional forecasts with BART-het for selected variables.

*Notes:* Posterior median alongside 50/68 percent credible sets. Restricted variables: consumer price inflation (CPIAUCSL), unemployment rate (UNRATE), 10-year government bond yields (GS10); unrestricted variables: Real GDP (GDPC1), industrial production (INDPRO), personal consumption expenditure inflation (PCECTPI), payroll employment (PAYEMS), federal funds rate (FEDFUNDS) and excess bond premium (EBP).


Financial conditions tighten initially but tend to improve subsequently, partially through a monetary easing response by the central bank as reflected in the policy rate. In addition, the assumed trajectory of the conditioning variables results in a modestly disinflationary episode that vanishes by 2027.

The conditional forecast distributions across model specifications are similar for most variables apart from output growth, industrial production, and employment. For BART-het, the magnitude of the peak contraction reduces by about half. This can be explained by noting that the algorithm decides to classify several observations (that are otherwise informative about directional movements of variables when assuming homoskedasticity) as outliers. We note that BART, due to the way how tree-based approaches fit data, is capable of dealing


(a) BART-het



(b) BART-hom



(c) BVAR-het



**Figure 2:** Unorthogonalized GIRF (UGIRF, scenario minus unconditional forecast) across model specifications for selected variables.

*Notes:* Posterior median alongside 50/68 percent credible sets; adverse and baseline scenarios in red and blue, respectively. Restricted variables: consumer price inflation (CPIAUCSL), unemployment rate (UNRATE), 10-year government bond yields (GS10); unrestricted variables: Real GDP (GDPC1), industrial production (INDPRO), personal consumption expenditure inflation (PCECTPI), payroll employment (PAYEMS), federal funds rate (FEDFUNDS) and excess bond premium (EBP).

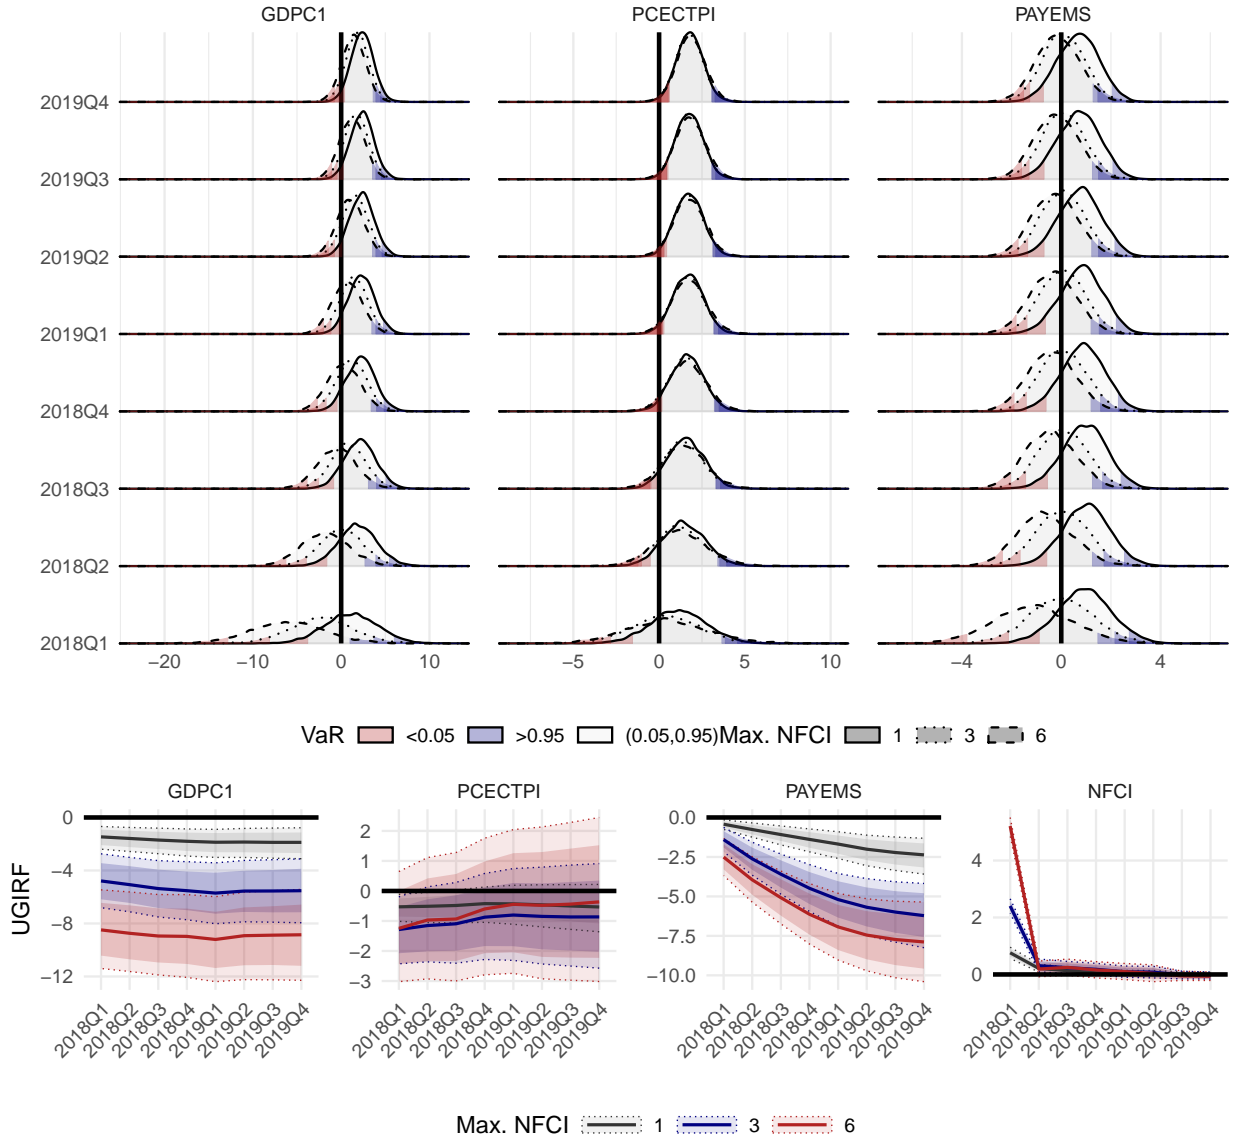
with outliers and heteroskedastic data features in the conditional mean function by design (see Huber *et al.*, 2023; Clark *et al.*, 2023, for discussions), even when assuming constant variances. But in this case several observations are classified as noise rather than signal. For linear models, the forecast trajectories are smoother, and an overshooting behavior is noticeable around 2027. Moreover, the response of the federal funds rate is somewhat more persistent, which may be due to nonlinearities arising at the effective lower bound which the BVAR fails to capture.

## 4.2. Financial Conditions in the US and Tail Risk Scenarios

In the next empirical application, we restrict our sample to 1976Q1–2017Q4 and consider the period from 2018Q1 until 2019Q1 as a laboratory to assess nonlinearities between economic variables and financial conditions. For this application we use BART-hom (since the pandemic observations are excluded). We investigate nonlinear patterns of macroeconomic risk, which, following the “growth-at-risk” approach of [Adrian \*et al.\* \(2019\)](#), is defined as the predictive quantiles of some variable of interest at a pre-defined probability level (in line with value-at-risk, VaR, in finance). We pick this period because the information set already contains the global financial crisis (and the model thus had the opportunity to learn from this severe financial episode), and because this “holdout sample” otherwise coincides with a comparatively eventless period. We impose hard constraints on the  $h = 1$  value of the National Financial Conditions Index (NFCI) and trace the effects of these scenarios on several macroeconomic variables.

The scenarios are defined to reflect an increase of the NFCI by approximately 1, 3 and 6 unconditional standard deviations (SDs, reflecting tighter financial conditions) in 2018Q1, i.e., in  $\mathcal{C}_1$ , which we implement by placing these values as hard restriction on the NFCI in that quarter. From 2018Q2 onward we leave the future unrestricted, i.e.,  $\mathcal{C}_h = \emptyset$  for  $h > 1$ . We investigate growth-at-risk (quantiles of real GDP), inflation-at-risk (quantiles of PCE inflation) and labor-at-risk (quantiles of growth in payroll employment) as our objects of interest (see [Adams \*et al.\*, 2021](#); [Pfarrhofer, 2022](#); [Botelho \*et al.\*, 2024](#); [Clark \*et al.\*, 2024](#); [Lopez-Salido and Loria, 2024](#), for related papers). The CF distributions (density estimates), for average growth rates  $\bar{\mathbf{y}}_{T+h} = \sum_{j=1}^h \mathbf{y}_{T+j}/h$ , are shown in Figure 3. In the lower panel, we show the difference between the conditional scenario distributions and the unconditional one, which yields the UGIRF (cumulated for all variables except the NFCI).

The different NFCI scenarios for 2018Q1 shift the predictive distributions. In all cases, the economy contracts which is reflected in a decrease of real GDP growth and



**Figure 3:** Conditional forecast distributions for selected variables and macroeconomic value-at-risk (VaR) and unorthogonalized GIRFs for different scenarios of financial stress.

*Notes:* “Max. NFCI” refers to the maximum value of the NFCI for the scenarios, with moderate (1), severe (3, comparable to the global financial crisis), and extreme stress (6). Sampling period 1976Q1 to 2017Q4, hard restriction applies 2018Q1. Posterior median alongside 50/68 percent credible set. Variables: Real GDP (GDPC1), personal consumption expenditure inflation (PCECTPI), payroll employment (PAYEMS), national financial conditions index (NFCI). Distribution of average growth rates from 2018Q1 until the quarter indicated on the y-axis in the upper panel. Cumulated UGIRFs for all variables except the NFCI.

payroll employment, and the simulated shock has a modestly disinflationary effect. While the upper tails of the distributions (upside risk) remain comparatively stable, downside risk as measured by the lower quantiles increases significantly for all considered variables (the

red shaded VaR  $< 0.05$  moves strongly leftwards), and there are some visible asymmetries. While the moderate NFCI scenario (max. NFCI = 1) results in growth-at-risk for the 5th percentile at about  $-5$  percent, the severe (max. NFCI = 3) and extreme (max. NFCI = 6) stress scenarios yield  $-7$  and  $-12.5$  percent, respectively. This finding is also present for inflation-at-risk and labor-at-risk.

The resulting predictive distributions exhibit non-Gaussian features, chief among them being heavy tails and skewness (see also Clark *et al.*, 2023, for a related but simplified scenario analysis in this context). In addition, there are hints of multimodality as the assumed values for the NFCI in 2018Q1 turn more extreme, which relates to the discussions in Adrian *et al.* (2021). These features can arise — even in one-step ahead predictions and for a single restricted period — due to, for example, the initial conditions of the nonlinear unconditional mean function at the forecast origin.

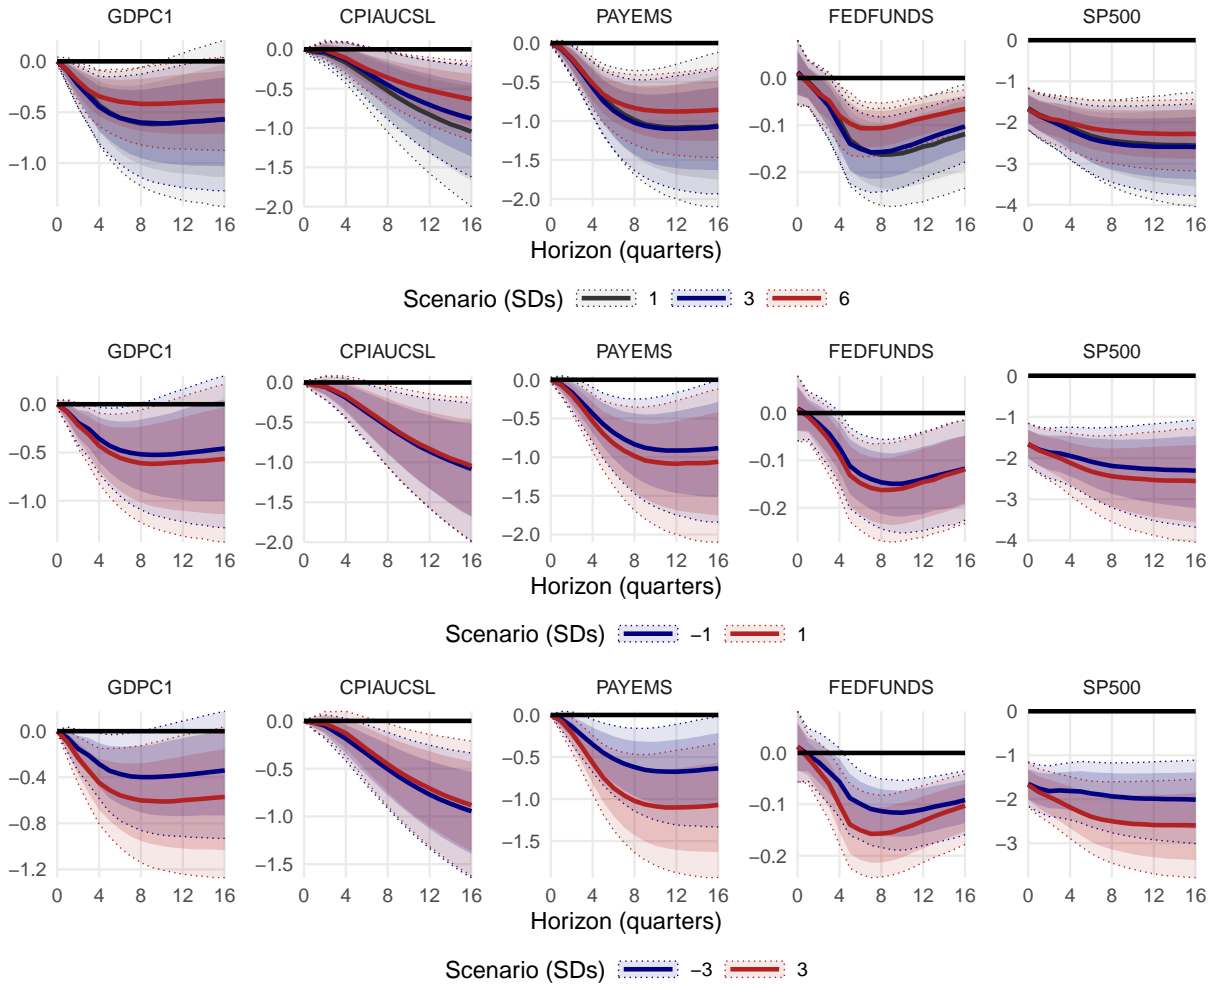
### 4.3. Spillovers and Spillbacks of US Financial Shocks

In our final application, we estimate the effect of a financial shock in the US and trace its effects through the domestic economy, but also capture spillovers and spillbacks to and from other economies. We use an adjusted dataset in this case, which drops several of the domestic indicators, but adds bilateral exchange rates alongside real GDP for the EA and the UK. To identify the financial shock, we place timing-restrictions on the contemporaneous impulse responses. This is operationalized with a specific ordering of the quantities in the vector  $\mathbf{y}_t$  — we structure this vector such that all slow moving domestic and foreign macroeconomic variables come first (which imposes zero restrictions on impact). These variables are then followed by the EBP, and all fast moving variables such as those capturing the financial economy. We then use a Cholesky decomposition of the form  $\Sigma = \mathbf{P}\mathbf{P}'$  where  $\mathbf{P}$  is lower triangular. That is,  $\mathbf{H}^{-1} = \mathbf{P}$ , and we orthogonalize the structural shocks different to our previous applications. The orthogonalized innovation of the EBP equation is interpreted as

the financial shock, similar to Gilchrist and Zakrajšek (2012); Barnichon *et al.* (2022). In a multicountry context, Huber *et al.* (2024) use an identification scheme identical to ours.

We use  $d_0 = 0$  and simulate different shock sizes and signs with  $d \in \{-3, -1, 1, 3, 6\}$ . In contrast with conventional linear frameworks, our approach allows to assess nonlinearities of higher-order responses with respect to different signs and magnitudes of a proportional shock impact with GIRFs. Such asymmetries and related nonlinearities have recently gained attention both in a VAR and local projection context, see, e.g., Mumtaz and Piffer (2025); Carriero *et al.* (2025); Forni *et al.* (2024); Hauzenberger *et al.* (2025a). While our framework allows to compute dynamic responses for each period in our sample, we focus on time averages in the results that follow. Further, we rescale all SGIRFs by computing  $\delta_{\tau,h}^{(d)}/d$  so that all responses shown below reflect a 1 SD financial shock (see also Gonçalves *et al.*, 2024; Kolesár and Plagborg-Møller, 2025). Note that for linear VARs, such scaling yields identical IRFs for all shock sizes. For nonlinear models, this is not necessarily the case and allows for a visual inspection of asymmetries. Selected variables are shown in Figure 4. The rows in the figure show different subsets of the same results, structured such that the shocks of different signs and sizes can be compared with ease.

Starting with the first row, we find that the size of the financial shock causes limited asymmetries in responses for the indicated variables. Financial shocks of different sizes rather symmetrically decrease real GDP and payroll employment. Interestingly, the effects of very large shocks increase slightly less than proportionally. Peak effects occur about two years after impact of the shock. In addition, the shock puts a persistent downward pressure on prices and leads to a decline in the federal funds rate which peaks at about  $-15$  basis points after around a year. Notably the federal funds rate does not react on impact, different to stock returns which immediately decline by about 1.5 percent during the quarter when the shock materializes. Qualitatively and in terms of magnitudes, these estimates are

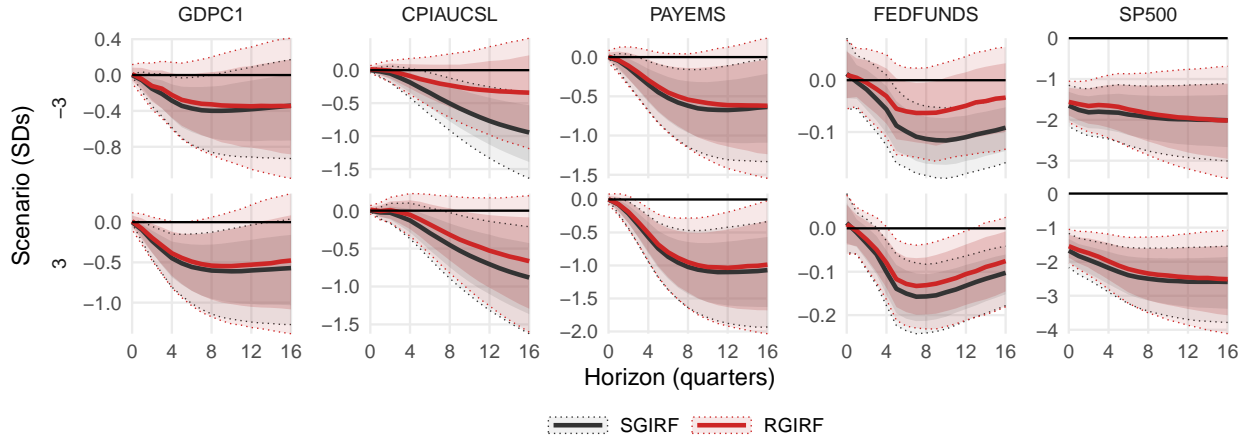


**Figure 4:** Structural GIRFs for selected variables to a financial shock in the US, comparing asymmetries due to size and sign of the shocks.

*Notes:* Posterior median alongside 50/68 percent credible sets. Cumulated responses for variables in differences and levels for all other variables. “Horizon” refers to periods after impact of the shock. Variables: Real GDP (GDPC1), consumer price inflation (CPIAUCSL), payroll employment (PAYEMS), federal funds rate (FEDFUNDS) and S&P500 index (SP500).

roughly in line with the previous literature. The US-based financial shock spills over to the other economies and leads to contractionary effects in terms of real GDP, see Appendix B.

Having established that the size of the financial shock does not seem to matter much, the second and third row zoom into sign asymmetries. The 1 SD US-based financial shock does not result in significant sign asymmetries (the posterior distributions overlap for the most part). Turning to the final row, this clearly differs for larger sized shocks of different



**Figure 5:** Structural and restricted GIRFs for selected variables.

*Notes:* The restricted case assumes that the US financial shock does not spill over to non-domestic variables. Posterior medians alongside 50/68 percent credible sets. Cumulated responses for variables in differences and levels for all other variables. “Horizon” refers to periods after impact of the shock. Variables: Real GDP (GDPC1), consumer price inflation (CPIAUCSL), payroll employment (PAYEMS), federal funds rate (FEDFUNDS) and S&P500 index (SP500).

signs. For these GIRFs that show the responses to a positive (adverse) and negative (benign) 3 SD shock, asymmetries are visible for most variables. In particular, we find that the negative effect on payroll employment is almost twice as large for adverse shocks, and the Federal Reserve responds more strongly to adverse financial shocks, as measured by the much stronger shift in the federal funds rate. These findings corroborate previous evidence (see, e.g., Forni *et al.*, 2024; Hauzenberger *et al.*, 2025a).

We next explore the role of international variables in the domestic transmission of the US shock. For this purpose, besides SGIRFs, we consider an alternative analysis where non-domestic transmission channels are switched off — we investigate how international channels affect the domestic transmission of shocks originating in the US. That is, we impose the restriction that foreign variables (real GDP in the EA and UK, and exchange rates) do not respond to the financial shock in the US in this counterfactual, thereby simulating a scenario where the financial shock is confined to the domestic economy without any real or financial spillovers (or spillbacks); see Breitenlechner *et al.* (2022) for a monetary application in this

context. We restrict all domestic shocks to their unconditional distribution, and use the non-domestic ones as driving shocks to compute the RGIRFs.

The results are displayed in Figure 5. The upper panels show the SGIRFs (those shown and discussed in the context of Figure 4) and “no spillovers” RGIRFs. Two key findings are worth reporting. First, for the most part, ruling out spillovers and spillbacks does not significantly alter the dynamic responses after an adverse shock. It is worth mentioning, however, that restricting the international transmission leads to slightly smaller effects on average. Second, international transmission channels appear to matter for inflation dynamics, and especially so for benign financial shocks. The response of inflation turns insignificant in this case, which is also associated with a less forceful action by the central bank as captured in more muted response of the policy rate.

## 5. CONCLUSIONS

This paper presents a unified methodology for conducting scenario analysis in multivariate macroeconomic settings, accommodating nonlinearities and unknown functional forms of conditional mean relationships. These methods are applicable to traditional nonlinear frameworks, such as variants of threshold or time-varying parameter models, but also to more recently developed models incorporating Bayesian machine learning. Our framework addresses some limitations of linear and parametric models in generating various types of counterfactual analyses and is suitable for large macroeconomic datasets.

The empirical applications, using Bayesian additive regression trees as an example of nonparametric modeling of the conditional mean function, underscore the role of nonlinearities in shaping macroeconomic dynamics. For instance, a scenario analysis based on Federal Reserve stress test assumptions reveals differences between linear and nonlinear models in forecasting economic contractions and recoveries. Similarly, in a growth-at-risk application we measure nonlinear macroeconomic risks under financial stress. Finally, an analysis of

financial spillovers reveals asymmetries in the transmission of shocks and the influence of international linkages on domestic outcomes.

## REFERENCES

- ADAMS PA, ADRIAN T, BOYARCHENKO N, AND GIANNONE D (2021), “Forecasting macroeconomic risks,” *International Journal of Forecasting* **37**(3), 1173–1191.
- ADRIAN T, BOYARCHENKO N, AND GIANNONE D (2019), “Vulnerable growth,” *American Economic Review* **109**(4), 1263–1289.
- (2021), “Multimodality in macrofinancial dynamics,” *International Economic Review* **62**(2), 861–886.
- ADRIAN T, GIANNONE D, LUCIANI M, AND WEST M (2025), “Scenario Synthesis and Macroeconomic Risk,” *arXiv* **2505.05193**.
- ALESSANDRI P, AND MUMTAZ H (2019), “Financial regimes and uncertainty shocks,” *Journal of Monetary Economics* **101**, 31–46.
- ANDERSSON MK, PALMQVIST S, AND WAGGONER DF (2010), “Density-conditional forecasts in dynamic multivariate models,” *Sveriges Riksbank Working Paper Series* **247**.
- ANDRIEU C, DOUCET A, AND HOLENSTEIN R (2010), “Particle Markov chain Monte Carlo methods,” *Journal of the Royal Statistical Society Series B* **72**(3), 269–342.
- ANTOLIN-DIAZ J, PETRELLA I, AND RUBIO-RAMÍREZ JF (2021), “Structural scenario analysis with SVARs,” *Journal of Monetary Economics* **117**, 798–815.
- ARIAS JE, RUBIO-RAMIREZ JF, AND SHIN M (2023), “Macroeconomic forecasting and variable ordering in multivariate stochastic volatility models,” *Journal of Econometrics* **235**(2), 1054–1086.
- BAÑBURA M, GIANNONE D, AND LENZA M (2015), “Conditional forecasts and scenario analysis with vector autoregressions for large cross-sections,” *International Journal of Forecasting* **31**(3), 739–756.
- BARNICHON R, MATTHES C, AND ZIEGENBEIN A (2022), “Are the effects of financial market disruptions big or small?” *Review of Economics and Statistics* **104**(3), 557–570.
- BAUMEISTER C, AND BENATI L (2013), “Unconventional Monetary Policy and the Great Recession: Estimating the Macroeconomic Effects of a Spread Compression at the Zero Lower Bound,” *International Journal of Central Banking* **31**.
- BAUMEISTER C, AND KILIAN L (2014), “Real-time analysis of oil price risks using forecast scenarios,” *IMF Economic Review* **62**(1), 119–145.

- BAUMEISTER C, AND PEERSMAN G (2013), “Time-varying effects of oil supply shocks on the US economy,” *American Economic Journal: Macroeconomics* **5**(4), 1–28.
- BERNANKE BS, GERTLER M, WATSON M, SIMS CA, AND FRIEDMAN BM (1997), “Systematic monetary policy and the effects of oil price shocks,” *Brookings Papers on Economic Activity* **1997**(1), 91–157.
- BOTELHO V, FORONI C, AND RENZETTI A (2024), “Labour at risk,” *European Economic Review* **170**, 104849.
- BREITENLECHNER M, GEORGIADIS G, AND SCHUMANN B (2022), “What goes around comes around: How large are spillbacks from US monetary policy?” *Journal of Monetary Economics* **131**, 45–60.
- CARRIERO A, CLARK TE, MARCELLINO M, AND MERTENS E (2024), “Addressing COVID-19 outliers in BVARs with stochastic volatility,” *Review of Economics and Statistics* 1–15.
- (2025), “Forecasting with shadow rate VARs,” *Quantitative Economics* **16**(3), 795–822.
- CHAN JC (2023), “Comparing stochastic volatility specifications for large Bayesian VARs,” *Journal of Econometrics* **235**(2), 1419–1446.
- CHAN JC, KOOP G, AND YU X (2024), “Large order-invariant Bayesian VARs with stochastic volatility,” *Journal of Business & Economic Statistics* **42**(2), 825–837.
- CHAN JC, PETTENUZZO D, POON A, AND ZHU D (2025), “Conditional Forecasts in Large Bayesian VARs with Multiple Equality and Inequality Constraints,” *Journal of Economic Dynamics and Control* **105061**.
- CERNIS T, HAUZENBERGER N, MUMTAZ H, AND PFARRHOFER M (2025), “A Bayesian Gaussian Process Dynamic Factor Model,” *arXiv* **2509.04928**.
- CERNIS T, KOOP G, TALLMAN E, AND WEST M (2024), “Decision synthesis in monetary policy,” *arXiv* **2406.03321**.
- CHIPMAN HA, GEORGE EI, AND McCULLOCH RE (2010), “BART: Bayesian additive regression trees,” *The Annals of Applied Statistics* **4**(1), 266–298.
- CIMADOMO J, GIANNONE D, LENZA M, MONTI F, AND SOKOL A (2022), “Nowcasting with large Bayesian vector autoregressions,” *Journal of Econometrics* **231**(2), 500–519.
- CLARK T, HUBER F, AND KOOP G (2025), “A Nonparametric Approach to Augmenting a Bayesian VAR with Nonlinear Factors,” *arXiv* **2508.13972**.
- CLARK TE, HUBER F, KOOP G, MARCELLINO M, AND PFARRHOFER M (2023), “Tail forecasting with multivariate Bayesian additive regression trees,” *International Economic Review* **64**(3), 979–1022.

- (2024), “Investigating growth-at-risk using a multicountry nonparametric quantile factor model,” *Journal of Business & Economic Statistics* **42**(4), 1302–1317.
- CONG Y, CHEN B, AND ZHOU M (2017), “Fast simulation of hyperplane-truncated multivariate normal distributions,” *Bayesian Analysis* **12**(4), 1017–1037.
- CRUMP RK, EUSEPI S, GIANNONE D, QIAN E, AND SBORDONE AM (2025), “A large Bayesian VAR of the United States economy,” *International Journal of Central Banking* **forthcoming**.
- DOAN T, LITTERMAN R, AND SIMS C (1984), “Forecasting and conditional projection using realistic prior distributions,” *Econometric Reviews* **3**(1), 1–100.
- ESSER J, MAIA M, PARNELL AC, BOSMANS J, VAN DONGEN H, KLAUSCH T, AND MURPHY K (2024), “Seemingly unrelated Bayesian additive regression trees for cost-effectiveness analyses in healthcare,” *arXiv* **2404.02228**.
- FISCHER MM, HAUZENBERGER N, HUBER F, AND PFARRHOFER M (2023), “General Bayesian time-varying parameter vector autoregressions for modeling government bond yields,” *Journal of Applied Econometrics* **38**(1), 69–87.
- FORNI M, GAMBETTI L, MAFFEI-FACCIOLI N, AND SALA L (2024), “Nonlinear transmission of financial shocks: Some new evidence,” *Journal of Money, Credit and Banking* **56**(1), 5–33.
- GALLANT AR, ROSSI PE, AND TAUCHEN G (1993), “Nonlinear dynamic structures,” *Econometrica* **61**(4), 871–907.
- GEWEKE J, AND AMISANO G (2010), “Comparing and evaluating Bayesian predictive distributions of asset returns,” *International Journal of Forecasting* **26**(2), 216–230.
- GILCHRIST S, AND ZAKRAJŠEK E (2012), “Credit spreads and business cycle fluctuations,” *American Economic Review* **102**(4), 1692–1720.
- GODSILL SJ, DOUCET A, AND WEST M (2004), “Monte Carlo smoothing for nonlinear time series,” *Journal of the American Statistical Association* **99**(465), 156–168.
- GONÇALVES S, HERRERA AM, KILIAN L, AND PESAVENTO E (2021), “Impulse response analysis for structural dynamic models with nonlinear regressors,” *Journal of Econometrics* **225**(1), 107–130.
- (2024), “State-dependent local projections,” *Journal of Econometrics* **244**(2), 105702.
- GOULET COULOMBE P (2024), “To bag is to prune,” *Studies in Nonlinear Dynamics & Econometrics* **forthcoming**.
- GOULET COULOMBE P, LEROUX M, STEVANOVIC D, AND SURPRENANT S (2022), “How is machine learning useful for macroeconomic forecasting?” *Journal of Applied Econometrics* **37**(5), 920–964.

- GOURIEROUX C, AND LEE Q (2023), “Nonlinear impulse response functions and local projections,” *arXiv* **2305.18145**.
- HAMILTON JD, AND HERRERA AM (2004), “Comment: “Oil shocks and aggregate macroeconomic behavior: the role of monetary policy”,” *Journal of Money, Credit and Banking* 265–286.
- HAUZENBERGER N, HUBER F, KLIEBER K, AND MARCELLINO M (2025a), “Machine learning the macroeconomic effects of financial shocks,” *Economics Letters* **250**, 112260.
- HAUZENBERGER N, HUBER F, AND KOOP G (2024), “Macroeconomic forecasting using BVARs,” in MP CLEMENTS, AND AB GALVÃO (eds.) “Handbook of Research Methods and Applications in Macroeconomic Forecasting,” chapter 2, 15–42, Cheltenham, UK/Northampton, USA: Edward Elgar Publishing.
- HAUZENBERGER N, HUBER F, MARCELLINO M, AND PETZ N (2025b), “Gaussian process vector autoregressions and macroeconomic uncertainty,” *Journal of Business & Economic Statistics* **43**(1), 27–43.
- HUBER F, KLIEBER K, MARCELLINO MG, ONORANTE L, AND PFARRHOFER M (2024), “Asymmetries in International Financial Spillovers,” *SSRN* **5054831**.
- HUBER F, KOOP G, ONORANTE L, PFARRHOFER M, AND SCHREINER J (2023), “Nowcasting in a pandemic using non-parametric mixed frequency VARs,” *Journal of Econometrics* **232**(1), 52–69.
- HUBER F, AND ROSSINI L (2022), “Inference in Bayesian additive vector autoregressive tree models,” *The Annals of Applied Statistics* **16**(1), 104–123.
- JAROCIŃSKI M (2010), “Conditional forecasts and uncertainty about forecast revisions in vector autoregressions,” *Economics Letters* **108**(3), 257–259.
- JORDÀ Ò (2005), “Estimation and inference of impulse responses by local projections,” *American Economic Review* **95**(1), 161–182.
- JORDÀ Ò, AND TAYLOR AM (2025), “Local projections,” *Journal of Economic Literature* **63**(1), 59–110.
- KILIAN L, AND LÜTKEPOHL H (2017), *Structural vector autoregressive analysis*, Cambridge University Press.
- KOLEŠÁR M, AND PLAGBORG-MØLLER M (2025), “Dynamic Causal Effects in a Nonlinear World: the Good, the Bad, and the Ugly,” *Journal of Business & Economic Statistics* **forthcoming**.
- KOOP G, PESARAN MH, AND POTTER SM (1996), “Impulse response analysis in nonlinear multivariate models,” *Journal of Econometrics* **74**(1), 119–147.
- LEEPER EM, AND ZHA T (2003), “Modest policy interventions,” *Journal of Monetary*

- Economics* **50**(8), 1673–1700.
- LIMA PA, CARVALHO CM, LOPES HF, AND HERREN A (2025), “Minnesota BART,” *arXiv* **2503.13759**.
- LINDSTEN F, JORDAN MI, AND SCHÖN TB (2014), “Particle Gibbs with Ancestor Sampling,” *Journal of Machine Learning Research* **15**(63), 2145–2184.
- LOPEZ-SALIDO D, AND LORIA F (2024), “Inflation at risk,” *Journal of Monetary Economics* **145**, 103570.
- LUCAS RE (1976), “Econometric policy evaluation: A critique,” *Carnegie-Rochester Conference Series on Public Policy* **1**, 19–46.
- MARCELLINO M, AND PFARRHOFER M (2024), “Bayesian nonparametric methods for macroeconomic forecasting,” in MP CLEMENTS, AND AB GALVÃO (eds.) “Handbook of Research Methods and Applications in Macroeconomic Forecasting,” chapter 5, 90–125, Cheltenham, UK/Northampton, USA: Edward Elgar Publishing.
- MCKAY A, AND WOLF CK (2023), “What can time-series regressions tell us about policy counterfactuals?” *Econometrica* **91**(5), 1695–1725.
- MEDEIROS MC, VASCONCELOS GF, VEIGA Á, AND ZILBERMAN E (2021), “Forecasting inflation in a data-rich environment: the benefits of machine learning methods,” *Journal of Business & Economic Statistics* **39**(1), 98–119.
- MUMTAZ H, AND PIFFER M (2025), “Impulse response estimation via flexible local projections,” *Journal of Money, Credit and Banking* **forthcoming**.
- PFARRHOFER M (2022), “Modeling tail risks of inflation using unobserved component quantile regressions,” *Journal of Economic Dynamics and Control* **143**, 104493.
- POTTER SM (2000), “Nonlinear impulse response functions,” *Journal of Economic Dynamics and Control* **24**(10), 1425–1446.
- RAMBACHAN A, AND SHEPHARD N (2021), “When do common time series estimands have nonparametric causal meaning?” *Mimeo* .
- SIMS CA, AND ZHA T (2006), “Does monetary policy generate recessions?” *Macroeconomic Dynamics* **10**(2), 231–272.
- WAGGONER DF, AND ZHA T (1999), “Conditional forecasts in dynamic multivariate models,” *Review of Economics and Statistics* **81**(4), 639–651.
- WEST M (2024), “Perspectives on constrained forecasting,” *Bayesian Analysis* **19**(4), 1013–1039.
- WEST M, AND HARRISON J (1997), *Bayesian forecasting and dynamic models*, Springer.

# Online Appendix: Scenario Analysis with Multivariate Bayesian Machine Learning Models

## A. TECHNICAL DETAILS

### A.1. Particle Gibbs with Ancestor Sampling

We run PGAS in each iteration  $m$  of our MCMC sampler, so uncertainty around model parameters is accounted for (see also our discussion in the context of unconditional forecasts in Section 2.1); we omit the corresponding index unless it is needed for clarity. A complete sweep requires the following steps. For horizon  $h = 1$ , we have:

- (i) The initial conditions are known and given by  $\mathbf{x}_{\tau+1} = (\mathbf{y}'_{\tau}, \dots, \mathbf{y}'_{\tau-p+1})'$ , that is,  $\boldsymbol{\mu}_{\tau+1} = \mathbf{F}(\mathbf{x}_{\tau+1})$ . We generate  $V - 1$  candidate particles,  $\mathbf{y}_{\tau+1}^{(v)}$ , from the unconditional predictive distribution in case  $\mathcal{C}_1 = \emptyset$ , and using (9) in case restrictions are present. The  $V$ th particle serves as the fixed reference trajectory, and is set to the previous MCMC draw at iteration  $m - 1$ , i.e.,  $\mathbf{y}_{\tau+1}^{(V)} = \mathbf{y}_{\tau+1}^{(m-1)}$ . For later reference in the context of  $h > 1$ , we define  $\mathbf{x}_{\tau+2}^{(v)} = (\mathbf{y}_{\tau+1}^{(v)'}, \mathbf{y}'_{\tau}, \dots)'$ ; we discuss the required ancestry tracking for generic horizons below.

- (ii) The next step is to compute weights based on the measurement equation. When no restrictions are imposed we trivially obtain equal weights — all particles are equally compatible because there are no “measurements.” In the presence of restrictions the unnormalized weights are generally given by  $\tilde{w}_1^{(v)} \propto \mathcal{N}(\mathbf{r}_1; \mathbf{R}_1 \boldsymbol{\mu}_{\tau+1}, \mathbf{R}_1 \boldsymbol{\Sigma} \mathbf{R}'_1 + \boldsymbol{\Omega}_1)$ , but since the initial conditions are fixed and  $\boldsymbol{\mu}_{\tau+1}$  is the same for all particles, we obtain equal weights in either case.

At each horizon  $h = 1, 2, \dots$ , we generally compute the normalized weights as

$$w_h^{(v)} = \frac{\exp\left(\log \tilde{w}_h^{(v)} - \max_j \log \tilde{w}_h^{(j)}\right)}{\sum_{j=1}^V \exp\left(\log \tilde{w}_h^{(j)} - \max_{j'} \log \tilde{w}_h^{(j')}\right)}.$$

For all subsequent horizons  $h > 1$ , the following steps are required:

- (i) We first resample ancestor indices,  $a_h^{(v)} \in \{1, \dots, V\}$ , for non-reference particles (for  $v = 1, \dots, V - 1$ ) from a categorical distribution,  $a_h^{(v)} \sim \text{Cat}(w_{h-1}^{(1)}, \dots, w_{h-1}^{(V)})$ . This step serves to retain particles with a high probability of having generated the measurements (in our case, restrictions) from the previous horizon.

For the reference particle  $v = V$ , the respective index is obtained with ancestor sampling. Specifically, we compute  $\pi_h^{(v)} = \log \tilde{w}_{h-1}^{(v)} + \log \mathcal{N}(\mathbf{y}_{\tau+h}^{(V)}; \mathbf{F}(\mathbf{x}_{\tau+h}^{(v)}), \Sigma)$ . We apply a weight-normalization analogous to the one above, to obtain  $\bar{w}_h^{(v)}$  using  $\pi_h^{(v)}$ . The ancestor index for the reference trajectory is sampled from  $a_h^{(V)} \sim \text{Cat}(\bar{w}_h^{(1)}, \dots, \bar{w}_h^{(V)})$ . This step reweights the candidate parents of the reference trajectory with their probability of generating the reference.

- (ii) In the next step, we propagate the non-reference particles based on their parent index to keep the lineage consistent also in view of the lag structures. We obtain  $\boldsymbol{\mu}_{\tau+h}^{(v)} = \mathbf{F}(\mathbf{x}_{\tau+h}^{(a_h^{(v)})})$ . If  $\mathcal{C}_h = \emptyset$ , we produce samples from the unconditional one-step ahead distributions; in case there are restrictions, we use (9) to draw particles that are consistent with these restrictions. For the reference particle, we again set  $\mathbf{y}_{\tau+h}^{(V)} = \mathbf{y}_{\tau+h}^{(m-1)}$ . The stacked vector is constructed by appending the new draw to the lagged components of its parent and dropping the previous last lag:  $\mathbf{x}_{\tau+h+1}^{(v)} = (\mathbf{y}_{\tau+h}^{(v)'}; \mathbf{y}_{\tau+h-1}^{(a_h^{(v)})'}; \dots; \mathbf{y}_{\tau+h-p+1}^{(a_h^{(v)})}')'$ .

(iii) The final step updates the weights in light of the restrictions given by the measurement equation. In case there are none, we again obtain equal weights. Otherwise, we have

$$\tilde{w}_h^{(v)} \propto \mathcal{N}(\mathbf{r}_h; \mathbf{R}_h \boldsymbol{\mu}_{\tau+h}^{(v)}, \mathbf{R}_h \boldsymbol{\Sigma} \mathbf{R}_h' + \boldsymbol{\Omega}_h).$$

When the maximum forecast horizon  $H$  is reached (i.e.,  $h = 1, 2, \dots, H$ ), it remains to generate a fully smoothed trajectory by tracing back through the stored ancestor indices. We draw one final particle index using  $v_H \sim \text{Cat}(w_H^{(1)}, \dots, w_H^{(V)})$  and then recursively for each  $h = H - 1, H - 2, \dots, 1$  trace its lineage, setting  $v_h = a_{h+1}^{(v_{h+1})}$ . A draw from the distribution  $p(\mathbf{y}_{\tau+1:\tau+H} \mid \mathcal{I}, \mathcal{C}_{1:H}, \boldsymbol{\Xi})$  is then given by the vector  $\mathbf{y}_{\tau+1:\tau+H}^{(m)} = (\mathbf{y}_{\tau+1}^{(v_1)'}, \dots, \mathbf{y}_{\tau+H}^{(v_H)'})'$ , which serves as the reference trajectory in the next sweep of the MCMC algorithm.

To compute GIRFs, we typically require the expectation of the respective scenario and baseline forecast distributions rather than a draw from them. The weights  $w_h^{(v)}$  relate to the forward filtering distributions, and do not encode information over the full forecast horizon, so they cannot be used without adjustment. To propagate information backwards, we define smoothing weights,  $s_h^{(v)}$ , which are initialized as  $s_H^{(v)} = w_H^{(v)}$ . For  $h = H - 1, \dots, 1$ , we compute  $\tilde{s}_h^{(v)} = \sum_{j=1}^V s_{h+1}^{(j)} \mathbb{I}(a_{h+1}^{(j)} = v)$  and normalize  $s_h^{(v)} = \tilde{s}_h^{(v)} / (\sum_{j=1}^V \tilde{s}_h^{(j)})$ . We use these smoothed weights to compute the required expectations at each horizon:

$$\mathbb{E}(\mathbf{y}_{t+h} \mid \mathcal{C}_{1:H}, \mathbf{x}_{\tau+1}, \mathcal{I}, \boldsymbol{\Xi}) \approx \sum_{v=1}^V s_h^{(v)} \mathbf{y}_{\tau+h}^{(v)}.$$

## A.2. Posterior Distributions and Sampling Algorithm

We may use the conditional distribution in (12) to update the trees equation-by-equation using the backfitting approach designed by Chipman *et al.* (2010); see also Esser *et al.* (2024). Here, one may define the vector of partial residuals

$$\tilde{y}_{is,t} = \left( y_{it} - \tilde{\mu}_{it} - \sum_{j \neq s} \ell_{ij}(\mathbf{x}_t \mid \mathcal{T}_{ij}, \mathbf{m}_{ij}) \right) \sim \mathcal{N}(\ell_{is}(\mathbf{x}_t \mid \mathcal{T}_{is}, \mathbf{m}_{is}), \varsigma_{it}^2),$$

conditioning on the fit of each of the  $S - 1$  trees except tree  $s$  and information in all but the  $i$ th equation. In full data notation,  $\tilde{\mathbf{y}}_{is} = (\tilde{y}_{is,1}, \dots, \tilde{y}_{is,T})'$ , this defines a conditionally Gaussian likelihood,  $p(\tilde{\mathbf{y}}_{is} \mid \mathcal{T}_{is}, \mathbf{m}_{is}, \bullet)$ , which can be marginalized analytically over the terminal node parameters  $\mathbf{m}_{is}$  (to keep the dimensionality of the inferential problem fixed).

Combining this conditional likelihood with the prior on the trees, and a suitable transition density (based on four distinct moves: grow a terminal node, prune a terminal node, change a splitting rule, swap a child/parent node), the trees are sampled using a standard accept/reject Metropolis-Hastings algorithm. These trees (and associated rules) partition the input space and we obtain a distinct set of observations for each terminal node. The posterior then takes the conventional Gaussian form for these parameters.

Updating all trees  $s = 1, \dots, S$ , across equations  $i = 1, \dots, n$ , yields an updated fit that can be used to compute the outlier-adjusted residuals  $\boldsymbol{\epsilon}_t / \sqrt{s_t} = \mathbf{y}_t - \mathbf{F}(\mathbf{x}_t)$ . The posterior of the constant part of the covariance matrix is then given by:

$$\boldsymbol{\Sigma} \mid \bullet \sim \mathcal{W}^{-1} \left( s_0 + T, \mathbf{S}_0 + \sum_{t=1}^T s_t^{-1} \boldsymbol{\epsilon}_t \boldsymbol{\epsilon}_t' \right)$$

The hierarchical parameters of the prior on the covariance matrix can be updated using:

$$a_i \mid \bullet \sim \mathcal{G}^{-1} \left( \frac{\nu + T}{2}, \frac{1}{A_i^2} + \nu \cdot \boldsymbol{\Sigma}_{[ii]}^{-1} \right),$$

where  $\boldsymbol{\Sigma}_{[ii]}^{-1}$  denotes the  $i$ th diagonal element of  $\boldsymbol{\Sigma}^{-1}$ . The outlier adjustment parameter  $s_t$  can be sampled, due to its discrete support, using the probabilities:

$$\Pr(s_t^{1/2} = 1 \mid \mathbf{y}_t, \mathbf{F}(\mathbf{x}_t), \boldsymbol{\Sigma}, \mathbf{p}) \propto \mathcal{N}(\mathbf{y}_t \mid \mathbf{F}(\mathbf{x}_t), \boldsymbol{\Sigma}) \cdot (1 - \mathbf{p}),$$

$$\Pr(s_t^{1/2} = \mathfrak{s} \mid \mathbf{y}_t, \mathbf{F}(\mathbf{x}_t), \boldsymbol{\Sigma}, \mathbf{p}) \propto \mathcal{N}(\mathbf{y}_t \mid \mathbf{F}(\mathbf{x}_t), \mathfrak{s}\boldsymbol{\Sigma}) \cdot (\mathbf{p}/(\bar{\mathfrak{s}} - 1)), \quad \text{for } \mathfrak{s} = 2, 3, \dots, \bar{\mathfrak{s}},$$

on a  $t$ -by- $t$  basis. The posterior distribution of the outlier probability is  $\mathbf{p} \mid \bullet \sim \mathcal{B}(a_p + T_o, b_p + T - T_o)$ , with the total number of observations classified as outliers denoted by  $T_o = \sum_{t=1}^T \mathbb{I}(s_t \neq 1)$ , where  $\mathbb{I}(\bullet)$  is an indicator function that yields 1 if its argument is true and 0 otherwise.

Note that all model parameters can be updated conditional on any restrictions via data augmentation (by appending the respective draws to the observed data in each MCMC sweep). This involves iteratively alternating between sampling from the (restricted) forecast distribution conditional on the parameters, and vice versa. See also Waggoner and Zha (1999), Bańbura *et al.* (2015) and Antolin-Diaz *et al.* (2021).

### A.3. Recursive substitution and structural (G)IRFs

Note that our model can be written in terms of its one-step ahead predictive distribution, taking the initial conditions  $\mathbf{x}_{\tau+1}$  as given, as  $\mathbf{y}_{\tau+1} = \mathbf{F}(\mathbf{x}_{\tau+1}) + \mathbf{H}^{-1}\mathbf{u}_{\tau+1}$  with  $\mathbf{u}_{\tau+1} \sim \mathcal{N}(\mathbf{0}_n, \mathbf{I}_n)$ . In case one desires to integrate over the marginal distributions of the other shocks (when using PGAS, one may simply specify the distribution of the structural shocks explicitly), one may randomly generate two distinct sets of structural shocks (see also Kilian and Lütkepohl, 2017):

$$\begin{aligned}\mathbf{u}_{j\tau}^{(s,m)} &= (u_{1\tau}^{(s,m)}, \dots, u_{j-1,\tau}^{(s,m)}, d_0 + d, u_{j+1,\tau}^{(s,m)}, \dots, u_{n\tau}^{(s,m)}) \\ \mathbf{u}_{j\tau}^{(b,m)} &= (u_{1\tau}^{(b,m)}, \dots, u_{j-1,\tau}^{(b,m)}, d_0, u_{j+1,\tau}^{(b,m)}, \dots, u_{n\tau}^{(b,m)}).\end{aligned}$$

When nonlinear impact effects are possible, one may also average over different baseline levels of the shock. These can be used to obtain the GIRF on impact:

$$\delta_{j\tau,1}^{(d,m)} = \mathbf{F}(\mathbf{x}_{\tau+1}) + \mathbf{H}^{-1}\mathbf{u}_{j\tau}^{(s,m)} - \mathbf{F}(\mathbf{x}_{\tau+1}) - \mathbf{H}^{-1}\mathbf{u}_{j\tau}^{(b,m)} = \mathbf{H}^{-1}(\mathbf{u}_{j\tau}^{(s,m)} - \mathbf{u}_{j\tau}^{(b,m)}).$$

When using the same sequence of random numbers for simulating  $\mathbf{s}$  and  $\mathbf{b}$  (see also the discussion in Appendix B), or when taking expectations in line with our definition of the structural GIRF in (11), we have a closed form expression for the impact response:

$$\boldsymbol{\delta}_{j\tau,1}^{(d,m)} = \mathbf{H}^{-1} \left( (d_0 + d) \cdot \mathbf{e}'_j - d_0 \cdot \mathbf{e}'_j \right) = d \cdot \mathbf{H}^{-1} \mathbf{e}'_j = d \cdot \boldsymbol{\beta}_0^{(m)},$$

for iteration  $m$  of our sampler. That is, we obtain a scaled version of the  $j$ th column of the impact matrix  $\mathbf{H}^{-1}$  as in standard VAR models. We thus assume with our additive error specification that the impact of the respective shock is constant over time, and shocks of different signs and sizes are (proportionally) introduced by setting  $d$  accordingly. This relates to scaled versions of the average causal effect for infinitesimally small shocks on impact,  $\lim_{d \rightarrow 0} \boldsymbol{\delta}_{j\tau,1}^{(d)}/d = \boldsymbol{\beta}_0^{(m)}$ , see also Kolesár and Plagborg-Møller (2025).

Next, using  $\boldsymbol{\beta}_0^{(\mathbf{s},m)} = \mathbf{H}^{-1} \mathbf{u}_{j\tau}^{(\mathbf{s},m)}$  and  $\boldsymbol{\beta}_0^{(\mathbf{b},m)} = \mathbf{H}^{-1} \mathbf{u}_{j\tau}^{(\mathbf{b},m)}$  to refer to the impact effects for the scenario and baseline forecasts, we can set  $\mathbf{x}_{\tau+2}^{(\mathbf{s},m)} = ((\mathbf{y}_{\tau+1} + \boldsymbol{\beta}_0^{(\mathbf{s},m)})', \mathbf{y}'_{\tau}, \dots, \mathbf{y}'_{\tau-p+2})'$  and  $\mathbf{x}_{\tau+2}^{(\mathbf{b},m)} = ((\mathbf{y}_{\tau+1} + \boldsymbol{\beta}_0^{(\mathbf{b},m)})', \mathbf{y}'_{\tau}, \dots, \mathbf{y}'_{\tau-p+2})'$ . As in Section 2, using (5), we can iterate forward and obtain  $\boldsymbol{\delta}_{\tau,2}^{(d,m)} = \mathbf{F}(\mathbf{x}_{\tau+2}^{(\mathbf{s},m)}) - \mathbf{F}(\mathbf{x}_{\tau+2}^{(\mathbf{b},m)})$ .<sup>1</sup> For higher-order responses, iterate forward, draw from the conditional predictive distributions, and reconfigure  $\mathbf{x}_{\tau+h}^{(\mathbf{s},m)}$  and  $\mathbf{x}_{\tau+h}^{(\mathbf{b},m)}$  accordingly at each horizon. The structural GIRF is given by  $\boldsymbol{\delta}_{\tau,h}^{(d,m)} = \mathbf{F}(\mathbf{x}_{\tau+h}^{(\mathbf{s},m)}) - \mathbf{F}(\mathbf{x}_{\tau+h}^{(\mathbf{b},m)})$ . This addresses future reduced form shocks to obtain the respective conditional expectation (implicitly via the recursive simulation of the scenario and baseline paths of the variables). Different to PGAS, this only requires a single realization instead of multiple particles.

**Comparison with linear VAR.** Under the assumption of Gaussian errors and when  $\mathbf{F}(\mathbf{y}_t) = \mathbf{A}\mathbf{y}_t$  is a linear function, where  $\mathbf{A}$  is an  $n \times n$  matrix of coefficients,  $\mathbf{y}_{t+h} = \mathbf{A}^h \mathbf{y}_t + \sum_{j=1}^h \mathbf{A}^{h-j} \boldsymbol{\epsilon}_{t+j}$ . Conditional on information up to time  $t$ , we obtain closed form expressions for, e.g.,  $\mathbb{E}(\mathbf{y}_{t+h}) = \mathbf{A}^h \mathbf{y}_t$  and  $\text{Var}(\mathbf{y}_{t+h}) = \sum_{j=1}^h \mathbf{A}^{h-j} \boldsymbol{\Sigma} \mathbf{A}^{h-j'}$ . In the nonlinear case, one

<sup>1</sup> Note that we may also compute other functions instead of the expectation, e.g., probabilities or quantiles, see also Gallant *et al.* (1993), and more recently, Jordà and Taylor (2025).

cannot apply expectations as straightforwardly and compute closed form multi-step ahead expressions, e.g., for the first moment, as  $\mathbb{E}(\tilde{\mathbf{F}}_{(h)}(\mathbf{y}_t)) \neq \mathbb{E}(\tilde{\mathbf{F}}_{(h)}(\mathbf{y}_t, \boldsymbol{\epsilon}_{t+1}, \boldsymbol{\epsilon}_{t+2}, \dots, \boldsymbol{\epsilon}_{t+h}))$ . The former expectation corresponds to a path which sets the innovations after time  $t$  to zero which is typically undesirable, see, e.g., [Potter \(2000\)](#) for a discussion.


Our approach yields standard IRFs when  $\mathbf{F}(\bullet)$  is linear — and when all future structural shocks are set equal to 0 (i.e., the model is iterated in expectations) or when the same randomly generated numbers are used for simulating the scenario and baseline distributions; otherwise they feature uncertainty but are equal in expectation. For instance, when  $\mathbf{F}$  is a linear mapping, we obtain the IRF at horizon  $h$  as  $d\mathbf{A}^h\boldsymbol{\beta}_0$ . That is, it can be obtained by simply projecting the shock impact forward using powers of  $\mathbf{A}$ . Since  $d$  factors out and the initial conditions cancel, the IRFs are time-invariant and proportional for different shock sizes and signs. In the more general nonlinear context, this is clearly not the case.

## B. DATA AND RESULTS

### B.1. Simulation-based results

**Conditioning on observables.** To assess the performance of our PGAS approach relative to closed form solutions, we provide a comparison with the precision sampler (PS) of [Chan \*et al.\* \(2025\)](#). Assuming a linear conditional mean function allows to contrast both algorithms one-to-one (we cannot use the PS for nonlinear implementations, for the reasons described in the paper). We use a homoskedastic linear VAR model with  $n = 5$  and  $p = 5$  using a conjugate Minnesota prior with hierarchically estimated hyperparameters, as in [Giannone \*et al.\* \(2015\)](#), to calibrate the parameters used for simulating artificial data. The data used for calibration features output, inflation, employment (GDPC1, CPIAUCSL, PAYEMS; all in annualized log-differences), interest rate (FEDFUNDS) and stock returns (SP500; log differences), reflecting key variables from our applications with real data.

We use the posterior median estimates of the dynamic coefficients and the covariance matrix, and simulate  $T = 1,000$  observations (to limit any effects of the initial conditions). We take the calibrated parameters as given and use the final  $p$  of the observations as initial conditions to compute 3,000 draws from the respective conditional forecast distributions. We impose three distinct restrictions at different horizons, to illustrate how our approach is capable of enforcing the restrictions over the full desired period: (i) the interest rate increases by one unconditional SD at  $h = 1$ , (ii) inflation is restricted to its unconditional mean for  $h = 9, \dots, 12$ , and (iii) output increases by one unconditional SD relative to its initial level at  $h = 20$ .



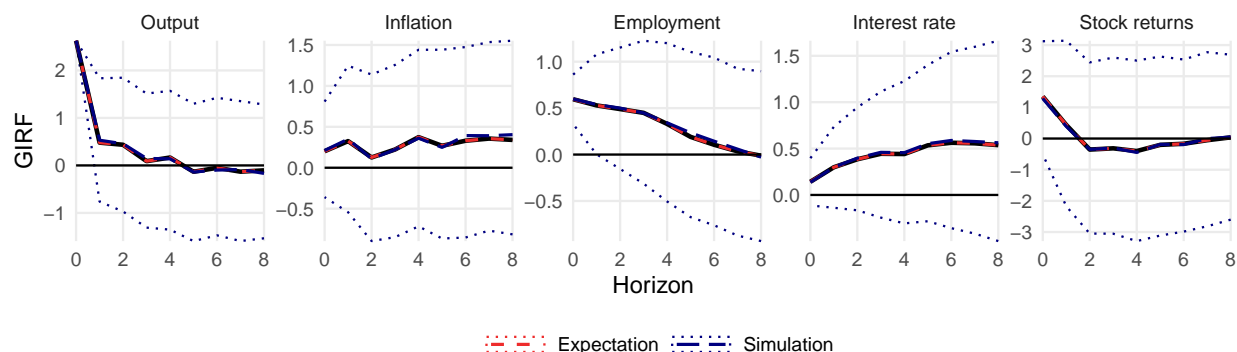
**Figure B.1:** Conditional forecast (CF) with precision sampling (PS) and particle Gibbs with ancestor sampling (PGAS).

*Notes:* We obtain 3,000 draws from the distribution of the conditional forecast. The black lines and grey shaded ribbon indicate the median and 68 percent credible set of the predictive distribution obtained from PS. The colored lines indicate draws from PGAS for different numbers of particles  $V \in \{5, 10, 25, 50\}$ .


The resulting CFs are shown in Figure B.1. The black lines and grey shaded ribbon indicate the median and 68 percent credible set of the predictive distribution obtained from PS. The colored lines indicate draws from PGAS for different numbers of particles for the same quantiles. Differences between implementations are muted, and as little as 5 particles are sufficient. On a 2020 Macbook Air M1, it takes about 5.5 seconds with our PS implementation to obtain 1,000 draws; PGAS runtimes in seconds are 9 (1.6 times as long as PS), 16.6 (3), 41 (7.5), 79.5 (14.5) for particles  $V \in \{5, 10, 25, 50\}$ .

**Conditioning on shocks: GIRF computation.** We use the same DGP as above, and treat the impact matrix  $H$  as well as the dynamic coefficients as given. This allows to compute the true impulse response function, which is shown in Figure B.2 as thick black line. We focus on the response of all variables to the first structural shock  $j = 1$ . Panel (a) shows the distribution of predictive draws when using different random numbers for the scenario and baseline forecast, while panel (b) uses an identical sequence of random numbers for both forecasts in each iteration. In the latter case, there is no MCMC uncertainty (see also Appendix A), so our discussion below focuses on panel (a).

(a) Distinct draws of random numbers (per iteration) for scenario and baseline forecast



(b) Identical draws of random numbers (per iteration) for scenario and baseline forecast



**Figure B.2:** Generalized impulse response function (IRF) computed with simulation-based methods for a single DGP.

*Notes:* When applicable we obtain 3,000 draws from the distribution of the GIRF. The thick black line is the true IRF, colored lines refer to different sampling-based estimation procedures; the thin dotted lines refer to the 68 percent credible set. When using PGAS we rely on  $V = 10$  particles.

We use PGAS with  $V = 10$  different particles, subject to distinct choices about the contemporaneous and future shock restrictions. For the variant labeled “Expectation,” we set the shock of interest equal to 1 and all others zero, while we impose these restrictions as approximately binding (with a variance of  $10^{-8}$ ). That is, we iterate in expectations, and obtain exactly the true IRF. For “Simulation,” we use  $\mathbf{r}_1^{(u)} = (1, 0, \dots, 0)'$  in  $\mathcal{C}_1^{(s)}$  and  $\mathbf{r}_1^{(u)} = \mathbf{0}_n$  in  $\mathcal{C}_1^{(b)}$ , with  $\mathbf{\Omega}_1 = \text{diag}(10^{-8}, 1, \dots, 1)$  in both cases.

## B.2. Dataset

Our dataset is roughly patterned after [Crump \*et al.\* \(2025\)](#). For reference, see also the data used in the [Dodd-Frank Act stress test](#). Variable codes, a brief description and the respective transformations are listed in [Table B.1](#).

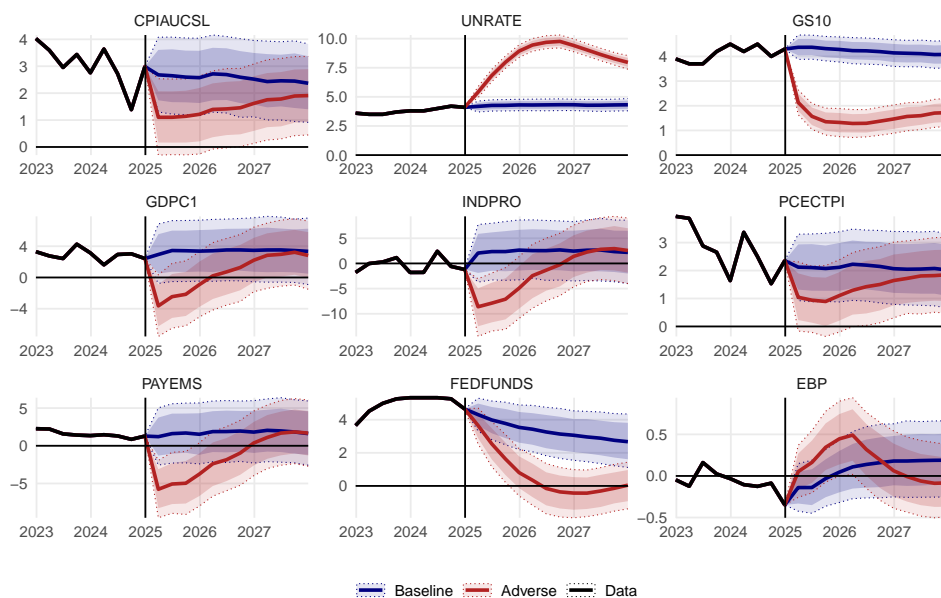
**Table B.1:** Variable codes, descriptions and transformation of our dataset.

Code	Description	$h(x_t)$	CF	GIRF
GDPC1	Real gross domestic product	1	✓	✓
PCECC96	Real personal consumption expenditure	1	✓	
PRFIx	Real private fixed investment (residential)	1	✓	
GCEC1	Real Government consumption and investment	1	✓	
RDI	Real disposable income	1	✓	
INDPRO	Industrial production	1	✓	
CPIAUCSL	Headline CPI	1	✓	✓
CPILFESL	Core CPI	1	✓	
PCECTPI	Headline PCE prices	1	✓	
PCEPILFE	Core PCE prices	1	✓	
HPI	House price index	1	✓	
HOUST	Housing starts	1	✓	
MR	Mortgage rate	0	✓	
PAYEMS	Payroll employment	1	✓	✓
UNRATE	Unemployment rate	0	✓	
FEDFUNDS	Federal funds rate	0	✓	✓
GS1	1-year treasury rate	0	✓	
GS10	10-year treasury rate	0	✓	✓
EBP	Excess bond premium	0	✓	✓
NFCI	National financial conditions index	0	✓	
OILPRICE <sub>x</sub>	Real crude oil prices (WTI)	2	✓	
SP500	S&P 500	2	✓	✓
EXUSUK <sub>x</sub>	US/UK foreign exchange rate	3	✓	✓
USDEUR	EU/US foreign exchange rate	3		✓
EARGDP	Real gross domestic product (EA)	1		✓
UKRGDP	Real gross domestic product (UK)	1		✓

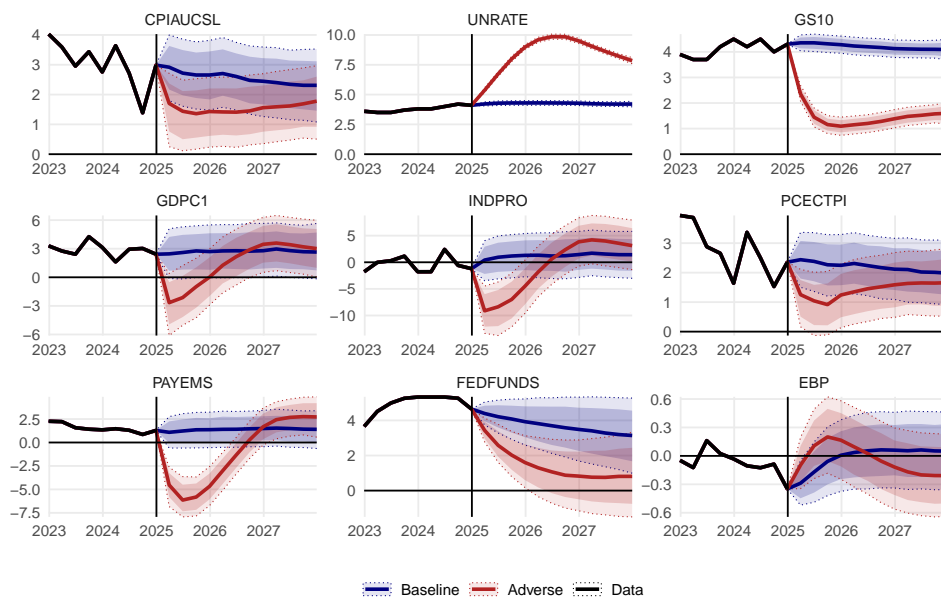
Notes: (0) no transformation  $h(x_t) = x_t$ ; (1) annualized log-differences  $h(x_t) = 400 \cdot \log(x_t/x_{t-1})$ ; (2) log-differences  $h(x_t) = 100 \cdot \log(x_t/x_{t-1})$ , (3) logarithm  $h(x_t) = \log(x_t)$ . Check marks indicate inclusion in the information sets for our applications (CF for Sections 4.1 and 4.2, GIRF for Section 4.3).

### B.3. Empirical results

(a) BART-hom

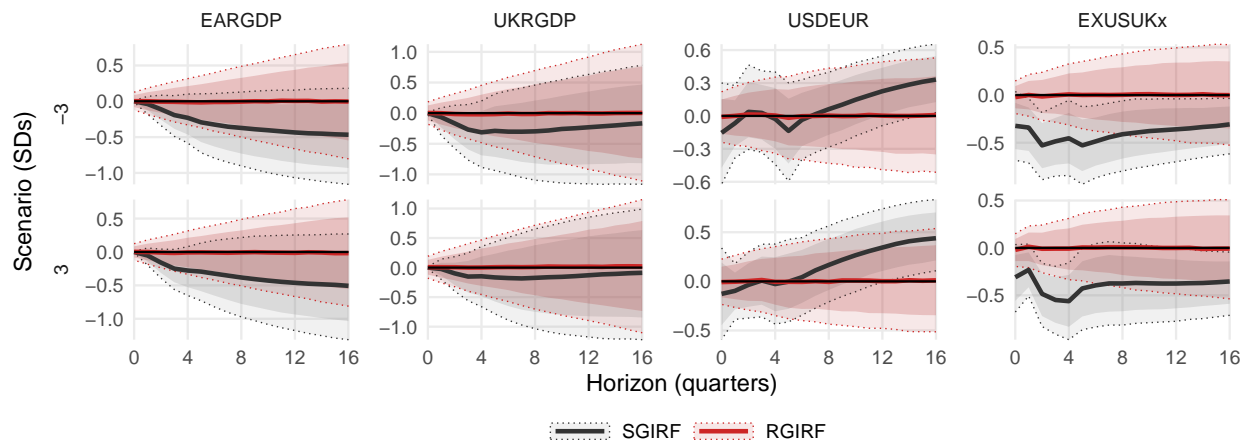


(b) BVAR-het



**Figure B.3:** Conditional forecasts for selected variables.

*Notes:* Posterior median alongside 50/68 percent credible sets. Restricted variables: consumer price inflation (CPIAUCSL), unemployment rate (UNRATE), 10-year government bond yields (GS10); unrestricted variables: Real GDP (GDPC1), industrial production (INDPRO), personal consumption expenditure inflation (PCECTPI), payroll employment (PAYEMS), federal funds rate (FEDFUNDS) and excess bond premium (EBP).



**Figure B.4:** Structural and restricted GIRFs for non-domestic variables.

*Notes:* The restricted case assumes that the US financial shock does not spill over to EA and UK real GDP (EARGDP, UKRGDP) and bilateral exchange rates (euro, USDEUR; British pound, EXUSUKx). Posterior medians alongside 50/68 percent credible sets. Cumulated responses for variables in differences and levels for all other variables.

## REFERENCES

- ANTOLIN-DIAZ J, PETRELLA I, AND RUBIO-RAMÍREZ JF (2021), “Structural scenario analysis with SVARs,” *Journal of Monetary Economics* **117**, 798–815.
- BAÑBURA M, GIANNONE D, AND LENZA M (2015), “Conditional forecasts and scenario analysis with vector autoregressions for large cross-sections,” *International Journal of Forecasting* **31**(3), 739–756.
- CHAN JC, PETTENUZZO D, POON A, AND ZHU D (2025), “Conditional Forecasts in Large Bayesian VARs with Multiple Equality and Inequality Constraints,” *Journal of Economic Dynamics and Control* **105061**.
- CHIPMAN HA, GEORGE EI, AND MCCULLOCH RE (2010), “BART: Bayesian additive regression trees,” *The Annals of Applied Statistics* **4**(1), 266–298.
- CRUMP RK, EUSEPI S, GIANNONE D, QIAN E, AND SBORDONE AM (2025), “A large Bayesian VAR of the United States economy,” *International Journal of Central Banking* **forthcoming**.
- ESSER J, MAIA M, PARNELL AC, BOSMANS J, VAN DONGEN H, KLAUSCH T, AND MURPHY K (2024), “Seemingly unrelated Bayesian additive regression trees for cost-effectiveness analyses in healthcare,” *arXiv* **2404.02228**.
- GALLANT AR, ROSSI PE, AND TAUCHEN G (1993), “Nonlinear dynamic structures,” *Econometrica* **61**(4), 871–907.
- GIANNONE D, LENZA M, AND PRIMICERI GE (2015), “Prior selection for vector autore-

- gressions,” *Review of Economics and Statistics* **97**(2), 436–451.
- JORDÀ Ò, AND TAYLOR AM (2025), “Local projections,” *Journal of Economic Literature* **63**(1), 59–110.
- KILIAN L, AND LÜTKEPOHL H (2017), *Structural vector autoregressive analysis*, Cambridge University Press.
- KOLESÁR M, AND PLAGBORG-MØLLER M (2025), “Dynamic Causal Effects in a Nonlinear World: the Good, the Bad, and the Ugly,” *Journal of Business & Economic Statistics* **forthcoming**.
- POTTER SM (2000), “Nonlinear impulse response functions,” *Journal of Economic Dynamics and Control* **24**(10), 1425–1446.
- WAGGONER DF, AND ZHA T (1999), “Conditional forecasts in dynamic multivariate models,” *Review of Economics and Statistics* **81**(4), 639–651.