Bocconi

unibocconi.it

WORKING PAPER
N. 224
JUNE 2024

Bayesian nonparametric methods for macroeconomic forecasting

Massimiliano Marcellino

Michael Pfarrhofer

Università Bocconi
BAFFI
Centre on Economics, Finance and Regulation

Via Röntgen 1 | 20136 Milano – Italia | Tel +39 02 5836.5908/5564
baffi@unibocconi.it | www.baffi.unibocconi.eu

Bayesian nonparametric methods for macroeconomic forecasting*

Massimiliano MARCELLINO[1] and Michael PFARRHOFER[2]

[1] *Bocconi University, BAFFI, IGIER* and *CEPR*
ORCID: 0000-0002-8737-2768
[2] *Vienna University of Economics and Business*
ORCID: 0000-0002-0168-688X

We review specification and estimation of multivariate Bayesian nonparametric models for forecasting (possibly large sets of) macroeconomic and financial variables. The focus is on Bayesian Additive Regression Trees and Gaussian Processes. We then apply various versions of these models for point, density and tail forecasting using datasets for the euro area and the US. The performance is compared with that of several variants of Bayesian VARs to assess the relevance of accounting for general forms of nonlinearities. We find that medium-scale linear VARs with stochastic volatility are tough benchmarks to beat. Some gains in predictive accuracy arise for nonparametric approaches, most notably for short-run forecasts of unemployment and longer-run predictions of inflation, and during recessionary or otherwise non-standard economic episodes.

JEL: C11, C32, C53
Keywords: United States, euro area, Bayesian Additive Regression Trees, Gaussian Processes, multivariate time series analysis, structural breaks

*We thank Marta Bańbura, Dario Caldara, Matteo Ciccarelli, Pablo Guerrón-Quintana and participants at the 2023 ECB Forecasting Conference for valuable comments and suggestions, and Marta Bańbura and Joan Paredes for providing us with the euro area real-time data. Marcellino acknowledges financial support from MUR- Prin 2022- Prot. 20227YZ9JK, financed by the European Union â Next Generation EU. Pfarrhofer acknowledges financial support from the Jubiläumsfonds of the Oesterreichische Nationalbank (OeNB, grant 18765).

1. INTRODUCTION

Predicting the future is notoriously difficult, even more so during crises when the realizations of economic variables are far from their average. In fact, econometric models are typically better at explaining and predicting values close to the average, particularly so in the case of linear models. Yet, reliable forecasts are particularly valuable during and after crisis time. Hence, the econometric literature has proposed a variety of sophisticated nonlinear/time-varying models. Examples include threshold and smooth transition models (see, e.g., Tong, 1990; Teräsvirta, 1994), Markov switching (MS, see, e.g., Hamilton, 1989) models, and double stochastic models (see, e.g., Nyblom, 1989).[1] While these are not without perils, in particular when used with short samples, specifications in this spirit can yield some gains, see e.g., Ferrara *et al.* (2015).

The situation got even worse during the Covid-19 period, with many macroeconomic time series exhibiting unprecedented shifts. Even sophisticated and typically well performing econometric models such as Bayesian vector autoregressions (BVARs) with stochastic volatility (SV), and possibly time varying parameters (TVPs), had troubles in tracking the tremendous fall and subsequent rebound in real activity and labor market indicators (see Carriero *et al.*, 202X, for the US). This contributed to the spread of an additional level of sophistication in the specification of econometric models for forecasting, with the adoption of a variety of machine learning (ML) or ML-inspired methods.

Examples include Goulet Coulombe *et al.* (2021) who assess whether and to what extent classical ML models can improve the forecasts, both in general and specifically during the onset of the Covid-19 pandemic. They focus on the UK economy that at the same time was also experiencing Brexit-related uncertainty. Huber *et al.* (2023) consider nowcasting GDP in the largest euro area (EA) economies, by combining a mixed frequency VAR with Bayesian Additive Regression Trees (BART). This extended the work of Huber and Rossini (2022) which was in turn inspired by Chipman *et al.* (2010). Further generalizations in this spirit, and applications to forecasting US variables are provided in Clark *et al.* (2023), which has a special focus on tail events. Indeed, regression trees (and more generally, random forests, see Breiman, 2001) are able to quickly adapt to extreme observations and to disentangle switches in the underlying regimes. They forecast well a variety of economic variables, see Masini *et al.* (2023) for a survey and Medeiros *et al.* (2021) for an application to forecasting US inflation.

[1] Primiceri (2005) popularized the use of time-varying parameters and stochastic volatility in macroeconometrics.

Another ML-inspired method is Gaussian process (GP) regression (see, Williams and Rasmussen, 2006, for an overview). Borrowing ideas from the literature on Bayesian linear VARs, Hauzenberger *et al.* (2021) assume, for each endogenous variable, a different nonlinear relationship with its own lags and with the lags of all the other variables. This particular approach can be viewed as a nonparametric alternative to Minnesota-type shrinkage. GPs are used to model nonlinearities in a flexible way, and have been found to yield some gains when forecasting US variables in an economic context. They are also related to neural networks, as they are universal approximators based on infinite mixtures of Gaussian distributions.[2]

Indeed another class of promising ML forecasting models are (classical or Bayesian) neural networks (NN), see, e.g., Hornik *et al.* (1989), Gu *et al.* (2021), Goulet Coulombe (2022), Hauzenberger *et al.* (2022).[3] Yet, NN are fairly different from the VAR models that are the workhorse of modern time series econometrics. Instead, BART and GP can be rather easily and efficiently adapted to provide a nonlinear and nonparametric extension of VAR models, see Clark *et al.* (2023) and Hauzenberger *et al.* (2021). Iterated multi-step ahead forecasts can be also computed, while the direct approach is typically used in single equation NN models. For these reasons, and for their good forecasting performance for economic and financial variables in previous studies, we focus on BART and GP models in this chapter.

A point worth making is that the flexible modeling of the conditional mean in BART and GP based specifications could make the error variances more stable than in linear models. On the other hand, the nonlinearities captured by BART and GP could be mainly due to large realizations of the errors and/or outlying observations. For a discussion of discriminating signal from noise during Covid-19 in BVARs estimated on US data, see Carriero *et al.* (202X). Hence, we also assess the role of assuming heteroskedastic features about the errors. More specifically, to address potentially time-varying variances in a flexible yet parsimonious manner, and to simplify estimation of very large nonparametric models, we introduce a factor stochastic volatility (FSV) structure on the multivariate errors. This makes equation-by-equation estimation feasible — and thus enables modeling a moderate to large number of variables jointly.

Empirically, overall we find improvements in predictive accuracy for BART and GPs. Heteroskedasticity does not appear to be as important as in the linear case for the nonparametric

[2] Specific choices of the kernels underlying Gaussian processes can produce a variety of neural network models, see Novak *et al.* (2018) for details.

[3] Further contributions to nonlinear and nonparametric times series analysis also make use of infinite mixtures, see Hirano (2002), Kalli and Griffin (2018), Billio *et al.* (2019), Jin *et al.* (2022).

implementations, and the size of the information set causes some differences in predictive performance. Both BART and GP perform particularly well for short-run forecasts of unemployment, and longer-run predictions of inflation. Some additional gains are masked when considering only the average performance over the full holdout sample; there are several periods when nonlinearities appear to be particularly useful. In fact, they help during recessions in some important cases such as downside-risk to GDP and upside-risk to unemployment, and the nonparametrics do well in predicting the recent surge of inflation.

The chapter is structured as follows. In Section 2 we describe model specification and estimation, and relate our discussions to nested simpler variants such as the BVAR. Section 3 illustrates the approaches by means of simple univariate examples using US data for unemployment, output and inflation. In Section 4 we apply the models for point, density and tail forecasting EA and US variables. Section 5 summarizes salient points and concludes. Technical details are gathered in an Appendix and additional empirical results are provided online as supplementary material.

2. MULTIVARIATE NONPARAMETRIC MODELS

Let $\boldsymbol{y}_t = (y_{1t}, \ldots, y_{nt})'$ be a standardized $n \times 1$ vector of endogenous variables, and $\boldsymbol{x}_t = (\boldsymbol{y}'_{t-1}, \ldots, \boldsymbol{y}'_{t-p})'$ the $k \times 1$ vector of lagged endogenous variables with $k = np$. We work with general multivariate models of the form:

$$\boldsymbol{y}_t = F(\boldsymbol{x}_t) + \boldsymbol{\epsilon}_t, \tag{1}$$

where $F(\boldsymbol{x}_t) = (f_1(\boldsymbol{x}_t), \ldots, f_n(\boldsymbol{x}_t))'$ collects a set of (possibly unknown) conditional mean functions $f_i(\boldsymbol{x}_t) : \mathbb{R}^k \to \mathbb{R}$ for $i = 1, \ldots, n$, such that $F(\boldsymbol{x}_t) : \mathbb{R}^k \to \mathbb{R}^n$. We rely on methods that allow for order-invariant estimation of the multivariate model equation-by-equation, i.e., in most cases, we assume the $f_i(\boldsymbol{x}_t)$'s to be conditionally independent across i. This provides substantial computational gains and flexibility.

The competing model specifications will be distinguished with respect to how we treat the function $F(\boldsymbol{x}_t)$ and what we assume about the error term $\boldsymbol{\epsilon}_t = (\epsilon_{1t}, \ldots, \epsilon_{nt})'$. In fact, the reduced form errors reflect the unpredictable component within our model framework, and assumptions about their correlation structure play a key role with respect to efficient estimation of multivariate nonparametric models.

2.1. *Conditionally independent equations*

Our implementation of the nonparametric conditional mean functions is based on imposing structure equation-by-equation. So we first will discuss how to transform the VAR into a system of (order-invariant) independent equations, conditional on auxiliary variables. In the spirit of FSV models, we decompose the reduced form errors as:

$$\boldsymbol{\epsilon}_t = \boldsymbol{L}\mathfrak{F}_t + \boldsymbol{\eta}_t, \tag{2}$$

where, $\mathfrak{F}_t = (\mathfrak{f}_{1t}, \dots, \mathfrak{f}_{qt})'$ is a $q \times 1$ vector of common static factors linked to the observed residuals via the $n \times q$ loadings matrix $\boldsymbol{L}$, and $\boldsymbol{\eta}_t = (\eta_{1t}, \dots, \eta_{nt})'$ is an $n \times 1$-vector of idiosyncratic shocks. In our baseline specification, we assume heteroskedastic factors and idiosyncratic innovations:

$$\mathfrak{F}_t \sim \mathcal{N}(\mathbf{0}_q, \boldsymbol{\Omega}_t), \quad \boldsymbol{\eta}_t \sim \mathcal{N}(\mathbf{0}_n, \boldsymbol{H}_t),$$
$$\boldsymbol{\Omega}_t = \text{diag}\left(\exp(\omega_{1t}), \dots, \exp(\omega_{qt})\right), \quad \boldsymbol{H}_t = \text{diag}\left(\exp(h_{1t}), \dots, \exp(h_{nt})\right).$$

The covariance matrix is thus given by:

$$\text{Var}(\boldsymbol{\epsilon}_t) \equiv \boldsymbol{\Sigma}_t = \boldsymbol{L}\boldsymbol{\Omega}_t\boldsymbol{L}' + \boldsymbol{H}_t,$$

where cross-variable correlations are encoded exclusively in $\boldsymbol{L}$, that is, joint dynamics are driven by a small number of $q \ll n$ shocks; see Kastner and Huber (2020) for an early implementation of FSV in the VAR context.[4] We will assign independent Gaussian priors on the elements of the loadings matrix, and equip these with a global-local shrinkage prior.

To see why this is useful from an algorithmic perspective, notice that $\tilde{\boldsymbol{y}}_t \equiv \boldsymbol{y}_t - \boldsymbol{L}\mathfrak{F}_t = F(\boldsymbol{x}_t) + \boldsymbol{\eta}_t$, where $\boldsymbol{\eta}_t$ are independent across equations due to the diagonal structure of $\boldsymbol{H}_t$. For

[4] We always use a minimum number of $q = 2$ factors irrespective of the size of the information set. Overfitting factor models carry less of a penalty than underfitting ones, see Chan (2023). We thus set the number of factors to the Ledermann bound, which is the largest positive solution q^* of the equation $(n - q^*)^2 \geq n + q^*$ and use a shrinkage prior to regularize the parameter space to alleviate overfitting concerns.

future reference, we also define a $T \times 1$-vector $\boldsymbol{h}_i = (\exp(h_{i1}), \ldots, \exp(h_{iT}))'$, which comprises the variance process for equations $i = 1, \ldots, n$, stacked over time $t = 1, \ldots, T$.

2.2. *Modeling the conditional mean*

Constant and time-varying parameter vector autoregressions

For the constant and TVP case, we simply assume that $F(\boldsymbol{x}_t) = \boldsymbol{A}_t \boldsymbol{x}_t$ is known to be a linear function (conditional on t), such that Eq. (1) turns into:

$$\boldsymbol{y}_t = \boldsymbol{A}_t \boldsymbol{x}_t + \boldsymbol{\epsilon}_t,$$

where $\boldsymbol{A}_t$ is an $n \times k$ matrix of time-varying autoregressive VAR coefficients. This specification is referred to as the TVP-VAR, as popularized by Primiceri (2005). The BVAR is obtained by assuming $\boldsymbol{A}_t = \boldsymbol{A}$ to be the same for all $t = 1, \ldots, T$, i.e., a constant set of VAR coefficients. Intercept terms, time trends, latent/observed factors or additional exogenous covariates can be introduced by augmenting the vector $\boldsymbol{x}_t$ accordingly. For instance, in our empirical work we always include an intercept term for the BVAR and TVP-VAR.

Priors. The linear model may be estimated via Bayesian methods by specifying a Gaussian prior on each coefficient of the matrix $\boldsymbol{A}$. A popular choice in this case is to regularize the parameter space via shrinkage priors, such as the Minnesota prior or global-local priors. We opt for the latter, and assume a horseshoe prior (HS, see Carvalho *et al.*, 2010, and Appendix A) on these coefficients.[5]

For the TVP case, we stack the coefficients column-wise in a vector $\boldsymbol{a}_t = \text{vec}(\boldsymbol{A}_t')$ and assume a random walk state equation:

$$\boldsymbol{a}_t = \boldsymbol{a}_{t-1} + \boldsymbol{\Theta}^{1/2} \boldsymbol{\varepsilon}_t, \quad \boldsymbol{\varepsilon}_t \sim \mathcal{N}(\boldsymbol{0}_{nk}, \boldsymbol{I}_{nk}), \quad \boldsymbol{\Theta}^{1/2} = \text{diag}\left(\theta_1^{1/2}, \ldots, \theta_{nk}^{1/2}\right).$$

To impose shrinkage, we rewrite this model in its non-centered parameterization which moves the square root of the state innovation variances in the diagonal matrix $\boldsymbol{\Theta}$ into the measurement equation (for details, see Frühwirth-Schnatter and Wagner, 2010; Bitto and Frühwirth-Schnatter, 2019). Let $\boldsymbol{a}_{it} = \boldsymbol{A}'_{[i\bullet]t}$ collect the parameters of the ith VAR equation and $\boldsymbol{\theta}_i^{1/2}$ is the

[5] Any prior which is Gaussian at the lowest level of the hierarchy can be (and many have been) used in this context, and we pick the horseshoe for its solid performance in many applications. For an overview of global-local shrinkage priors, see Cadonna *et al.* (2020); an empirical comparison for VARs is provided by Cross *et al.* (2020).

k-vector subsetting the associated state innovations on the diagonal of $\boldsymbol{\Theta}^{1/2}$. Using the transformation $\boldsymbol{a}_{it} = \boldsymbol{a}_i + \text{diag}(\boldsymbol{\theta}_i^{1/2})\tilde{\boldsymbol{a}}_{it}$, which splits the TVPs into a constant initial state and time-varying part, we may write:

$$y_{it} = \boldsymbol{x}_t' \left(\boldsymbol{a}_i + \text{diag}\left(\boldsymbol{\theta}_i^{1/2}\right)\tilde{\boldsymbol{a}}_{it}\right) + \epsilon_{it}, \quad \tilde{\boldsymbol{a}}_{it} \sim \mathcal{N}(\tilde{\boldsymbol{a}}_{it-1}, \boldsymbol{I}_k), \quad \tilde{\boldsymbol{a}}_{i0} = \boldsymbol{0}_k,$$

which opens up various avenues for efficient Bayesian estimation. For this purpose, we assume a HS prior on the constant part of the parameters, $\boldsymbol{a}_i$, and the square root of the state innovation variances, $\boldsymbol{\theta}_i^{1/2}$ for each equation $i = 1, \ldots, n$, centering the prior model on a constant parameter variant to avoid overfitting (a similar specification has been used by Huber *et al.*, 2021).[6]

This establishes a specific state space representation, and combining the likelihood with suitable priors allows for devising a straightforward Gibbs sampling algorithm augmented with a forward-filtering backward-sampling step for drawing the TVPs if applicable. The functional relationship between the $\boldsymbol{y}_t$'s and the $\boldsymbol{x}_t$'s, may change over time, but conditional on each point in time, it is assumed to be known and linear. We break this assumption next.

Bayesian Additive Regression Trees

One of the goals of this chapter is to discuss and review methods capable of inferring unknown (and potentially nonlinear) conditional mean functions without imposing a large number of perhaps restrictive assumptions. The first approach we consider in this regard is based on regression trees. In particular, we use Bayesian Additive Regression Trees, often referred to by the acronym BART, as developed by Chipman *et al.* (2010). In essence, regression trees are step-functions which partition the input space provided by $\boldsymbol{x}_t$. The terminal node parameters that are associated with these partitions (the "leaves" of the tree) provide the fitted values for the output variable values.

When using BART, we approximate the unknown functions $f_i(\boldsymbol{x}_t)$ equation-by-equation using a sum of individual regression tree functions $\ell_{is}(\boldsymbol{x}_t|\mathcal{T}_{is}, \boldsymbol{\mu}_{is})$:

$$f_i(\boldsymbol{x}_t) \approx \sum_{s=1}^{S} \ell_{is}(\boldsymbol{x}_t|\mathcal{T}_{is}, \boldsymbol{\mu}_{is}).$$

[6] Consider the conditional prior for the $j = 1, \ldots, nk$, coefficient at time t, $p(a_{jt}|a_{jt-1}, \Theta_j) \sim \mathcal{N}(a_{jt-1}, \Theta_j)$. For large Θ_j, it allows for large shifts, while if $\Theta_j \rightarrow 0$, this specification collapses to a constant parameter model. In addition, shrinkage is imposed on the constant part (initial state) of the parameters, similar to the BVAR without TVPs.

Here, S denotes the number of trees, $\mathcal{T}_{is}$ are tree structures, and $\boldsymbol{\mu}_{is}$ are tree-specific terminal node parameters. Rather than having a single complex tree, the idea of BART is to have many simplistic trees. In this respect, BART is in fact related to the popular "classical" ML approach Random Forest. Each of the trees will be pruned heavily via our priors, such that individually they only explain a fraction of the variance of the dependent variable. But by having a combination of many simplistic trees, we are able to approximate virtually any nonlinear relationship between the dependent variable and the predictors.

Priors. BART uses many trees and we have many predictors, which requires introducing some sort of regularization. In our algorithm, the splitting variables are estimated alongside the thresholds and the terminal node parameters, on which we specify priors. In fact, we do not specify priors directly on the trees, but rather in the context of a tree-generating stochastic process which features three aspects. Let $\alpha \in (0, 1)$ and $\beta \in \mathbb{R}^+$. The first ingredient is to define the probability of a node at depth $d = 1, \dots,$ being nonterminal, which we set to $\alpha/(1 + d)^\beta$. This implies that the probability of having additional child nodes decreases the more complex (or deep, encoded in d) the respective tree becomes. Second, we use a discrete uniform prior over the splitting variables, which implies that each of them is equally likely *a priori* to be selected as determining partitions of the input space. Third, all thresholds within the splitting rules are assigned a uniform prior over the range of the relevant splitting variables. These are the assumptions about $\mathcal{T}_{is}$.

Next, we define the prior on the terminal node parameters, which act as the fitted values conditional on the partitioned input space. Let $\#\mathrm{TN}_{is}$ denote the number of terminal node parameters of tree s. Here, we use independent Gaussian priors that are symmetric across trees, $\mu_{is,l} \sim \mathcal{N}(0, v_i)$, for $l = 1, \dots, \#\mathrm{TN}_{is}$, and identical for all terminal nodes in equation i. The prior variance is chosen in a data-driven way:

$$\sqrt{v_i} = \frac{\max(\boldsymbol{y}_i) - \min(\boldsymbol{y}_i)}{2\gamma\sqrt{S}},$$

which puts most prior mass on the observed range of values, where $\boldsymbol{y}_i = (y_{i1}, \dots, y_{iT})'$.

Regularization of the trees arises from two main sources. First, it is due to how the probability of nodes being nonterminal is specified. We use the default values $\alpha = 0.95$ and $\beta = 2$ taken from Chipman *et al.* (2010). These choices work well across a large number of datasets (as demonstrated in the original paper), and favor simplistic over complicated trees (the prior probability of 2 to 4 terminal nodes is 0.92). Second, the tightness of the prior on the terminal node parameters increases with S. That is, shrinkage on the output values becomes tighter the more trees there are. It is worth noting, however, that these priors do not rule out larger and more complex trees in case they are required. We set the number of trees to $S = 250$.[7]

Algorithmically, it suffices to note here that the trees $\mathcal{T}_{is}$ are drawn marginally of the terminal node parameters, using a Metropolis-Hastings algorithm with a transition distribution defined by four distinct probabilistic moves: growing a terminal node, $\Pr(\texttt{grow}) = 0.25$; pruning a pair of terminal nodes, $\Pr(\texttt{prune}) = 0.25$; change a splitting rule, $\Pr(\texttt{change}) = 0.4$; swap parent and child node, $\Pr(\texttt{swap}) = 0.1$.

Example. To gain intuition how this procedure approximates the unknown functions, consider a single tree in stacked notation for a scalar dependent variable $\boldsymbol{y} = (y_1, \ldots, y_T)'$, a single predictor $\boldsymbol{x} = (x_1, \ldots, x_T)'$ and shocks $\boldsymbol{\epsilon} = (\epsilon_1, \ldots, \epsilon_T)'$:

$$\boldsymbol{y} = \ell(\boldsymbol{x}|\mathcal{T}, \boldsymbol{\mu}) + \boldsymbol{\epsilon}, \quad E(\boldsymbol{y}|\boldsymbol{x}) = \ell(\boldsymbol{x}|\mathcal{T}, \boldsymbol{\mu}) = \sum_{s=1}^{\#\text{TN}} \mu_s \mathbb{I}(\boldsymbol{x} \in \mathcal{S}_s).$$

The tree is fully defined by sets $\mathcal{T} = \{\mathcal{S}_s\}_{s=1}^{\#\text{TN}}$, formed by partitioning $\boldsymbol{x}$ via splitting rules: $\{\boldsymbol{x} \leq c\}$ or $\{\boldsymbol{x} > c\}$ with associated number of #TN terminal nodes μ_s. This simplistic example is visualized in Figure 1. Here we show the case of a single predictor, where the splitting rules amount to splitting the real line into non-overlapping intervals that define the sets $\mathcal{S}_s$. With two predictors, it produces partitioned areas on $\mathbb{R}^2$; for three predictors cuboids on $\mathbb{R}^3$, and so on. Moving to BART from this trivial example is as simple as adding an additional summation operator in front of the tree function, and introducing several additional indexes. And this consideration already hints at the increased flexibility that arises from having many rather than a single tree.

The tree in Figure 1 can be interpreted as follows. At the initial split, the rule is $x < -0.06$. In case it is fulfilled, one moves down the left-hand side, otherwise right. The next (non-terminal) splitting rule is $x < -0.68$; if it is fulfilled, one moves down left and reaches another rule. Otherwise, one moves right and reaches a terminal node containing 9% of the observations, with associated parameter $\mu_s = -0.85$. This value represents the assigned output, or

[7] It is worth mentioning that differences stemming from varying the number of trees are muted as long as this number is not set too small. In earlier related research, we found predictive performance typically to increase up to 100 or 150 trees. It subsequently plateaus when using even more trees — so our choice strikes a balance between having too few trees, and using needlessly many of them.

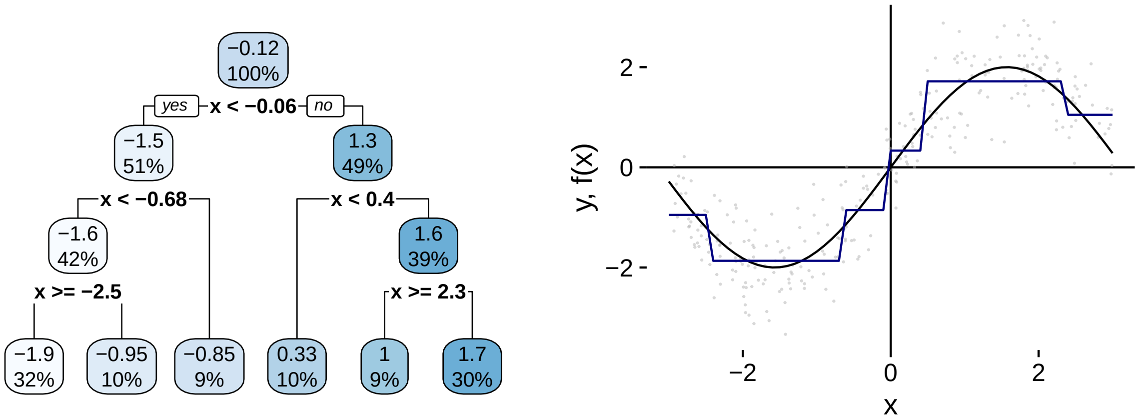

Fig. 1: How a regression tree partitions the (input) parameter space. Based on a single tree, i.e., $S = 1$. The left panel shows the tree, the right panel the conditional mean function.

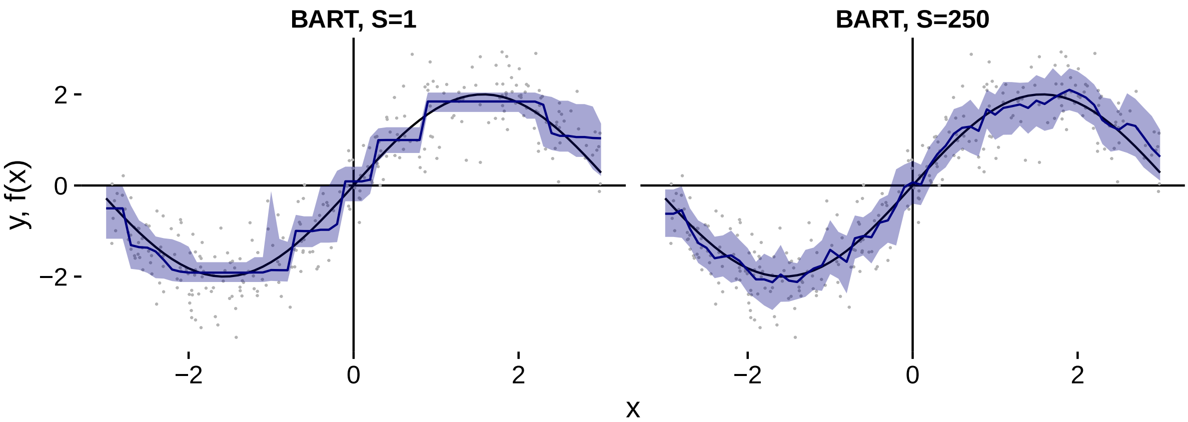

Fig. 2: Varying the number of trees. The blue shaded are is the 99% posterior credible set of the estimate, the black line indicates the true conditional mean function.

fitted value. The resulting piece-wise linear partitions of the input space give rise to the conditional mean function shown in the right panel. Rather than a single draw, however, Figure 2 shows the posterior distribution of BART with a single tree ($S = 1$), and our default specification with many trees ($S = 250$). Exploring the posterior this showcases that BART is capable of approximating virtually any functional relationship via using a sum of trees.

Gaussian process regression

The second approach to nonparametric inference is based on GP regression, see Williams and Rasmussen (2006) for a textbook treatment: "a GP is defined as a collection of random variables, any finite number of which have a joint Gaussian distribution." In essence, the idea in this context is to place priors directly on the conditional mean function. For our purposes, we set the mean of the GP itself to zero, such that it is exclusively specified by its variance:

$$f_i(\boldsymbol{x}_t) \sim \mathcal{GP}\left(0, \mathcal{K}_{\vartheta_i}(\boldsymbol{x}_t, \boldsymbol{x}_t)\right), \tag{3}$$

where $\mathcal{K}_{\vartheta_i}$ denotes a suitable covariance function. We will work with a simple distance-based covariance function, reflecting a specific notion of distance between periods of the input space across $t = 1, \ldots, T$; in general, we have $\text{Cov}\left(f_i(\boldsymbol{x}_t), f_i(\boldsymbol{x}_{\tilde{t}})\right) = \mathcal{K}_{\vartheta_i}\left(\boldsymbol{x}_t, \boldsymbol{x}_{\tilde{t}}\right)$, where $\mathcal{K}_{\vartheta_i}\left(\boldsymbol{x}_t, \boldsymbol{x}_{\tilde{t}}\right)$ is referred to as a kernel which depends on a set of tuning parameters $\boldsymbol{\vartheta}_i$. Many such kernels are available, which allows to express different beliefs about the respective functional relationships. We discuss a specific kernel to be used in our applications below, but note that the framework generalizes trivially to alternative choices.

Define $\mathcal{K}_{\vartheta_i}(\boldsymbol{X}, \boldsymbol{X}')$ with typical $(t, \tilde{t})$ element $\mathcal{K}_{\vartheta_i}(\boldsymbol{x}_t, \boldsymbol{x}_{\tilde{t}})$, i.e., it captures the covariances of the function values across all input vector combinations. Further, we stack $\tilde{\boldsymbol{y}}_i = (\tilde{y}_{i1}, \ldots, \tilde{y}_{iT})'$, $\boldsymbol{X} = (\boldsymbol{x}_1, \ldots, \boldsymbol{x}_T)'$ and $\boldsymbol{f}_i = (f_i(\boldsymbol{x}_1), \ldots, f_i(\boldsymbol{x}_T))'$, and obtain:

$$\boldsymbol{f}_i \sim \mathcal{N}\left(\boldsymbol{0}_T, \mathcal{K}_{\vartheta_i}(\boldsymbol{X}, \boldsymbol{X}')\right), \quad \tilde{\boldsymbol{y}}_i \sim \mathcal{N}\left(\boldsymbol{0}_T, \mathcal{K}_{\vartheta_i}(\boldsymbol{X}, \boldsymbol{X}') + \text{diag}(\boldsymbol{h}_i)\right).$$

In other words, Eq. (3) transforms into a multivariate Gaussian distribution when conditioning on the data. Note that the representation for $\tilde{\boldsymbol{y}}_i$ above is conditional on $\boldsymbol{L}\mathfrak{F}_t$ for notational simplicity, resulting in a diagonal white noise component with variances over time stored in $\boldsymbol{h}_i$. Intuitively, the distance-based covariances capture similarity of observations in the input space that will transmit to similarities in the output space of functional values.

Priors. The specific kernel we choose is the squared-exponential kernel. This is a popular choice since it directly relates to a Bayesian linear regression model with infinitely many basis functions, among several other useful features and properties. Let us first define our notion of distance, with $d_{t\tilde{t}} = ||\boldsymbol{x}_t - \boldsymbol{x}_{\tilde{t}}||^2$ capturing the Euclidean norm between two different periods in the input space.

The squared-exponential kernel is defined as:

$$\mathcal{K}_{\vartheta_i}\left(\boldsymbol{x}_t, \boldsymbol{x}_{\tilde{t}}\right) = \xi_i \times \exp\left(-\frac{l_i d_{t\tilde{t}}}{2}\right),$$

with two hyperparameters collected in $\boldsymbol{\vartheta}_i = (\xi_i, l_i)'$. The parameter ξ_i regulates the unconditional variance of the prior, since for $\boldsymbol{x}_t = \boldsymbol{x}_{\tilde{t}}$ which yields $d_{t\tilde{t}} = 0$ one obtains $\mathcal{K}_{\boldsymbol{\vartheta}_i}(\boldsymbol{x}_t, \boldsymbol{x}_{\tilde{t}}) = \xi_i$. The parameter l_i is referred to as the inverse length-scale, which in simple terms governs the sensitivity of the output values to varying distance in the input space.

The parameters in $\boldsymbol{\vartheta}_i$ may be fixed and treated as tuning parameters to be optimized via cross-validation. We take a different route and assign independent Gamma priors, and infer them alongside all other parameters when running our algorithm.

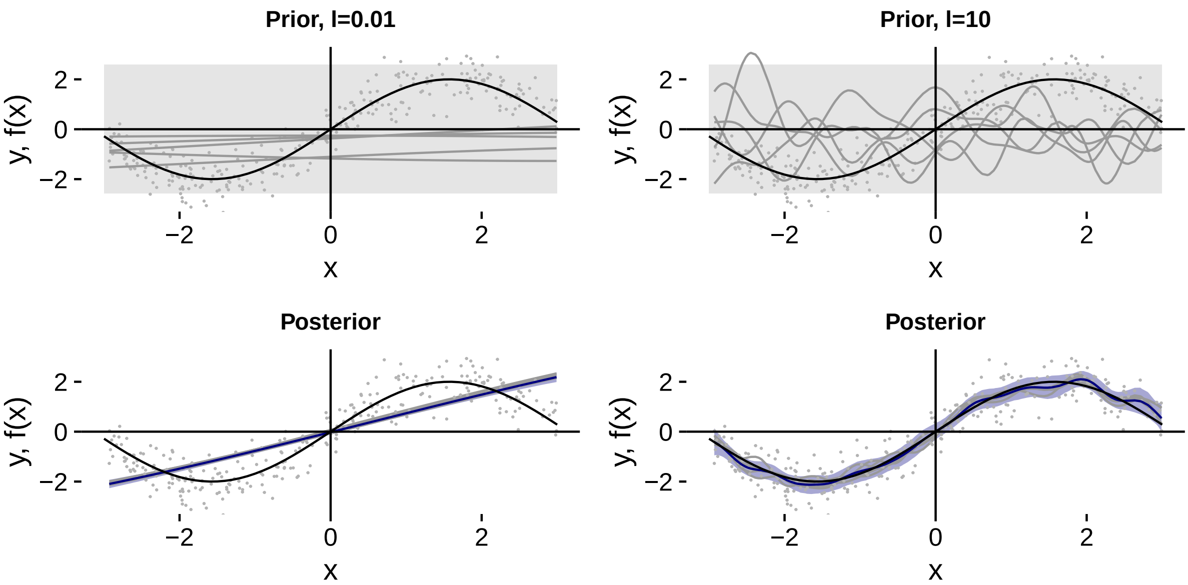

Fig. 3: Visualizing the GP prior and resulting posterior. The blue shaded area is the 99% posterior credible set of the estimate, the black line indicates the true conditional mean function. The grey shaded area refers to the unconditional variance ξ of the GP prior; grey lines refer to five random draws from this prior varying the inverse length scale l.

Example. Figure 3 provides a visualization of our GP prior, varying the inverse length-scale l_i and fixing $\xi_i = 1$. The upper panels indicate the implied priors, while the bottom panels show the resulting posterior distribution. The grey shaded area refers to 95% prior mass over the possible space of functions, and is the same for both cases of the inverse length-scale. To visualize the role of the inverse length-scale, we draw five realizations from these priors (dark grey lines). On the left-hand side, we have a short inverse length-scale (and conversely, a long length-scale). This results in very slow-moving prior functions. In other words, a substantial shift in the input space is necessary to produce a noteworthy change in the output space. Assessing the associated posterior indicates that a very short inverse length-scale produces an almost linear relationship. By contrast, a long inverse length-scale allows for rapid changes of the function values. This is also reflected in the posterior, which exhibits some overfitting.

2.3. *Modeling the conditional variances*

As stated above, we rely on a FSV model for the reduced form errors due to a number of computational and statistical advantages, alongside its established solid empirical performance (see, e.g., Chan, 2023).[8] So far we have not discussed specifics about how we model the drifting volatilities of the factors and the idiosyncratic errors. The natural logarithm of the diagonal elements of the respective covariance matrices are assumed to follow independent AR(1) processes:

$$h_{it} = \mu_{ih} + \phi_{ih}(h_{it-1} - \mu_{ih}) + \varsigma_{ih}\zeta_{it,h}, \quad \zeta_{it,h} \sim \mathcal{N}(0,1), \quad \text{for } i = 1, \ldots, n,$$
$$\omega_{jt} = \phi_{j\omega}\omega_{jt-1} + \varsigma_{j\omega}\zeta_{jt,\omega}, \quad \zeta_{jt,\omega} \sim \mathcal{N}(0,1), \quad \text{for } j = 1, \ldots, q,$$

Note that we normalize the unconditional mean of the factor-specific volatility processes to zero to pin down the scale of the factors. We refer to this heteroskedastic normal specification of the errors as `SV`.[9] In homoskedastic specifications, we assume that $\omega_{jt} = 0$ for $j = 1, \ldots, q$, and $h_{it} = h_i$ for $i = 1, \ldots, n$, and all $t = 1, \ldots, T$. We assign weakly informative independent inverse gamma priors on $\exp(h_i)$, and indicate this specification as `hom`.

For the BVAR, to reflect large variance shocks such as during the Covid-19 pandemic, we also consider the following alternative specification using t-distributed idiosyncratic shocks:

$$\eta_{it} \sim t_{\nu_i}(0, \exp(h_{it})),$$

where ν_i refers to the degrees of freedom which we estimate under a uniform prior. This t-distributed error specification is referred to as `SV-t` and assumes the same independent AR(1) processes as written above on h_{it}. Rewriting this distribution as a scale mixture allows for a representation similar to Carriero *et al.* (202X). [10]

[8] In this framework, the sign of the factors is not econometrically identified. This poses no further issues when interest is on forecasting, since we merely use the FSV model as a tool to estimate the full covariance matrix of the system. Several avenues to achieve econometric identification are available if necessary (see Aguilar and West, 2000; Chan *et al.*, 2022).

[9] Note that nonparametric methods for estimating variances are available. We refrain from including such approaches in this chapter, but note that our baseline framework could trivially be extended to also allow for nonparametric treatments of the conditional variance, see, e.g., Clark *et al.* (2023).

[10] Empirically, we considered this specification of the errors also for BART and GP but got similar results to the case of normal errors.

2.4. *Estimation and prediction*

We use a Bayesian approach for posterior and predictive inference. In particular, our basic setup allows for a comparatively straightforward sampling algorithm. Our implementation uses mostly conditional Gibbs sampling steps, augmented with occasional MH updates when the conditional posterior distributions are not available in closed form. For the SV processes and TVPs, we rely on standard filtering and smoothing methods for state space models (see Chan *et al.*, 2020, for a textbook treatment). The codes are implemented in the software `R`.

Details about priors, posteriors and the implementation of the algorithm are provided in Appendix A. In the following, we summarize key steps and discuss how to compute higher-order forecasts.

Sampling algorithm. The key aspect of the FSV model from an algorithmic perspective is that we can break up the full multivariate system into independent equations. Recall that we earlier defined $(\boldsymbol{y}_t - \boldsymbol{L}\mathfrak{F}_t) \equiv \tilde{\boldsymbol{y}}_t = F(\boldsymbol{x}_t) + \boldsymbol{\eta}_t$, where $\boldsymbol{\eta}_t$ has a diagonal covariance matrix. Conditional on $\boldsymbol{L}\mathfrak{F}_t$, we thus may work with the univariate equation $\tilde{y}_{it} = f_i(\boldsymbol{x}_t) + \eta_{it}$ and update the $f_i(\boldsymbol{x}_t)$'s equation-by-equation for $i = 1, \ldots, n$. This yields a draw for the full vector $F(\boldsymbol{x}_t)$, which subsequently can be used to update any hyperparameters related to the conditional mean (e.g., the shrinkage parameters for the BVAR, or the hyperparameters of the kernel).

Conversely, notice that $(\boldsymbol{y}_t - F(\boldsymbol{x}_t)) \equiv \boldsymbol{\epsilon}_t = \boldsymbol{L}\mathfrak{F}_t + \boldsymbol{\eta}_t$. This auxiliary representation is a simple linear Gaussian regression model with (potentially) heteroskedastic errors. Treating the latent factors $\mathfrak{F}_t$ as conditionally observed implies that the loadings $\boldsymbol{L}$ can be drawn variable-by-variable from their Gaussian posterior. Having updated the loadings, one may stack the reduced form errors and update the full history of the factors $\mathfrak{F}_t$ jointly in one block.

Depending on the respective applicable assumptions about the factor and idiosyncratic volatilities, these can be sampled factor-by-factor and variable-by-variable, using either the textbook quantities for homoskedastic errors, or a standard algorithm for sampling stochastic volatilities.

Predictive Simulation. For the BVAR and TVP-VAR versions of our models, the moments of the h-step ahead predictive distribution are available in closed-form, and can be derived by iterating the endogenous vector and covariance matrix forward in time. This is not the case for the nonparametric approaches due to the innate nonlinearity, but we may use simulation-based methods to produce a Monte Carlo sample from the predictive distribution.

To set up this procedure, let $\boldsymbol{x}_T = (\boldsymbol{y}'_T, \ldots, \boldsymbol{y}'_{T-p+1})'$ denote the predictors at the end of the sample. Also note that $\boldsymbol{\epsilon}_{T+1} \sim \mathcal{N}(\mathbf{0}_n, \boldsymbol{L}\boldsymbol{\Omega}_{T+1}\boldsymbol{L}' + \boldsymbol{H}_{T+1})$, where the covariance matrices can be updated by iterating forward their independent AR(1) state equations. We denote a draw from these errors with $\tilde{\boldsymbol{\epsilon}}_{T+1}$. Depending on the conditional mean function, we obtain either the predictive location $\tilde{y}_{iT+1} = f_i(\boldsymbol{x}_T)$ or distribution from which we can sample. These equation-specific predictive draws are collected in $\tilde{\boldsymbol{y}}_{T+1} = (\tilde{y}_{1T+1}, \ldots, \tilde{y}_{nT+1})'$, which in conjunction with the errors can be used to form a draw $\hat{\boldsymbol{y}}_{T+1} = \tilde{\boldsymbol{y}}_{T+1} + \tilde{\boldsymbol{\epsilon}}_{T+1}$ from the multivariate system. For the two-step ahead prediction, we condition on this draw such that $\boldsymbol{x}_{T+1} = (\hat{\boldsymbol{y}}'_{T+1}, \boldsymbol{y}'_T, \ldots, \boldsymbol{y}'_{T-p+2})'$ configures the input space, which allows to obtain $\hat{\boldsymbol{y}}_{T+2} = \tilde{\boldsymbol{y}}_{T+2} + \tilde{\boldsymbol{\epsilon}}_{T+2}$.

Higher-order forecasts, $\hat{\boldsymbol{y}}_{T+h}$, for $h = 1, 2, \ldots$, are obtained analogously by iteratively simulating forward and re-configuring the input vector at each additional horizon. Applying this procedure for each sweep of the algorithm integrates out all sources of predictive uncertainty.

3. ILLUSTRATIONS REVISITING OKUN'S LAW AND THE PHILLIPS CURVE

Before proceeding with the multivariate forecasting models, we illustrate the several types of nonlinearities that may arise from our nonparametric approaches using simplified examples. We revisit two prominent (single-equation) macroeconomic classics: Okun's Law (OL) and two variants of the Phillips Curve (PC).[11] In particular, we intend to showcase commonalities and differences between BART and GP, since the two take very different approaches to achieve a potentially nonlinear relationship between a dependent variable and one or more predictors. As a rough gauge of goodness of fit, we compute the coefficient of determination (R^2) and the deviance information criterion (DIC). They are displayed in Table 1.

For these illustrations, we use US data for the unemployment rate u_t, its natural rate u_t^*, inflation π_t, inflation expectations $\mathbb{E}_t(\pi_{t+1})$, and the logs of real GDP y_t and potential output y_t^*. These measures allow us to define the unemployment gap $u_t^g = u_t - u_t^*$ and the output gap $y_t^g = y_t - y_t^*$. The data are quarterly and taken from the Real-Time Dataset for Macroeconomists (GDP, unemployment and inflation), the Survey of Professional Forecasters (inflation expectations) and FRED (natural rate of unemployment and potential output).

[11] Note that the purpose of these examples is purely to visualize how the nonparametric approaches fit data, and we make no claim of "causal" or structural identification. This would require more intricate econometric procedures on top of the nonparametric approaches.

Table 1: Goodness of fit for each illustrative example across model specifications.

	OL					PC_1					PC_2			
	BLR	BART	GP(e)	GP(f) l = 1	GP(f) l = 0.1	BLR	BART	GP(e)	GP(f) l = 1	GP(f) l = 0.1	BLR	BART	GP(e)	GP(f) l = 1
R^2	75.8	**80.8**	77.2	78.0	77.2	3.9	**14.2**	1.7	6.0	1.6	28.8	**40.8**	26.9	33.2
DIC	231.5	**223.1**	232.6	243.8	232.4	393.8	**372.5**	397.2	401.9	395.8	326.9	**292.0**	342.4	355.5

Note: Coefficient of determination R^2 and deviance information criterion (DIC) for Okun's law (OL) and the two Phillips curves (PC). Bayesian linear regression (BLR), Bayesian additive regression trees (BART) and Gaussian Process (GP). The letter in parentheses for GP refers to whether the inverse length-scale l is fixed (f) to the indicated value, or estimated (e).

We consider three distinct specifications. The first two, $\texttt{OL}$ and $\texttt{PC}_1$, are bivariate nonlinear regressions, allowing for comparatively straightforward visualizations. The third, $\texttt{PC}_2$, serves to illustrate a multiple nonlinear regression; it nests $\texttt{PC}_1$ in the case of expected inflation entering linearly with a coefficient of 1:

$$u_t^g = f(y_t^g) + \epsilon_t, \tag{OL}$$

$$\pi_t - \mathbb{E}_t(\pi_{t+1}) = f(u_t^g) + \epsilon_t, \tag{$\texttt{PC}_1$}$$

$$\pi_t = f(\mathbb{E}_t(\pi_{t+1}), u_t^g) + \epsilon_t. \tag{$\texttt{PC}_2$}$$

All three specifications are estimated with SV. We normalize each variable prior to estimation such that the charts can be interpreted like standardized regression coefficients. Figures 4 and 5 show the function values when fitting the estimated relationships over a grid of all possible combinations of the input variables. That is, these charts are closely related to so-called partial dependence plots, a tool designed to ease interpretation of the output of black-box machine learning methods.[12]

We start with $\texttt{OL}$ in the upper panels. The first chart on the far left indicates the Bayesian linear regression (BLR) model; it estimates a straight line with tight credible sets. By contrast, when the function is unknown, BART and GP exhibit wider credible sets. BART and the GP variants ($l = 0.1$ in red, $l = 1$ in blue, estimated l in green) agree on the shape of the function, which somewhat flattens for large positive values of the output gap. There are some (barely significant) differences for very negative values among the nonparametrics. A further key difference between them arises from the different assumptions about how to estimate the conditional mean function; clearly, the GP variants are smooth, whereas BART results in somewhat

[12] For a textbook review of such tools and devices, see e.g., Hastie *et al.* (2009).

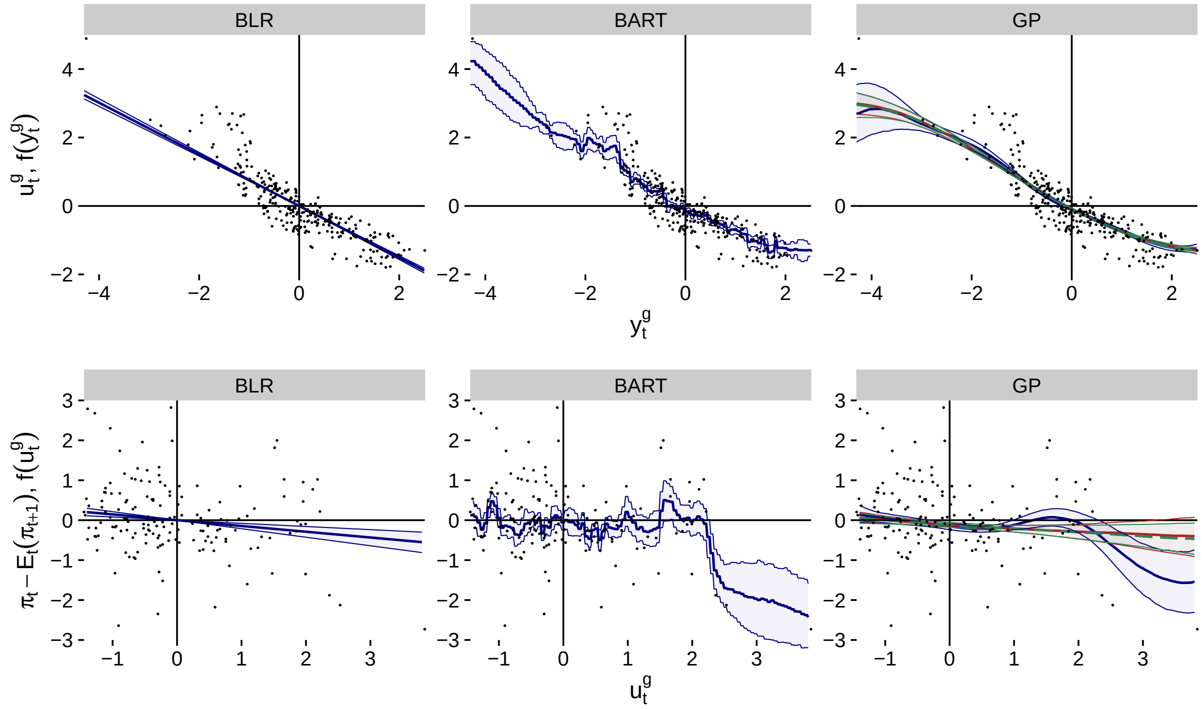

Fig. 4: Conditional mean functions for `OL` (upper panels) and `PC`$_1$ (lower panels). Bayesian linear regression (BLR), Bayesian additive regression trees (BART, $S = 250$) and Gaussian process (GP, $\xi = 1$; $l = 0.1$ in red, $l = 1$ in blue, estimated l in green). Points mark observations.

rougher movements and piece-wise almost constant fitted values. However, linearity appears to be a reasonable approximation. This is also evident when comparing the estimated inverse length-scale to the fixed length-scale versions: a comparatively small estimate for l results, suggesting a very slow-moving conditional mean function.

A different picture emerges for `PC`$_1$, in the lower panels of Figure 4. Again, by construction, BLR shows a linear, weakly negative association between the inflation differential and the unemployment gap. Prominent nonlinearities are visible when the conditional mean function is approximated with BART. These estimates more or less coincide with the looser of the two GP(f) specifications. The small inverse length-scale case (in red), on the other hand, results in an estimate which is rather similar to the linear one, as is the case for the estimated version (in green). Overall, this procedure detects a rather flat PC over most of the range of the unemployment gap, and a somewhat more negative slope for high values. A particularly interesting finding appears for the BART case. Zooming into some of the partitions of the input space — e.g., u_t^g in the intervals $(-1.2, -0.8)$, $(-0.2, 0.5)$, or $(1.0, 1.5)$ — points towards piece-wise approximately linear PCs, shifted conditional on the magnitude of the unemployment gap.

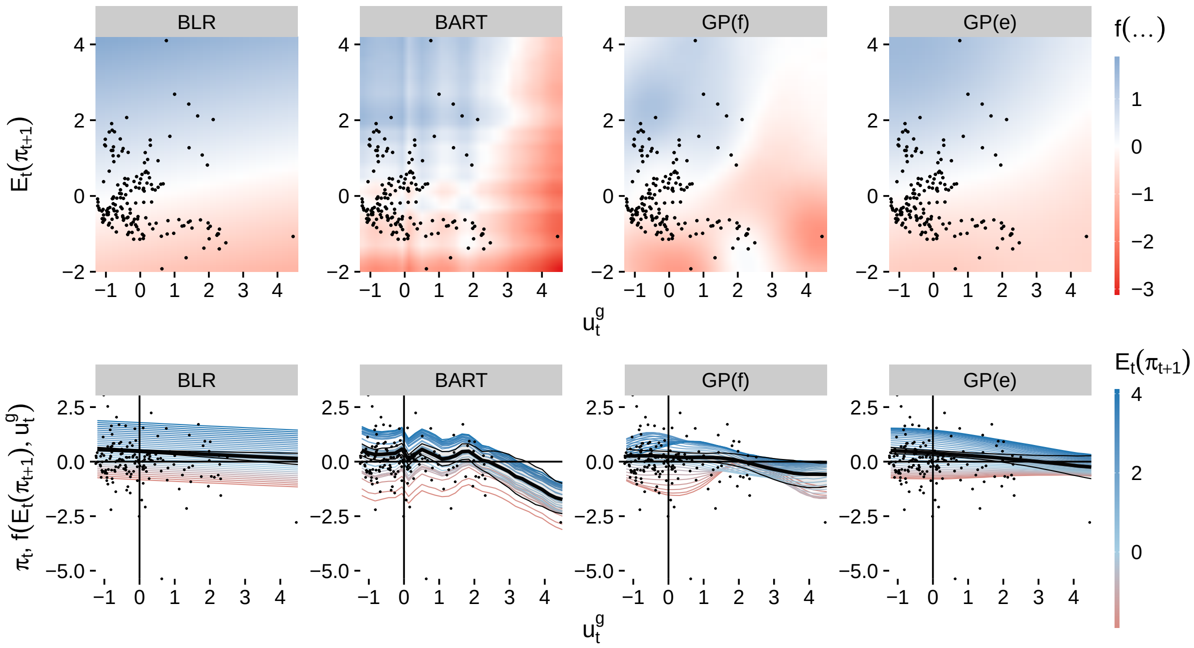

Fig. 5: Posterior surface (upper panels) and marginal conditional mean functions (lower panels) for $\texttt{PC}_2$. Bayesian linear regression (BLR), Bayesian additive regression trees (BART, $S = 250$) and the Gaussian process (GP, $\xi = 1$, $l = 1$ **f**ixed or l **e**stimated).

Finally, we turn to the more sophisticated case of $\texttt{PC}_2$, with two right-hand side variables which adds a dimension to our estimates of the conditional mean function. The posteriors are visualized in Figure 5 (upper panels), with the multiple BLR now by construction fitting a plane, and more complex surfaces arising from BART and GP. The lower panels condition on various levels of expected inflation, indicated using the color scale ranging from red (low) to blue (high). The black line with associated grey shaded area refers to the posterior median and 68 percent credible set averaging across all input combinations.[13]

The findings from $\texttt{PC}_1$ generalize to this case; comparing the lower panels of Figure 5 to the lower panels of Figure 4, the estimated marginal functional relationship is virtually identical. However, adding potential nonlinearities via using inflation expectations as an additional covariate (rather than fixing its linear coefficient to 1 as in the $\texttt{PC}_1$ example) results in a more nuanced and potentially heterogeneous estimate of the PC. This is also reflected in more favorable measures of fit for $\texttt{PC}_2$ over $\texttt{PC}_1$. Abstracting from the location shift, varying magnitudes of inflation expectations change the shape of the PC, particularly in the case of GP(f), and to a lesser extent, also for BART and GP(e). This finding is also clearly visible in the upper panels,

[13]In principle, such analyses could also be conducted for the multivariate (higher-dimensional) case. However, it is worth mentioning that adding predictors (let alone additional equations in a multivariate system) increases computational complexity significantly, and producing such charts is computationally prohibitive.

with more red in the top right quadrant of the BART and GP charts when compared to BLR. Again, it is worth noting that BART appears to produce a slightly more noisy fit; but on average, the two nonparametric methods tend to agree, at least qualitatively, about the shape of the nonlinear relationship. BART is often close to a loose GP prior.

Considering the measures of fit provided in Table 1, it is worth mentioning that BART consistently produces the highest R^2 and lowest DIC. GP(e) is usually rather close to GP(f) with $l = 0.1$, and BLR indicates competitive metrics in all cases, again pointing towards linearity being a good approximation. How can these patterns be explained? First, note that the DIC is an imperfect in-sample measure of model fit, and may in fact disproportionally reward overfitting models. In light of the high R^2 value and from eyeballing the charts, this is likely what happens. Second, a high in-sample R^2 is not necessarily indicative of good predictive performance. In fact, this is rarely the case: high in-sample R^2 often coincides with low out-of-sample R^2, pointing towards overfitting issues. Good predictive performance in spite of in-sample overfitting is a somewhat unique feature of particular versions of regression trees (see, e.g., the discussion in Goulet Coulombe, 2020).

Summarizing, this simplistic illustration provides two lessons (we again note that these in-sample results must be viewed and interpreted with caution). First, linear models are sensible choices in many cases, at least for the examples discussed above. Second, nonparametric approaches may discover intricate and useful nonlinearities, but depending on hyperparameter choices, they may also overfit the data. Next, we turn to out-of-sample predictive inference, a more natural and informative means of model comparison.

4. FORECASTS FOR THE EA AND THE US

4.1. *Data and forecast designs*

Euro Area. For the EA forecast exercise, we rely on the real-time dataset used in Banbura *et al.* (2021). The history of vintages starts in 2001Q1 and ends in 2021Q4. Since the final vintage is comprised of data up to 2021Q3, this results in a maximum holdout sample from 2001Q2 until 2021Q3 (we use the final available vintage for evaluating our real-time forecasts) for one-step-ahead forecasts. For h-step-ahead forecasts, the holdout (evaluation period) shortens by the initial h quarters due to having a fixed date/vintage as our initial training sample. All series start in 1980Q2 and we use $p = 4$ lags with the small and medium information sets. We account for ragged edges by imputing the missing values from their joint multivariate conditional distribution online within our algorithm.

United States. An often used dataset for the purpose of evaluating competing forecast models, particularly in multivariate time series analysis, is FRED-QD, which was compiled and described by McCracken and Ng (2020). We consider this dataset in addition to the EA for two reasons. First, to also employ a well-known dataset, and to see whether lessons from the EA carry over to the US. Second, to have a longer training sample and somewhat larger information set for the medium-sized specifications. Here, we adopt a pseudo out-of-sample forecast design using an expanding window. The initial training sample ranges from 1959Q2 to 1984Q4. The final available observation is 2022Q4, which implies that our forecast evaluation period ranges from 1985Q1 until that quarter. We again use $p = 4$ lags with small and medium information sets.

Although we use differently sized information sets (`Small` and `Medium`), we evaluate the forecast performance for three focus variables. In particular, we investigate the predictive performance of our competing models for real GDP growth, headline inflation and the change in unemployment rates. The full dataset for the EA and the US are summarized in Table 2.

4.2. *Competing models and predictive losses*

The models are differentiated mainly with respect to how the conditional mean functions $F(\boldsymbol{x}_t)$ are estimated. Here, we consider the BVAR, TVP-VAR, BART and GP as potential variants. Moreover, we consider various assumptions about the errors. In particular, we have homoskedastic (labeled `hom`) and SV factors and idiosyncratic shocks for all versions of the conditional mean. For the BVAR, we also consider t-distributed errors (labeled `SV-t`). Third, we estimate all models using two differently sized information sets, a small one with just the target variables plus a short-term interest rate, and a medium-sized set with 14 (EA) and 19 (US) series.

Due to the practical relevance, we provide an overview of estimation times in Table 3. The gray shaded cell displays seconds per 1000 predictive draws, and all other models are shown as ratios to this benchmark which is the small BVAR with homoskedastic errors. The indicated values are averages over the holdout, to also reflect the effect of an increasing number of training observations. The forecast exercise was run on a high performance computing cluster with an Intel `E5-2650v3` 2.3 GHz processor and 256GB memory. For instance, in the case of the

Table 2: Datasets, abbreviations, transformation codes and variable inclusion.

Euro area	Code	Transform	S	M
Real Gross Domestic Product	YER	$100 \cdot \Delta \log(x)$	✓	✓
Harmonized Index of Consumer Prices	HICSA	$100 \cdot \Delta \log(x)$	✓	✓
Unemployment Rate	URX	Δx	✓	✓
Short-term interest rate	STN	Δx	✓	✓
Private consumption, real	PCR	$100 \cdot \Delta \log(x)$		✓
Total investment, real	ITR	$100 \cdot \Delta \log(x)$		✓
GDP deflator	YED	$100 \cdot \Delta \log(x)$		✓
HICP excluding energy and food	HEFSA	$100 \cdot \Delta \log(x)$		✓
Long-term interest rate	LTN	Δx		✓
Total employment	LNN	$100 \cdot \Delta \log(x)$		✓
Compensation per employee	CEX	$100 \cdot \Delta \log(x)$		✓
ESI	ESI	$100 \cdot \Delta \log(x)$		✓
Price of oil in EUR	POE	$100 \cdot \Delta \log(x)$		✓
Nominal effective exchange rate	EERB	$100 \cdot \Delta \log(x)$		✓
United States				
Real Gross Domestic Product	GDPC1	$100 \cdot \Delta \log(x)$	✓	✓
Consumer Price Index	CPIAUCSL	$100 \cdot \Delta \log(x)$	✓	✓
Unemployment Rate	UNRATE	Δx	✓	✓
Federal Funds Rate	FEDFUNDS	Δx	✓	✓
Personal Consumption Expenditures	PCECC96	$100 \cdot \Delta \log(x)$		✓
Gross Private Domestic Investment	GPDIC1	$100 \cdot \Delta \log(x)$		✓
GDP Deflator	GDPCTPI	$100 \cdot \Delta \log(x)$		✓
Producer Price Index	PPIACO	$100 \cdot \Delta \log(x)$		✓
Industrial Production Index	INDPRO	$100 \cdot \Delta \log(x)$		✓
1-Year Treasury Constant Maturity Rate	GS1	Δx		✓
5-Year Treasury Constant Maturity Rate	GS5	Δx		✓
Nonfarm All Employees	PAYEMS	$100 \cdot \Delta \log(x)$		✓
Nonfarm Hours All Persons	HOANBS	$100 \cdot \Delta \log(x)$		✓
Nonfarm Compensation Per Hour	COMPRNFB	$100 \cdot \Delta \log(x)$		✓
Capacity Utilization: Manufacturing	CUMFNS	x		✓
Real Crude Oil Prices WTI	OILPRICEx	$100 \cdot \Delta \log(x)$		✓
US/UK Foreign Exchange Rate	EXUSUKx	$100 \cdot \Delta \log(x)$		✓
Real M2 Money Stock	M2REAL	$100 \cdot \Delta \log(x)$		✓
S&P 500	SP500	$100 \cdot \Delta \log(x)$		✓

EA, estimating the model and producing 1000 draws from the predictive distribution for the small BVAR takes about 25 seconds. Doing the same for the TVP-VAR with a medium-sized information set takes 54 times as long. For further reference, on a 2020 Macbook Air M1 it takes about 5 seconds per 1000 draws to estimate and predict with the homoskedastic BVAR using the small EA information set.[14]

To assess predictive performance, we rely on distinct loss functions that are designed to measure accuracy along several dimensions. The predictive distributions of our nonparametric models may feature heavy tails, skewness, or even multiple modes — particularly so for higher-order forecasts. Besides overall measures of the adequacy of our predictions, we thus aim at painting a more nuanced picture by also focusing on specific parts of the predictive distribu-

[14] Our empirical results are based on 11000 iterations of the algorithm, discarding the initial 2000 as burnin and considering each third of the remaining 9000 draws. Thus, the samples used for inference comprise 3000 draws.

Table 3: Estimation times over the holdout.

		BVAR hom	BVAR SV	BVAR SV−t	TVP SV	BART hom	BART SV	GP hom	GP SV
EA	S	24.49	1.09	1.10	1.80	1.19	1.28	2.36	2.47
EA	M	2.38	2.72	2.75	54.32	3.19	3.55	8.14	8.53
US	S	26.45	1.10	1.11	1.90	1.18	1.10	4.48	4.61
US	M	5.35	5.88	5.93	218.13	5.92	5.45	26.14	26.68

Note: Required time for obtaining 1000 draws from the predictive distribution. Values are averages across all holdout samples. Rows refer to the small (S) and medium (M) information sets. The gray shaded cells display seconds per 1000 draws, all other models are shown as ratios to this benchmark.

tions. In this context, we first define the Quantile Score (QS, see, e.g., Giacomini and Komunjer, 2005) for quantile $\tau \in (0, 1)$:

$$\text{QS}_{\tau,it} = 2(y_{it} - Q_{\tau,it})(\tau - \mathbb{I}(y_{it} \leq Q_{\tau,it})),$$

which is based on the tick loss function; $Q_{\tau,it}$ is the τth quantile of the distribution of variable $i = 1, \ldots, n$, and y_{it} the realization of the forecasted variable. Note that for $\tau = 0.5$, the QS collapses to the mean absolute error (MAE); the QS thus generalizes this common point forecast metric to generic quantiles.

The QS allows to comment on predictive accuracy at any desired quantile. As a lower-dimensional summary across quantiles, we follow Gneiting and Ranjan (2011) and define variants of the quantile-weighted continuous ranked probability score (CRPS). The CRPS is a proper scoring rule and given by:

$$\text{CRPS}_{m,it} = \frac{1}{J-1} \sum_{j=1}^{J-1} \mathfrak{w}_m(\tau_j) \text{QS}_{\tau_j,it}, \quad \tau_j = \frac{j}{J},$$

where the weights $\mathfrak{w}_m(\tau_j)$ define which part of the predictive distribution shall be targeted. The trivial case, $\mathfrak{w}_m(\tau_j) = 1$, is labeled `CRPS` in our results. The CRPS is a standard metric for density forecast performance and can be interpreted in the original scale of the data. An alternative representation is given by:

$$\text{CRPS}_{it} = \int_{-\infty}^{\infty} (\mathcal{F}(z) - \mathbb{I}(y_{it} \leq z))^2 dz,$$

where $\mathcal{F}(z)$ denotes the cumulative distribution function of the predictive density, see Gneiting and Raftery (2007).

With an eye to assessing the performance of our models regarding tail events and in non-normal times, in addition to the non-weighted CRPS we consider weighting schemes that target downside-risk (left tail, CRPS-L) and upside-risk (right tail, CRPS-R) respectively: $\mathfrak{w}_{\mathtt{L}}(\tau_j) = (1 - \tau_j)^2$ and $\mathfrak{w}_{\mathtt{R}}(\tau_j) = \tau_j^2$. As a rough gauge of significance of the predictive premia, we conduct Diebold and Mariano (DM, 1995) tests for all predictive losses.

4.3. *Forecast results*

Tables 4 and 5 summarize our results for the EA and US, respectively. To provide additional discussions with respect to distinct economic phases, we evaluate the predictive losses also for subsamples. In particular, we consider four splits: recessions (Re), expansions (Ex), pre Covid-19 (Pre) and post Covid-19 (Post).[15] All results are benchmarked (as ratios) relative to the medium BVAR with homoskedastic errors, whose grey shaded row displays raw predictive losses. The best overall specification is indicated in bold, the best specification by size is in italics (in case no model outperforms the benchmark, this marking is omitted). Levels of statistical significance for the DM-test, testing for equal versus greater predictive accuracy relative to the benchmark are: ′ (10 percent), ° (5 percent) and * (1 percent).

A helicopter tour

We first provide an overview of general patterns that are apparent for both the EA and the US across Tables 4 and 5. Overall, there is little to be gained from moving beyond linearity when the focus is on predicting GDP growth. There are some minor gains, most noteworthy for two-year ahead point forecasts in the US. If there are gains for the nonparametrics, however, most of them are either negligible due to tiny improvements in magnitude, statistically insignificant, or result for less economically important metrics (e.g., upward-risk to GDP as captured by CRPS-R).

For both economies, we find improvements in predictive accuracy for headline inflation and the unemployment rate, to a varying extent. Starting with inflation, particularly BART exploiting the medium sized information set indicates gains of up to 5% (EA) and 14% (US),

[15]Recession dates are taken from the NBER Business Cycle Dating Committee and the Euro Area Business Cycle Network. The pre and post Covid-19 split is 2019Q4, such that this quarter is the final one in the Pre sub-sample. The forecasts are allocated to subsamples based on the quarter for which the they were made, not when they were produced.

Table 4: Predictive loss metrics for the EA.

S

		CRPS			CRPS−L			CRPS−R			MAE		
	Horizon	1	4	8	1	4	8	1	4	8	1	4	8
YER	BVAR hom	1.27	1.01	1.01	1.23	1.01	1.00	1.34	1.01	1.01	1.31	1.01	1.01
	BVAR SV	1.03	1.01	0.99	1.05	1.02	0.99	1.01	1.00	1.00	1.03	1.01	1.00
	BVAR SV−t	1.02	1.00	0.99°	1.04	1.01	1.00	**1.00**	1.00	**0.99***	*1.02*	1.01	0.99'
	TVP SV	1.53	1.01	**0.99'**	1.72	1.01	*0.99*	1.34	1.00	0.99	1.56	1.00	**0.99**
	BART hom	1.18	1.01	1.00	1.15	1.01	0.99	1.23	1.02	1.00	1.20	1.01	1.00
	BART SV	1.03	1.01	1.00	1.05	1.02	1.01	1.02	1.00	1.00	1.04	1.01	1.01
	GP hom	*1.01*	1.00	1.00	*1.02*	*1.00*	1.00	1.01	1.01	1.01	1.03	0.99	1.01
	GP SV	1.05	**0.99**	1.00	1.06	1.00	1.00	1.04	*0.99*	1.01	1.04	**0.99**	1.00
HICSA	BVAR hom	1.08	1.11	1.10	1.09	1.11	1.12	1.06	1.09	1.09	1.08	1.11	1.12
	BVAR SV	1.07	1.10	1.11	1.09	1.12	1.14	1.04	1.07	1.08	1.07	1.10	1.13
	BVAR SV−t	1.08	1.13	1.10	1.09	1.14	1.11	1.07	1.11	1.10	1.08	1.14	1.12
	TVP SV	1.08	1.12	1.11	1.11	1.14	1.15	1.05	1.09	1.08	1.08	1.13	1.13
	BART hom	1.05	1.03	1.07	1.05	1.03	1.09	1.05	1.03	1.06	1.05	1.04	1.07
	BART SV	1.04	1.03	*1.05*	1.05	1.02	1.05	1.03	*1.03*	*1.06*	1.04	1.03	*1.05*
	GP hom	*1.00*	*1.01*	1.07	*1.02*	*0.98*	*1.04*	**0.97**	1.03	1.09	*0.98*	*1.01*	1.06
	GP SV	1.02	1.03	1.09	1.04	1.01	1.05	1.00	1.05	1.14	1.01	1.03	1.08
URX	BVAR hom	1.07	0.98	0.97	1.06	0.96	0.97	1.07	0.99	0.97	1.10	1.00	0.99
	BVAR SV	1.03	1.00	0.96	1.01	1.00	0.98	1.05	1.00	0.94	1.04	1.03	1.00
	BVAR SV−t	1.03	1.00	0.96	1.01	0.99	0.96	1.05	1.00	0.96	1.04	1.02	1.00
	TVP SV	1.00	0.96	**0.94'**	0.97	0.94	**0.94**	1.03	0.98	**0.93**	1.01	0.97	**0.97**
	BART hom	0.97	0.97	1.01	0.97	0.98	1.06	0.97	0.96	0.97	0.98	1.00	1.04
	BART SV	**0.93'**	0.96	0.99	0.93	0.97	1.03	**0.93°**	0.96	0.96	**0.93'**	0.98	1.03
	GP hom	0.94	0.97	1.00	0.92	0.98	1.05	0.95	0.96	0.97	0.95	1.01	1.05
	GP SV	0.93'	**0.93'**	0.98	**0.92**	*0.94*	1.01	0.94'	**0.92**	0.96	0.93'	**0.94**	1.03

M

		CRPS			CRPS−L			CRPS−R			MAE		
	Horizon	1	4	8	1	4	8	1	4	8	1	4	8
YER		0.61	0.55	0.68	0.19	0.16	0.21	0.16	0.16	0.18	0.70	0.64	0.77
	BVAR SV	1.07	1.01	1.01	1.10	1.02	1.00	1.06	**0.99**	1.01	1.03	1.01	1.01
	BVAR SV−t	1.07	1.00	1.01	1.09	1.02	1.00	1.05	0.99	1.02	1.03	1.02	1.01
	TVP SV	1.12	1.02	1.02	1.16	1.03	1.01	1.09	1.01	1.03	1.09	1.02	1.02
	BART hom	1.04	1.01	1.00	1.03	1.00	1.00	1.04	1.01	1.01	1.05	1.01	1.00
	BART SV	1.43	1.07	1.07	1.35	1.08	1.09	1.57	1.06	1.06	1.37	1.06	1.05
	GP hom	1.04	1.00	*0.99*	1.03	0.99	0.98	1.05	1.01	1.01	1.04	0.99	*0.99*
	GP SV	1.08	*1.00*	1.00	1.08	**0.99**	**0.98**	1.08	1.01	1.02	1.05	*0.99*	1.00
HICSA		0.20	0.21	0.23	0.06	0.06	0.07	0.06	0.06	0.07	0.28	0.29	0.32
	BVAR SV	1.02	1.02	0.99	1.04	1.02	0.99	1.02	1.01	0.99	1.02	1.02	0.98
	BVAR SV−t	1.02	1.01	1.00	1.04	1.02	0.99	1.01	1.00	1.00	1.02	1.01	1.01
	TVP SV	1.05	1.06	1.07	1.08	1.08	1.07	1.02	1.04	1.06	1.05	1.06	1.08
	BART hom	**0.98**	**0.96**	**0.96**	**0.99**	**0.96**	**0.98**	*0.97*	**0.97**	**0.96**	**0.95**	**0.95'**	**0.95**
	BART SV	1.01	1.00	1.02	1.04	0.99	1.02	1.00'	1.01	1.02	1.00'	0.99	1.01
	GP hom	1.14	1.24	1.27	1.07	1.08	1.14	1.22	1.40	1.39	1.12	1.26	1.32
	GP SV	1.25	1.28	1.27	1.13	1.12	1.16	1.38	1.43	1.38	1.23	1.31	1.32
URX		0.12	0.14	0.15	0.03	0.04	0.04	0.04	0.05	0.05	0.16	0.18	0.20
	BVAR SV	1.03	1.01	0.98	1.01	1.00	0.98	1.06	1.02	0.99	1.03	1.01	0.99
	BVAR SV−t	1.03	1.00	*0.98*	1.01	0.98	*0.97*	1.05	1.01	0.99	1.03	1.01	1.00
	TVP SV	1.00	*0.96*	0.98	0.95	**0.93**	0.97	1.04	0.99	1.00	0.98	*0.95*	*0.98*
	BART hom	*0.96'*	0.99	1.02	*0.94'*	0.99	1.05	*0.98*	*0.99*	0.99	*0.98*	1.00	1.02
	BART SV	1.01'	1.03°	1.00	0.97°	1.03'	1.03	1.05	1.04°	0.98	1.02'	1.03°	1.00
	GP hom	1.05	1.04	1.03	1.03	1.07	1.09	1.07	1.01	0.98	1.10	1.11	1.10
	GP SV	1.12	1.05	1.04	1.11	1.10	1.11	1.12	1.02	*0.98*	1.17	1.12	1.10

Notes: Real-time data, holdout period from 2001Q2 to 2021Q3. Continuous ranked probability score (CRPS), quantile weighted CRPS for left tail (CRPS-L) and right tail (CRPS-R), mean absolute error (MAE). DM tests: ′ (10 percent), ° (5 percent) and * (1 percent).

depending on the horizon and loss. The performance of GPs deteriorates when moving from the small to medium-sized information set. Its performance with the small information set in the US, however, is noteworthy, with gains of about 10% for one-quarter and one-year ahead predictions of upside-risk. Predictive premia overall appear to be slightly larger for longer-horizon forecasts.

Additional sizable gains over the benchmark, between 7 to 11% for the EA and US respectively, are present in the context of one-quarter ahead (and in some cases, one-year ahead) unemployment forecasts for the small information set. BART and GP show very similar out-of-sample metrics in this case, with the GPs having a slight edge when considering the results for the US. When assessing longer-horizons, these improvements vanish, and the strong performance of the TVP-VAR is worth mentioning for the EA.

The gains for unemployment are present mainly for the small-scale nonparametric model variants. On average, moving from the small to the large information set does not change much for BART in most cases, although there are some improvements particularly for inflation. The GP results for headline inflation and unemployment, however, are worse for the large dataset (albeit they are not significantly worse than the benchmark), which points towards overfitting issues. By contrast, as expected and in line with the previous forecasting literature, larger linear

Table 5: Predictive loss metrics for the US.

S

Variable	Model	CRPS 1	CRPS 4	CRPS 8	CRPS−L 1	CRPS−L 4	CRPS−L 8	CRPS−R 1	CRPS−R 4	CRPS−R 8	MAE 1	MAE 4	MAE 8
GDPC1	BVAR hom	1.08	1.04	1.03	1.06	1.03	1.02	1.09	1.05	1.04	1.15	1.08	1.04
GDPC1	BVAR SV	1.02	1.05	1.04	1.05	1.05	1.03	0.99	1.05	1.05	1.09	1.07	1.03
GDPC1	BVAR SV−t	1.04	1.06	1.04	1.07	1.06	1.03	1.01	1.06	1.05	1.11	1.09	1.03
GDPC1	TVP SV	1.00	1.04	1.02	1.05	1.05	1.02	0.96	1.02	1.03	1.07	1.04	0.99
GDPC1	BART hom	1.01	1.01	0.98*	1.00	1.00	0.99°	1.02	1.01	0.97*	1.07	1.03	0.97*
GDPC1	BART SV	0.99	1.02	1.00	1.03	1.03	1.01	0.95	1.00	1.00	1.04	1.01	0.98'
GDPC1	GP hom	1.02	1.01	1.00	1.04	1.01	1.00	1.00	1.00	1.00	1.12	1.03	0.98'
GDPC1	GP SV	1.02	1.02	1.02	1.06	1.03	1.01	0.97	1.02	1.02	1.13	1.03	0.98'
CPIAUCSL	BVAR hom	1.06	1.06	1.06	1.04	1.05	1.04	1.08	1.07	1.09	1.06	1.05	1.06
CPIAUCSL	BVAR SV	1.04	1.04	1.07	1.06	1.08	1.09	1.01	1.00	1.06	1.05	1.05	1.06
CPIAUCSL	BVAR SV−t	1.06	1.04	1.07	1.07	1.08	1.09	1.05	1.01	1.06	1.07	1.07	1.06
CPIAUCSL	TVP SV	1.04	1.00	1.02	1.06	1.05	1.06	1.01	0.96	1.00	1.05	1.02	1.03
CPIAUCSL	BART hom	1.01	0.98	1.12	1.04	0.99	1.06	0.97	0.97	1.18	1.03	0.99	1.15
CPIAUCSL	BART SV	0.99	0.93'	1.01	1.03	0.97	0.99	0.95	0.89°	1.04	1.02	0.94	1.01
CPIAUCSL	GP hom	0.97	0.93	0.98	1.02	0.97	0.94	0.91'	0.89'	1.03	0.96	0.94	0.98
CPIAUCSL	GP SV	0.97	0.92'	0.97	1.03	0.96	0.95	0.90'	0.87'	1.00	0.95	0.93	0.95
UNRATE	BVAR hom	1.14	1.03	1.02	1.19	1.04	1.03	1.10	1.03	1.02	1.16	1.03	1.02
UNRATE	BVAR SV	1.03	1.06	1.04	1.05	1.07	1.06	1.03	1.06	1.05	1.03	1.00	0.99
UNRATE	BVAR SV−t	1.05	1.07	1.04	1.08	1.07	1.06	1.04	1.07	1.04	1.04	1.01	0.99
UNRATE	TVP SV	1.02	1.07	1.06	1.03	1.09	1.08	1.03	1.07	1.05	1.02	1.01	1.01
UNRATE	BART hom	0.97	1.01	1.01	0.99	1.03	1.01	0.95°	1.00	1.02	0.97	1.02	1.02
UNRATE	BART SV	0.95°	1.00	1.03	0.94°	1.00	1.01	0.96°	1.00	1.04	0.97	1.01	1.03
UNRATE	GP hom	0.95	0.99	1.01	0.94	1.01	1.01	0.96	0.98	1.02	0.97	1.01	1.01
UNRATE	GP SV	0.89'	0.99	1.02	0.82'	1.01	1.02	0.95'	0.98	1.02	0.96	1.00	1.01

M

Variable	Model	CRPS 1	CRPS 4	CRPS 8	CRPS−L 1	CRPS−L 4	CRPS−L 8	CRPS−R 1	CRPS−R 4	CRPS−R 8	MAE 1	MAE 4	MAE 8
GDPC1	BVAR hom	0.44	0.45	0.46	0.13	0.13	0.14	0.13	0.13	0.14	0.53	0.54	0.57
GDPC1	BVAR SV	0.97	1.02	1.02	1.01	1.04	1.03	0.95'	1.01	1.02	1.00	1.01	0.99°
GDPC1	BVAR SV−t	0.97	1.03	1.03	1.01	1.04	1.02	0.93'	1.02	1.03	1.01	1.02	1.00
GDPC1	TVP SV	0.98	1.03	1.04	1.01	1.05	1.04	0.96	1.02	1.05	1.02	1.02	1.01
GDPC1	BART hom	1.02	1.01	0.99	1.03	1.01	1.00	1.00	1.00	0.99	1.01	1.00	0.97*
GDPC1	BART SV	0.98	1.02	1.01	1.02	1.04	1.02	0.93	1.00	1.01	1.01	0.99	0.96*
GDPC1	GP hom	1.04	1.04	1.03	1.07	1.05	1.03	1.03	1.05	1.04	1.00	0.99	0.96*
GDPC1	GP SV	0.99	1.05	1.04	1.04	1.07	1.05	0.96	1.04	1.05	1.00	0.99	0.96*
CPIAUCSL	BVAR hom	0.26	0.31	0.33	0.08	0.09	0.09	0.07	0.09	0.10	0.35	0.41	0.44
CPIAUCSL	BVAR SV	0.97'	0.95°	0.98	1.01	1.01	1.02	0.92*	0.90°	0.95'	0.97'	0.96*	0.96*
CPIAUCSL	BVAR SV−t	0.97	0.96'	0.99	1.01	1.01	1.04	0.93°	0.91°	0.96	0.98	0.97	0.97
CPIAUCSL	TVP SV	0.97	0.97	1.00	1.01	1.04	1.07	0.93°	0.91°	0.95	0.99	0.98	0.98
CPIAUCSL	BART hom	0.95	0.92°	0.98	0.98	0.93°	0.97	0.93'	0.91°	0.98	0.95	0.91°	0.99
CPIAUCSL	BART SV	0.96	0.89*	0.95	0.99	0.93*	0.95	0.92'	0.86°	0.95	0.96	0.89*	0.93'
CPIAUCSL	GP hom	1.16	1.14	1.12	1.08	1.02	1.01	1.26	1.26	1.22	1.08	1.14	1.15
CPIAUCSL	GP SV	1.14	1.12	1.13	1.05	1.03	1.03	1.25	1.23	1.23	1.08	1.12	1.15
UNRATE	BVAR hom	0.22	0.24	0.25	0.06	0.06	0.07	0.07	0.08	0.08	0.26	0.29	0.29
UNRATE	BVAR SV	0.98	1.06	1.07	0.98	1.07	1.08	1.01	1.06	1.07	0.96	1.00	1.02
UNRATE	BVAR SV−t	1.02	1.07	1.07	1.06	1.08	1.08	1.00	1.07	1.06	0.97	1.02	1.02
UNRATE	TVP SV	0.96'	1.05	1.06	0.95'	1.07	1.08	0.98	1.05	1.06	0.94	1.00	0.99
UNRATE	BART hom	0.99	1.01	1.02	0.97	1.01	1.01	1.00	1.01	1.02	0.98	1.00	1.02
UNRATE	BART SV	0.95	1.04	1.05	0.86	1.04	1.04	1.02	1.04	1.05	1.00	1.02	1.03
UNRATE	GP hom	1.09	1.06	1.07	1.07	1.06	1.06	1.12	1.07	1.09	1.04	1.05	1.07
UNRATE	GP SV	1.00	1.07	1.10	0.91	1.07	1.09	1.09	1.09	1.11	1.03	1.03	1.06

Horizon

Notes: Pseudo out-of-sample scheme, holdout period from 1985Q1 to 2022Q4, expanding window. Continuous ranked probability score (CRPS), quantile weighted CRPS for left tail (CRPS-L) and right tail (CRPS-R), mean absolute error (MAE). DM tests: ′ (10 percent), ° (5 percent) and * (1 percent).

BVARs tend to improve forecasts, particularly so for the US. BART does not appear to suffer to this extent from the issue, in line with our discussion about regression trees and in-sample versus out-of-sample overfitting in Section 3. In the next subsection, we zoom into specifics and find that the weaker performance of medium-scale GPs is driven by specific economic periods and varies across the predictive distribution.

On a general note, however, nonlinearities tend to become less important to improve predictions as the size of the information set increases. The following argument is in line with the flexibility provided by TVPs becoming less important for larger models, see, e.g., Huber *et al.* (2021). When there is risk of, for instance, omitted variables in small-scale models, the nonparametric features may offset such misspecification. This results in the margins between linear and nonlinear models for the small information set being slightly bigger on average. These margins in many cases can reduce (or even vanish) when moving to the medium-sized information set. The additional flexibility does not improve forecasts when there is more information to be exploited, and linearity appears to be a good approximation for larger-scale models, particularly so for the EA.

Commenting on the necessity of heteroskedastic features, it is worth noting that linear models tend to improve when adding SV (at least SVs never significantly hurt predictive accuracy), particularly in the US. This is a well-known result (see, e.g., Clark, 2011). For the EA, adding SVs is not as crucial, likely due to the shorter available time period. By contrast, there are overall no systematic patterns in favor or against SVs for BART or GP. This less important role of SVs has also been described in Clark *et al.* (2023), and can be explained by the notion that the flexible conditional mean functions, particularly for BART, can already address particular forms of heteroskedastic patterns in the data. In many cases, the homoskedastic variant is very close to the heteroskedastic one, but usually it does not hurt to add SV. There is an exception to this, however, namely inflation forecasts with medium-sized models in the EA.

With respect to accuracy in different parts of the distributions, we find that the CRPS often looks like an average of CRPS-L and CRPS-R. Typically, these two tail forecast metrics are very similar, although there are some noteworthy exceptions. An example are inflation forecasts for the US, where BART shows more sizable gains when the focus is on upward-risk (CRPS-R). Moreover, we find that the relative metrics for overall density forecasts, measured by the CRPS, very often agree with those for point forecasts, the MAE. For this comparison, we observe an exception for two-year ahead forecasts of US GDP, where the nonparametric models yield modest gains of up to 4% for point forecasts but they do not improve density forecasts.

Performance over time

Next we investigate whether the forecast performance differs across economic phases and distinct periods, and if so, when. This sheds light on particular circumstances which render nonlinearities more important, and thus the discussed nonparametric methods particularly useful. To economize on space, we show only the one-quarter ahead forecasts in Tables 6 and 7 due to the practical relevance of this horizon. Many patterns we describe are similar across all horizons we consider, but we note that there are some differences, particularly for short-horizon forecasts. Tables of these additional results are provided in the Online Appendix.

Starting with GDP, we find that averaging over the full holdout sample masks some noteworthy gains from using GPs during recessions in both the EA and the US. In particular, we find overall improvements for density forecasts of about 5% with the medium versions. More importantly, they show even more sizable improvements when considering CRPS-L, of up to almost 10% lower values. In fact, CRPS-L tracks the lower tail of the distribution of GDP, and is thus closely linked to measuring growth-at-risk. And such downside-risk to GDP is naturally

Table 6: One-year ahead predictive loss metrics for the EA.

S

		CRPS					CRPS−L					CRPS−R					MAE				
		FS	Re	Ex	Pre	Post	FS	Re	Ex	Pre	Post	FS	Re	Ex	Pre	Post	FS	Re	Ex	Pre	Post
YER	BVAR hom	1.01	1.01	1.01	1.03	0.99	1.01	1.00	1.01	1.02	1.00	1.01	1.02	1.01	1.04	0.99	1.01	1.02	1.01	1.02	1.00
YER	BVAR SV	1.01	1.01	1.01	1.01	1.00	1.02	1.02	1.02	1.03	1.01	1.00	1.00	1.00	0.99	1.00	1.01	1.00	1.01	1.02	1.00
YER	BVAR SV−t	1.00	1.00	*1.01*	1.01	1.00	1.01	1.01	1.02	1.01	1.01	1.00	0.99°	1.00	**0.99**	1.00	1.01	0.99°	1.03	1.02	1.00
YER	TVP SV	1.01	1.00	1.01	1.01	1.00	1.01	1.01	1.02	1.02	1.01	1.00	0.99'	1.00	0.99	1.00	1.00	0.99	1.01	1.01	1.00
YER	BART hom	1.01	1.01	1.01	1.02	1.00	1.01	1.00	*1.01*	1.01	1.00	1.02	1.02	1.01	1.04	1.00	1.01	1.02	1.00	1.01	1.01
YER	BART SV	1.01	1.00	1.01	1.00	1.02	1.02	1.01	1.03	1.01	1.03	1.00	1.00	1.01	1.00	1.01	1.01	1.00	1.01	1.00	1.01
YER	GP hom	1.00	*0.97*	1.02	1.02	*0.99*	*1.00*	*0.95'*	1.05	1.01	*0.98*	1.01	1.00	1.01	1.03	0.99	0.99	0.99	1.00	0.98	1.00
YER	GP SV	**0.99**	0.97'	1.01	*1.00*	0.99	1.00	0.96'	1.05	*1.00*	1.00	*0.99*	**0.98**	*0.99*	1.01	*0.97*	**0.99**	*0.98*	**1.00**	**0.98**	*1.00*
HICSA	BVAR hom	1.11	1.15	1.09	1.11	1.07	1.11	1.13	1.11	1.12	1.06	1.09	1.16	1.07	1.09	1.07	1.11	1.18	1.09	1.12	1.07
HICSA	BVAR SV	1.10	1.17	1.08	1.11	1.05	1.12	1.15	1.11	1.13	1.06	1.07	1.17	1.05	1.08	1.04	1.10	1.20	1.08	1.11	1.06
HICSA	BVAR SV−t	1.13	1.14	1.12	1.13	1.08	1.14	1.10	1.15	1.15	1.08	1.11	1.20	1.08	1.11	1.08	1.14	1.18	1.13	1.15	1.08
HICSA	TVP SV	1.12	1.18	1.10	1.13	1.05	1.14	1.16	1.14	1.16	1.04	1.09	1.19	1.06	1.09	1.06	1.13	1.21	1.11	1.14	1.07
HICSA	BART hom	1.03	0.99	1.04	1.04	0.95	1.03	0.97	1.05	1.04	0.97	1.03	*1.01*	1.03	1.05	0.94	1.04	1.03	1.04	1.05	0.95
HICSA	BART SV	1.03	*0.98*	1.04	1.04	0.93	1.02	*0.94*	1.05	1.04	0.91	*1.03*	1.02	1.03	*1.04*	0.94	1.03	*1.01*	1.03	1.04	0.93
HICSA	GP hom	*1.01*	1.01	*1.00*	*1.03*	0.87	*0.98*	0.98	*0.99*	*1.00*	0.89	1.03	1.05	*1.02*	1.06	0.84	*1.01*	1.02	*1.01*	*1.03*	0.92
HICSA	GP SV	1.03	1.04	1.03	1.06	**0.85**	1.01	1.00	1.01	1.03	**0.85**	1.05	1.09	1.04	1.10	*0.83*	1.03	1.07	1.02	1.06	**0.88**
URX	BVAR hom	0.98	0.98	0.98	1.00	0.89	0.96	1.01	0.94	0.99	0.83	0.99	0.96	1.01	1.00	0.95	1.00	1.00	1.01	1.02	0.91
URX	BVAR SV	1.00	0.96	1.01	1.02	0.91	1.00	1.01	0.99	1.03	0.86	1.00	0.94'	1.03	1.01	0.96	1.03	1.00	1.04	1.05	0.93
URX	BVAR SV−t	1.00	0.95°	1.02	1.02	0.87	0.99	0.99	0.99	1.03	0.83	1.00	0.94°	1.04	1.02	0.92	1.02	0.98'	1.04	1.04	0.92
URX	TVP SV	0.96	0.96	*0.96*	0.98	**0.84**	0.94	1.01	*0.91*	0.98	**0.75**	0.98	0.94'	0.99	0.99	**0.92**	0.97	0.98	0.97	1.00	**0.84**
URX	BART hom	0.97	0.91'	1.00	0.97	0.97	0.98	0.95	0.99	0.98	0.94	0.96	0.89'	1.01	0.96	1.00	1.00	0.94	1.03	1.00	0.98
URX	BART SV	0.96	0.90'	0.99	0.97	0.95	0.97	0.94	0.98	0.98	0.91	0.96	0.89'	1.00	0.96	0.98	0.98	0.92	1.01	0.99	0.93
URX	GP hom	0.97	0.88'	1.01	0.97	0.96	0.98	0.95	1.00	0.99	0.92	0.96	0.84'	1.03	0.95	0.99	1.01	0.93	1.04	1.01	0.99
URX	GP SV	**0.93'**	*0.86'*	0.96	**0.93**	0.92	*0.94*	*0.93*	0.94	**0.95**	0.87	**0.92**	*0.81'*	*0.98*	**0.91**	0.97	**0.94**	*0.90*	*0.96*	**0.95**	0.89

M

		CRPS					CRPS−L					CRPS−R					MAE				
		FS	Re	Ex	Pre	Post	FS	Re	Ex	Pre	Post	FS	Re	Ex	Pre	Post	FS	Re	Ex	Pre	Post
YER	BVAR hom	0.55	1.37	0.39	0.29	3.19	0.16	0.51	0.09	0.10	0.81	0.16	0.29	0.14	0.08	1.03	0.64	1.63	0.45	0.37	3.47
YER	BVAR SV	1.01	1.00	1.01	1.01	1.00	1.02	1.01	1.03	1.00	1.04	**0.99**	1.00	**0.99**	1.02	0.97	1.01	1.00	1.02	1.01	1.01
YER	BVAR SV−t	1.00	0.98'	1.02	1.01	1.00	1.02	0.97°	1.07	0.99	1.05	0.99	1.00	0.99	1.03	**0.97**	1.02	0.99	1.03	1.01	1.03
YER	TVP SV	1.02	1.01	1.03	1.03	1.01	1.03	1.01	1.05	1.02	1.04	1.01	1.01	1.01	1.05	0.99	1.02	1.01	1.02	1.02	1.02
YER	BART hom	1.01	1.00	1.01	1.02	0.99	1.00	0.99°	1.02	1.01	0.99	1.01	1.01	1.01	1.03	0.99	1.01	1.00	1.01	1.01	1.00
YER	BART SV	1.07	1.01	1.05	1.04	1.00	1.08	1.02	1.04	1.03	1.03	1.06	1.01	1.06	1.04	0.98	1.06	1.01	1.04	1.03	1.00
YER	GP hom	1.00	**0.95***	1.03	1.02	**0.98'**	0.99	**0.92***	1.07	1.00	**0.98**	1.01	1.01	1.02	1.05	0.99	0.99	0.98*	1.00	0.99	**1.00**
YER	GP SV	*1.00*	0.95*	1.03	1.00	0.99	**0.99**	0.93*	1.05	**0.99**	0.99	1.01	*0.99*	1.01	1.03	0.99	*0.99*	**0.97***	*1.00*	*0.98*	1.00
HICSA	BVAR hom	0.21	0.29	0.19	0.20	0.34	0.06	0.09	0.06	0.06	0.08	0.06	0.08	0.06	0.06	0.11	0.29	0.37	0.28	0.27	0.47
HICSA	BVAR SV	1.02	0.97	1.03	1.01	1.05	1.02	0.96	1.04	1.02	1.05	1.01	0.99	1.01	1.00	1.04	1.02	0.98	1.03	1.01	1.05
HICSA	BVAR SV−t	1.01	1.00	1.01	1.01	1.02	1.02	1.00	1.03	1.02	1.06	1.00	1.02	0.99	1.00	0.98	1.01	1.00	1.01	1.00	1.06
HICSA	TVP SV	1.06	0.98	1.09	1.07	1.02	1.08	0.98	1.12	1.09	1.03	1.04	0.99	1.06	1.05	1.01	1.06	0.98	1.08	1.06	1.05
HICSA	BART hom	**0.96**	**0.92**	**0.97**	**0.96**	0.98	**0.96**	**0.92**	**0.97**	**0.95**	0.97	**0.97**	**0.93'**	**0.98**	**0.96**	0.98	**0.95'**	**0.93**	**0.95**	**0.94**	0.97
HICSA	BART SV	1.00	0.95	1.01	1.00	0.95	0.99	0.94	1.00	0.99	*0.95*	1.01	0.96	1.01	1.01	0.95	0.99	0.96	0.99	0.99	0.96
HICSA	GP hom	1.24	1.06	1.30	1.30	*0.89*	1.08	0.94	1.13	1.10	0.95	1.40	1.21	1.45	1.52	**0.82**	1.26	1.14	1.29	1.31	*0.95*
HICSA	GP SV	1.28	1.10	1.33	1.34	0.91	1.12	0.98	1.17	1.14	0.98	1.43	1.25	1.48	1.55	0.83	1.31	1.18	1.34	1.36	0.99
URX	BVAR hom	0.14	0.28	0.11	0.12	0.27	0.04	0.06	0.03	0.03	0.07	0.05	0.10	0.03	0.04	0.09	0.18	0.35	0.15	0.17	0.32
URX	BVAR SV	1.01	1.02	1.00	1.01	0.99	1.00	1.02	0.99	1.00	0.97	1.02	1.03	1.01	1.02	1.02	1.01	1.02	1.01	1.01	1.01
URX	BVAR SV−t	1.00	1.01	0.99	0.99	1.02	0.98	1.04	0.95'	0.98	0.97	1.01	1.00	1.02	1.00	1.07	1.01	1.03	1.01	1.00	1.08
URX	TVP SV	*0.96*	1.05	**0.92'**	*0.98*	*0.90*	**0.93**	1.07	**0.88°**	*0.95*	*0.82*	0.99	1.05	**0.96**	1.00	*0.97*	*0.95*	1.04	**0.90°**	*0.96*	*0.88°*
URX	BART hom	0.99	0.98	1.00	1.00	0.96	0.99	0.99	0.99	1.00	0.93	*0.99*	0.98	1.00	*1.00*	0.98	1.00	0.99	1.01	1.02	0.94
URX	BART SV	1.03°	0.98	1.02'	1.04'	0.94	1.03'	0.98	1.02	1.04	0.90	1.04°	0.98	1.03°	1.04°	0.98	1.03°	0.98	1.03'	1.04	0.92
URX	GP hom	1.04	0.80°	1.16	1.06	0.96	1.07	0.92'	1.13	1.10	0.92	1.01	0.73°	1.18	1.02	1.00	1.11	0.88'	1.21	1.12	1.02
URX	GP SV	1.05	**0.80°**	1.19	1.07	0.96	1.10	**0.91**	1.17	1.13	0.91	1.02	**0.73°**	1.19	1.02	1.00	1.12	**0.87'**	1.23	1.13	1.03

a key issue during recessions, making these gains particularly relevant. Interestingly, this appears to relate specifically to how GPs estimate the conditional mean functions, because these improvements are non-existent for BART.

Dissecting the comparatively weak full sample performance for the medium-scale GP variants (for unemployment, and even more so, inflation) reveals that these relatively worse metrics are mainly due to expansionary (i.e., non-recessionary) economic phases. In addition, it is worth noting that CRPS-L and CRPS-R differ substantially especially for the expansionary sub-sample. The poor overall density forecast performance is in large parts due to the right tail of the respective predictive distributions during expansions. Zooming into these sub-samples

Table 7: One-year ahead predictive loss metrics for the US.

S

		CRPS					CRPS−L					CRPS−R					MAE				
		FS	Re	Ex	Pre	Post	FS	Re	Ex	Pre	Post	FS	Re	Ex	Pre	Post	FS	Re	Ex	Pre	Post
GDPC1	BVAR hom	1.04	1.02	1.05	1.04	1.04	1.03	1.02	1.03	1.02	1.04	1.05	1.01	1.06	1.05	1.05	1.08	1.02	1.12	1.09	1.06
GDPC1	BVAR SV	1.05	1.03	1.06	1.02	1.12	1.05	1.06	1.04	1.01	1.13	1.05	0.99'	1.07	1.01	1.14	1.07	1.01	1.11	1.07	1.07
GDPC1	BVAR SV−t	1.06	1.03	1.08	1.03	1.13	1.06	1.06	1.06	1.03	1.13	1.06	1.00	1.08	1.03	1.15	1.09	1.01	1.13	1.09	1.09
GDPC1	TVP SV	1.04	1.01	1.05	1.00	1.12	1.05	1.04	1.06	1.01	1.14	1.02	0.97*	1.04	0.98	1.13	1.04	0.99	1.06	1.03	1.05
GDPC1	BART hom	1.01	0.99	*1.01*	1.00	*1.01*	*1.00*	1.00	*1.00*	**0.99**	*1.02*	1.01	0.99'	1.02	1.01	*1.00*	1.03	0.99	1.05	1.03	1.02
GDPC1	BART SV	1.02	1.01	1.02	*0.99*	1.08	1.03	1.04	1.02	1.00	1.10	1.00	0.97*	1.02	*0.98'*	1.07	*1.01*	0.98'	*1.02*	*1.01*	1.02
GDPC1	GP hom	*1.01*	*0.97°*	1.03	1.00	1.02	1.01	*0.97'*	1.04	1.00	1.03	**1.00**	0.97°	*1.01*	1.00	1.00	1.03	**0.97'**	1.06	1.03	1.02
GDPC1	GP SV	1.02	0.98	1.05	1.00	1.08	1.03	0.98	1.06	0.99	1.10	1.02	**0.96°**	1.04	0.99	1.07	1.03	0.97'	1.06	1.03	*1.01*
CPIAUCSL	BVAR hom	1.06	1.04	1.06	1.08	0.97	1.05	1.04	1.05	1.05	1.04	1.07	1.02	1.08	1.11	0.92	1.05	1.05	1.06	1.07	1.00
CPIAUCSL	BVAR SV	1.04	1.03	1.04	1.05	0.97	1.08	1.06	1.09	1.08	1.10	1.00	0.98	1.00	1.03	0.86	1.05	1.04	1.06	1.06	1.04
CPIAUCSL	BVAR SV−t	1.04	*0.99*	1.06	1.06	0.97	1.08	*1.01*	1.10	1.07	1.10	1.01	0.97	1.02	1.05	0.87	1.07	1.01	1.08	1.07	1.04
CPIAUCSL	TVP SV	1.00	1.04	0.99	1.03	0.89	1.05	1.08	1.04	1.05	1.06	0.96	1.00	0.95	1.02	0.75	1.02	1.05	1.02	1.03	0.98
CPIAUCSL	BART hom	0.98	1.00	0.97	1.01	0.86	0.99	1.02	0.97	0.99	0.95	0.97	0.97	0.97	1.02	0.79	0.99	1.01	0.98	1.01	0.89
CPIAUCSL	BART SV	0.93'	0.99	0.91'	0.96	0.79	0.97	1.03	0.96	0.98	0.93	0.89°	*0.94*	0.87'	*0.94'*	0.68	0.94	*0.98*	0.93	0.96	0.85
CPIAUCSL	GP hom	0.93	1.01	0.91	0.97	0.78	0.97	1.01	0.95	0.97	0.94	0.89'	1.00	0.87'	0.96	0.64	0.94	1.00	0.92	0.96	0.87
CPIAUCSL	GP SV	*0.92'*	0.99	*0.90'*	*0.96*	*0.75*	*0.96*	1.02	*0.94*	*0.97*	*0.91*	*0.87'*	0.96	*0.85'*	0.94'	*0.61*	*0.93*	0.99	*0.92*	*0.95*	*0.84*
UNRATE	BVAR hom	1.03	1.00	1.06	1.05	1.01	1.04	1.01	1.05	1.07	1.01	1.03	0.99°	1.08	1.04	1.02	1.03	1.01	1.05	1.04	1.02
UNRATE	BVAR SV	1.06	1.01	1.09	1.01	1.11	1.07	0.99	1.10	1.01	1.13	1.06	1.03	1.10	1.00	1.12	**1.00**	1.00	**0.99**	1.01	**0.98**
UNRATE	BVAR SV−t	1.07	1.02	1.10	1.01	1.12	1.07	1.00	1.11	1.01	1.14	1.07	1.03	1.11	1.02	1.13	1.01	1.01	1.00	1.02	0.99
UNRATE	TVP SV	1.07	1.01	1.11	1.01	1.13	1.09	0.99	1.13	1.03	1.15	1.07	1.03	1.11	1.00	1.14	1.01	1.00	1.02	1.02	1.01
UNRATE	BART hom	1.01	0.97*	1.04	1.00	1.02	1.03	0.97*	1.05	1.03	1.03	1.00	0.96*	1.03	0.98	1.02	1.02	0.97*	1.05	1.01	1.04
UNRATE	BART SV	1.00	0.99	*1.00*	0.99	*1.01*	**1.00**	0.97*	*1.01*	*0.99*	*1.00*	1.00	1.01	**0.99**	0.99	*1.01*	1.01	0.99°	1.03	1.00	1.02
UNRATE	GP hom	**0.99**	*0.96°*	1.02	0.96	1.02	1.01	**0.96***	1.03	1.00	1.02	0.98	*0.96'*	1.00	0.93'	1.03	1.01	**0.96°**	1.05	0.99	1.04
UNRATE	GP SV	0.99	0.96'	1.02	**0.96**	1.03	1.01	0.96°	1.04	1.00	1.03	**0.98**	0.96	0.99	**0.93'**	1.03	1.00	0.97'	1.02	*0.99*	1.02

M

		CRPS					CRPS−L					CRPS−R					MAE				
		FS	Re	Ex	Pre	Post	FS	Re	Ex	Pre	Post	FS	Re	Ex	Pre	Post	FS	Re	Ex	Pre	Post
GDPC1	BVAR hom	0.45	1.44	0.33	0.34	1.72	0.13	0.51	0.09	0.10	0.54	0.13	0.33	0.11	0.11	0.46	0.54	1.79	0.39	0.42	1.93
GDPC1	BVAR SV	1.02	1.01	1.03	**0.98'**	1.12	1.04	1.03	1.04	1.00	1.11	1.01	0.98*	1.03	0.97*	1.14	1.01	1.00	1.01	1.00	1.02
GDPC1	BVAR SV−t	1.03	1.02	1.04	0.99	1.13	1.04	1.04	1.05	*1.00*	1.13	1.02	0.98*	1.03	0.97*	1.14	1.02	1.00	1.03	1.02	1.02
GDPC1	TVP SV	1.03	1.01	1.04	1.00	1.11	1.05	1.03	1.06	1.01	1.12	1.02	0.98*	1.03	0.98	1.12	1.02	1.00	1.03	1.02	1.02
GDPC1	BART hom	1.01	0.99°	1.02	1.00	1.02	1.01	0.98°	1.04	1.01	1.02	1.00	1.00	1.01	1.00	1.02	1.00	0.99	1.00	0.98	1.03
GDPC1	BART SV	1.02	1.00	1.03	0.98	1.10	1.04	1.02	1.05	1.01	1.11	1.00	*0.98**	1.02	**0.96°**	1.11	0.99	0.99'	**0.99**	0.98°	1.02
GDPC1	GP hom	1.04	**0.95***	1.09	1.06	1.01	1.05	**0.91***	1.15	1.07	1.00	1.05	1.00	1.06	1.06	1.02	0.99	*0.97°*	1.00	0.98	1.02
GDPC1	GP SV	1.05	0.97*	1.09	1.03	1.10	1.07	0.96°	1.15	1.06	1.11	1.04	0.98'	1.06	1.01	1.11	**0.99**	0.97°	0.99	**0.97'**	1.02
CPIAUCSL	BVAR hom	0.31	0.62	0.27	0.27	0.71	0.09	0.21	0.08	0.08	0.19	0.09	0.15	0.08	0.08	0.23	0.41	0.80	0.36	0.36	0.92
CPIAUCSL	BVAR SV	0.95°	0.95°	0.96'	0.97°	0.87	1.01	0.97	1.02	1.01	1.01	0.90°	0.92*	0.90°	0.94*	0.75	0.96*	0.96°	0.96°	0.96*	0.96
CPIAUCSL	BVAR SV−t	0.96'	0.93'	0.97	0.99	0.83	1.01	0.96	1.03	1.02	0.98	0.91°	0.91*	0.91'	0.96°	0.70	0.97	0.95'	0.98	0.99	0.92
CPIAUCSL	TVP SV	0.97	0.95*	0.98	0.99	0.87	1.04	0.97'	1.06	1.03	1.05	0.91°	0.94*	0.90'	0.96°	0.71	0.98	0.97°	0.98	0.98	0.97
CPIAUCSL	BART hom	0.92°	0.94°	0.91°	0.93°	0.86	**0.93°**	0.97'	**0.91'**	**0.93'**	0.89	0.91°	0.90°	0.91'	0.93°	0.84	0.91°	0.93°	0.90°	0.92°	0.87
CPIAUCSL	BART SV	**0.89***	**0.93°**	**0.88°**	**0.92***	0.76	0.93*	0.97	0.92°	0.94°	0.90	**0.86°**	**0.89°**	**0.85°**	**0.91***	0.63	**0.89***	**0.92°**	**0.88***	**0.90***	0.84
CPIAUCSL	GP hom	1.14	0.95	1.19	1.23	0.73	1.02	**0.94'**	1.04	1.05	**0.84**	1.26	0.98	1.33	1.42	0.65	1.14	0.97	1.18	1.22	0.79
CPIAUCSL	GP SV	1.12	0.97	1.17	1.22	**0.71**	1.03	0.98	1.04	1.06	0.88	1.23	0.98	1.28	1.40	**0.56**	1.12	0.98	1.16	1.20	**0.78**
UNRATE	BVAR hom	0.24	0.97	0.15	0.13	1.45	0.06	0.19	0.05	0.04	0.37	0.08	0.37	0.04	0.04	0.46	0.29	1.10	0.19	0.17	1.55
UNRATE	BVAR SV	1.06	1.01	1.10	0.98'	1.15	1.07	0.99°	1.11	**0.98'**	1.17	1.06	1.03	1.10	0.98	1.15	1.00	1.00	1.01	1.00	1.01
UNRATE	BVAR SV−t	1.07	1.01	1.11	0.98	1.16	1.08	0.99°	1.12	0.98	1.18	1.07	1.03	1.11	0.98	1.17	1.02	1.00	1.03	1.00	1.05
UNRATE	TVP SV	1.05	1.01	1.08	*0.97°*	1.14	1.07	0.99'	1.11	0.99	1.15	1.05	1.02	1.07	*0.95**	1.16	*1.00*	1.00	*1.00*	**0.98**	1.03
UNRATE	BART hom	1.01	0.99*	1.02	1.02	1.00	1.01	1.00	1.01	1.01	1.00	1.01	0.98*	1.04	1.02	**1.00**	1.00	0.99'	1.01	1.01	1.00
UNRATE	BART SV	1.04	1.00	1.06	1.00	1.08	1.04	*0.98**	1.07	1.00	1.08	1.04	1.02	1.06	0.99	1.09	1.02	0.99	1.03	1.02	1.01
UNRATE	GP hom	1.06	**0.94***	1.15	1.12	**1.00**	1.06	1.00	1.08	1.12	**0.99**	1.07	**0.92***	1.23	1.14	1.00	1.05	*0.97**	1.11	1.08	1.01
UNRATE	GP SV	1.07	0.96*	1.16	1.07	1.08	1.07	0.99'	1.11	1.07	1.08	1.09	0.94*	1.24	1.08	1.09	1.03	0.97*	1.08	1.05	1.02

actually suggests that we need to redeem the medium-sized GPs, since they are indeed the best performing models for unemployment forecasts during recessions, with improvements for the EA of about 20% lower CRPSs, and almost 30% lower CRPS-R, which measure the economically important accuracy for predicting upward-risk to unemployment. A similar pattern is observable for the US, albeit with smaller magnitudes.

Turning to inflation specifically, the most striking finding is the consistent solid performance of BART across all sample splits in both economies. The homoskedastic version is superior in the EA, while the SVs contribute modest further improvements for the US. We again observe somewhat better relative metrics in recessions when compared to expansions for the EA, especially for the medium-sized information set. Interestingly, this is not the case for the US, where some of the gains are larger in expansions than recessions by a few percentage points. This suggests that nonlinearities are also present in recovery periods from economic downturns, and these are best picked up by BART. For instance, we observe between 12 to 15% decreases in the predictive losses (depending on the specific loss), relative to the benchmark for BART-SV in expansions.

Last, we distinguish between the pre- and post-pandemic periods. Most importantly, when assessing the raw predictive losses we find them to be an order of magnitude larger when comparing the latter to the former period. We begin with listing several patterns with respect to the BVARs. Looking at the relative performance metrics on aggregate for GDP and unemployment, we find that they are roughly stable for the EA. In the US, the relative numbers increase modestly. This is different for inflation in the US, where particularly the BVAR-SV and SV-t variants show a better relative performance benchmarked against the homoskedastic one. Especially the t-distributed errors appear to offer some gains. When turning to the nonparametrics, there are several improvements for post-pandemic GDP and unemployment forecasts for the EA (but not the US). But these improvements pale in comparison to those for inflation, especially so in the US. The recent inflation surge is picked up satisfactorily by both BART and the GP, with between 12 to 15% improvements over the benchmark depending on the loss function in the EA, and between 20 to almost 40% in the US. It is worth mentioning that the largest relative gains in magnitude arise for upside-risk to inflation, which is indeed arguably the most relevant metric given the context of observed inflation dynamics during this period.

Odds and ends

While we have found many commonalities between the EA and the US, there are some differences which deserve explanations. The EA aggregates are averages of a large and heterogeneous block of individual countries, and these averages perhaps smooth out country-level dynamics that could be exploited by the nonparametric models (as in, e.g., Huber *et al.*, 2023). Further, publication lags in information set, i.e., for the euro area, imply that the one-step ahead forecast is *de facto* already a multi-step ahead forecast due to the ragged edges. It is also worth mentioning that there are three recessions in the EA (the global financial crisis, sovereign debt crisis and pandemic, with the latter being not a conventional recession), and four in the US. In general, the evaluation period for the US is much longer, and the ratio between recession/expansion quarters is much smaller in the US. This indeed also has implications with respect to the DM tests (e.g., the lack of significance for some of the subsamples is explained by the small number of observations used to construct the test statistic). Further, the recent high-inflation period is only partly captured in the EA dataset.

The EA as an average of many individual countries, alongside the much longer sampling period in the US, and potential timing heterogeneities arising from the real-time versus pseudo out-of-sample forecast evaluation schemes explain some of the differences between the two economies we discussed above.

5. CLOSING REMARKS

In this chapter we have reviewed specification and estimation of multivariate Bayesian nonparametric models for forecasting macroeconomic and financial variables. To flexibly model the conditional mean functions, we focus on BART and GPs. Efficient estimation is enabled by relying on a FSV specification of the reduced form errors, which also addresses heteroskedastic patterns and co-movement across a potentially large set of time series of interest.

How these methods fit data, and which kinds of patterns they are capable of detecting, is illustrated with single-equation examples using US data for inflation, unemployment and output. We subsequently apply small and medium sized multivariate models for point, density and tail forecasting using a EA and US dataset. Various metrics of predictive accuracy are compared to variants of BVARs and TVP-VARs, and particularly the medium-scale BVAR-SV is an often used and very capable benchmark.

We find some gains in predictive accuracy for the nonparametric approaches. Most notably, they perform well for short-run forecasts of unemployment (especially GPs), and longer-run predictions of inflation (with BART being particularly strong). There is some evidence for differentials in performance over the business cycle, and for non-standard economic phases such as during and after the Covid-19 pandemic. Nonlinearities as captured by BART and the GPs are particularly helpful during recessions, and they do well in predicting the recent surge of inflation.

REFERENCES

Aguilar O, and West M (2000), "Bayesian dynamic factor models and portfolio allocation," *Journal of Business & Economic Statistics* **18**(3), 338–357.

Banbura M, Brenna F, Paredes J, and Ravazzolo F (2021), "Combining Bayesian VARs with survey density forecasts: Does it pay off?" *ECB Working Paper* **2543**.

Billio M, Casarin R, and Rossini L (2019), "Bayesian nonparametric sparse VAR models," *Journal of Econometrics* **212**(1), 97–115.

Bitto A, and Frühwirth-Schnatter S (2019), "Achieving shrinkage in a time-varying parameter model framework," *Journal of Econometrics* **210**(1), 75–97.

Breiman L (2001), "Random forests," *Machine Learning* **45**, 5–32.

Cadonna A, Frühwirth-Schnatter S, and Knaus P (2020), "Triple the gamma–a unifying shrinkage prior for variance and variable selection in sparse state space and TVP models," *Econometrics* **8**(2), 20.

Carriero A, Clark TE, Marcellino M, and Mertens E (202X), "Addressing COVID-19 outliers in BVARs with stochastic volatility," *Review of Economics and Statistics* **in-press**, 1–38.

Carvalho CM, Polson NG, and Scott JG (2010), "The horseshoe estimator for sparse signals," *Biometrika* **97**(2), 465–480.

Chan JC (2023), "Comparing stochastic volatility specifications for large Bayesian VARs," *Journal of Econometrics* **235**(2), 1419–1446.

Chan JC, Eisenstat E, and Yu X (2022), "Large Bayesian VARs with factor stochastic volatility: Identification, order invariance and structural analysis," *arXiv* **2207.03988**.

Chan JC, Koop G, Poirier DJ, and Tobias JL (2020), *Bayesian Econometric Methods*, Cambridge: Cambridge University Press, 2 edition.

Chipman HA, George EI, and McCulloch RE (2010), "BART: Bayesian additive regression trees," *The Annals of Applied Statistics* **4**(1), 266–298.

Clark TE (2011), "Real-time density forecasts from Bayesian vector autoregressions with stochastic volatility," *Journal of Business & Economic Statistics* **29**(3), 327–341.

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.

Cross JL, Hou C, and Poon A (2020), "Macroeconomic forecasting with large Bayesian VARs: Global-local priors and the illusion of sparsity," *International Journal of Forecasting* **36**(3), 899–915.

Diebold FX, and Mariano RS (1995), "Comparing predictive accuracy," *Journal of Business & Economic Statistics* **13**, 253–263.

Ferrara L, Marcellino M, and Mogliani M (2015), "Macroeconomic forecasting during the Great Recession: The return of non-linearity?" *International Journal of Forecasting* **31**(3), 664–679.

Frühwirth-Schnatter S, and Wagner H (2010), "Stochastic model specification search for Gaussian and partial non-Gaussian state space models," *Journal of Econometrics* **154**(1), 85–100.

Giacomini R, and Komunjer I (2005), "Evaluation and combination of conditional quantile forecasts," *Journal of Business & Economic Statistics* **23**(4), 416–431.

Gneiting T, and Raftery AE (2007), "Strictly proper scoring rules, prediction, and estimation," *Journal of the American Statistical Association* **102**(477), 359–378.

Gneiting T, and Ranjan R (2011), "Comparing density forecasts using threshold-and quantile-weighted scoring rules," *Journal of Business & Economic Statistics* **29**(3), 411–422.

Goulet Coulombe P (2020), "To bag is to prune," *arXiv* **2008.07063**.

——— (2022), "A Neural Phillips Curve and a Deep Output Gap," *arXiv* **2202.04146**.

Goulet Coulombe P, Marcellino M, and Stevanović D (2021), "Can machine learning catch the Covid-19 recession?" *National Institute Economic Review* **256**, 71–109.

Gu S, Kelly B, and Xiu D (2021), "Autoencoder asset pricing models," *Journal of Econometrics* **222**(1), 429–450.

Hamilton JD (1989), "A new approach to the economic analysis of nonstationary time series and the business cycle," *Econometrica* **57**(2), 357–384.

Hastie T, Tibshirani R, and Friedman JH (2009), *The elements of statistical learning: Data mining, inference, and prediction*, volume 2, New York, NY: Springer.

Hauzenberger N, Huber F, Klieber K, and Marcellino M (2022), "Enhanced Bayesian Neural Networks for Macroeconomics and Finance," *arXiv* **2211.04752**.

Hauzenberger N, Huber F, Marcellino M, and Petz N (2021), "Gaussian process vector autoregressions and macroeconomic uncertainty," *arXiv* **2112.01995**.

Hirano K (2002), "Semiparametric Bayesian inference in autoregressive panel data models," *Econometrica* **70**(2), 781–799.

Hornik K, Stinchcombe M, and White H (1989), "Multilayer feedforward networks are universal approximators," *Neural Networks* **2**(5), 359–366.

Huber F, Koop G, and Onorante L (2021), "Inducing sparsity and shrinkage in time-varying parameter models," *Journal of Business & Economic Statistics* **39**(3), 669–683.

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.

Jin X, Maheu JM, and Yang Q (2022), "Infinite Markov pooling of predictive distributions," *Journal of Econometrics* **228**(2), 302–321.
Kalli M, and Griffin JE (2018), "Bayesian nonparametric vector autoregressive models," *Journal of Econometrics* **203**(2), 267–282.
Kastner G, and Huber F (2020), "Sparse Bayesian vector autoregressions in huge dimensions," *Journal of Forecasting* **39**(7), 1142–1165.
Makalic E, and Schmidt DF (2015), "A simple sampler for the horseshoe estimator," *IEEE Signal Processing Letters* **23**(1), 179–182.
Masini RP, Medeiros MC, and Mendes EF (2023), "Machine learning advances for time series forecasting," *Journal of Economic Surveys* **37**(1), 76–111.
McCracken M, and Ng S (2020), "FRED-QD: A quarterly database for macroeconomic research," *NBER Working Paper* **26872**.
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.
Novak R, Xiao L, Lee J, Bahri Y, Yang G, Hron J, Abolafia DA, Pennington J, and Sohl-Dickstein J (2018), "Bayesian deep convolutional networks with many channels are gaussian processes," *arXiv* **1810.05148**.
Nyblom J (1989), "Testing for the constancy of parameters over time," *Journal of the American Statistical Association* **84**(405), 223–230.
Omori Y, Chib S, Shephard N, and Nakajima J (2007), "Stochastic volatility with leverage: Fast and efficient likelihood inference," *Journal of Econometrics* **140**(2), 425–449.
Primiceri GE (2005), "Time varying structural vector autoregressions and monetary policy," *The Review of Economic Studies* **72**(3), 821–852.
Teräsvirta T (1994), "Specification, estimation, and evaluation of smooth transition autoregressive models," *Journal of the American Statistical Association* **89**(425), 208–218.
Tong H (1990), *Non-linear time series: A dynamical system approach*, Oxford University Press.
Williams CK, and Rasmussen CE (2006), *Gaussian processes for machine learning*, volume 2, MIT press Cambridge, MA.

A. ECONOMETRIC APPENDIX

A.1. *Details about priors*

On the dynamic coefficients of the BVAR (and the constant part of the TVPs in the non-centered parameterization, when applicable), we assume a Gaussian prior of the form:

$$\boldsymbol{a}_i \sim \mathcal{N}(\underline{\boldsymbol{a}_i}, \underline{\boldsymbol{V}_{ai}}),$$

where $\underline{\boldsymbol{V}_{ai}} = \tau_{ai}^2 \cdot \text{diag}(\lambda_{a,i1}^2, \dots, \lambda_{a,ik}^2)$; in our application we set $\boldsymbol{a}_i = \boldsymbol{0}_k$, but note that a variant of the canonical Minnesota prior could be established by setting specific elements of the prior mean to 1 rather than 0. On the square root of the state innovations we use a Gaussian prior, which implies independent marginal Gamma priors on the state innovations:

$$\boldsymbol{\theta}_i^{1/2} \sim \mathcal{N}(\boldsymbol{0}_k, \underline{\boldsymbol{V}_{\theta i}}) \iff \theta_{ij} \sim \mathcal{G}\left(\frac{1}{2}, \frac{1}{2\tau_{\theta i}^2 \lambda_{\theta,ij}^2}\right), \text{ for } j = 1, \dots, k,$$

where $\underline{\boldsymbol{V}_{\theta i}} = \tau_{\theta i}^2 \cdot \text{diag}(\lambda_{\theta,i1}^2, \dots, \lambda_{\theta,ik}^2)$. A very similar Gaussian prior is assumed for the factor loadings matrix, with the key difference that the global shrinkage parameter in this case applies across equations $i = 1, \ldots, n$:

$$\boldsymbol{l}_i \sim \mathcal{N}(\mathbf{0}_q, \underline{\boldsymbol{V}_{li}}),$$

where $\underline{\boldsymbol{V}_{li}} = \tau_l^2 \cdot \text{diag}(\lambda_{l,i1}^2, \ldots, \lambda_{l,iq}^2)$.

The τ's refer to the global shrinkage parameter of the HS prior, while the λ's indicate local scalings. The horseshoe is obtained by assuming independent half-Cauchy priors on all of them. We briefly present a generic version of the horseshoe that applies to all parameters where it is used discussed. Let $\boldsymbol{\beta} = (\beta_1, \ldots, \beta_d)$ denote a generic d-dimensional vector of parameters, and assume a prior $\beta_j \sim \mathcal{N}(\underline{\beta_j}, \tau^2\lambda_j^2)$ with global shrinkage parameter $\tau \sim \mathcal{C}^+(0, 1)$ and local scalings $\lambda_j \sim \mathcal{C}^+(0, 1)$ for $j = 1, \ldots, d$.

In particular, we exploit the auxiliary representation of the HS, see Makalic and Schmidt (2015), which enables straightforward Gibbs updates:

$$\lambda_j^2 | e_j \sim \mathcal{G}^{-1}(1/2, 1/e_j), \quad \tau^2 | f \sim \mathcal{G}^{-1}(1/2, 1/f); \quad f, e_j \sim \mathcal{G}^{-1}(1/2, 1)$$

which yields the following posterior distributions:

$$\lambda_j^2 | \bullet \sim \mathcal{G}^{-1}\left(1, \frac{1}{e_j} + \frac{(\beta_j - \underline{\beta_j})^2}{2\tau^2}\right), \quad \tau^2 | \bullet \sim \mathcal{G}^{-1}\left(\frac{d+1}{2}, \frac{1}{f} + \sum_{j=1}^{d} \frac{(\beta_j - \underline{\beta_j})^2}{2\lambda_j^2}\right),$$
$$e_j | \bullet \sim \mathcal{G}^{-1}\left(1, 1 + \tau_j^{-2}\right), \quad f | \bullet \sim \mathcal{G}^{-1}\left(1, 1 + \tau^{-2}\right)$$

On the parameters governing the SVs, we use independent vague priors for the unconditional means, $\mu_{ih} \sim \mathcal{N}(0, 10)$. In addition, we impose stationarity via the prior using transformed beta distributed priors, $(\phi_{ih} + 1)/2, (\phi_{j\omega} + 1)/2 \sim \mathcal{B}(5, 1.5)$. On the variances, we assume $\varsigma_{ih}^2, \varsigma_{j\omega}^2 \sim \mathcal{G}(1/2, 1/2)$, which implies a standard normal prior on $\varsigma_{ih}, \varsigma_{j\omega}$. The priors on the initial states, h_{i0}, ω_{j0} are assumed to follow the unconditional distribution of the AR(1) processes.

The kernel used for the GP regression features two tuning parameters collected in $\boldsymbol{\vartheta}_i = (\xi_i, l_i)'$. Here, we use equation-specific independent Gamma priors, $p(\boldsymbol{\vartheta}_i) = p(\xi_i)p(l_i)$; and thus assume $\xi_i \sim \mathcal{G}(a_{\xi i}, b_{\xi i})$ and $l_i \sim \mathcal{G}(a_{li}, b_{li})$. Given the positive support of these parameters, we use a log-normal transition density in a random walk MH algorithm to sample them.

A.2. *Sampling the conditional mean functions*

Constant and time-varying parameter vector autoregressions

Constant Parameters. Writing the ith equation in full data matrices, using $\tilde{\boldsymbol{y}}_i = (\tilde{y}_{i1}, \ldots, \tilde{y}_{iT})'$, $\boldsymbol{X} = (\boldsymbol{x}_1, \ldots, \boldsymbol{x}_T)'$, $\boldsymbol{\eta}_i = (\eta_{i1}, \ldots, \eta_{iT})'$ and the variances of the idiosyncratic error term $\boldsymbol{H}_i = \text{diag}(\boldsymbol{h}_i)$ we have:

$$\tilde{\boldsymbol{y}}_i = \boldsymbol{X}\boldsymbol{a}_i + \boldsymbol{\eta}_i, \quad \boldsymbol{\eta}_i \sim \mathcal{N}(\boldsymbol{0}_T, \boldsymbol{H}_i),$$

and the posterior of takes a textbook Gaussian form, $\boldsymbol{a}_i|\bullet \sim \mathcal{N}(\overline{\boldsymbol{a}_i}, \boldsymbol{V}_{ai})$, with moments:

$$\boldsymbol{V}_{ai} = \left(\boldsymbol{X}'\boldsymbol{H}_i^{-1}\boldsymbol{X} + \underline{\boldsymbol{V}_{ai}}^{-1}\right)^{-1}, \quad \overline{\boldsymbol{a}_i} = \boldsymbol{V}_{ai}(\boldsymbol{X}'\boldsymbol{H}_i^{-1}\tilde{\boldsymbol{y}}_i + \underline{\boldsymbol{V}_{ai}}^{-1}\underline{\boldsymbol{a}_i}).$$

Time Varying Parameters. For sampling the TVPs, we rely on the following representation of equation i of the full system:

$$(\tilde{y}_{it} - \boldsymbol{x}_t'\boldsymbol{a}_i) = \boldsymbol{x}_t'\text{diag}\left(\boldsymbol{\theta}_i^{1/2}\right)\tilde{\boldsymbol{a}}_{it} + \eta_{it},$$

which is the measurement equation of a standard conditionally Gaussian state space model, enabling Kalman filter based updates. Then, conditional on a draw of the TVPs, we may define $\tilde{\boldsymbol{x}}_t = (\boldsymbol{x}_t \odot \tilde{\boldsymbol{a}}_{it})$, such that we have an equivalent representation:

$$\tilde{y}_{it} = \boldsymbol{x}_t'\boldsymbol{a}_i + \tilde{\boldsymbol{x}}_t'\boldsymbol{\theta}_i^{1/2} + \eta_{it}.$$

This is again a conditionally Gaussian regression model, and textbook posteriors apply for sampling the conditional mean parameters $\boldsymbol{a}_i$ and square roots of the state innovations $\boldsymbol{\theta}_i^{1/2}$ in one block.

Bayesian Additive Regression Trees

Each tree in BART is sampled conditional on the remaining $S-1$ trees. In this context, define partial residuals $\check{\boldsymbol{y}}_{is} = \boldsymbol{y}_i - \sum_{j\neq s} \ell_{ij}(\boldsymbol{X}|\mathcal{T}_{ij}, \boldsymbol{\mu}_{ij})$ excluding the sth tree. Chipman *et al.* (2010) show how to analytically marginalize the conditional posterior with respect to the terminal node parameters $\boldsymbol{\mu}_{ij}$, which keeps the dimensionality of the inference problem fixed:

$$p(\mathcal{T}_{ij}|\check{\boldsymbol{y}}_{is}, \bullet) \propto p(\mathcal{T}_{ij}) \int p(\check{\boldsymbol{y}}_{is}|\mathcal{T}_{ij}, \boldsymbol{\mu}_{ij}, \bullet) p(\boldsymbol{\mu}_{ij}|\mathcal{T}_{ij}, \bullet) d\boldsymbol{\mu}_{ij}.$$

The posterior $p(\mathcal{T}_{ij}|\check{\boldsymbol{y}}_{is}, \bullet)$ can be used in an MH-updating scheme, based on transition density $q(\mathcal{T}_{ij}^{(*)}|\mathcal{T}_{ij}^{(c)})$ which produces a candidate tree $\mathcal{T}_{ij}^{(*)}$ conditional on the current (c) state of the tree $\mathcal{T}_{ij}^{(c)}$. The resulting acceptance probability has a standard form and is given by:

$$\min\left(\frac{p(\mathcal{T}_{ij}^{(*)}|\check{\boldsymbol{y}}_{is}, \bullet)}{p(\mathcal{T}_{ij}^{(c)}|\check{\boldsymbol{y}}_{is}, \bullet)} \frac{q(\mathcal{T}_{ij}^{(c)}|\mathcal{T}_{ij}^{(*)})}{q(\mathcal{T}_{ij}^{(*)}|\mathcal{T}_{ij}^{(c)})}, 1\right).$$

Conditional on the tree structures, sampling the terminal node parameters is trivial. In particular, subject to the partitions of the input space we obtain distinct observations for each terminal node. The posterior takes the form of an intercept-only regression model.

Gaussian process regression

For a generic set of training observations in $\boldsymbol{X}$ and test observations $\boldsymbol{X}^*$, the joint distribution of the target $\tilde{\boldsymbol{y}}_i$ and the function values is given by:

$$\begin{pmatrix} \tilde{\boldsymbol{y}}_i \\ \boldsymbol{f}_i^* \end{pmatrix} \sim \mathcal{N}\left(\boldsymbol{0}, \begin{bmatrix} \mathcal{K}_{\boldsymbol{\vartheta}_i}(\boldsymbol{X}, \boldsymbol{X}') + \boldsymbol{H}_i & \mathcal{K}_{\boldsymbol{\vartheta}_i}(\boldsymbol{X}, \boldsymbol{X}^{*\prime}) \\ \mathcal{K}_{\boldsymbol{\vartheta}_i}(\boldsymbol{X}^*, \boldsymbol{X}') & \mathcal{K}_{\boldsymbol{\vartheta}_i}(\boldsymbol{X}^*, \boldsymbol{X}^{*\prime}) \end{bmatrix}\right)$$

where $\boldsymbol{H}_i = \text{diag}(\boldsymbol{h}_i)$. In particular, this expression gives rise both to the in-sample distribution of the conditional mean function, and the predictive distribution. In general, we have $\boldsymbol{f}_i^* \sim \mathcal{N}(\overline{\boldsymbol{f}_i}, \boldsymbol{V}_{fi})$, with moments:

$$\begin{aligned} \overline{\boldsymbol{f}_i} &= \mathcal{K}_{\boldsymbol{\vartheta}_i}(\boldsymbol{X}^*, \boldsymbol{X}') \left(\mathcal{K}_{\boldsymbol{\vartheta}_i}(\boldsymbol{X}, \boldsymbol{X}') + \boldsymbol{H}_i\right)^{-1} \tilde{\boldsymbol{y}}_i, \\ \boldsymbol{V}_{fi} &= \mathcal{K}_{\boldsymbol{\vartheta}_i}(\boldsymbol{X}^*, \boldsymbol{X}^{*\prime}) - \mathcal{K}_{\boldsymbol{\vartheta}_i}(\boldsymbol{X}^*, \boldsymbol{X}') \left(\mathcal{K}_{\boldsymbol{\vartheta}_i}(\boldsymbol{X}, \boldsymbol{X}') + \boldsymbol{H}_i\right)^{-1} \mathcal{K}_{\boldsymbol{\vartheta}_i}(\boldsymbol{X}, \boldsymbol{X}^{*\prime}). \end{aligned}$$

We will also need an expression for the logarithm of the conditional likelihood, which can be used to optimize the tuning parameters encoded in $\boldsymbol{\vartheta}_i$ via a MH-updating scheme. Note that $\tilde{\boldsymbol{y}}_i \sim \mathcal{N}(\boldsymbol{0}, \mathcal{K}_{\boldsymbol{\vartheta}_i}(\boldsymbol{X}, \boldsymbol{X}') + \boldsymbol{H}_i)$, and $p(\tilde{\boldsymbol{y}}_i|\bullet)$ is thus the density of a zero mean multivariate Gaussian with covariance matrix $\mathcal{K}_{\boldsymbol{\vartheta}_i}(\boldsymbol{X}, \boldsymbol{X}') + \boldsymbol{H}_i)$.

To update the GP tuning parameters via a MH-step, note that the acceptance probabilities for candidate draws $\boldsymbol{\vartheta}_i^*$ using a transition density $q(\boldsymbol{\vartheta}_i^{(*)}|\boldsymbol{\vartheta}_i^{(c)})$ are given by:

$$\min\left(\frac{p(\tilde{\boldsymbol{y}}_i|\boldsymbol{\vartheta}_i^{(*)},\bullet)p(\boldsymbol{\theta}_i^{(*)})}{p(\tilde{\boldsymbol{y}}_i|\boldsymbol{\vartheta}_i^{(c)},\bullet)p(\boldsymbol{\theta}_i^{(c)})}\frac{q(\boldsymbol{\vartheta}_i^{(c)}|\boldsymbol{\vartheta}_i^{(*)})}{q(\boldsymbol{\vartheta}_i^{(*)}|\boldsymbol{\vartheta}_i^{(c)})},1\right).$$

A.3. *Additional sampling steps*

Sampling the Factors. Recall that the FSV model of Eq. (2), which is given by $\boldsymbol{\epsilon}_t = \boldsymbol{L}\boldsymbol{\mathfrak{F}}_t + \boldsymbol{\eta}_t$. In stacked notation, using $Tn\times 1$-vectors $\boldsymbol{\epsilon} = (\boldsymbol{\epsilon}_1', \dots, \boldsymbol{\epsilon}_T')'$, $\boldsymbol{\eta} = (\boldsymbol{\eta}_1', \dots, \boldsymbol{\eta}_T')'$, and $\boldsymbol{\mathfrak{F}} = (\boldsymbol{\mathfrak{F}}_1', \dots, \boldsymbol{\mathfrak{F}}_T')'$, $\boldsymbol{\Omega} = \text{bdiag}(\boldsymbol{\Omega}_1, \dots, \boldsymbol{\Omega}_T)$, $\boldsymbol{H} = \text{bdiag}(\boldsymbol{H}_1, \dots, \boldsymbol{H}_T)$, we may write:

$$\boldsymbol{\epsilon} = (\boldsymbol{I}_T \otimes \boldsymbol{L})\boldsymbol{\mathfrak{F}} + \boldsymbol{\eta}, \quad \boldsymbol{\mathfrak{F}} \sim \mathcal{N}(\boldsymbol{0}_{Tq}, \boldsymbol{\Omega}), \quad \boldsymbol{\eta} \sim \mathcal{N}(\boldsymbol{0}_{Tn}, \boldsymbol{H}).$$

The operator bdiag($\bullet$) refers to stacking matrices, and outputs a block diagonal matrix. This is just a big linear regression model, and the joint posterior of the factors arises as a Gaussian distribution $\boldsymbol{\mathfrak{F}}|\bullet \sim \mathcal{N}\left(\overline{\boldsymbol{\mathfrak{F}}}, \boldsymbol{V}_{\mathfrak{F}}\right)$, with textbook moments:

$$\boldsymbol{V}_{\mathfrak{F}} = \left((\boldsymbol{I}_T \otimes \boldsymbol{L}')\boldsymbol{H}^{-1}(\boldsymbol{I}_T \otimes \boldsymbol{L}) + \boldsymbol{\Omega}^{-1}\right)^{-1}, \quad \overline{\boldsymbol{\mathfrak{F}}} = \boldsymbol{V}_{\mathfrak{F}}\left((\boldsymbol{I}_T \otimes \boldsymbol{L}')\boldsymbol{H}^{-1}\boldsymbol{\epsilon}\right).$$

Sampling the Loadings. The loadings are sampled variable-by-variable. Let $\boldsymbol{y}_i = (y_{i1}, \dots, y_{iT})'$ denote variable i stacked over time, $\boldsymbol{\epsilon}_i = (\epsilon_{i1}, \dots, \epsilon_{iT})'$, $\boldsymbol{\eta}_i = (\eta_{i1}, \dots, \eta_{iT})'$ and $\tilde{\boldsymbol{\mathfrak{F}}}$ is the $T \times q$ full data matrix of factors, $\boldsymbol{l}_i = \boldsymbol{L}_{i\bullet}'$ the loadings associated with the ith equation and $\boldsymbol{H}_i = \text{diag}(\boldsymbol{h}_i)$, such that we may write:

$$(\boldsymbol{y}_i - \boldsymbol{f}_i) \equiv \boldsymbol{\epsilon}_i = \tilde{\boldsymbol{\mathfrak{F}}}\boldsymbol{l}_i + \boldsymbol{\eta}_i, \quad \boldsymbol{\eta}_i \sim \mathcal{N}\left(\boldsymbol{0}_T, \boldsymbol{H}_i\right).$$

The posterior of the loadings is Gaussian, $\boldsymbol{l}_i|\bullet \sim \mathcal{N}\left(\overline{\boldsymbol{l}}_i, \boldsymbol{V}_{li}\right)$, with moments:

$$\boldsymbol{V}_{li} = \left(\tilde{\boldsymbol{\mathfrak{F}}}'\boldsymbol{H}_i^{-1}\tilde{\boldsymbol{\mathfrak{F}}} + \underline{\boldsymbol{V}_{li}}^{-1}\right)^{-1}, \quad \overline{\boldsymbol{l}}_i = \boldsymbol{V}_{li}(\tilde{\boldsymbol{\mathfrak{F}}}'\boldsymbol{H}_i^{-1}\boldsymbol{\epsilon}_i).$$

Sampling the Volatilities. The factors and idiosyncratic shocks can be written independently as $\mathfrak{F}_{qt} = \exp(\omega_{jt}/2)\zeta_{\mathfrak{F},qt}$ and $\eta_{it} = \exp(h_{it}/2)\zeta_{\eta,it}$ with $\zeta_{\bullet t} \sim \mathcal{N}(0,1)$. Squaring and taking logs moves the SVs into the conditional mean of these measurement equations, which then feature a log χ^2 error term. Using a 10-component Gaussian distribution as an approximation, as suggested in Omori *et al.* (2007), allows for standard filtering and smoothing algorithms to draw the SVs factor-by-factor and equation-by-equation.

For the homoskedastic case, the variances of the factors are normalized to 1, and under inverse Gamma priors $\exp(h_i) \sim \mathcal{G}^{-1}(a_0, b_0)$ on the idiosyncratic shocks, we obtain textbook posteriors: $\exp(h_i)|\bullet \sim \mathcal{G}^{-1}(a_0 + T/2, b_0 + \sum_{t=1}^{T} \eta_{it}^2/2)$.

The extension to t-distributed errors is achieved by introducing auxiliary variables $\gamma_{it} \sim \mathcal{G}^{-1}(\nu_i/2, \nu_i/2)$. In particular, these can be used to write $\eta_{it} \sim t_{\nu_i}\left(0, \exp(h_{jt})\right)$ conditionally as $\eta_{it}|\gamma_{it} \sim \mathcal{N}\left(0, \gamma_{it} \exp(h_{jt})\right)$, where the marginal distributions coincide. This enables efficient conditional Gibbs updates via data augmentation.

Online Appendix: Bayesian nonparametric methods for macroeconomic forecasting

OA. FORECAST RESULTS BY SUBSAMPLE

OA.1. *Euro area: Results by forecast horizon*

Table OA.1: One-quarter ahead predictive loss metrics for the EA.

S

		CRPS					CRPS−L					CRPS−R					MAE				
		FS	Re	Ex	Pre	Post	FS	Re	Ex	Pre	Post	FS	Re	Ex	Pre	Post	FS	Re	Ex	Pre	Post
YER	BVAR hom	1.27	1.02	1.53	1.05	1.43	1.23	1.02	1.56	1.06	1.35	1.34	1.03	1.52	1.05	1.54	1.31	1.03	1.59	1.05	1.52
YER	BVAR SV	1.03	1.03	1.02	1.04	1.02	1.05	1.04	1.05	1.07	1.03	1.01	1.01	1.00	0.99	1.02	1.03	1.02	1.04	1.04	1.02
YER	BVAR SV−t	1.02	1.03	1.01	1.04	1.01	1.04	1.04	1.04	1.08	1.02	1.00	1.01	0.99	0.99	1.01	1.02	1.02	1.02	1.05	1.00
YER	TVP SV	1.53	1.02	2.06	1.04	1.87	1.72	1.04	2.82	1.07	2.17	1.34	1.00	1.54	0.99	1.59	1.56	1.01	2.11	1.04	1.99
YER	BART hom	1.18	1.02	1.35	1.05	1.27	1.15	1.02	1.35	1.05	1.22	1.23	1.03	1.35	1.05	1.36	1.20	1.02	1.38	1.04	1.34
YER	BART SV	1.03	1.01	1.05	1.02	1.04	1.05	1.03	1.08	1.04	1.05	1.02	0.99	1.03	1.00	1.03	1.04	1.00	1.07	1.03	1.04
YER	GP hom	1.01	0.96°	1.07	1.01	1.01	1.02	0.98	1.10	1.05	1.00	1.01	0.95°	1.04	0.97	1.03	1.03	0.95°	1.10	1.01	1.04
YER	GP SV	1.05	0.99	1.10	1.02	1.06	1.06	1.00	1.15	1.05	1.06	1.04	0.98	1.07	0.99	1.07	1.04	0.98	1.09	1.03	1.04
HICSA	BVAR hom	1.08	1.10	1.07	1.05	1.24	1.09	1.12	1.07	1.07	1.19	1.06	1.07	1.06	1.02	1.29	1.08	1.13	1.07	1.06	1.19
HICSA	BVAR SV	1.07	1.11	1.06	1.05	1.18	1.09	1.13	1.08	1.08	1.14	1.04	1.07	1.03	1.01	1.21	1.07	1.12	1.05	1.06	1.14
HICSA	BVAR SV−t	1.08	1.10	1.07	1.06	1.17	1.09	1.12	1.08	1.08	1.13	1.07	1.06	1.07	1.03	1.21	1.08	1.10	1.07	1.07	1.11
HICSA	TVP SV	1.08	1.10	1.07	1.06	1.17	1.11	1.12	1.10	1.10	1.16	1.05	1.07	1.04	1.02	1.18	1.08	1.12	1.07	1.07	1.13
HICSA	BART hom	1.05	1.05	1.05	1.05	1.06	1.05	1.07	1.05	1.06	1.04	1.05	1.03	1.05	1.04	1.08	1.05	1.08	1.05	1.05	1.04
HICSA	BART SV	1.04	1.04	1.04	1.05	0.99	1.05	1.04	1.05	1.06	1.00	1.03	1.03	1.04	1.05	0.98	1.04	1.06	1.04	1.05	0.99
HICSA	GP hom	1.00	1.08	0.97	1.05	0.72	1.02	1.11	0.99	1.06	0.79	0.97	1.04	0.96	1.04	0.64	0.98	1.05	0.97	1.03	0.77
HICSA	GP SV	1.02	1.10	0.99	1.07	0.74	1.04	1.11	1.01	1.07	0.83	1.00	1.08	0.98	1.07	0.65	1.01	1.08	0.98	1.05	0.78
URX	BVAR hom	1.07	1.09	1.06	1.06	1.11	1.06	1.09	1.06	1.04	1.13	1.07	1.09	1.07	1.07	1.10	1.10	1.09	1.10	1.08	1.18
URX	BVAR SV	1.03	1.08	1.01	1.07	0.91	1.01	1.07	0.98	1.07	0.82	1.05	1.10	1.02	1.07	0.98	1.04	1.08	1.02	1.07	0.93
URX	BVAR SV−t	1.03	1.07	1.01	1.07	0.91	1.01	1.07	0.99	1.07	0.83	1.05	1.08	1.04	1.07	0.98	1.04	1.07	1.03	1.07	0.94
URX	TVP SV	1.00	1.07	0.97	1.03	0.92	0.97	1.07	0.94	1.01	0.86	1.03	1.09	1.00	1.04	0.98	1.01	1.07	0.98	1.02	0.95
URX	BART hom	0.97	0.97	0.97	0.93°	1.10	0.97	0.98	0.97	0.93°	1.11	0.97	0.97	0.97	0.93°	1.09	0.98	0.98	0.98	0.93°	1.16
URX	BART SV	0.93'	0.93°	0.93	0.94°	0.90	0.93	0.93°	0.93	0.96	0.83	0.93°	0.93'	0.93'	0.92*	0.96	0.93'	0.91°	0.93	0.92°	0.94
URX	GP hom	0.94	0.87'	0.96	0.97	0.82	0.92	0.88°	0.94	0.99	0.73	0.95	0.87'	0.99	0.96	0.91	0.95	0.89°	0.98	0.97	0.89
URX	GP SV	0.93'	0.88'	0.95	0.95	0.86	0.92	0.90'	0.93	0.98	0.75	0.94'	0.88'	0.97	0.93'	0.96	0.93'	0.86°	0.96	0.94	0.91

M

		CRPS					CRPS−L					CRPS−R					MAE				
		FS	Re	Ex	Pre	Post	FS	Re	Ex	Pre	Post	FS	Re	Ex	Pre	Post	FS	Re	Ex	Pre	Post
YER	BVAR hom	0.61	1.96	0.36	0.27	4.24	0.19	0.75	0.09	0.09	1.31	0.16	0.38	0.12	0.07	1.12	0.70	2.19	0.42	0.34	4.50
YER	BVAR SV	1.07	1.01	1.14	1.00	1.12	1.10	1.02	1.22	1.01	1.15	1.06	0.99°	1.10	0.99	1.11	1.03	1.00	1.05	1.02	1.03
YER	BVAR SV−t	1.07	0.98°	1.16	0.98'	1.13	1.09	0.99'	1.26	0.98	1.17	1.05	0.97*	1.10	0.97°	1.11	1.03	0.97*	1.09	1.00	1.06
YER	TVP SV	1.12	1.01	1.24	1.02	1.19	1.16	1.02	1.38	1.03	1.25	1.09	0.99°	1.15	1.01	1.15	1.09	1.00	1.18	1.03	1.14
YER	BART hom	1.04	1.03	1.04	1.07	1.01	1.03	1.02	1.04	1.07	1.00	1.04	1.05	1.04	1.07	1.02	1.05	1.04	1.05	1.06	1.04
YER	BART SV	1.43	1.05	1.75	1.08	1.52	1.35	1.05	1.69	1.10	1.39	1.57	1.04	1.85	1.06	1.74	1.37	1.05	1.62	1.10	1.45
YER	GP hom	1.04	1.00	1.07	1.05	1.03	1.03	0.99'	1.10	1.05	1.02	1.05	1.04	1.06	1.07	1.04	1.04	1.02	1.05	1.01	1.06
YER	GP SV	1.08	1.02	1.14	1.06	1.09	1.08	1.01	1.20	1.06	1.10	1.08	1.04	1.11	1.07	1.09	1.05	1.03	1.06	1.02	1.07
HICSA	BVAR hom	0.20	0.31	0.17	0.18	0.37	0.06	0.10	0.05	0.05	0.10	0.06	0.08	0.05	0.05	0.12	0.28	0.39	0.25	0.25	0.53
HICSA	BVAR SV	1.02	0.98	1.04	1.01	1.09	1.04	0.98	1.06	1.02	1.14	1.02	0.99	1.02	1.01	1.06	1.02	0.98	1.03	1.00	1.11
HICSA	BVAR SV−t	1.02	0.96°	1.04	1.01	1.10	1.04	0.96	1.08	1.03	1.13	1.01	0.97'	1.02	0.99	1.09	1.02	0.95°	1.04	1.01	1.09
HICSA	TVP SV	1.05	0.99	1.07	1.04	1.09	1.08	1.00	1.12	1.06	1.22	1.02	0.99	1.03	1.03	0.98	1.05	0.99	1.06	1.04	1.09
HICSA	BART hom	0.98	0.95	0.99	0.99	0.94	0.99	0.97	1.00	1.00	0.94	0.97	0.94	0.98	0.98	0.95	0.95	0.94	0.96	0.96	0.91
HICSA	BART SV	1.01	0.94	1.03	1.02	0.93	1.04	0.95	1.05	1.04	0.96	1.00'	0.93	1.01	1.00	0.91'	1.00'	0.93	1.00	1.01	0.89'
HICSA	GP hom	1.14	1.01	1.18	1.20	0.81	1.07	0.93	1.13	1.11	0.87	1.22	1.14	1.24	1.32	0.75	1.12	1.08	1.13	1.17	0.89
HICSA	GP SV	1.25	1.06	1.31	1.32	0.90	1.13	0.99	1.19	1.16	0.97	1.38	1.19	1.44	1.50	0.81	1.23	1.15	1.26	1.28	0.97
URX	BVAR hom	0.12	0.23	0.10	0.10	0.34	0.03	0.05	0.03	0.03	0.09	0.04	0.08	0.03	0.03	0.11	0.16	0.28	0.13	0.13	0.40
URX	BVAR SV	1.03	1.07	1.02	1.03	1.03	1.01	1.05	0.99	1.03	0.94	1.06	1.08	1.04	1.04	1.10	1.03	1.06	1.02	1.02	1.06
URX	BVAR SV−t	1.03	1.02	1.04	1.02	1.08	1.01	1.03	1.00	1.01	1.02	1.05	1.01	1.07	1.02	1.14	1.03	1.02	1.03	1.00	1.14
URX	TVP SV	1.00	1.03	0.98	0.98	1.07	0.95	1.04	0.92'	0.96'	0.95	1.04	1.04	1.04	0.99	1.17	0.98	1.02	0.97	0.97'	1.05
URX	BART hom	0.96'	0.99	0.95'	0.98	0.91	0.94'	1.00	0.91'	0.96'	0.87	0.98	0.99	0.97	0.98	0.95	0.98	0.99	0.97	0.99	0.94
URX	BART SV	1.01'	1.00	0.99'	1.02'	0.91	0.97°	1.00	0.93°	1.00°	0.81	1.05	1.00	1.04	1.04	1.00	1.02'	0.99	1.01'	1.03'	0.95
URX	GP hom	1.05	0.98	1.09	1.12	0.85	1.03	1.11	1.01	1.11	0.80	1.07	0.89	1.15	1.12	0.91	1.10	1.05	1.12	1.15	0.92
URX	GP SV	1.12	0.96	1.18	1.20	0.87	1.11	1.08	1.12	1.22	0.80	1.12	0.89	1.23	1.18	0.93	1.17	1.02	1.23	1.24	0.93

Table OA.2: Two-year ahead predictive loss metrics for the EA.

S

		CRPS					CRPS−L					CRPS−R					MAE				
		FS	Re	Ex	Pre	Post	FS	Re	Ex	Pre	Post	FS	Re	Ex	Pre	Post	FS	Re	Ex	Pre	Post
YER	BVAR hom	1.01	1.00	1.01	1.02	1.00	1.00	1.00	1.01	1.01	1.00	1.01	1.01	1.01	1.03	1.00	1.01	1.01	1.01	1.01	1.00
YER	BVAR SV	0.99	0.99	1.00	**0.98°**	1.00	0.99	0.99	*1.00*	**0.97'**	1.01	1.00	1.00	1.00	0.98°	1.00	1.00	0.99°	1.00	1.00	1.00
YER	BVAR SV−t	0.99°	0.99	**0.99°**	0.98*	1.00	1.00	0.99	1.00	0.98	1.01	**0.99***	**0.99'**	**0.98°**	**0.97***	*1.00*	0.99'	0.99'	**0.99**	0.99	**1.00**
YER	TVP SV	**0.99'**	0.98*	1.00	0.98'	1.00	*0.99*	0.98°	1.01	0.98	1.00	0.99	0.99	1.00	0.98'	1.00	**0.99**	*0.99°*	0.99	*0.98*	1.00
YER	BART hom	1.00	0.99	1.00	1.00	1.00	0.99	0.99	1.00	0.99	1.00	1.00	1.00	1.00	1.01	1.00	1.00	1.00	1.00	1.00	1.00
YER	BART SV	1.00	1.00	1.01	1.00	1.01	1.01	1.00	1.03	1.00	1.01	1.00	1.00	1.00	1.00	1.00	1.01	1.00	1.02	1.01	1.00
YER	GP hom	1.00	0.98*	1.03	1.01	*1.00*	1.00	0.97*	1.05	1.00	*0.99*	1.01	1.00	1.01	1.02	1.00	1.01	1.00	1.02	1.01	1.00
YER	GP SV	1.00	*0.98**	1.03	1.01	1.00	1.00	*0.97**	1.05	1.00	1.00	1.01	1.00	1.02	1.02	1.00	1.00	0.99	1.01	1.00	1.00
HICSA	BVAR hom	1.10	0.99	1.13	1.12	1.03	1.12	0.97	1.17	1.12	1.10	1.09	1.02	1.10	1.11	0.98	1.12	1.00	1.16	1.14	1.06
HICSA	BVAR SV	1.11	1.00	1.14	1.13	1.03	1.14	0.98	1.20	1.15	1.12	1.08	1.02	1.09	1.11	0.97	1.13	1.03	1.16	1.15	1.06
HICSA	BVAR SV−t	1.10	1.01	1.13	1.13	1.00	1.11	0.98	1.16	1.12	1.05	1.10	1.04	1.11	1.13	0.97	1.12	1.02	1.14	1.14	1.04
HICSA	TVP SV	1.11	*0.96*	1.16	1.13	1.03	1.15	0.94	1.22	1.16	1.10	1.08	*0.99*	1.10	1.10	0.97	1.13	*0.96*	1.17	1.14	1.05
HICSA	BART hom	1.07	1.01	1.09	1.09	1.02	1.09	0.99	1.12	1.09	1.10	1.06	1.02	*1.07*	1.09	0.96	1.07	1.00	1.09	1.08	1.04
HICSA	BART SV	*1.05*	0.98	*1.08*	*1.06*	1.02	1.05	0.95	1.09	1.06	1.04	*1.06*	1.01	1.07	*1.07*	1.00	*1.05*	0.97	*1.07*	*1.05*	*1.03*
HICSA	GP hom	1.07	0.99	1.09	1.08	*0.99*	*1.04*	0.95	*1.08*	*1.04*	1.07	1.09	1.03	1.11	1.14	*0.94*	1.06	0.98	1.08	1.06	1.04
HICSA	GP SV	1.09	0.97	1.12	1.10	1.02	1.05	**0.92**	1.09	1.05	*1.02*	1.14	1.05	1.16	1.17	1.02	1.08	0.97	1.12	1.09	1.05
URX	BVAR hom	0.97	0.93'	0.99	0.97	0.99	0.97	0.98	0.96	0.97	0.98	0.97	0.89°	1.02	0.96	1.00	0.99	0.96	1.00	0.99	1.00
URX	BVAR SV	0.96	0.88'	1.01	0.95	1.01	0.98	0.97	0.99	0.98	1.02	0.94	0.82°	1.02	0.93	1.01	1.00	0.94'	1.03	1.00	1.01
URX	BVAR SV−t	0.96	0.90	1.00	0.96	1.00	0.96	0.97	0.96	0.96	0.96	0.96	0.86	1.03	0.95	1.03	1.00	0.95	1.03	1.00	1.03
URX	TVP SV	**0.94'**	0.90'	*0.96*	**0.93**	*0.96*	**0.94**	1.01	*0.92*	**0.94**	*0.96*	**0.93**	0.84'	*0.99*	**0.92**	**0.97**	**0.97**	0.96	*0.97*	**0.97**	*0.96*
URX	BART hom	1.01	0.82°	1.11	1.01	1.03	1.06	*0.89'*	1.13	1.06	1.05	0.97	0.77°	1.09	0.96	1.01	1.04	0.87'	1.12	1.04	1.06
URX	BART SV	0.99	0.84	1.08	1.00	0.99	1.03	0.91	1.09	1.04	0.99	0.96	0.78'	1.08	0.96	0.99	1.03	0.89	1.09	1.03	0.99
URX	GP hom	1.00	0.80°	1.12	1.01	0.99	1.05	0.90'	1.11	1.05	1.02	0.97	0.73°	1.12	0.97	0.97	1.05	0.88'	1.13	1.05	1.03
URX	GP SV	0.98	*0.80'*	1.09	0.99	0.98	1.01	0.89'	1.07	1.02	0.98	0.96	*0.73'*	1.10	0.95	0.98	1.03	*0.87'*	1.10	1.03	0.97

M

		CRPS					CRPS−L					CRPS−R					MAE				
		FS	Re	Ex	Pre	Post	FS	Re	Ex	Pre	Post	FS	Re	Ex	Pre	Post	FS	Re	Ex	Pre	Post
YER	BVAR hom	0.68	2.06	0.39	0.30	4.30	0.21	0.79	0.09	0.10	1.30	0.18	0.40	0.14	0.08	1.17	0.77	2.32	0.44	0.38	4.57
YER	BVAR SV	1.01	1.00	1.01	1.01	1.00	1.00	1.00	1.01	1.00	1.01	1.01	1.01	1.01	1.03	1.00	1.01	1.01	1.01	1.01	1.00
YER	BVAR SV−t	1.01	0.99	1.03	1.02	1.00	1.00	0.99	1.04	1.00	1.00	1.02	1.01	1.02	1.04	1.00	1.01	1.01	1.02	1.02	1.00
YER	TVP SV	1.02	1.01	1.03	1.04	1.01	1.01	1.00	1.04	1.01	1.01	1.03	1.03	1.03	1.07	1.01	1.02	1.02	1.02	1.03	1.01
YER	BART hom	1.00	1.00	1.01	1.01	1.00	1.00	0.99	1.01	1.00	1.00	1.01	1.00	1.01	1.02	1.00	1.00	1.00	1.00	1.01	1.00
YER	BART SV	1.07	1.23	1.06	1.00	1.01	1.09	1.25	1.07	1.03	1.02	1.06	1.19	1.05	*0.97*	1.01	1.05	1.20	1.03	**0.97**	1.00
YER	GP hom	*0.99*	**0.96***	1.03	1.01	**0.99**	0.98	**0.94***	1.06	0.99	**0.98**	1.01	*0.99°*	1.02	1.04	**0.99**	*0.99*	**0.98***	1.01	0.99	*1.00*
YER	GP SV	1.00	0.96*	1.03	1.00	0.99	**0.98**	0.95*	1.05	*0.98*	0.99	1.02	1.00	1.02	1.04	1.00	1.00	0.99*	1.01	1.00	1.00
HICSA	BVAR hom	0.23	0.31	0.21	0.21	0.47	0.07	0.10	0.06	0.06	0.13	0.07	0.08	0.07	0.06	0.15	0.32	0.40	0.30	0.29	0.65
HICSA	BVAR SV	0.99	0.99	0.99	1.00	0.95	0.99	0.99	0.99	1.00	**0.91**	0.99	1.00	0.99	0.99	0.98	0.98	1.01	0.97	0.98	**0.96**
HICSA	BVAR SV−t	1.00	1.00	1.00	1.01	0.96	0.99	0.98	1.00	1.01	0.92	1.00	1.03	1.00	1.01	1.00	1.01	1.01	1.00	1.01	0.97
HICSA	TVP SV	1.07	0.96	1.10	1.10	0.97	1.07	*0.93*	1.13	1.10	0.94	1.06	1.01	1.08	1.08	0.99	1.08	1.01	1.10	1.10	0.99
HICSA	BART hom	**0.96**	**0.94**	**0.97**	**0.95**	1.01	**0.98**	0.94	**0.99**	**0.96**	1.02	**0.96**	**0.94**	**0.96**	**0.95**	1.00	**0.95**	**0.92**	**0.95**	**0.93**	1.01
HICSA	BART SV	1.02	1.08	1.02	1.01	1.00	1.02	1.06	1.03	1.01	1.01	1.02	1.11	1.01	1.01	0.99	1.01	1.07	1.00	1.00	1.00
HICSA	GP hom	1.27	1.14	1.31	1.35	0.92	1.14	1.02	1.18	1.16	1.03	1.39	1.29	1.41	1.54	0.81	1.32	1.21	1.35	1.38	1.04
HICSA	GP SV	1.27	1.16	1.31	1.36	**0.92**	1.16	1.05	1.19	1.18	1.04	1.38	1.30	1.40	1.53	**0.79**	1.32	1.21	1.35	1.38	1.03
URX	BVAR hom	0.15	0.30	0.12	0.14	0.27	0.04	0.07	0.03	0.04	0.07	0.05	0.11	0.04	0.05	0.09	0.20	0.38	0.16	0.19	0.32
URX	BVAR SV	0.98	1.01	0.97	0.98	1.00	0.98	1.04	0.95	0.98	0.99	0.99	1.01	0.98	0.98	1.02	0.99	1.03	0.97	0.99	0.99
URX	BVAR SV−t	*0.98*	1.00	0.97	*0.98*	0.98	*0.97*	1.06	0.93	*0.97*	**0.92**	0.99	0.97	0.99	0.98	1.03	1.00	1.04	0.98	1.00	1.00
URX	TVP SV	0.98	1.10	**0.92**	0.99	**0.96**	0.97	1.16	**0.88**	0.98	0.93	1.00	1.07	**0.95**	1.00	*0.99*	*0.98*	1.12	**0.91**	*0.99*	**0.92**
URX	BART hom	1.02	0.94	1.06	1.02	1.02	1.05	0.95	1.09	1.05	1.02	0.99	0.94	1.02	0.99	1.01	1.02	0.95	1.05	1.02	1.02
URX	BART SV	1.00	0.99	1.06	0.99	1.01	1.03	0.97	1.09	1.03	1.00	0.98	0.99	1.03	*0.96*	1.02	1.00	0.98	1.05	1.00	0.99
URX	GP hom	1.03	0.74°	1.19	1.04	1.03	1.09	**0.86°**	1.20	1.10	1.05	0.98	0.66°	1.19	0.98	1.00	1.10	0.81°	1.24	1.10	1.09
URX	GP SV	1.04	**0.73°**	1.21	1.04	1.04	1.11	0.86°	1.23	1.12	1.07	*0.98*	**0.64°**	1.19	0.97	1.01	1.10	**0.81°**	1.25	1.10	1.12

OA.2. *United States: Results by forecast horizon*

Table OA.3: One-quarter ahead predictive loss metrics for the US.

S

		CRPS					CRPS−L					CRPS−R					MAE				
		FS	Re	Ex	Pre	Post	FS	Re	Ex	Pre	Post	FS	Re	Ex	Pre	Post	FS	Re	Ex	Pre	Post
GDPC1	BVAR hom	1.08	0.99	1.13	1.04	1.18	1.06	1.01	1.10	1.01	1.16	1.09	0.96°	1.14	1.05	1.20	1.15	0.96'	1.25	1.12	1.23
	BVAR SV	1.02	0.99	1.03	1.00	1.05	1.05	1.01	1.07	0.99	1.16	0.99	0.97°	1.00	1.00	0.97	1.09	0.97	1.15	1.09	1.10
	BVAR SV−t	1.04	0.96	1.07	1.00	1.11	1.07	0.98	1.12	1.00	1.21	1.01	**0.93°**	1.03	0.99	1.04	1.11	**0.93°**	1.20	1.10	1.13
	TVP SV	1.00	0.97	1.02	0.98	1.05	1.05	0.99	1.08	0.99	1.16	0.96	0.95*	0.97	0.96*	0.97	1.07	0.94*	1.13	1.06	1.08
	BART hom	1.01	1.01	1.02	1.04	0.96	1.00	1.02	**0.99**	1.01	0.99	1.02	0.98	1.03	1.05	0.93	1.07	0.99	1.11	1.11	0.99
	BART SV	0.99	1.02	0.97	1.00	0.96	1.03	1.04	1.02	0.99	1.10	0.95	1.00	0.93	0.99	0.84	1.04	1.01	1.06	1.07	0.98
	GP hom	1.02	0.96	1.05	1.05	0.97	1.04	**0.97**	1.08	1.04	1.03	1.00	0.95'	1.02	1.04	0.90	1.12	0.96'	1.20	1.17	1.00
	GP SV	1.02	0.97	1.04	1.04	0.98	1.06	0.97	1.12	1.03	1.13	0.97	0.97	0.98	1.02	0.86	1.13	0.97	1.21	1.17	1.02
CPIAUCSL	BVAR hom	1.06	1.04	1.06	1.03	1.21	1.04	1.04	1.03	1.02	1.13	1.08	1.02	1.09	1.03	1.30	1.06	1.03	1.06	1.03	1.21
	BVAR SV	1.04	1.04	1.03	1.01	1.14	1.06	1.08	1.05	1.05	1.13	1.01	0.97	1.01	0.98	1.16	1.05	1.06	1.05	1.03	1.15
	BVAR SV−t	1.06	1.05	1.07	1.04	1.21	1.07	1.08	1.07	1.06	1.16	1.05	1.02	1.06	1.02	1.26	1.07	1.06	1.07	1.04	1.22
	TVP SV	1.04	1.03	1.04	0.99	1.27	1.06	1.07	1.05	1.03	1.21	1.01	**0.96**	1.02	0.95'	1.32	1.05	1.01	1.06	1.00	1.29
	BART hom	1.01	1.12	0.98	0.98	1.14	1.04	1.13	1.01	1.02	1.13	0.97	1.11	0.94	0.94	1.15	1.03	1.16	1.00	1.01	1.13
	BART SV	0.99	1.08	0.97	0.98	1.09	1.03	1.09	1.02	1.02	1.10	0.95	1.07	0.93	**0.92'**	1.09	1.02	1.12	1.00	1.01	1.09
	GP hom	0.97	1.16	0.92'	0.98	0.89	1.02	1.17	0.98	1.03	1.01	0.91'	1.16	0.87°	0.94	0.76	0.96	1.17	0.91'	0.98	0.87
	GP SV	0.97	1.17	**0.92'**	0.98	0.91	1.03	1.18	0.99	1.03	1.05	**0.90'**	1.17	**0.85°**	0.93	0.77	0.95	1.17	**0.90'**	0.97	0.87
UNRATE	BVAR hom	1.14	0.98	1.24	1.05	1.21	1.19	0.97	1.28	1.05	1.30	1.10	0.99	1.19	1.04	1.14	1.16	0.97	1.27	1.09	1.23
	BVAR SV	1.03	0.99	1.06	1.00	1.06	1.05	0.98'	1.07	1.01	1.08	1.03	0.99	1.07	0.98	1.08	1.03	0.98	1.05	1.06	1.00
	BVAR SV−t	1.05	0.97	1.11	1.00	1.09	1.08	0.95°	1.13	1.00	1.15	1.04	0.98	1.09	0.99	1.08	1.04	0.96°	1.09	1.06	1.02
	TVP SV	1.02	0.98	1.05	0.99	1.05	1.03	0.97°	1.05	1.01	1.04	1.03	0.99	1.06	0.97	1.08	1.02	0.97'	1.05	1.05	1.00
	BART hom	0.97	0.97	0.98	1.00	0.95'	0.99	0.96'	1.01	1.03	0.96	0.95°	0.98	0.94°	0.97	0.94°	0.97	0.96	0.97	1.01	0.93'
	BART SV	0.95°	0.99	0.92°	1.00	0.91*	0.94°	0.99	0.91°	1.00	0.88°	0.96°	1.00	0.92°	0.99	0.93*	0.97	0.99	0.96	1.03	0.91°
	GP hom	0.95	0.98	0.92	1.01	0.90'	0.94	0.95	0.93	1.03	0.85'	0.96	1.00	0.92°	0.98	0.94'	0.97	0.96	0.98	1.06	0.90'
	GP SV	**0.89'**	0.98	**0.83'**	1.00	**0.81'**	**0.82'**	0.96	0.76	1.03	**0.65'**	**0.95'**	1.00	**0.90***	0.97	0.93°	0.96	0.97	0.96	1.05	**0.88°**

M

		CRPS					CRPS−L					CRPS−R					MAE				
		FS	Re	Ex	Pre	Post	FS	Re	Ex	Pre	Post	FS	Re	Ex	Pre	Post	FS	Re	Ex	Pre	Post
GDPC1	BVAR hom	0.44	1.34	0.34	0.32	1.82	0.13	0.47	0.09	0.09	0.55	0.13	0.31	0.11	0.10	0.50	0.53	1.68	0.40	0.41	2.02
	BVAR SV	0.97	0.98	0.97	0.94*	1.04	1.01	0.98	1.02	**0.94***	1.14	0.95'	0.99	0.93'	0.93*	0.98	1.00	0.99'	1.01	0.99'	1.04
	BVAR SV−t	**0.97**	**0.96'**	0.98	**0.94***	1.04	1.01	0.97	1.04	0.96'	1.11	**0.93'**	0.94*	0.93	**0.91***	0.99	1.01	0.93°	1.05	**0.98**	1.07
	TVP SV	0.98	0.99	0.98	0.96*	1.03	1.01	0.98	1.03	0.95°	1.13	0.96	1.00	0.94	0.96*	0.95	1.02	0.99	1.04	1.03	1.00
	BART hom	1.02	1.07	0.99	1.05	**0.95**	1.03	1.07	1.01	1.06	**0.99**	1.00	1.07	0.98	1.04	0.90	1.01	1.07	0.98	1.03	**0.96**
	BART SV	0.98	1.06	**0.94**	0.98	0.96	1.02	1.05	1.00	0.98	1.10	0.93	1.08	**0.89**	0.98	0.83	1.01	1.07	0.99	1.03	0.98
	GP hom	1.04	1.02	1.05	1.09	0.95	1.07	0.99	1.12	1.11	0.99	1.03	1.06	1.02	1.08	0.91	**1.00**	1.04	**0.98**	1.01	0.96
	GP SV	0.99	1.01	0.99	1.02	0.95	1.04	0.98	1.08	1.01	1.09	0.96	1.06	0.93	1.02	**0.82**	1.00	1.02	0.99	1.02	0.96
CPIAUCSL	BVAR hom	0.26	0.48	0.23	0.23	0.51	0.08	0.17	0.07	0.07	0.16	0.07	0.12	0.07	0.07	0.15	0.35	0.61	0.32	0.32	0.71
	BVAR SV	0.97'	1.04	0.95°	0.98	0.89°	1.01	1.07	0.99	1.02	0.95	0.92*	1.01	0.90*	0.93°	0.82°	0.97'	1.06	0.95*	0.98	0.93'
	BVAR SV−t	0.97	1.02	0.96'	0.99	0.89'	1.01	1.04	1.00	1.02	0.94	0.93°	1.02	0.91°	0.95	0.81°	0.98	1.05	0.97'	0.99	0.92'
	TVP SV	0.97	1.05	0.96'	0.99	0.90	1.01	1.06	1.00	1.02	0.99	0.93°	1.03	0.91°	0.95'	0.79'	0.99	1.06	0.97	0.99	0.94
	BART hom	**0.95**	1.06	0.93'	**0.96**	0.91	**0.98**	1.07	**0.95**	**0.99**	0.90	0.93'	1.05	0.90'	0.93	0.92	**0.95**	1.09	0.92°	**0.95**	0.92
	BART SV	0.96	1.06	0.93'	0.98	0.83'	0.99	1.04	0.97	1.00	0.92	0.92'	1.09	0.89°	0.95	0.73'	0.96	1.09	0.93'	0.98	0.87
	GP hom	1.16	1.23	1.14	1.23	**0.76'**	1.08	1.20	1.04	1.13	**0.82'**	1.26	1.27	1.26	1.36	**0.69'**	1.08	1.28	1.03	1.13	0.79
	GP SV	1.14	1.22	1.12	1.20	0.79	1.05	1.16	1.02	1.09	0.86	1.25	1.33	1.23	1.34	0.71'	1.08	1.28	1.04	1.14	**0.79**
UNRATE	BVAR hom	0.22	0.81	0.15	0.11	1.55	0.06	0.16	0.05	0.03	0.40	0.07	0.31	0.04	0.03	0.50	0.26	0.92	0.19	0.14	1.73
	BVAR SV	0.98	0.98°	0.99	0.93*	1.03	0.98	0.98*	0.98	0.95*	1.00	1.01	0.98	1.03	0.90*	1.08	0.96	0.98*	0.95	0.98°	0.95
	BVAR SV−t	1.02	**0.95***	1.07	**0.91***	1.11	1.06	**0.95***	1.11	**0.93***	1.17	1.00	**0.95°**	1.03	0.89*	1.08	0.97	**0.95***	0.98	**0.95°**	0.98
	TVP SV	0.96'	0.96°	0.96	0.92*	0.99	0.95'	0.97*	0.94'	0.96*	0.93	0.98	0.96'	1.01	**0.88***	1.07	**0.94**	0.96*	0.93	0.96°	0.93
	BART hom	0.99	1.06	0.94	1.09	0.91'	0.97	1.07	0.93	1.09	0.87'	1.00	1.05	0.96	1.10	**0.93°**	0.98	1.06	**0.93**	1.07	0.89°
	BART SV	0.95	1.07	0.87'	1.05	0.86'	0.86	1.10	**0.76'**	1.08	0.68'	1.02	1.05	0.99	1.03	1.01	1.00	1.10	0.95	1.12	0.90°
	GP hom	1.09	1.11	1.08	1.31	0.91'	1.07	1.16	1.04	1.29	0.89	1.12	1.08	1.15	1.35	0.94'	1.04	1.14	0.98	1.19	0.90'
	GP SV	1.00	1.07	0.96	1.18	0.86	0.91	1.15	0.81	1.20	0.66	1.09	1.03	1.13	1.17	1.02	1.03	1.12	0.98	1.18	0.89'

Table OA.4: Two-year ahead predictive loss metrics for the US.

S

Variable	Model	CRPS FS	CRPS Re	CRPS Ex	CRPS Pre	CRPS Post	CRPS−L FS	CRPS−L Re	CRPS−L Ex	CRPS−L Pre	CRPS−L Post	CRPS−R FS	CRPS−R Re	CRPS−R Ex	CRPS−R Pre	CRPS−R Post	MAE FS	MAE Re	MAE Ex	MAE Pre	MAE Post
GDPC1	BVAR hom	1.03	1.02	1.03	1.04	1.01	1.02	1.02	1.02	1.03	1.00	1.04	1.02	1.04	1.05	1.01	1.04	1.02	1.05	1.05	1.02
GDPC1	BVAR SV	1.04	1.02	1.05	1.03	1.07	1.03	1.04	1.03	1.02	1.06	1.05	1.00	1.07	1.04	1.09	1.03	1.01	1.04	1.04	1.01
GDPC1	BVAR SV−t	1.04	1.01	1.05	1.02	1.07	1.03	1.03	1.03	1.01	1.07	1.05	1.00	1.07	1.04	1.09	1.03	1.01	1.03	1.03	1.01
GDPC1	TVP SV	1.02	0.99	1.04	1.00	1.07	1.02	1.00	1.04	1.00	1.05	1.03	0.98	1.04	1.00	1.11	0.99	0.99	0.99	0.99	1.00
GDPC1	BART hom	**0.98***	_0.98_	**0.97***	**0.97***	_0.99_	**0.99°**	_0.98_	**0.99**	**0.98°**	1.00	**0.97***	0.98	**0.97***	**0.96***	_0.99'_	_0.97*_	_0.98_	_0.96*_	_0.96*_	0.99
GDPC1	BART SV	1.00	1.00	1.01	0.99'	1.04	1.01	1.01	1.01	1.00	1.04	1.00	_0.98°_	1.01	0.98*	1.06	0.98'	0.99	0.98	0.98'	1.00
GDPC1	GP hom	1.00	0.99	1.01	1.00	1.00	1.00	0.98	1.01	1.01	_1.00_	1.00	1.00	1.00	1.00	1.00	0.98'	0.99	0.97	0.98	_0.99_
GDPC1	GP SV	1.02	0.99	1.03	1.01	1.04	1.01	0.99	1.03	1.00	1.03	1.02	0.99	1.03	1.01	1.06	0.98'	0.99	0.97'	0.97'	1.00
CPIAUCSL	BVAR hom	1.06	0.96	1.08	1.09	0.94	1.04	0.94	1.07	1.04	1.01	1.09	1.02	1.10	1.14	0.89	1.06	0.98	1.08	1.08	0.99
CPIAUCSL	BVAR SV	1.07	**0.93**	1.10	1.08	1.01	1.09	**0.91**	1.14	1.09	1.09	1.06	0.99	1.07	1.08	0.97	1.06	**0.96'**	1.08	1.06	1.05
CPIAUCSL	BVAR SV−t	1.07	0.96	1.09	1.09	1.00	1.09	0.94	1.14	1.09	1.10	1.06	1.02	1.06	1.09	0.93	1.06	0.98	1.07	1.06	1.04
CPIAUCSL	TVP SV	1.02	0.94	1.04	1.05	0.90	1.06	0.93	1.09	1.07	1.00	1.00	_0.98_	1.01	_1.05_	0.83	1.03	0.96	1.04	1.05	0.92
CPIAUCSL	BART hom	1.12	1.09	1.13	1.20	_0.79_	1.06	1.06	1.06	1.10	_0.86_	1.18	1.14	1.19	1.29	0.74	1.15	0.99	1.19	1.22	0.82
CPIAUCSL	BART SV	1.01	1.06	1.00	1.07	0.79	0.99	1.06	0.97	1.01	0.88	1.04	1.07	1.03	1.12	_0.72_	1.01	0.97	1.02	1.06	0.83
CPIAUCSL	GP hom	0.98	1.05	0.97	1.03	0.80	**0.94**	1.03	**0.92'**	0.96	0.87	1.03	1.09	1.02	1.10	0.75	0.98	1.01	0.97	1.01	_0.82_
CPIAUCSL	GP SV	_0.97_	1.04	_0.95_	_1.00_	0.82	0.95	1.04	0.92'	**0.96**	0.89	_1.00_	1.05	_0.99_	1.06	0.77	_0.95_	1.00	_0.94_	_0.97_	0.85
UNRATE	BVAR hom	1.02	1.00	_1.03_	1.04	**1.00**	1.03	1.02	1.03	1.06	_0.99_	_1.02_	0.99	_1.04_	1.03	_1.00_	1.02	1.01	1.02	1.02	1.01
UNRATE	BVAR SV	1.04	1.01	1.07	1.02	1.07	1.06	1.00	1.08	1.02	1.09	1.05	1.01	1.08	1.02	1.07	_0.99_	1.01	_0.98_	**0.98**	1.00
UNRATE	BVAR SV−t	1.04	1.01	1.07	1.01	1.07	1.06	1.01	1.08	1.03	1.09	1.04	1.01	1.07	_1.01_	1.07	0.99	1.01	0.98	1.00	**0.99**
UNRATE	TVP SV	1.06	1.01	1.10	1.04	1.08	1.08	1.01	1.12	1.08	1.09	1.05	1.01	1.10	1.02	1.09	1.01	1.01	1.01	1.01	1.01
UNRATE	BART hom	1.01	_0.98*_	1.04	1.02	1.01	1.01	**0.99***	1.02	1.01	1.01	1.02	_0.97*_	1.07	1.03	1.01	1.02	**0.98***	1.04	1.02	1.01
UNRATE	BART SV	1.03	1.00	1.05	1.03	1.02	1.01	0.99°	1.02	_1.00_	1.03	1.04	1.01	1.07	1.05	1.02	1.03	1.00	1.06	1.04	1.02
UNRATE	GP hom	_1.01_	0.98	1.03	1.01	1.01	_1.01_	0.99	_1.02_	1.01	1.01	1.02	0.98	1.05	1.02	1.01	1.01	0.98	1.03	1.01	1.01
UNRATE	GP SV	1.02	0.98'	1.05	_1.01_	1.03	1.02	0.99'	1.04	1.01	1.03	1.02	0.98	1.06	1.02	1.03	1.01	0.99	1.02	1.01	1.01

M

Variable	Model	CRPS FS	CRPS Re	CRPS Ex	CRPS Pre	CRPS Post	CRPS−L FS	CRPS−L Re	CRPS−L Ex	CRPS−L Pre	CRPS−L Post	CRPS−R FS	CRPS−R Re	CRPS−R Ex	CRPS−R Pre	CRPS−R Post	MAE FS	MAE Re	MAE Ex	MAE Pre	MAE Post
GDPC1	BVAR hom	0.46	1.44	0.34	0.35	1.75	0.14	0.51	0.09	0.10	0.54	0.14	0.33	0.11	0.11	0.47	0.57	1.78	0.42	0.44	1.97
GDPC1	BVAR SV	1.02	1.01	1.03	1.00	1.08	1.03	1.03	1.03	1.01	1.07	1.02	0.99*	1.03	0.99	1.10	0.99°	1.00	0.98°	0.99°	1.00
GDPC1	BVAR SV−t	1.03	1.02	1.03	1.00	1.08	1.02	1.04	1.00	_1.00_	1.07	1.03	1.00	1.05	1.00	1.11	1.00	1.01	0.99	1.00	1.00
GDPC1	TVP SV	1.04	1.01	1.06	1.03	1.08	1.04	1.02	1.05	1.02	1.08	1.05	1.00	1.07	1.03	1.09	1.01	1.01	1.00	1.01	1.00
GDPC1	BART hom	_0.99_	0.98*	_1.00_	0.99	0.99	_1.00_	0.98*	1.02	1.00	1.00	_0.99_	1.00	_0.99_	0.99	0.99	0.97*	0.99	0.96*	0.97°	1.00
GDPC1	BART SV	1.01	1.00	1.02	_0.99_	1.06	1.02	1.01	1.03	1.01	1.06	1.01	**0.98***	1.02	_0.98°_	1.07	0.96*	0.99°	0.95*	0.96*	0.99
GDPC1	GP hom	1.03	**0.96°**	1.06	1.05	**0.99**	1.03	**0.92***	1.11	1.05	**0.99**	1.04	1.01	1.04	1.06	**0.99**	0.96*	0.98	**0.95***	**0.95***	1.00
GDPC1	GP SV	1.04	0.97°	1.08	1.04	1.05	1.05	0.96°	1.12	1.05	1.05	1.05	1.00	1.07	1.04	1.07	**0.96***	**0.98'**	0.95*	0.95*	**0.99**
CPIAUCSL	BVAR hom	0.33	0.57	0.30	0.29	0.76	0.09	0.20	0.08	0.08	0.20	0.10	0.13	0.10	0.09	0.24	0.44	0.75	0.40	0.39	0.99
CPIAUCSL	BVAR SV	0.98	0.96°	0.99	0.99	0.96	1.02	0.96	1.04	1.02	1.04	0.95'	0.98	0.95'	**0.96**	0.92	0.96*	0.97°	0.96°	0.96°	0.97
CPIAUCSL	BVAR SV−t	0.99	0.98'	0.99	1.00	0.94	1.04	0.98	1.05	1.04	1.01	0.96	0.99	0.95	0.97	0.91	0.97	0.98	0.97	0.98	0.93
CPIAUCSL	TVP SV	1.00	_0.94'_	1.02	1.02	0.95	1.07	_0.94'_	1.10	1.07	1.04	0.95	**0.97**	0.95	0.97	0.90	0.98	_0.96_	0.98	0.98	0.96
CPIAUCSL	BART hom	0.98	1.05	0.96	1.00	0.88	0.97	1.06	0.95	0.99	0.89	0.98	1.03	0.97	1.01	0.88	0.99	0.99	0.99	1.01	0.88
CPIAUCSL	BART SV	**0.95**	1.02	**0.93**	**0.97**	0.87	_0.95_	1.03	_0.93'_	_0.96_	0.90	**0.95**	1.01	**0.94**	0.97	0.85	**0.93'**	0.98	**0.91**	**0.94**	0.87
CPIAUCSL	GP hom	1.12	1.07	1.13	1.20	**0.74**	1.01	1.02	1.00	1.05	**0.83**	1.22	1.15	1.23	1.36	0.69	1.15	1.06	1.17	1.23	**0.78**
CPIAUCSL	GP SV	1.13	1.07	1.15	1.22	0.76	1.03	1.02	1.03	1.06	0.87	1.23	1.16	1.25	1.37	**0.67**	1.15	1.07	1.17	1.23	0.80
UNRATE	BVAR hom	0.25	0.96	0.16	0.14	1.46	0.07	0.19	0.05	0.04	0.39	0.08	0.36	0.04	0.04	0.45	0.29	1.08	0.19	0.18	1.53
UNRATE	BVAR SV	1.07	1.01	1.12	1.01	1.13	1.08	0.99	1.12	**0.99**	1.17	1.07	1.02	1.13	1.02	1.13	1.02	1.00	1.03	1.02	1.01
UNRATE	BVAR SV−t	1.07	1.02	1.10	1.02	1.11	1.08	1.00	1.11	1.01	1.14	1.06	1.03	1.10	1.03	1.11	1.02	1.01	1.03	1.01	1.04
UNRATE	TVP SV	1.06	1.02	1.09	1.01	1.11	1.08	1.01	1.12	1.04	1.13	1.06	1.03	1.09	**1.00**	1.12	**0.99**	1.02	**0.97°**	_0.98_	_1.00_
UNRATE	BART hom	1.02	0.99*	1.04	1.02	1.01	1.01	0.99*	1.02	1.01	1.01	1.02	0.99°	1.05	1.03	1.01	1.02	0.99*	1.05	1.03	1.02
UNRATE	BART SV	1.05	1.01	1.08	1.03	1.07	1.04	_0.99°_	1.07	1.00	1.09	1.05	1.02	1.09	1.05	1.06	1.03	1.00	1.05	1.03	1.01
UNRATE	GP hom	1.07	**0.96°**	1.15	1.14	1.00	1.06	1.01	1.08	1.13	**0.99**	1.09	**0.93***	1.25	1.15	1.02	1.07	_0.98'_	1.13	1.10	1.03
UNRATE	GP SV	1.10	0.97*	1.19	1.12	1.07	1.09	1.00	1.13	1.10	1.08	1.11	0.95*	1.28	1.15	1.07	1.06	0.98°	1.12	1.10	1.02