

# INFERENCE IN BAYESIAN ADDITIVE VECTOR AUTOREGRESSIVE TREE MODELS

BY FLORIAN HUBER<sup>1,a</sup> AND LUCA ROSSINI<sup>2,b</sup>

<sup>1</sup>*Department of Economics, University of Salzburg, [florian.huber@sbg.ac.at](mailto:florian.huber@sbg.ac.at)*

<sup>2</sup>*Department of Economics, Management and Quantitative Methods, University of Milan, [luca.rossini@unimi.it](mailto:luca.rossini@unimi.it)*

Vector autoregressive (VAR) models assume linearity between the endogenous variables and their lags. This assumption might be overly restrictive and could have a deleterious impact on forecasting accuracy. As a solution we propose combining VAR with Bayesian additive regression tree (BART) models. The resulting Bayesian additive vector autoregressive tree (BAVART) model is capable of capturing arbitrary nonlinear relations between the endogenous variables and the covariates without much input from the researcher. Since controlling for heteroscedasticity is key for producing precise density forecasts, our model allows for stochastic volatility in the errors. We apply our model to two datasets. The first application shows that the BAVART model yields highly competitive forecasts of the U.S. term structure of interest rates. In a second application we estimate our model using a moderately sized Eurozone dataset to investigate the dynamic effects of uncertainty on the economy.

**1. Introduction.** In macroeconomics and finance, most models commonly employed to study the transmission of economic shocks or produce predictions assume linearity in their parameters and are fully parametric. One prominent example is the vector autoregressive (VAR) model that is extensively used in central banks and academia (see Doan, Litterman and Sims (1984), Litterman (1986), Sims (1980), Sims and Zha (1998)). In normal times, and with macroeconomic relations remaining stable, this linearity assumption might describe the data well. In turbulent times, however, key transmission channels often change, and more flexibility could be necessary. Ignoring such changes or failing to effectively control for outliers could translate into weak out-of-sample forecasts and potentially has adverse effects on the estimation of impulse responses.

The linearity assumption has been subject to substantial criticism in the literature. For instance, Granger and Terasvirta (1993) show that macroeconomic and financial quantities depend nonlinearly on each other and thus assuming linearity might be overly restrictive. As a solution, researchers increasingly rely on nonlinear models which feature time-varying parameters. These models either allow for gradually evolving coefficients or feature a rather low number of structural breaks. All these models have in common that, within each point in time, the relationship between the endogenous and explanatory variables is linear and deciding on the specific law of motion is an important modeling decision.

Another strand of the literature proposes Bayesian nonparametric time series models in order to relax the linearity assumption. In particular, several recent papers (see, e.g., Bassetti, Casarin and Leisen (2014), Billio, Casarin and Rossini (2019), Kalli and Griffin (2018)) propose novel methods that assume the transition densities to be nonparametric. These techniques are characterized by featuring an infinite dimensional parameter space that is flexibly adjusted to the complexity of the data at hand.<sup>1</sup> Within the nonparametric paradigm there

---

Received August 2020; revised April 2021.

*Key words and phrases.* Bayesian additive regression trees, BAVART, decision trees, nonparametric regression, vector autoregressions.

<sup>1</sup>For a review on Bayesian nonparametric methods, see Hjort et al. (2010).

has been a number of popular competing approaches that make use of machine learning techniques, such as boosting (Freund and Schapire (1997), Friedman (2001)), bagging and random forests (Breiman (2001)), decision trees and Mondrian forests (Lakshminarayanan, Roy and Teh (2014), Roy and Teh (2009)) and Bayesian additive regression tree (BART) models (Chipman, George and McCulloch (2002), Chipman, George and McCulloch (2010), Gramacy and Lee (2008), Linero (2018)).

In this paper our focus is on BART models and applying them to multivariate time series data. BART has been successfully applied for dealing with model uncertainty (Hernández et al. (2018)), sparse regression problems (Linero (2018)) and spatial models as well as for limited dependent variable models (Krueger, Bansal and Buddhavarapu (2020)). Moreover, in several recent studies BART has been shown to yield precise forecasts and improve upon several competing models from the statistics and machine learning literature (see He, Yalov and Hahn (2019), Kapelner and Bleich (2015), Linero (2018), Prüser (2019), Waldmann (2016), Huber et al. (2021)). BART has also been successfully used for carrying out causal effect estimation (see, e.g., Dorie et al. (2019), Green and Kern (2012), Hill (2011), Kern et al. (2016), Hahn, Murray and Carvalho (2020)).

The strong empirical performance of BART for forecasting and causal inference gives rise to the main contribution of the paper. We aim to bridge the literature on BART (see Chipman, George and McCulloch (2010)) with the literature on VAR models. BART, being a flexible nonparametric regression approach, allows for unveiling nonlinear relations between a set of endogenous and explanatory variables without needing much input from the researcher. Intuitively speaking, it models the conditional mean of the regression model by summing over a large number of trees which are, by themselves, constrained through a regularization prior. The resulting individual trees will take a particularly simple form and can thus be classified as “weak learners”. Each of these simple trees only explains a small fraction of the variation in the response variable while a large sum of them is capable of describing the data extremely well.

It is precisely this intuition upon which we build when we generalize these techniques to a multivariate setting. More precisely, we assume that a potentially large dimensional vector of endogenous variables is determined by its lagged values. As opposed to VAR models which model this relationship using a linear function, we assume the precise functional form to be unknown. This function is then estimated using BART. The resulting model, labeled the Bayesian additive vector autoregressive tree (BAVART) model, is a highly flexible variant that can be used for forecasting and impulse response analysis.

Estimation and inference are carried out using Markov chain Monte Carlo (MCMC) techniques. Since the error covariance matrix is nondiagonal, we propose methods that allow for equation-by-equation estimation of the multivariate model. These techniques imply that model estimation scales well in high dimensions and permits estimation of huge dimensional models. Another novel feature of our approach is that it allows for flexibly handling heteroscedasticity. We control for time-variation in the error variances by proposing a stochastic volatility specification. This feature is crucial for producing precise density forecasts. To produce higher-order forecasts and impulse response functions (IRFs), we develop techniques that enable us to sample from the predictive distribution and the posterior distribution of the IRFs. This proves to be another important contribution of the paper which is closely related to the literature on estimating nonlinear impulse response functions (see, e.g., Barnichon and Matthes (2018), Plagborg-Møller (2019)).

We illustrate our BAVART model using two empirical applications. The first deals with forecasting the U.S. term structure of interest rates. Since financial data often exhibit nonlinear features, our BAVART model might be well suited for dealing with interest rates at differing maturities. We consider two ways of modeling the yield curve. The first models a

panel of yields simultaneously along the lines of [Carriero, Kapetanios and Marcellino \(2012\)](#), whereas the second approach fits a Nelson–Siegel (NS) model in the spirit of [Diebold and Li \(2006\)](#) but assumes the three NS factors to follow a BAVART model. Within each of these classes, we consider several linear and nonlinear competing models. Since interest not only centers on one-step-ahead predictive distributions but also on multistep-ahead forecasts, we provide algorithms to simulate from the relevant predictive distributions. The findings show that, for one-month-ahead point forecasts, our proposed BAVART model with NS factors yields highly competitive forecasts for bonds with maturities greater than five years. For higher-order forecasts, these accuracy gains become smaller. When the full predictive distribution is considered, jointly modeling the yields without imposing a factor structure gives rise to more precise density predictions. As opposed to point forecasts, we find that using BART on the conditional mean helps in improving density forecasts for maturities below seven years and both one-step-ahead and three-steps-ahead predictions.

In a second application, we apply the BAVART model to macroeconomic data for the Eurozone. Instead of using our model to produce forecasts, we analyze the dynamic effects of uncertainty on the Eurozone economy. As opposed to the existing literature, which deals with the macroeconomic effects of uncertainty using nonlinear models (see, e.g., [Aastveit, Natvik and Sola \(2017\)](#), [Alessandri and Mumtaz \(2019\)](#), [Caggiano, Castelnuovo and Groshenny \(2014\)](#), [Crespo Cuaresma, Huber and Onorante \(2020\)](#), [Ferrara and Guérin \(2018\)](#), [Jackson, Kliesen and Owyang \(2020\)](#), [Mumtaz and Theodoridis \(2018\)](#), [Paccagini and Colombo \(2020\)](#), [Caggiano, Castelnuovo and Nodari \(2021\)](#)), our approach remains agnostic on the precise form of nonlinearities and infers these in a data-based manner. This constitutes a substantial advantage since there is considerable uncertainty with respect to selecting the appropriate type of nonlinearities in multivariate time series models.

To assess how macroeconomic reactions change in our nonlinear and nonparametric framework if interest rates are at the zero lower bound, we simulate dynamic responses under the restriction that short-term interest rates are zero and not allowed to react to uncertainty shocks. Our findings indicate that increases in economic uncertainty have negative effects on real activity. Specifically, we observe increases in unemployment rates, declining consumption levels, and a drop in prices. By contrast, financial market quantities display adverse reactions with declines in stock prices and decreasing short-term interest rates. These findings are, in general, consistent with established findings in the literature and thus show that our BAVART model can be used to carry out meaningful structural inference.

The remainder of this paper is organized as follows. [Section 2](#) discusses the BART model in the context of the homoscedastic regression model, while [Section 3](#) extends this method to the VAR case and proposes the BAVART specification and how we control for heteroscedasticity. [Section 4](#) presents the results of the term structure forecasting exercise, while [Section 5](#) investigates the relationship between uncertainty and macroeconomic dynamics in the Eurozone. Finally, the last section summarizes our findings and concludes the paper.

**2. Bayesian additive regression tree models.** In this section we briefly review BART.<sup>2</sup> Let  $\mathbf{y} = (y_1, \dots, y_T)'$  denote the  $T$ -dimensional response vector and  $\mathbf{X} = (\mathbf{x}_1, \dots, \mathbf{x}_T)'$  be a  $T \times K$ -dimensional matrix of exogenous variables with  $\mathbf{x}_t$  being the  $K$  covariates in time  $t$ . We assume that  $\mathbf{y}$  depends on  $\mathbf{X}$  through a potentially nonlinear function  $f : \mathbb{R}^{T \times K} \rightarrow \mathbb{R}^T$  as follows:

$$\mathbf{y} = f(\mathbf{X}) + \boldsymbol{\varepsilon}, \quad \boldsymbol{\varepsilon} \sim \mathcal{N}(\mathbf{0}_T, \sigma^2 \mathbf{I}_T),$$

<sup>2</sup>For an extensive introduction to BART models, see [Hill, Linero and Murray \(2020\)](#).

where  $\sigma^2$  denotes the error variance and the function  $f$  is generally not known. BART approximates the function  $f$  by summing over  $N$  (which is a large number) regression trees,

$$(1) \quad f(\mathbf{X}) \approx \sum_{j=1}^N g(\mathbf{X}|\mathcal{T}_j, \mathbf{m}_j).$$

In Eq. (1), the function  $g(\mathbf{X}|\mathcal{T}_j, \mathbf{m}_j)$  corresponds to a single tree model with  $\mathcal{T}_j$  denoting the tree structure associated with the  $j$ th binary tree,  $\mathbf{m}_j = (\mu_{j1}, \dots, \mu_{jb_j})'$  is the vector of terminal node parameters associated with  $\mathcal{T}_j$  and  $b_j$  are the leaves of the  $j$ th tree. In what follows, and in consistency with Chipman, George and McCulloch (2010), we set  $N = 250$  in all our empirical applications.

These binary trees are constructed by considering splitting rules of the form  $\{\mathbf{X} \in \mathcal{A}_{jk}\}$  or  $\{\mathbf{X} \notin \mathcal{A}_{jk}\}$  with  $\mathcal{A}_{jk}$  being a partition set. These rules typically depend on selected columns of  $\mathbf{X}$ , denoted as  $\mathbf{X}_{\bullet j}$  ( $j = 1, \dots, K$ ) and a threshold  $c$ . The set  $\mathcal{A}_{jk}$  is then defined by splitting the predictor space, according to  $\{\mathbf{X}_{\bullet j} \leq c\}$  or  $\{\mathbf{X}_{\bullet j} > c\}$ .

The step function  $g$  is constant over the elements of  $\mathcal{A}_{jk}$ ,

$$g(\mathbf{X}|\mathcal{T}_j, \mathbf{m}_j) = \mu_{jk}, \quad \text{if } \mathbf{X} \in \mathcal{A}_{jk}, k = 1, \dots, b_j.$$

Hence, the set  $\mathcal{A}_{jk}$  defines a tree-specific unique partition of the covariate space such that the function  $g$  returns a specific value  $\mu_{jk}$  for specific values of  $\mathbf{x}_t$ .

To avoid overfitting, the trees are encouraged to be small (i.e., take a particularly simple form) and the terminal node parameters to be shrunk to zero. If the first tree,  $g(\mathbf{x}|\mathcal{T}_1, \mathbf{m}_1)$ , is a weak learner and fitted in a reasonable way, the corresponding tree structure will be very simple and elements in  $\mathbf{m}_1$  will be pushed toward zero. This implies that the first tree will explain a small fraction of the variation in  $\mathbf{y}$ . Subtracting  $g(\mathbf{X}|\mathcal{T}_1, \mathbf{m}_1)$  from  $\mathbf{y}$  yields a new conditional model with transformed  $\tilde{\mathbf{y}} = \mathbf{y} - g(\mathbf{X}|\mathcal{T}_1, \mathbf{m}_1)$ , and then the next tree will be fitted with  $\mathbf{y}$  being replaced by  $\tilde{\mathbf{y}}$ . This procedure is repeated for a sufficiently large number  $N$  of trees until the fit of the additive model becomes reasonably good.

We illustrate BART using a simple example. In this simple example we first focus on a single regression tree model (i.e.,  $N = 1$ ). The case of several regression trees is considered afterward.

EXAMPLE 1. Consider the regression tree  $g(\mathbf{X}|\mathcal{T}_j, \mathbf{m}_j)$  in Figure 1. We include two covariates in  $\mathbf{x}_t = (x_{1t}, x_{2t})$ . In the figure each rectangle each refers to a splitting rule, and we henceforth label this a node. At the top node (also called root node or simply root), we have

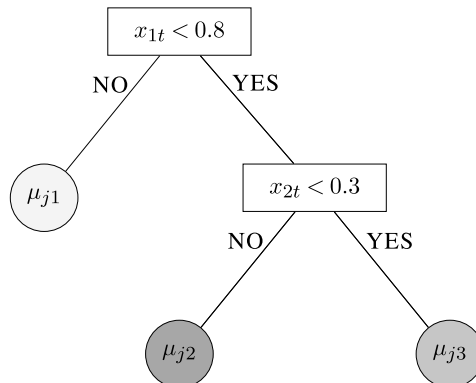


FIG. 1. Example of binary regression tree,  $g(\mathbf{x}|\mathcal{T}_j, \mathbf{m}_j)$ , with internal nodes labeled by their splitting rules and leaf nodes labeled with the corresponding parameters  $\mu_{jk}$ , with  $j = 1$  and  $k = 1, 2, 3$ .

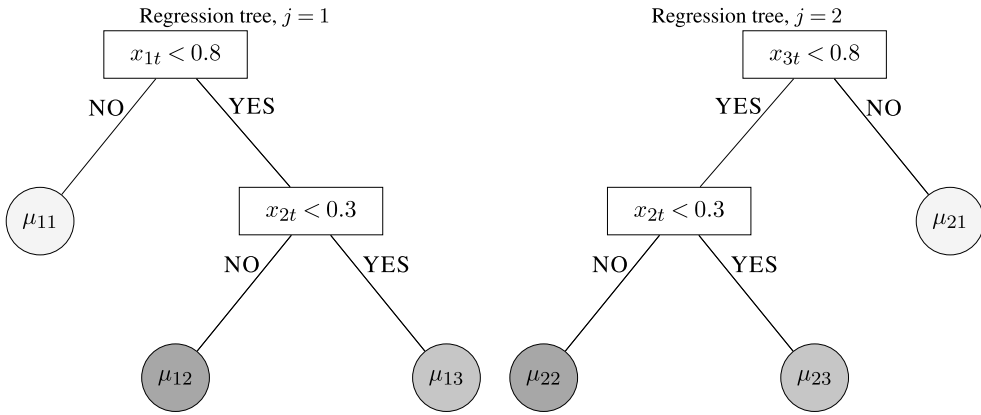


FIG. 2. Example of a sum of regression tree,  $g(\mathbf{x}|\mathcal{T}_j, \mathbf{m}_j)$ , with internal nodes labeled by their splitting rules and leaf nodes labeled with the corresponding parameters  $\mu_{jk}$ , where  $j = 1, 2$  and  $k = 1, 2, 3$ .

the condition that asks whether  $x_{1t} < 0.8$ . If this condition is not true, then we arrive at the terminal node,  $\mu_{j1}$  and set  $\mathbb{E}(y_t) = \mu_{j1}$ . By contrast, if the condition is fulfilled, we move to the next node. The corresponding splitting rule states that if  $x_{2t} \geq 0.3$ , we reach a terminal node with  $\mathbb{E}(y_t) = \mu_{j2}$ , whereas if  $x_{2t} < 0.3$ , we set  $\mathbb{E}(y_t) = \mu_{j3}$ . Using a single tree thus yields three possible values for the conditional mean,  $\mu_{j1}$ ,  $\mu_{j2}$  and  $\mu_{j3}$  and implies abrupt shifts between them. If  $y_t$  is evolving smoothly over time, using a single tree would thus lead to a rather simplistic conditional mean structure.

EXAMPLE 2. Since the tree model in the first example is too simplistic, our second example deals with a more flexible conditional mean structure. Let us consider a sum of regression trees with  $N = 2$  trees and 3 covariates, depicted in Figure 2. Within a given tree the same intuition as in Example 1 applies. However, since we now use two trees, we gain more flexibility in modeling the conditional mean of  $y_t$ . This is because the conditional mean at a given point in time is equal to the sum of the terminal node parameters for the two trees, and the combination of the two parameters depends on the specific set of decision rules across both trees. This allows for a substantially richer mean structure and increased flexibility, as opposed to a single tree model.

The second example shows how flexibility is increased by adding more trees and illustrates how BART handles nonlinearities in a flexible manner. In particular, each regression tree is a simple stepwise function, and when we sum over the different regression trees, we gain flexibility. The resulting additive model essentially allows for approximating nonlinearities without prior assumptions on the specific form of the nonlinearities.

### 3. BART models for multivariate time series analysis.

3.1. *Bayesian additive vector autoregressive tree models.* In this section we generalize the model outlined in the previous section to the multivariate case. Consider an  $M$ -dimensional vector of endogenous variables  $\mathbf{y}_t = (y_{1t}, \dots, y_{Mt})'$ . Stacking the rows yields a  $T \times M$  matrix  $\mathbf{Y} = (\mathbf{y}_1, \dots, \mathbf{y}_T)'$ . We assume that  $\mathbf{Y}$  depends on a  $T \times K$  matrix  $\mathbf{X} = (\mathbf{X}_1, \dots, \mathbf{X}_T)'$  with each  $\mathbf{X}_t = (\mathbf{y}'_{t-1}, \dots, \mathbf{y}'_{t-p})'$  being a  $K (= PM)$ -dimensional vector of lagged endogenous variables. The BAVART model is then given by

$$\mathbf{y}_t = F(\mathbf{X}_t) + \boldsymbol{\varepsilon}_t,$$

or, in terms of full-data matrices,

$$(2) \quad \mathbf{Y} = F(\mathbf{X}) + \boldsymbol{\varepsilon},$$

with  $\boldsymbol{\varepsilon} = (\boldsymbol{\varepsilon}_1, \dots, \boldsymbol{\varepsilon}_T)'$  denoting a  $T \times M$  matrix of shocks with typical  $t$ th row  $\boldsymbol{\varepsilon}_t \sim \mathcal{N}(\mathbf{0}_M, \boldsymbol{\Sigma})$ . For the moment, we assume that the error variance-covariance matrix  $\boldsymbol{\Sigma}$  is time-invariant.

In Eq. (2),  $F$  is defined in terms of equation-specific functions  $f_j(\mathbf{X})$ :

$$F(\mathbf{X}) = (f_1(\mathbf{X}), \dots, f_M(\mathbf{X}))'.$$

Similarly to the standard BART specification, we approximate each  $f_j$  ( $j = 1, \dots, M$ ) through a sum of  $N$  regression trees,

$$f_j(\mathbf{X}) \approx \sum_{k=1}^N g_{jk}(\mathbf{X}|\mathcal{T}_{jk}, \mathbf{m}_{jk}).$$

Here,  $g_{jk}$  ( $k = 1, \dots, N$ ) denotes an equation-specific step function with arguments  $\mathcal{T}_{jk}$  and  $\mathbf{m}_{jk}$ . As before, the individual tree structures  $\mathcal{T}_{jk}$  are associated with a  $b_{jk}$ -dimensional vector  $\mathbf{m}_{jk} = (\mu_{jk,1}, \dots, \mu_{jk,b_{jk}})'$  of terminal node coefficients associated with  $b_{jk}$ , denoting the number of leaves per tree in equation  $j$ . Notice that both the tree structures and the terminal node parameters are now specific to equation  $j$ . The main difference to the model, illustrated in Section 2, is that we approximate  $M$  different functions  $f_j$ . This implies that if certain elements in  $\mathbf{y}_t$  depend linearly on  $\mathbf{X}_t$ , then our flexible multivariate specification can pick this up.

In what follows, we will estimate the BAVART model by exploiting its structural form. Using the structural form of the VAR to speed up computation has been done in several recent papers (see, e.g., [Carriero, Clark and Marcellino \(2019\)](#), [Huber, Koop and Onorante \(2021\)](#)).<sup>3</sup> This implies that Eq. (2) can be written as follows:

$$(3) \quad \mathbf{Y} = F(\mathbf{X}) + \boldsymbol{\varepsilon} \mathbf{A}'_0,$$

whereby  $\mathbf{A}_0$  denotes a  $M \times M$ -dimensional lower triangular matrix with  $\text{diag}(\mathbf{A}_0) = (1, \dots, 1)'$  and  $\boldsymbol{\varepsilon} = (\boldsymbol{\varepsilon}_1, \dots, \boldsymbol{\varepsilon}_T)'$  is a  $T \times M$  matrix of orthogonal shocks with  $\boldsymbol{\varepsilon}_t \sim \mathcal{N}(\mathbf{0}_M, \mathbf{H})$ .  $\mathbf{H}$  denotes a  $M \times M$ -dimensional diagonal matrix with the variances on its main diagonal. Hence, the covariance matrix is  $\boldsymbol{\Sigma} = \mathbf{A}_0 \mathbf{H} \mathbf{A}'_0$ .

Conditional on  $\mathbf{A}_0$ , this form permits equation-by-equation estimation since the shocks are independent. This leads to enormous computational gains. The  $j$ th  $> 1$  equation can be written as

$$(4) \quad \mathbf{y}_{\bullet j} = \sum_{k=1}^N g_{jk}(\mathbf{X}|\mathcal{T}_{jk}, \mathbf{m}_{jk}) + \sum_{l=1}^{j-1} a_{jl} \boldsymbol{\varepsilon}_{\bullet l} + \boldsymbol{\varepsilon}_{\bullet j}.$$

Here,  $\mathbf{y}_{\bullet j}$ ,  $\boldsymbol{\varepsilon}_{\bullet l}$  and  $\boldsymbol{\varepsilon}_{\bullet j}$  refers to the  $j$ th or  $l$ th column of  $\mathbf{Y}$ ,  $\boldsymbol{\varepsilon}$  and  $\boldsymbol{\varepsilon}$ , respectively.  $a_{jl}$  denotes the  $(j, l)$ th element of  $\mathbf{A}_0$ .

Notice that Eq. (4) is a generalized additive model that consists of a nonparametric part ( $\sum_{k=1}^N g_{jk}(\mathbf{X}|\mathcal{T}_{jk}, \mathbf{m}_{jk})$ ) and a regression part ( $\sum_{l=1}^{j-1} a_{jl} \boldsymbol{\varepsilon}_{\bullet l}$ ). For  $j = 1$ , the model reduces to a standard BART specification without the regression part.

---

<sup>3</sup>The algorithm proposed in [Carriero, Clark and Marcellino \(2019\)](#) is sampling from an “inexact” full conditional posterior distribution (see, [Carriero et al. \(2021\)](#)). As a robustness check, we have recomputed several of the results in Sections 4 and 5 using an algorithm that replaces the shocks in the previous  $j - 1$  equations with the actual contemporaneous values in  $\mathbf{y}_t$  and thus sample the trees in the structural representation. The corresponding results are similar and we expect that the differences which arise are mainly driven by the fact that all priors are specified on structural as opposed to reduced-form parameters.

The key idea behind this formulation is that, for a sufficiently large number of trees, we approximate nonlinear relations between  $\mathbf{y}_t$  and its lags while allowing for linear relations between the contemporaneous values of  $\boldsymbol{\varepsilon}_t$ . These linear relations determine the covariances and are of vital importance for the identification of structural shocks.

3.2. *Allowing for heteroscedasticity.* Up to this point, we assumed the error variance to be constant. If the researcher wishes to relax this assumption, several feasible options exist. In a recent paper, [Pratola et al. \(2020\)](#) propose a combination of a standard additive BART model with a multiplicative BART specification to flexibly control for heteroscedasticity. But modeling heteroscedasticity with BART implies that we need some information on how the volatility of economic shocks depends on additional covariates (which potentially differ from  $\mathbf{X}$ ). As a simple yet flexible solution, we adopt a standard stochastic volatility (SV) model. SV models are frequently used in macroeconomics and finance and have a proven track record for forecasting applications (see, e.g., [Clark \(2011\)](#), [Clark and Ravazzolo \(2015\)](#)).

Our SV specification assumes that  $\mathbf{H}$  is time-varying,

$$\mathbf{H}_t = \text{diag}(e^{h_{1t}}, \dots, e^{h_{Mt}}),$$

with the time-varying variances  $e^{h_{jt}}$ . We assume that the  $h_{jt}$ 's follow an AR(1) model:

$$(5) \quad h_{jt} = c_j + \rho_j(h_{jt-1} - c_j) + \sigma_{jh}v_{jt}, \quad v_{jt} \sim \mathcal{N}(0, 1),$$

$$(6) \quad h_{j0} \sim \mathcal{N}\left(c_j, \frac{\sigma_{jh}^2}{1 - \rho_j^2}\right).$$

Here,  $c_j$  is the unconditional mean,  $\rho_j$  is the persistence parameter and  $\sigma_{jh}^2$  is the error variance of the log-volatility process. This specification essentially implies that, if  $\rho_j$  is close to one, the log-volatilities evolve smoothly over time and tend to be persistent.

For macroeconomic and financial data, using SV has been shown to greatly increase forecast accuracy. Moreover, it is worth noting that using SV entails a much more flexible error distribution than the one used in, for example, [Chipman, George and McCulloch \(2010\)](#). Hence, coupling the BAVART model with a stochastic volatility component yields a model which allows for flexible adjustments of the conditional mean while also being flexible on the error variances. In our empirical work, the models considered feature stochastic volatility of this form.

3.3. *The prior and posterior simulation.* The Bayesian approach calls for the specification of suitable priors over the parameters of the model. Here, we mainly follow the different strands of the literature our approach combines. In particular, we focus on the priors associated with the trees  $\mathcal{T}_{jk}$  and the terminal node parameters  $\mathbf{m}_{jk}$  for  $j = 1, \dots, M$  and  $k = 1, \dots, N$ . We assume that the hyperparameters across priors are the same for each equation. On the covariances we use the horseshoe prior ([Carvalho, Polson and Scott \(2009\)](#)).

For each equation  $j$ , the joint prior structure is given by

$$p((\mathcal{T}_{j1}, \mathbf{m}_{j1}), \dots, (\mathcal{T}_{jN}, \mathbf{m}_{jN}), c_j, \rho_j, \sigma_j^2, \mathbf{a}_j) = p((\mathcal{T}_{j1}, \mathbf{m}_{j1}), \dots, (\mathcal{T}_{jN}, \mathbf{m}_{jN})) \\ \times p(c_j) \times p(\rho_j) \times p(\sigma_{j,h}^2) \times p(\mathbf{a}_j).$$

At this level the prior implies independence between  $p((\mathcal{T}_{j1}, \mathbf{m}_{j1}), \dots, (\mathcal{T}_{jN}, \mathbf{m}_{jN}))$  and the remaining model parameters. [Chipman, George and McCulloch \(2010\)](#) further introduce the

following factorization:

$$p((T_{j1}, \mathbf{m}_{j1}), \dots, (T_{jN}, \mathbf{m}_{jN})) = \prod_{k=1}^N p(\mathbf{m}_{jk} | \mathcal{T}_{jk}) \times p(\mathcal{T}_{jk}),$$

$$p(\mathbf{m}_{jk} | \mathcal{T}_{jk}) = \prod_{q=1}^{b_{jq}} p(\mu_{jk,q} | \mathcal{T}_{jk}),$$

implying that the tree components  $(\mathcal{T}_{jk}, \mathbf{m}_{jk})$  are independent of each other and of the remaining parameters while the prior on  $\mathbf{m}_{jk}$  depends on the tree structure. This prior structure allows for integrating out  $\mathbf{m}_{jk}$  from the posterior of the trees and thus greatly simplifies computation.

The prior on the tree structure is specified along the lines suggested in Chipman, George and McCulloch (1998) and Chipman, George and McCulloch (2010). Specifically, instead of constructing a prior on the trees directly, we construct a tree generating stochastic process that consists of three steps to grow trees. Let  $s = 0$  be the first iteration of this tree generating process. Then, the following steps for constructing trees are used:

Start with the trivial tree that consists of a single terminal node (i.e., its root) which will be denoted by  $\eta_{jk}$  for each equation  $j$  and  $k = 1, \dots, N$ .

Step 1. Split the terminal node  $\eta_{jk}$  with probability

$$p_{\text{SPLIT}}(\eta_{jk}, \mathcal{T}_{jk}^{(s)}) = \alpha(1+d)^{-\beta}, \quad \alpha \in (0, 1), \quad \beta \geq 0, \quad d = 0, 1, 2, \dots,$$

where  $d$  denotes the depth of the tree and  $(\alpha, \beta)$  are prior hyperparameters. These parameters are set as follows. In detail, a node at depth  $d$  spawns children with  $p_{\text{SPLIT}}(\eta_{jk}, \mathcal{T}_{jk}^{(s)})$ , and, as the tree grows,  $d$  increases while the prior probability decreases, making it less likely that a node spawns children. This implies that the probability that a given node becomes a bottom node increases, and thus a penalty on tree complexity is introduced. We set  $\alpha = 0.95$  and  $\beta = 2$ , implying that the probability that a given node is nonterminal decreases quadratically if the trees become more complicated (i.e., for increasing levels of  $d$ ).

Step 2. If the current node is split, we introduce a splitting rule  $\kappa_{jk}$  drawn from the distribution function  $p_{\text{RULE}}(\kappa_{jk} | \eta_{jk}, \mathcal{T}_{jk}^{(s)})$  and use this rule to spawn a left and right children node. In particular, the rule  $p_{\text{RULE}}(\kappa_{jk} | \eta_{jk}, \mathcal{T}_{jk}^{(s)})$  is set such that a given ‘‘splitting’’ covariate  $\mathbf{X}_{\bullet,j}$  is chosen with probability  $1/K$ . As described above, conditional on a chosen covariate, one needs to estimate a threshold. The prior on this threshold is, in the absence of substantial information, assumed to be uniformly distributed over the range of  $\mathbf{X}_{\bullet,j}$ .

Step 3. Once we obtain all terminal nodes (i.e., no node is split anymore), we will label this new tree  $\mathcal{T}_{jk}^{(s+1)}$  and return to step (2).

On the terminal node parameter  $\mu_{jk,q}$ , we use a conjugate Gaussian prior distribution  $\mathcal{N}(0, \sigma_{\mu}^2)$ , where  $\sigma_{\mu}^2$  is set as follows:

$$\sigma_{\mu}^2 = \frac{\mathfrak{R}_j}{2\tilde{s}\sqrt{N}},$$

with  $\mathfrak{R}_j$  denoting the range of the endogenous variable in equation  $j$ , and  $\tilde{s}$  denotes the number of prior standard deviations. This hyperparameter is set equal to 2, implying that  $\mu_{jk,q}$  will place around 95% prior mass on the range of  $\mathbf{y}_{\bullet,j}$ . One key property of this prior is that it increases with  $\mathfrak{R}_j$ . This implies that if outliers in  $\mathbf{Y}$  are observed, the range increases, and the prior becomes looser. By contrast, it decreases in the number of trees, implying that if  $N$  is large, the terminal node parameters associated with a given tree will be strongly pushed



toward zero. This is consistent with the notion that we aim to use a composite model of many weak learners, as opposed to having a small to moderate number of complex trees.

For the free elements in  $A_0$ , we introduce a horseshoe prior on each element of  $\mathbf{a}_j = (a_{j1}, \dots, a_{jj-1})'$ . This prior consists of local scaling parameters  $\tau_{jl}$ , which are specific to each covariance parameter  $a_{jl}$ , and a global shrinkage parameter  $\lambda$  which pushes all covariances toward zero. The horseshoe prior on the covariances is given by

$$\begin{aligned} a_{jl} | \tau_{jl}, \lambda^2 &\sim \mathcal{N}(0, \tau_{jl}^2 \lambda^2), \\ \tau_{jl} &\sim \mathcal{C}^+(0, 1), \\ \lambda &\sim \mathcal{C}^+(0, 1), \end{aligned}$$

where  $j = 1, \dots, M; l = 1, \dots, j - 1$  and  $\mathcal{C}^+$  denote the half-Cauchy distribution.

The prior specification on the parameters of the log-volatility equation follows the setup proposed in [Kastner and Frühwirth-Schnatter \(2014\)](#). In details we use a zero mean Gaussian prior with variance  $10^2$  on the unconditional mean  $c_j$ , a Beta prior on the (transformed) persistence parameter,  $\frac{\rho_j+1}{2} \sim \mathcal{B}(25, 5)$  and a Gamma prior on the error variance of the log-volatility process  $\sigma_{j,h}^2 \sim \mathcal{G}(1/2, 1/2)$ .

These priors can be combined with the likelihood to yield a joint posterior distribution over the coefficients and latent states in the model. To simulate from this joint posterior distribution, we adopt an MCMC algorithm. This algorithm simulates all quantities in an equation-specific manner. Since all steps necessary are standard, we only provide an overview. The [Appendix](#) provides more detail on sampling the trees. Here, it suffices to say that we sample all quantities related to the trees (i.e.,  $\mathcal{T}_{jk}$ ,  $\mathbf{m}_{jk}$  for all  $j, k$ ) using the algorithm outlined in [Chipman, George and McCulloch \(2010\)](#). The latent states and the coefficients associated with the state equation of the log-volatilities are obtained through the efficient algorithm discussed in [Kastner and Frühwirth-Schnatter \(2014\)](#). Conditional on the trees and log-volatilities, the posterior of  $\mathbf{a}_j$  is multivariate Gaussian and takes a standard form since the resulting conditional model is a linear regression model with heteroscedastic shocks. Finally, the parameters of the horseshoe prior are simulated using techniques outlined in [Makalic and Schmidt \(2015\)](#) which involve sampling from inverse Gamma distributions (conditional on introducing auxiliary shrinkage parameters) only.

**4. Forecasting the term structure of interest rates.** In this section we illustrate the predictive capabilities of the BAVART model. After providing an overview of the dataset, the model specification and the design of the forecasting exercise adopted, our focus will be on how well our approach works when applied to US yield curve data.

*4.1. Data overview, model specification and design of the forecasting exercise.* In this section we briefly discuss the two datasets adopted. For our forecasting application we use data on the nominal yield curve which is close to the dataset proposed in [Gürkaynak, Sack and Wright \(2007\)](#). These are downloaded from the website of the Federal Reserve Board (<https://www.federalreserve.gov/data/nominal-yield-curve.htm>) and range from June 1961 to December 2019. The maturities included are one, three, five, seven, 10 and 15 years, and we will consider changes in the interest rates. We will then use the period from June 1961 to June 2006 as our initial training sample. This allows us to compute forecast distributions for July 2006. We then move on to expand this estimation period by one month, until we reach the end of the full sample. This procedure yields a sequence of 160 predictive distributions, which we then evaluate using the observed outcomes.

We consider two approaches of modeling the term structure of interest rates. Due to its empirical success, we use the three-factor Nelson–Siegel (NS) model, originally proposed in

Nelson and Siegel (1987) and combined with a VAR state evolution equation on the factors in Diebold and Li (2006). The NS model assumes that the yield at maturity  $\varpi$ ,  $i_t(\varpi)$  depends on three latent factors, as follows:

$$(7) \quad i_t(\varpi) = \mathcal{L}_t + \frac{1 - e^{-\gamma\varpi}}{\gamma\varpi} \mathcal{S}_t + \left( \frac{1 - e^{-\gamma\varpi}}{\gamma\varpi} - e^{-\gamma\varpi} \right) \mathcal{C}_t,$$

whereby  $\mathcal{L}_t$ ,  $\mathcal{S}_t$  and  $\mathcal{C}_t$  denote a level, slope and curvature factor, respectively. Moreover,  $\gamma$  is a parameter which shapes the factor loadings. Consistent with Diebold and Li (2006), we set  $\gamma = 0.0609$ . This value maximizes the factor loadings on  $\mathcal{C}_t$ . The three factors are then obtained on a  $t$ -by- $t$  basis using OLS estimation. The NS model sets  $\mathbf{y}_t = (\mathcal{L}_t, \mathcal{S}_t, \mathcal{C}_t)'$  and assumes it to evolve according to a multivariate dynamic model (such as our BAVART specification). Predictions of  $\mathbf{y}_t$  are then mapped back using Eq. (7).

The second approach follows Carriero, Kapetanios and Marcellino (2012) and models the  $M = 7$  yields directly in a VAR. This approach is thus less parsimonious but also more flexible since it essentially allows for maturity-specific idiosyncrasies.

For these two approaches we benchmark the BAVART model against several linear and nonlinear competing models. The first model we consider is a time-varying parameter (TVP) VAR with a normal-gamma shrinkage prior on the state innovation variances (for the exact specification; see Huber, Koop and Onorante (2021)). Since the period of the zero lower bound (ZLB) is a dominant source of nonlinearities in yield curve data, we estimate nonlinear models that assume a dependence between the parameters of the model and the lagged short-term interest rate (in our case the one-year yield). This gives rise to the second model which is an interacted VAR (IVAR) that introduces interaction effects between the first lag of the short-term interest rate and the remaining model parameters (see, e.g., Caggiano, Castelnuovo and Pellegrino (2017)), whereas the third model is a smooth transition VAR (STVAR) which assumes that coefficients evolve slowly between two regimes (see Gefang and Strachan (2010), Auerbach and Gorodnichenko (2012)). Fourth, we consider a Bayesian threshold VAR (TVAR) which assumes that coefficients change if the (first lag) of the short-term interest rate passes a threshold to be estimated (see Alessandri and Mumtaz (2017), Huber and Zörner (2019)). All these models are estimated using a horseshoe prior on the VAR coefficients and benchmarked against a constant parameter NS -VAR equipped with a Minnesota prior and stochastic volatility. Moreover, we include  $P = 2$  lags of the endogenous variables in all models considered.

4.2. *Forecasting results.* Before discussing the results of the forecast exercise, the question on how to compute the predictive distribution naturally arises. Producing one-step-ahead forecasts is computationally easy, since, conditional on the tree structures, splitting values and splitting covariates, it is straightforward to compute a prediction for  $\mathbf{y}_{T+1}$ . More precisely, (and with a slight abuse of notation), the one-step-ahead predictive distribution is given by

$$p(\mathbf{y}_{T+1} | \mathbf{y}_{1:T}) = \int p(\mathbf{y}_{T+1} | \mathbf{y}_{1:T}, \mathbf{\Xi}) p(\mathbf{\Xi} | \mathbf{y}_{1:T}) d\mathbf{\Xi},$$

where  $\mathbf{y}_{1:T}$  denotes the full history of  $\mathbf{y}_t$  and  $\mathbf{\Xi}$  is a generic notation that summarizes all parameters and latent states (i.e., the tree structures, terminal nodes, log-volatilities etc.). This integral is solved numerically through Monte Carlo integration.

The conditional density  $p(\mathbf{y}_{T+1} | \mathbf{y}_{1:T}, \mathbf{\Xi})$  is

$$y_{T+1} | \mathbf{y}_{1:T}, \mathbf{\Xi} \sim \mathcal{N}(F(\mathbf{X}_{T+1}), \tilde{\Sigma}_{T+1|T}),$$

with  $\tilde{\Sigma}_{T+1|T}$  is a random draw obtained by using Eq. (6) to predict the log-volatilities and using these predictions to form  $\mathbf{H}_{T+1|T}$ . Higher order forecasts are then computed iteratively

by first simulating from the one-step-ahead predictive distributions to obtain a prediction for  $\mathbf{y}_{T+1}$ . We label this draw  $\tilde{\mathbf{y}}_{T+1}$ .  $\tilde{\Sigma}_{T+2|T}$  is again computed by exploiting Eq. (6). This allows us to draw  $\tilde{\mathbf{y}}_{T+2}$  from

$$\tilde{\mathbf{y}}_{T+2} \sim \mathcal{N}(F(\tilde{\mathbf{X}}_{T+2}), \tilde{\Sigma}_{T+2|T}).$$

We repeat this procedure until we have a prediction  $\mathbf{y}_{T+o}$  with  $o$  denoting the desired forecast horizon. In what follows, we will consider  $o = 1$  and *three-steps-ahead* forecasts.

The point (i.e., median forecasts) and density predictions are then evaluated by using relative mean squared forecast errors (MSFEs) and the continuous ranked probability score (CRPS), respectively. The MSFEs allow us to gauge the quality of point forecasts while the CRPS is used to also take into account higher-order moments of the predictive distribution. All results are benchmarked against the NS-Minnesota VAR.

The findings of our forecasting exercise are summarized in heatmaps depicted in Figures 3 and 4. If a given model is performing worse than the benchmark VAR, the corresponding cell will become darker (marking a relative MSE/CRPS score well below, or above unity).

Our forecasting horse race draws a rich picture of relative model performance. We consider models of different sizes (namely, the NS variants and the ones that use the selected yields exclusively), priors, assumptions on the conditional mean and forecast horizons. Besides, we consider both point and density forecasts.

Starting with one-month-ahead point forecasting accuracy, depicted in Figure 3, we observe that the set of competing models improves upon the benchmark VAR. But these improvements are rather small and mostly concentrated toward the short-end (i.e., maturities

	One-month-ahead						One-quarter-ahead					
NS-BART	+1.02	+0.97	+0.96	+0.95	+0.93	+0.91	+0.98	+0.99	+0.99	+0.99	+0.99	+0.99
NS-TVP	+0.93	+0.91	+0.96	+0.98	+0.99	+0.97	+0.95	+0.99	+0.99	+0.99	+1.00	+1.00
NS-IVAR	+1.04	+1.00	+1.00	+1.01	+1.01	+1.01	+1.02	+1.01	+1.00	+1.00	+1.00	+1.00
NS-STVAR	+1.01	+0.98	+1.00	+1.01	+1.02	+1.02	+1.02	+1.01	+1.01	+1.00	+1.01	+1.01
NS-TVAR	+0.88	+0.91	+0.96	+0.99	+1.00	+1.01	+1.03	+1.01	+1.00	+1.00	+1.00	+1.00
BART	+1.11	+1.03	+0.99	+1.00	+1.01	+1.02	+0.99	+0.99	+0.98	+0.99	+0.99	+0.98
TVP	+0.96	+0.95	+0.99	+1.00	+1.02	+1.04	+0.99	+1.05	+1.02	+1.00	+1.00	+0.99
IVAR	+0.96	+0.97	+1.00	+1.02	+1.03	+1.08	+0.99	+1.00	+1.00	+1.00	+1.00	+1.01
STVAR	+1.02	+1.00	+0.99	+1.01	+1.02	+1.04	+0.99	+0.99	+0.99	+0.98	+0.98	+0.99
TVAR	+0.92	+0.93	+0.96	+1.01	+1.03	+1.04	+1.00	+0.98	+0.98	+0.98	+0.98	+0.98
	1y	3y	5y	7y	10y	15y	1y	3y	5y	7y	10y	15y

FIG. 3. Mean squared forecast errors relative to the NS-Minnesota VAR over the hold-out period: July 2006 to December 2019.

shorter than seven years) of the yield curve. For this yield curve segment there is no clear picture emerging on whether using the NS factor model or jointly modeling the different maturities improves upon the other. For some models the NS variant yields slightly smaller MSE ratios than the unrestricted model (e.g., our BAVART model or NS-TVP), whereas, for other models, unrestricted modeling of the maturities yields more precise point forecasts (e.g., the IVAR and the TVAR).

The relative performance of our proposed BART-based model increases with the maturity. For shorter maturities we find that BAVART performs well but in most cases is outperformed by one of the competing models (such as the NS-TVAR for one-year bonds and the NS-TVP for the three-year bonds). When we consider bonds with maturities larger or equal than five years, the performance of the NS-BART model improves appreciably. These improvements reach 7% for 10-year government bonds and 9% for 15-year government bonds. Interestingly, these relative MSE ratios are always smaller than the ones we observe for the unrestricted BAVART model. We conjecture that the stronger performance of BAVART for bonds with longer maturities is mainly driven by the fact that these time series display more variation, especially after the global financial crisis. This is the period where the U.S. Federal Reserve introduced unconventional monetary policy measures, such as quantitative easing to push down the slope of the yield curve.

When we focus on the one-quarter-ahead forecasts, we observe a great deal of light gray cells on the right panel of Figure 3, indicating only small differences in predictive accuracy vis-à-vis the BVAR benchmark. Most models feature MSE ratios close to unity with accuracy gains ranging from around two to five percent (in the case of the NS-TVP model and for the one-year-ahead bond yield).

Considering only point forecasts imply that we do not factor in how well a given model predicts higher order moments of the predictive distribution. We now turn to discuss the accuracy of density forecasts. Figure 4 shows that several of our competing models yield density forecasts which are slightly worse than the ones of the benchmark VAR for both forecast horizons. There are two exceptions to this pattern. First, we find that both versions of the BAVART and the TVP VAR yield precise density predictions for one- up to five-year yields (with the TVP VAR being better for one-year and three-year yields and the NS-BAVART outperforming all competitors for five-year yields). Especially for the short-end of the yield curve, these improvements are sizable (around 25% for the unrestricted TVP-VAR and 14% for the unrestricted BAVART). We conjecture that this stems from the fact that one-year yields display little variation during the period of the ZLB. In such a situation the most flexible models in our pool (i.e., BAVART and the TVP VAR) quickly adjust both the parameters and the error variances and thus yields tighter predictive intervals centered around values close to zero. The other models lack this flexibility since all parameters in the system either change or display no/little change. Even though the signal variable is the short-term interest rate, this potentially is a limitation and thus could negatively impact density forecasting accuracy.

Second, when we focus on three-months-ahead forecasts, a similar picture emerges. The unrestricted BAVART and the TVP model work well at the short end of the yield (with sizable gains for one- and three-year yields in both cases). For higher maturities and in contrast to the results shown in Figure 3, we find relative CRPSs close to one, indicating that all models included have a hard time beating the benchmark VAR with SV.

To sum up, our forecasting exercise shows that the BAVART model works remarkably well for predicting the U.S. term structure of interest rates. When we focus on point forecasts, the predictive gains of the BAVART model increase with the maturity of the bond. For density predictions this story is reversed, indicating that both variants of the BAVART work well at the short-end of the yield curve.

	One-month-ahead						One-quarter-ahead					
NS-BART	+1.09	+1.01	+1.00	+0.99	+0.98	+0.98	+1.08	+1.02	+1.01	+1.00	+1.00	+1.00
NS-TVP	+0.90	+0.93	+0.96	+1.00	+1.00	+0.98	+0.92	+0.94	+0.96	+0.98	+0.99	+0.99
NS-IVAR	+1.41	+1.10	+1.05	+1.03	+1.02	+1.03	+1.31	+1.08	+1.04	+1.03	+1.02	+1.03
NS-STVAR	+1.58	+1.20	+1.12	+1.10	+1.09	+1.10	+1.52	+1.20	+1.13	+1.11	+1.10	+1.11
NS-TVAR	+1.59	+1.19	+1.12	+1.09	+1.09	+1.10	+1.51	+1.20	+1.12	+1.10	+1.09	+1.10
BART	+0.86	+0.95	+0.98	+0.98	+1.00	+1.03	+0.88	+0.91	+0.96	+0.98	+1.00	+1.02
TVP	+0.75	+0.91	+0.99	+1.07	+1.10	+1.09	+0.78	+0.92	+0.98	+1.03	+1.06	+1.04
IVAR	+1.37	+1.08	+1.04	+1.03	+1.03	+1.04	+1.31	+1.08	+1.03	+1.02	+1.02	+1.02
STVAR	+1.53	+1.18	+1.12	+1.10	+1.09	+1.10	+1.49	+1.17	+1.11	+1.09	+1.08	+1.10
TVAR	+1.52	+1.17	+1.11	+1.09	+1.09	+1.10	+1.46	+1.16	+1.10	+1.09	+1.08	+1.09
	1y	3y	5y	7y	10y	15y	1y	3y	5y	7y	10y	15y

FIG. 4. *Continuous ranked probability scores relative to the NS-Minnesota VAR over the hold-out period: July 2006 to December 2019.*

**5. The effects of macroeconomic uncertainty on the Eurozone.** We now turn to our application based on Eurozone data. The next subsection provides a brief introduction to the dataset used while Section 5.2 shows some in-sample results. Section 5.2.1 deals with the question of how macroeconomic uncertainty impacts the Eurozone economy with a particular emphasis on the zero lower bound.

**5.1. Data overview.** In the second application we apply our BAVART model to analyze the effects of uncertainty shocks on macroeconomic outcomes. Several recent papers analyze this question using U.S. datasets (Bloom (2009), Caggiano, Castelnuovo and Pellegrino (2017), Jurado, Ludvigson and Ng (2015), Ramey and Zubairy (2017), Carriero, Clark and Marcellino (2018)). Instead of focusing on U.S. data, we apply our approach to monthly Eurozone data that ranges from January 1999 to January 2019. We opt for this dataset because the Eurozone economy underwent structural changes over this estimation period, multiple recessions and, in general, a challenging environment since the time series we consider are rather short.

The dataset we use is a selection of time series from the popular Euro Area Real Time Database (EA-RTD). We consider a medium-scale model that includes six macroeconomic and financial time series. These six time series are augmented by an uncertainty index (abbreviated as UNC) which is estimated using the approach outlined in Jurado, Ludvigson and Ng (2015). Apart from the uncertainty indicator, we include the (log) level Dow Jones Euro Stoxx 50 price index (DJE50), HICP inflation (C\_OV), unemployment rate (UNETO), three-month Euribor (EUR3M), industrial production (XCONS) and the yield on Euro area

TABLE 1

*Number of times a given covariate shows up across splitting rules in a medium-scale VAR with one lag*

	UNC <sub>t</sub>	DJE50 <sub>t</sub>	EUR3M <sub>t</sub>	C_OV <sub>t</sub>	UNETO <sub>t</sub>	XCONS <sub>t</sub>	10Y <sub>t</sub>
UNC <sub>t-1</sub>	52	44	44	38	41	45	44
DJE50 <sub>t-1</sub>	46	56	44	32	39	43	47
EUR3M <sub>t-1</sub>	49	50	62	55	40	47	45
C_OV <sub>t-1</sub>	47	48	48	60	43	50	50
UNETO <sub>t-1</sub>	47	46	27	38	60	49	47
XCONS <sub>t-1</sub>	45	43	43	30	45	51	47
10Y <sub>t-1</sub>	43	42	33	37	41	43	51

10-year benchmark government bonds (10Y). We include one lag of the endogenous variables.<sup>4</sup>

5.2. *In-sample results.* In this section we illustrate key features of our BART-based approach using a subset of the dataset discussed in the previous section. Variable importance is gauged by considering the posterior median of the number of times a given quantity shows up in a splitting rule. These frequencies are shown in Table 1. The columns refer to the different equations and the rows to the corresponding covariates.

A simple inspection of the main diagonal elements of the table reveals that the first, own lag of a given variable within the corresponding equation shows up frequently. This finding holds for all equations and resembles key results of the literature on Bayesian VARs which states that the AR(1) term explains most variation in  $y_t$ . However, here it is worth stressing that our model is far from sparse, and also the lags of other variables seem to play a role in determining the dynamics of a given endogenous variable. The lags of other quantities in a given equation are also often included as splitting variables, indicating that we can not tell a simple story about some few elements in  $x_t$  exclusively shaping the dynamics of  $y_t$ .

5.2.1. *Impulse responses to an uncertainty shock and the role of the zero lower bound.* In this section we illustrate how the BAVART model can be used to carry out structural inference. There is a broad body of literature dealing with the macroeconomic effects of uncertainty using nonlinear models (Alessandri and Mumtaz (2019), Crespo Cuaresma, Huber and Onorante (2020), Ferrara and Guérin (2018), Mumtaz and Theodoridis (2018), Paccagini and Colombo (2020), Caggiano, Castelnovo and Nodari (2021)). These papers all rely on parametric approaches and make rather strong assumptions on the nature of nonlinearities. Our BAVART model, by contrast, introduces almost no restrictions and thus remains agnostic on the specific form of nonlinearities in the transmission mechanisms.

Because of the highly nonlinear nature of our model, we resort to generalized impulse response (GIRF) functions (see Koop, Pesaran and Potter (1996)). These impulse responses are computed as follows. Let  $\epsilon_t$  denote the  $M$  structural shocks, and  $s_j$  denotes a  $M \times 1$  selection vector that equals 1 in the  $j$ th position. Hence,  $\epsilon_{jt} = s'_j \epsilon_t$  yields the  $j$ th structural shock. The GIRF to the  $j$ th structural shock is then defined as the difference between the forecast which assumes  $\epsilon_{jt} = 1$  (while setting the other shocks to zero) and the unconditional forecast (i.e., with  $\epsilon = \mathbf{0}_M$ ) for a given forecast horizon. The posterior distributions of the GIRFs are

<sup>4</sup>Using a single lag allows for simple inspection of several features of the model. Including more lags is straightforward, but lags larger than one only rarely show up in the splitting rules and the impulse responses look very similar to the ones reported in the main body of the text.

computed by repeating this procedure during MCMC sampling. The structural shock is computed by setting  $\mathbf{H}_t$  equal to the unconditional mean, that is,  $\mathbf{H}_t = \text{diag}(e^{c^1}, \dots, e^{c^M})$  and assuming that the uncertainty factor is ordered first.

To assess whether BART uncovers implicit nonlinear relations consistent with the literature quoted above, we consider two experiments. The first assumes that the economy is hit by a one standard deviation uncertainty shock. The dynamic reactions of  $\mathbf{y}_t$  are left unrestricted. This implies that the central bank is able to react to increases in uncertainty by lowering interest rates.

The second experiment asks how the reactions in  $\mathbf{y}_t$  change if the economy is stuck at the zero lower bound (ZLB) and the central bank can not use conventional monetary policy tools. This second experiment is interesting because our BAVART model can use the information that the short-term interest rate is zero so that different branches of a given tree are effectively ruled out. For instance, consider a simple tree which has a root node with a decision rule that splits the observations according to whether the short-term interest rate is below one percent or greater than one percent. If we introduce the restriction that the interest rate is zero, the second branch of the tree (the above one percent interest rate branch) will play no role in constructing the impulse responses. Hence, only different configurations of the covariate space, which are consistent with interest rates close to zero, are being considered. The restriction that the short-term interest rate is at the ZLB is easily incorporated by zeroing out the forecast of the short rate.

Figure 5 presents the dynamic responses to a one standard deviation uncertainty shock. The gray shaded area is the marginal posterior distribution of the IRFs, with solid black lines denoting the 16(25)th and 75(84)th percentiles. The dotted black lines represent the 16th and 84th credible intervals of the impulse responses, based on zeroing out the reaction of the short-term interest rate.

In general, we find that the IRFs closely resemble the ones reported in the literature (see, among many others, Bloom (2009), Jurado, Ludvigson and Ng (2015), Mumtaz and Theodoridis (2018)). Focusing on the reactions of stock markets (DJE50), there is some limited evidence that equity prices decline. If interest rates are stuck at the ZLB, these declines are somewhat more pronounced and appear to peak slightly later.

Considering the dynamic responses of short-term interest rates (EUR3M) reveal that, after a month, short rates decrease appreciably. This decline in interest rates peaks after about two months and reaches around 100 basis points. In economic terms the negative reaction of short-term interest rates is likely induced by expansionary monetary policy measures undertaken by

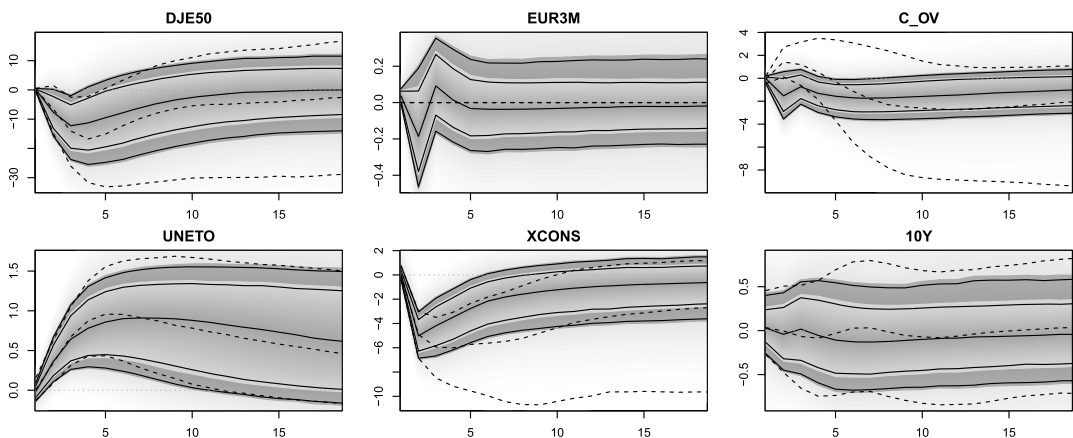


FIG. 5. *Dynamic responses to a one-standard deviation uncertainty shock. The dashed gray lines represent 16th and 84th posterior credible intervals.*

the central bank. Notice that, in the presence of the ZLB, we assume no interest rate reaction at the short-end of the yield curve.

Turning to inflation reactions ( $C_{OV}$ ) suggest that prices tend to fall after about six months. This decline in inflation is consistent with a negative demand channel which implies that firms lower prices in response to a decline in aggregate demand (see Bloom (2009)). When the economy is stuck at the ZLB, we observe somewhat stronger but insignificant inflation reactions. These responses are consistent with Caggiano, Castelnuovo and Pellegrino (2017) who also find more pronounced but insignificant inflation responses if short rates hit the ZLB.

The unemployment rate (UNETO), with a lag of around two to three months, increases and remains elevated for several months before turning insignificant. This increase in the unemployment rate is consistent with Jurado, Ludvigson and Ng (2015) and Carriero, Clark and Marcellino (2018) who, for U.S. data, find similar unemployment responses. One interesting finding is that the unemployment responses in the presence of the ZLB are similar in magnitude but tend to peter out faster and thus are more short-lived than the ones if interest rates are allowed to react freely. It is noteworthy that the impulse responses are very similar over the first four to five months and then depart from each other.

When we consider industrial production (XCONS), we observe a sluggish decline which peaks at around minus five percent in the second month and then, after around six months the reactions of real activity turn insignificant. Interestingly, we do not observe a rebound in real activity arising from a “wait-and-see” mechanism reported in, for example, Bloom (2009). Gieseck and Largent (2016), for Eurozone data, find similar responses for GDP which are also rather short-lived and do not display a substantial real activity overshoot. When we assume that the ZLB is binding, the reactions of industrial production become more pronounced, which is consistent with other findings who report that, if interest rates hit zero, real activity reacts stronger to uncertainty shocks (Caggiano, Castelnuovo and Nodari (2021)). Finally, long-term interest rates display no statistically significant reaction throughout the impulse response horizon.

To sum up, a key take away from this exercise is that our BAVART model is capable of producing meaningful impulse responses which are consistent with the literature. Without introducing any assumptions on the specific form of nonlinearities but in the presence of the ZLB, our BAVART specification yields impulse responses that are consistent with papers that assess how the uncertainty, real activity nexus, changes if the central bank is constrained by the ZLB.

**6. Conclusions.** VAR models assume that the lagged dependent variables influence the contemporaneous values in a linear fashion. In this paper we relax this assumption by blending the literature on BART models and VARs. The resulting BAVART model can handle arbitrary nonlinear relations between the endogenous and the exogenous variables. Our proposed model is, moreover, capable of handling stochastic volatility in the shocks. As opposed to existing models, which make strong assumptions on the nature of nonlinearities, our model remains agnostic and allows to estimate these forms in a data-based manner. To make the model operational, we briefly discuss Bayesian estimation but also show how to compute multistep-ahead forecasts and generalized impulse responses.

We illustrate our approach using two topical applications. In the first application we apply the model to the U.S. term structure of interest rates. Using several linear and nonlinear competing models and different ways of modeling the yield curve, we show that our BAVART model yields precise point and density forecasts. The point forecasting accuracy differences become larger with the maturity of a given bond. An opposite picture emerges when the full predictive distribution is used: in that case, forecast gains can be mostly found at the short-end of the yield curve.



The second application deals with the effects of uncertainty on the Eurozone economy. To investigate the role of the ZLB on interest rates, we consider impulse responses which restrict interest rate reactions to zero and compare these to their unrestricted counterpart. The findings indicate that uncertainty has a detrimental effect on macroeconomic outcomes. Unemployment increases, prices fall, stock markets decline and industrial production drops markedly. In general, these reactions are similar to other findings in the literature which mainly focus on U.S. data. When we assume that interest rates are stuck at the ZLB, real activity responses become somewhat more elevated. Especially for output, we observe much stronger responses if the central bank is not able to react adequately.

There are many possible avenues for further research and possible applications of our model. For instance, the model can be used to track asymmetries in the transmission of economic shocks. Or it could be applied to high-frequency financial data, such as daily stock returns and then combined with a heavy-tailed error distribution to produce precise density forecasts. From an econometric perspective, BART could be used to model time-variation in regression coefficients and thus generalize TVP regressions which assume a random walk evolution on the latent states.

## APPENDIX: POSTERIOR APPROXIMATION OF THE TREES

In this appendix, we provide details on the MCMC algorithm used to simulate from the joint posterior distribution of the trees and the terminal node parameters. The remaining quantities take well known forms and are very easy to simulate using textbook results for the Gaussian linear regression model.

As stated in Section 3.3 and following [Carriero, Clark and Marcellino \(2019\)](#), [Koop, Korobilis and Pettenuzzo \(2019\)](#) and [Kastner and Huber \(2020\)](#), we rely on a structural representation of the model that entails equation-by-equation estimation.

We simulate from the conditional posterior of  $(\mathcal{T}_{jk}, \mathbf{m}_{jk})$  by conditioning on all trees except the  $k$ th one. This is achieved by using the Bayesian backfitting strategy discussed in [Chipman, George and McCulloch \(2010\)](#). More precisely, the likelihood function depends on  $(\mathcal{T}_{jk}, \mathbf{m}_{jk})$  through the partial residuals

$$\mathbf{R}_{jk} = \mathbf{y}_{\bullet j} - \sum_{l=1}^{j-1} a_{jl} \mathbf{e}_{\bullet l} - \sum_{s \neq k} g_{js}(\mathbf{X} | \mathcal{T}_{js}, \mathbf{m}_{js}).$$

The algorithm proposed in [Chipman, George and McCulloch \(2010\)](#) then proceeds by integrating out  $\mathbf{m}_{jk}$  and then drawing  $\mathcal{T}_{jk}$  with Metropolis–Hastings (MH) algorithm, detailed in [Chipman, George and McCulloch \(1998\)](#).

This step is implemented by generating a candidate tree  $\mathcal{T}_{jk}^*$  from a proposal distribution  $q(\mathcal{T}_{jk}, \mathcal{T}_{jk}^*)$ , and then accept the proposed value with probability

$$\alpha(\mathcal{T}_{jk}, \mathcal{T}_{jk}^*) = \min \left\{ 1, \frac{q(\mathcal{T}_{jk}^*, \mathcal{T}_{jk}) p(\mathbf{R}_{jk} | \mathbf{X}, \mathcal{T}_{jk}^*, \mathbf{M}_j) p(\mathcal{T}_{jk}^*)}{q(\mathcal{T}_{jk}, \mathcal{T}_{jk}^*) p(\mathbf{R}_{jk} | \mathbf{X}, \mathcal{T}_{jk}, \mathbf{M}_j) p(\mathcal{T}_{jk})} \right\}.$$

The first term is the ratio of the probability of how the previous tree moves to the new tree against the probability of how the new tree moves to the previous one. The second term is the likelihood ratio of the new tree against the previous tree, and the last term denotes the likelihood ratio of the new against the previous tree under the prior distribution.

The proposal distribution  $q(\mathcal{T}_{jk}, \mathcal{T}_{jk}^*)$  features four local steps ([Chipman, George and McCulloch \(1998\)](#)). The first step grows the tree by splitting a node into two different nodes. This step is chosen with probability 0.25. The second step combines two nonterminal nodes into a single terminal node. We select this step with probability 0.25 as well. The third step

swaps splitting rules between two terminal nodes (with probability 0.4), and the final step changes the splitting rule of a single nonterminal node (with probability 0.1).

After sampling the tree structure we can easily simulate from the posterior distribution of the terminal nodes. Since the prior is conjugate, the resulting posterior will also be a Gaussian distribution that takes a standard form and does not explicitly depend on the tree structure but only on the subset of elements in the residuals allocated to a specific terminal node.

**Acknowledgments.** We would like to thank two anonymous reviewers, an Associate Editor as well as the Handling Editor for numerous comments that greatly improved the paper.

**Funding.** Florian Huber gratefully acknowledges funding from the Austrian Science Fund (FWF): ZK 35.

## REFERENCES

- AASTVEIT, K. A., NATVIK, G. J. and SOLA, S. (2017). Economic uncertainty and the influence of monetary policy. *J. Int. Money Financ.* **76** 50–67.
- ALESSANDRI, P. and MUMTAZ, H. (2017). Financial conditions and density forecasts for US output and inflation. *Rev. Econ. Dyn.* **24** 66–78.
- ALESSANDRI, P. and MUMTAZ, H. (2019). Financial regimes and uncertainty shocks. *J. Monet. Econ.* **101** 31–46.
- AUERBACH, A. J. and GORODNICHENKO, Y. (2012). Measuring the output responses to fiscal policy. *Am. Econ. J. Econ. Policy* **4** 1–27.
- BARNICHON, R. and MATTHES, C. (2018). Functional approximation of impulse responses. *J. Monet. Econ.* **99** 41–55.
- BASSETTI, F., CASARIN, R. and LEISEN, F. (2014). Beta-product dependent Pitman–Yor processes for Bayesian inference. *J. Econometrics* **180** 49–72. MR3188911 <https://doi.org/10.1016/j.jeconom.2014.01.007>
- BILLIO, M., CASARIN, R. and ROSSINI, L. (2019). Bayesian nonparametric sparse VAR models. *J. Econometrics* **212** 97–115. MR3994009 <https://doi.org/10.1016/j.jeconom.2019.04.022>
- BLOOM, N. (2009). The impact of uncertainty shocks. *Econometrica* **77** 623–685. MR2531358 <https://doi.org/10.3982/ECTA6248>
- BREIMAN, L. (2001). Random forests. *Mach. Learn.* **45** 5–32.
- CAGGIANO, G., CASTELNUOVO, E. and GROSHENNY, N. (2014). Uncertainty shocks and unemployment dynamics in US recessions. *J. Monet. Econ.* **67** 78–92.
- CAGGIANO, G., CASTELNUOVO, E. and NODARI, G. (2021). Uncertainty and monetary policy in good and bad times. *J. Appl. Econometrics*. To appear.
- CAGGIANO, G., CASTELNUOVO, E. and PELLEGRINO, G. (2017). Estimating the real effects of uncertainty shocks at the zero lower bound. *Eur. Econ. Rev.* **100** 257–272.
- CARRIERO, A., CLARK, T. E. and MARCELLINO, M. (2018). Measuring uncertainty and its impact on the economy. *Rev. Econ. Stat.* **100** 799–815.
- CARRIERO, A., CLARK, T. E. and MARCELLINO, M. (2019). Large Bayesian vector autoregressions with stochastic volatility and non-conjugate priors. *J. Econometrics* **212** 137–154. MR3994011 <https://doi.org/10.1016/j.jeconom.2019.04.024>
- CARRIERO, A., KAPETANIOS, G. and MARCELLINO, M. (2012). Forecasting government bond yields with large Bayesian vector autoregressions. *J. Bank. Financ.* **36** 2026–2047.
- CARRIERO, A., CHAN, J. C. C., CLARK, T. E. and MARCELLINO, M. (2021). Corrigendum to: Large Bayesian vector autoregressions with stochastic volatility and non-conjugate priors. Working paper.
- CARVALHO, C. M., POLSON, N. G. and SCOTT, J. G. (2009). Handling sparsity via the horseshoe. In *Artificial Intelligence and Statistics* 73–80. PMLR.
- CHIPMAN, H. A., GEORGE, E. I. and MCCULLOCH, R. E. (1998). Bayesian CART model search. *J. Amer. Statist. Assoc.* **93** 935–948.
- CHIPMAN, H. A., GEORGE, E. I. and MCCULLOCH, R. E. (2002). Bayesian treed models. *Mach. Learn.* **48** 299–320.
- CHIPMAN, H. A., GEORGE, E. I. and MCCULLOCH, R. E. (2010). BART: Bayesian additive regression trees. *Ann. Appl. Stat.* **4** 266–298. MR2758172 <https://doi.org/10.1214/09-AOAS285>
- CLARK, T. E. (2011). Real-time density forecasts from Bayesian vector autoregressions with stochastic volatility. *J. Bus. Econom. Statist.* **29** 327–341. MR2848507 <https://doi.org/10.1198/jbes.2010.09248>

- CLARK, T. E. and RAVAZZOLO, F. (2015). Macroeconomic forecasting performance under alternative specifications of time-varying volatility. *J. Appl. Econometrics* **30** 551–575. MR3358636 <https://doi.org/10.1002/jae.2379>
- CRESPO CUARESMA, J. C., HUBER, F. and ONORANTE, L. (2020). Fragility and the effect of international uncertainty shocks. *J. Int. Money Financ.* **108** 102151.
- DIEBOLD, F. X. and LI, C. (2006). Forecasting the term structure of government bond yields. *J. Econometrics* **130** 337–364. MR2211798 <https://doi.org/10.1016/j.jeconom.2005.03.005>
- DOAN, T., LITTERMAN, R. and SIMS, C. (1984). Forecasting and conditional projection using realistic prior distributions. *Econometric Rev.* **3** 1–100.
- DORIE, V., HILL, J., SHALIT, U., SCOTT, M. and CERVONE, D. (2019). Automated versus do-it-yourself methods for causal inference: Lessons learned from a data analysis competition. *Statist. Sci.* **34** 43–68. MR3938963 <https://doi.org/10.1214/18-STS667>
- FERRARA, L. and GUÉRIN, P. (2018). What are the macroeconomic effects of high-frequency uncertainty shocks? *J. Appl. Econometrics* **33** 662–679. MR3850753 <https://doi.org/10.1002/jae.2624>
- FREUND, Y. and SCHAPIRE, R. E. (1997). A decision-theoretic generalization of on-line learning and an application to boosting. *J. Comput. System Sci.* **55** 119–139. MR1473055 <https://doi.org/10.1006/jcss.1997.1504>
- FRIEDMAN, J. H. (2001). Greedy function approximation: A gradient boosting machine. *Ann. Statist.* **29** 1189–1232. MR1873328 <https://doi.org/10.1214/aos/1013203451>
- GEFANG, D. and STRACHAN, R. (2010). Nonlinear impacts of international business cycles on the U.K.—a Bayesian smooth transition VAR approach. *Stud. Nonlinear Dyn. Econom.* **14** Art. 2, 33. MR2585173 <https://doi.org/10.2202/1558-3708.1677>
- GIESECK, A. and LARGENT, Y. (2016). The impact of macroeconomic uncertainty on activity in the euro area. *Rev. Econ.* **67** 25–52.
- GRAMACY, R. B. and LEE, H. K. H. (2008). Bayesian treed Gaussian process models with an application to computer modeling. *J. Amer. Statist. Assoc.* **103** 1119–1130. MR2528830 <https://doi.org/10.1198/016214508000000689>
- GRANGER, C. W. and TERASVIRTA, T. (1993). Modelling non-linear economic relationships. OUP Catalogue.
- GREEN, D. P. and KERN, H. L. (2012). Modeling heterogeneous treatment effects in survey experiments with Bayesian additive regression trees. *Public Opin. Q.* **76** 491–511.
- GÜRKAYNAK, R. S., SACK, B. and WRIGHT, J. H. (2007). The US treasury yield curve: 1961 to the present. *J. Monet. Econ.* **54** 2291–2304.
- HAHN, P. R., MURRAY, J. S. and CARVALHO, C. M. (2020). Bayesian regression tree models for causal inference: Regularization, confounding, and heterogeneous effects (with discussion). *Bayesian Anal.* **15** 965–1056. Includes comments and discussions by 25 discussants and a rejoinder by the authors. MR4154846 <https://doi.org/10.1214/19-BA1195>
- HE, J., YALOV, S. and HAHN, P. R. (2019). XBART: Accelerated Bayesian additive regression trees. In *The 22nd International Conference on Artificial Intelligence and Statistics* 1130–1138. PMLR.
- HERNÁNDEZ, B., RAFTERY, A. E., PENNINGTON, S. R. and PARNELL, A. C. (2018). Bayesian additive regression trees using Bayesian model averaging. *Stat. Comput.* **28** 869–890. MR3766048 <https://doi.org/10.1007/s11222-017-9767-1>
- HILL, J. L. (2011). Bayesian nonparametric modeling for causal inference. *J. Comput. Graph. Statist.* **20** 217–240. Supplementary material available online. MR2816546 <https://doi.org/10.1198/jcgs.2010.08162>
- HILL, J., LINERO, A. and MURRAY, J. (2020). Bayesian additive regression trees: A review and look forward. *Annu. Rev. Stat. Appl.* **7** 251–278. MR4104193 <https://doi.org/10.1146/annurev-statistics-031219-041110>
- HJORT, N. L., HOLMES, C., MÜLLER, P. and WALKER, S. G. (2010). *Bayesian Nonparametrics*. Cambridge Univ. Press, Cambridge.
- HUBER, F., KOOP, G. and ONORANTE, L. (2021). Inducing sparsity and shrinkage in time-varying parameter models. *J. Bus. Econom. Statist.* **39** 669–683. MR4272927 <https://doi.org/10.1080/07350015.2020.1713796>
- HUBER, F. and ZÖRNER, T. O. (2019). Threshold cointegration in international exchange rates: A Bayesian approach. *Int. J. Forecast.* **35** 458–473.
- HUBER, F., KOOP, G., ONORANTE, L., PFARRHOFER, M. and SCHREINER, J. (2021). Nowcasting in a pandemic using non-parametric mixed frequency VARs. *J. Econometrics*. To appear.
- JACKSON, L. E., KLIESEN, K. L. and OWYANG, M. T. (2020). The nonlinear effects of uncertainty shocks. *Stud. Nonlinear Dyn. Econom.* **24** 20190024, 19. MR4145962 <https://doi.org/10.1515/snde-2019-0024>
- JURADO, K., LUDVIGSON, S. C. and NG, S. (2015). Measuring uncertainty. *Am. Econ. Rev.* **105** 1177–1216.
- KALLI, M. and GRIFFIN, J. E. (2018). Bayesian nonparametric vector autoregressive models. *J. Econometrics* **203** 267–282. MR3770826 <https://doi.org/10.1016/j.jeconom.2017.11.009>
- KAPELNER, A. and BLEICH, J. (2015). Prediction with missing data via Bayesian additive regression trees. *Canad. J. Statist.* **43** 224–239. MR3353381 <https://doi.org/10.1002/cjs.11248>

- KASTNER, G. and FRÜHWIRTH-SCHNATTER, S. (2014). Ancillarity-sufficiency interweaving strategy (ASIS) for boosting MCMC estimation of stochastic volatility models. *Comput. Statist. Data Anal.* **76** 408–423. MR3209449 <https://doi.org/10.1016/j.csda.2013.01.002>
- KASTNER, G. and HUBER, F. (2020). Sparse Bayesian vector autoregressions in huge dimensions. *J. Forecast.* **39** 1142–1165. MR4161021 <https://doi.org/10.1002/for.2680>
- KERN, H. L., STUART, E. A., HILL, J. and GREEN, D. P. (2016). Assessing methods for generalizing experimental impact estimates to target populations. *J. Res. Educ. Eff.* **9** 103–127. <https://doi.org/10.1080/19345747.2015.1060282>
- KOOP, G., KOROBILIS, D. and PETTENUZZO, D. (2019). Bayesian compressed vector autoregressions. *J. Econometrics* **210** 135–154. MR3944767 <https://doi.org/10.1016/j.jeconom.2018.11.009>
- KOOP, G., PESARAN, M. H. and POTTER, S. M. (1996). Impulse response analysis in nonlinear multivariate models. *J. Econometrics* **74** 119–147. MR1409037 [https://doi.org/10.1016/0304-4076\(95\)01753-4](https://doi.org/10.1016/0304-4076(95)01753-4)
- KRUEGER, R., BANSAL, P. and BUDDHAVARAPU, P. (2020). A new spatial count data model with Bayesian additive regression trees for accident hot spot identification. *Accident Anal. Prev.* **144** 105623. <https://doi.org/10.1016/j.aap.2020.105623>
- LAKSHMINARAYANAN, B., ROY, D. M. and TEH, Y. W. (2014). Mondrian forests: Efficient online random forests. In *Advances in Neural Information Processing Systems* (Z. Ghahramani, M. Welling, C. Cortes, N. D. Lawrence and K. Q. Weinberger, eds.) **27** 3140–3148. Curran Associates.
- LINERO, A. R. (2018). Bayesian regression trees for high-dimensional prediction and variable selection. *J. Amer. Statist. Assoc.* **113** 626–636. MR3832214 <https://doi.org/10.1080/01621459.2016.1264957>
- LITTERMAN, R. B. (1986). Forecasting with Bayesian vector autoregressions—five years of experience. *J. Bus. Econom. Statist.* **4** 25–38.
- MAKALIC, E. and SCHMIDT, D. F. (2015). A simple sampler for the horseshoe estimator. *IEEE Signal Process. Lett.* **23** 179–182.
- MUMTAZ, H. and THEODORIDIS, K. (2018). The changing transmission of uncertainty shocks in the U.S. *J. Bus. Econom. Statist.* **36** 239–252. MR3790211 <https://doi.org/10.1080/07350015.2016.1147357>
- NELSON, C. R. and SIEGEL, A. F. (1987). Parsimonious modeling of yield curves. *J. Bus.* **60** 473–489.
- PACCAGINI, A. and COLOMBO, V. (2020). The asymmetric effect of uncertainty shocks. CAMA Working Paper, 72/2020.
- PLAGBORG-MØLLER, M. (2019). Bayesian inference on structural impulse response functions. *Quant. Econ.* **10** 145–184. MR3915253 <https://doi.org/10.3982/QE926>
- PRATOLA, M. T., CHIPMAN, H. A., GEORGE, E. I. and MCCULLOCH, R. E. (2020). Heteroscedastic BART via multiplicative regression trees. *J. Comput. Graph. Statist.* **29** 405–417. MR4116052 <https://doi.org/10.1080/10618600.2019.1677243>
- PRÜSER, J. (2019). Forecasting with many predictors using Bayesian additive regression trees. *J. Forecast.* **38** 621–631. MR4021441 <https://doi.org/10.1002/for.2587>
- RAMEY, V. A. and ZUBAIRY, S. (2017). Government spending multipliers in good times and in bad: Evidence from us historical data. *J. Polit. Econ.* **126** 850–901.
- ROY, D. M. and TEH, Y. W. (2009). The mondrian process. In *Advances in Neural Information Processing Systems* (D. Koller, D. Schuurmans, Y. Bengio and L. Bottou, eds.) **21** 1377–1384. Curran Associates.
- SIMS, C. A. (1980). Macroeconomics and reality. *Econometrica* **48** 1–48.
- SIMS, C. A. and ZHA, T. (1998). Bayesian methods for dynamic multivariate models. *Internat. Econom. Rev.* **39** 949–968.
- WALDMANN, P. (2016). Genome-wide prediction using Bayesian additive regression trees. *Genet. Sel. Evol.* **48** 1–12.