LET THE TREE DECIDE: FABART. A NON-PARAMETRIC FACTOR MODEL FOR NONLINEAR OIL SHOCK TRANSMISSION

2026

BANCO DE ESPAÑA
Eurosistema

Documentos de Trabajo
N.º 2627

Eoghan O'Neill and Sofia Velasco

LET THE TREE DECIDE: FABART. A NON-PARAMETRIC FACTOR MODEL FOR NONLINEAR OIL SHOCK TRANSMISSION

LET THE TREE DECIDE: FABART. A NON-PARAMETRIC FACTOR MODEL FOR NONLINEAR OIL SHOCK TRANSMISSION [*]

Eoghan O'Neill

UNIVERSITY COLLEGE DUBLIN

Sofia Velasco [**]

BANCO DE ESPAÑA

(*) The authors thank the Society for Nonlinear Dynamics and Econometrics (SNDE) for awarding this paper the Time Series Award. We are grateful to Matteo Ciccarelli, Juan J. Dolado, Florian Huber, Lutz Kilian, Karin Klieber, Lorenzo Mori, Haroon Mumtaz and Luca Rossini, as well as to seminar participants at the 2025 European Winter Meeting of the Econometric Society, the 2025 Banque de France AI Methods Conference, CFE-CMStatistics 2025, the 16th Workshop in Time Series Econometrics, Zaragoza, the 33rd Symposium of the Society for Nonlinear Dynamics and Econometrics and the anonymous referee, for helpful comments and suggestions.
Generative AI tools (OpenAI GPT and Google Gemini) were used to assist with code review and editing. The authors take full responsibility for all analyses, results and conclusions. The views expressed in this paper are those of the authors only and do not necessarily reflect those of Banco de España or the European System of Central Banks (ESCB). Replication code will be made available on GitHub.
(**) sofia.velasco@bde.es. An earlier version of this paper was circulated on arXiv as Velasco (2025) "Let the Tree Decide: FABART. A Non-Parametric Factor Model".

Documentos de Trabajo. N.º 2627

September 2026

https://doi.org/10.53479/44248

The Working Paper Series seeks to disseminate original research in economics and finance. All papers have been anonymously refereed. By publishing these papers, the Banco de España aims to contribute to economic analysis and, in particular, to knowledge of the Spanish economy and its international environment.

The opinions and analyses in the Working Paper Series are the responsibility of the authors and, therefore, do not necessarily coincide with those of the Banco de España or the Eurosystem.

The Banco de España disseminates its main reports and most of its publications via the Internet at the following website: http://www.bde.es.

Reproduction for educational and non-commercial purposes is permitted provided that the source is acknowledged.

© BANCO DE ESPAÑA, Madrid, 2026

ISSN: 1579-8666 (online edition)

Abstract

The question of how oil supply news shocks transmit to real activity, financial conditions and regional labor markets is back at the center of the macroeconomic research agenda. To answer this question, we introduce the Factor Bayesian Additive Regression Tree (FABART) model, a nonlinear factor-augmented vector autoregression model, and apply it to a large U.S. macro-financial dataset with externally identified oil supply news shocks. The framework incorporates a flexible nonparametric measurement equation that allows nonlinear transmission to emerge from the data without imposing a specific functional form on the asymmetry. We find that adverse oil supply news shocks generate stronger and more persistent contractions in real activity than the expansions associated with favorable shocks of comparable magnitude, with especially pronounced differences in industrial production, interest rates and equity prices. Employment responses are highly heterogeneous across U.S. states, with substantially stronger contractions in manufacturing-intensive regions than in energy-producing states. Across shock magnitudes, nonlinearities arise mainly between very small and moderate oil-price movements: small shocks generate weak and imprecisely estimated responses, while even moderate shocks produce economically meaningful effects on industrial production and regional employment. Larger shocks, however, do not systematically generate proportionally stronger responses across variables and shock signs.

Keywords: Bayesian FAVAR, nonlinear models, non-parametric techniques, oil price shocks, factor models.

JEL classification: C11, C32, E32, Q43.

Resumen

La transmisión de las noticias acerca de las perturbaciones sobre la oferta de petróleo a la actividad económica, las condiciones financieras y los mercados laborales regionales vuelve a ser una cuestión central en la investigación macroeconómica. Para abordarla, desarollamos un modelo no lineal de vectores autorregresivos aumentados con factores, el modelo FABART (Factor Bayesian Additive Regression Tree, por sus siglas en inglés), y lo aplicamos a un amplio conjunto de datos macroeconómicos y financieros de Estados Unidos, junto con noticias acerca de los *shocks* sobre la oferta de petróleo identificados mediante instrumentos externos. El modelo incorpora una ecuación de observación no paramétrica y flexible que permite que las no linealidades en el mecanismo de transmisión emerjan de los datos, sin imponer una forma funcional específica para las asimetrías. Los resultados muestran que las noticias adversas sobre perturbaciones de oferta de petróleo dan lugar a contracciones de la actividad económica más intensas y persistentes que las expansiones asociadas a noticias favorables de magnitud comparable. Estas asimetrías son especialmente pronunciadas en la producción industrial, los tipos de interés y las cotizaciones bursátiles. Asimismo, la respuesta del empleo presenta una notable heterogeneidad entre los distintos estados de Estados Unidos, con contracciones significativamente más intensas en las regiones con un mayor peso de la industria manufacturera que en los estados productores de energía. En cuanto al tamaño de los *shocks,* las no linealidades se manifiestan principalmente al comparar perturbaciones de muy pequeña magnitud con las moderadas. Mientras que los *shocks* de pequeña magnitud generan respuestas débiles y estimadas con escasa precisión, los *shocks* moderados producen efectos económicamente relevantes sobre la producción industrial y el empleo regional. En cambio, las perturbaciones de mayor magnitud no generan, de forma sistemática, respuestas proporcionalmente más intensas, ni entre las distintas variables ni en función del signo del *shock*.

Palabras clave: modelo de factores dinámicos bayesiano, modelos no lineales, técnicas no paramétricas, *shocks* al precio del petróleo, modelos de factores.

Códigos JEL: C11, C32, E32, Q43.

1 Introduction

Estimating the macroeconomic effects of oil supply news shocks requires addressing two closely related challenges. First, the structural shock must be credibly identified. Second, the propagation mechanism must be flexible enough to capture nonlinear transmission without imposing asymmetry by construction. This paper proposes a framework designed to jointly address these issues. We combine externally identified oil supply news shocks with a large macro-financial information set and a flexible nonlinear transmission mechanism in which asymmetries are estimated rather than imposed. Specifically, we develop a Factor Bayesian Additive Regression Tree (FABART) model that embeds Bayesian Additive Regression Trees into a factor-augmented vector autoregression (FAVAR) framework. The factor structure summarizes information from a large set of macroeconomic, financial, and regional variables, while the nonparametric measurement equation allows the mapping between latent factors and observables to vary flexibly across the support of the data.

First, the factor structure allows the model to incorporate a large macro-financial information set, helping overcome the informational deficiencies of low-dimensional oil-market VARs highlighted in the recent literature. Second, unlike many conventional nonlinear VAR approaches, the FABART framework does not require the researcher to specify ex ante the functional form of nonlinear transmission through threshold effects, regime-switching mechanisms, transformed regressors, or other restrictive nonlinear parameterizations. Instead, nonlinear responses emerge flexibly from the relationship between latent factors and observables through the nonparametric measurement equation. This distinction is important because it separates the economic question of whether nonlinear transmission exists from the econometric assumptions used to characterize it.

We apply FABART to study how externally identified oil supply news shocks propagate through the U.S. economy.[1] Recent geopolitical tensions have renewed interest in the macroeconomic consequences of oil supply disruptions. Recent evidence from analyses of the 2026 Iran War illustrates that geopolitical oil supply disturbances can still generate persistent effects on inflation and broader macroeconomic conditions (Kilian et al., 2026). More generally, oil supply shocks and the associated movements in oil prices have

[1] We use the external instrument proposed by Känzig (2021), which exploits high-frequency changes in oil futures prices around OPEC production announcements to identify oil supply news shocks. Within the structural oil-market literature, unexpected supply disruptions are distinguished from anticipated supply news shocks. The former are flow supply shocks, whereas the latter are identified as storage-demand shocks because they operate through inventory accumulation rather than contemporaneous production changes (Kilian and Murphy, 2014). Subsequent work has refined the construction and interpretation of these high-frequency proxies (Kilian, 2024; Degasperi, 2026) and examined the econometric conditions under which they provide valid external instruments in SVARs (Miranda-Agrippino and Ricco, 2023; Mori and Peersman, 2024).

long been recognized as an important source of macroeconomic fluctuations, affecting real activity, inflation, and labor-market dynamics (Hamilton, 2003; Kilian, 2008; Herrera and Karaki, 2015). Structural models distinguishing oil supply, aggregate demand, and oil-specific demand shocks show that these disturbances have markedly different macroeconomic consequences (Kilian, 2009). Even after conditioning on a common oil supply shock, however, important questions remain about the magnitude, persistence, and symmetry of its transmission.

A central challenge in this literature is the credible identification of structural oil shocks. Structural oil-market VARs address this challenge by distinguishing oil supply, flow-demand, and storage-demand shocks using economically motivated identifying restrictions and information on inventories (Kilian, 2009; Kilian and Murphy, 2014; Kilian and Zhou, 2023). More recently, Mori and Peersman (2024) show that conventional oil-market VARs remain informationally deficient because they omit macro-financial information that predicts both oil-market shocks and the external instrument of Känzig (2021). As a result, augmenting the information set materially alters the estimated transmission of oil supply news shocks, suggesting that informational deficiencies may persist even when identification relies on externally identified shocks.

This problem is part of a broader limitation of structural VAR analysis. Structural VARs are used to identify macroeconomic shocks and trace their transmission through real and financial variables. Yet empirical VARs are typically estimated using a relatively small number of carefully selected variables, implicitly assuming that the researcher observes the relevant information set. When important information is omitted, the identified shocks may fail to capture the structural disturbances perceived by economic agents and the resulting impulse responses may be distorted. This concern extends well beyond oil markets and represents a broader issue in structural macroeconometrics. A growing literature shows that expanding the information set can materially alter the estimated effects of identified macroeconomic shocks; see, for example, Forni and Gambetti (2010, 2014); Juvenal and Petrella (2015); Caldara and Herbst (2019); Miranda-Agrippino and Ricco (2023); Mori and Peersman (2024).

Once the structural oil shock has been identified, the remaining question concerns its transmission. Many nonlinear econometric approaches require the researcher to specify ex ante the functional form governing nonlinearities. While such specifications can be appropriate, valid inference depends on correct model specification. A large literature following Mork (1989), Hamilton (2003), and Hamilton (2011) argues that positive oil shocks generate substantially larger macroeconomic effects than negative shocks. At the same time, Kilian and Vigfusson (2011a) and Kilian and Zhou (2023) show that commonly used nonlinear oil VAR specifications may yield inconsistent estimation and invalid impulse-

response inference, potentially generating spurious evidence of sign asymmetry. More generally, Gonçalves et al. (2021) emphasize that valid impulse-response analysis requires estimation and inference procedures consistent with the underlying nonlinear model. Recent work has therefore proposed more flexible approaches, including BART-based local projections (Mumtaz and Piffer, 2022), semiparametric local projections (Goncalves et al., 2026), and related methods that allow asymmetries to emerge from the data rather than being imposed through restrictive functional forms (David et al., 2026). Our approach follows this direction by embedding flexible nonlinear transmission within a nonlinear factor model rather than estimating impulse responses equation by equation.

The dynamic FABART model introduced in this paper is an example of a broad class of dynamic nonlinear dimension reduction methods (Yalcin and Amemiya, 2001) and nonlinear state-space models. Moran (2019) introduced a static BART-based factor model in the context of dimension reduction. Bobeica et al. (2025) combine BART with a VAR and latent factors in a nonlinear multivariate framework. FABART differs from Moran (2019) and Bobeica et al. (2025) in two respects. First, it models the latent factors dynamically through a VAR transition equation, allowing them to capture persistent macroeconomic dynamics. Second, structural identification is imposed on a linear approximation to the nonlinear measurement equation while the nonlinear mapping between latent factors and observables is learned nonparametrically. By contrast, Bobeica et al. (2025) decompose the shock-loading function into linear and nonlinear components and impose sign, zero, and magnitude restrictions on the linear component. Clark et al. (2025) instead augment a Bayesian VAR with nonlinear factors constructed as sum-of-tree functions of lagged observables, with outcome-specific loadings. Their factors are therefore static nonlinear transformations of observed predictors, whereas FABART specifies sum-of-tree functions of dynamic latent factors with a linear transition equation. Other approaches to nonlinear dynamic factor models include Gaussian Processes (Chernis et al., 2025) and deep neural networks (Andreini et al., 2020).[2] An alternative approach is a nonlinear VAR with Bayesian trees and no latent factors, i.e. a Bayesian Additive Vector Autoregression Tree model (Huber and Rossini, 2022).

Our empirical application augments the oil-market variables of Känzig (2021) with macroeconomic indicators, financial variables, uncertainty measures, and state-level labor-market data. In particular, incorporating regional employment variables allows us to study heterogeneous transmission mechanisms associated with sectoral specialization, energy dependence, and labor-market composition. This dimension is particularly relevant because oil shocks are often believed to propagate through sectoral reallocation and heterogeneous

[2]While Chernis et al. (2025) recover the latent factors using Particle Gibbs with Ancestor Sampling (PGAS), FABART relies on a linear approximation of the BART measurement equation for latent factor recovery.

employment adjustment (Davis and Haltiwanger, 2001; Herrera and Karaki, 2015; Hamilton and Owyang, 2012).

Our main findings can be summarized as follows. First, augmenting the information set with financial variables materially changes the estimated transmission of oil supply news shocks, producing sharper contractions in real activity, stronger financial-market responses, and weaker inflationary effects. Second, we find strong evidence of sign asymmetry. Adverse oil supply news shocks generate substantially larger contractions than favorable shocks generate expansions, particularly in employment and labor-market variables. Third, the effects are highly heterogeneous across U.S. states. Manufacturing-intensive states exhibit substantially stronger employment contractions, while oil-producing states display partially offsetting dynamics following oil-price increases. Fourth, size also plays an important role in the transmission, with responses changing materially when moving from small to moderate oil-price movements.

These results relate to recent contributions studying oil-market transmission using externally identified oil supply news shocks based on the instrument of Känzig (2021). The first finding reinforces the importance of informational sufficiency emphasized by Mori and Peersman (2024): expanding the information set with macro-financial and regional variables materially alters the estimated propagation of oil supply news shocks. Our evidence complements the findings of Forni et al. (2025), who use a nonlinear Proxy-SVAR framework and emphasize the role of uncertainty and financial conditions in amplifying the transmission of adverse shocks. At the same time, our approach is designed to address the concerns raised by Kilian and Vigfusson (2011a, 2017): rather than imposing a particular functional form for nonlinear transmission, the model allows nonlinear responses to emerge from a flexible nonparametric measurement equation estimated using externally identified oil supply news shocks and a large information set. The size results speak to a related but distinct form of nonlinearity, consistent with the argument in Caravello and Martinez-Bruera (2024) that sign and size asymmetries need not coincide empirically.

The results show that accounting jointly for informational sufficiency and flexible nonlinear transmission materially changes the estimated effects of oil supply news shocks. More broadly, the findings suggest that several apparently conflicting results in the oil-shock literature may reflect differences in identification, information sets, and nonlinear specifications rather than fundamentally different economic mechanisms. In this sense, FABART provides a framework for revisiting the asymmetry debate while minimizing the risk that estimated nonlinear responses are driven by restrictive functional-form assumptions rather than by the data themselves.

The remainder of the paper is organized as follows. Section 2 presents the model, esti-

mation strategy, and identification approach. Section 3 discusses the empirical findings. Section 4 concludes.

2 FABART: A Non-Parametric Dynamic Factor Model

In this section, we present the empirical framework used to extract latent factors from a large information set without imposing a parametric functional form on the relationship between observables and factors, identify oil supply news shocks, and study their potentially nonlinear transmission to macroeconomic variables.

Our approach combines two strands of literature. The first relates to dynamic factor models and FAVARs (Bernanke et al., 2005; Stock and Watson, 2005; Forni et al., 2009; Charnavoki and Dolado, 2014), which provide a parsimonious way to summarize large datasets and trace the effects of structural disturbances on a broad set of macroeconomic variables. The second relates to Bayesian additive regression trees (BART) (Chipman et al., 2010), which offer a flexible nonparametric way of capturing nonlinear, state-dependent, and asymmetric relationships without requiring a specific functional form. Combining these two approaches, the Factor Bayesian Additive Regression Tree (FABART) model preserves the dimension-reduction advantages of factor models while letting the data determine how latent factors map into the large set of observables.

2.1 Model

FABART is designed to capture nonlinearities in the measurement equation of a dynamic factor model. Related contributions introduce nonlinear factor structures through parametric or semiparametric specifications, including quantile-based factor models and nonlinear loading functions (Chen et al., 2021, 2026; Korobilis and Schröder, 2024). We instead use BART to approximate the mapping from latent factors to observables nonparametrically. This places the model closer to recent work that combines factor structures with flexible nonparametric components, including Bobeica et al. (2025), Hauzenberger et al. (2023), Chernis et al. (2025), and Clark et al. (2024). Introducing this flexible measurement equation allows the model to summarize high-dimensional datasets while letting nonlinear, asymmetric, and state-dependent relationships between factors and observables be learned from the data. [3]

Let $X_{i,t}$, with $i = 1, \ldots, N$ and $t = 1, \ldots, T$, denote a large panel of N observable variables

[3] MC experiments in the online Supplemental Appendix demonstrate point and interval forecasting performance, and the accuracy of impulse response estimates.

in T time periods. We assume these variables are driven by a lower-dimensional vector of J factors, $Y_t = (F_t^1, \ldots, F_t^J)'$, where the first factor is observed and corresponds to real oil prices in the application, $F_t^1 \equiv Z_t$. The remaining factors, $F_t^2, \ldots, F_t^J$, are latent. The loading of the observed factor in its own measurement equation is normalized to one.[4] [5]

To capture a flexible and potentially nonlinear relationship between the latent factors and the high-dimensional dataset $X_t = [X_{1,t}, X_{2,t}, \ldots, X_{N,t}]'$, we postulate the following measurement equation:

$$X_t = \mathcal{F}(Y_t) + \varepsilon_t, \quad \varepsilon_t \sim \mathcal{N}(0, \Sigma) \tag{1}$$

where $\mathcal{F}(Y_t) = (f_1(Y_t), \ldots, f_N(Y_t))'$ is an N-dimensional vector of unknown functions. Each $f_i(\cdot)$ is estimated nonparametrically using BART, which allows for heterogeneous, nonlinear, and asymmetric dependencies without strong parametric assumptions. The covariance matrix of the errors is diagonal, $\Sigma = \text{diag}(0, \sigma_2^2, \sigma_3^2, \ldots, \sigma_N^2)$

The dynamics of the latent factor vector Y_t follow a VAR(L) process (the transition equation):

$$Y_t = \sum_{l=1}^{L} \Phi_l Y_{t-l} + u_t, \quad u_t \sim \mathcal{N}(0, \Omega) \quad , \tag{2}$$

where and Φ_l are the autoregressive coefficient matrices in our application $L = 6$. No restrictions are placed on the covariance matrix Ω.

Each function $f_i(Y_t)$ in the measurement equation is approximated using a sum-of-trees representation:

$$f_i(Y_t) \approx \sum_{s=1}^{S} g_{is}(Y_t \mid \tau_{is}, \mu_{is}) \tag{3}$$

where $g_{is}(\cdot)$ is a regression tree with structure τ_{is} and terminal node values μ_{is}. The total number of trees S is fixed and controls the flexibility of the BART approximation.

Following Chipman et al. (2010), each tree partitions the predictor space using binary splits. Each rule selects one factor dimension, k, and a threshold c, which determine whether an observation lies in the left or right subtree. Since the predictors are the factors observed at time t, a binary split is written directly in terms of the contemporaneous factor vector Y_t as:

$$Y_{k,t} \leq c \quad \text{or} \quad Y_{k,t} > c \quad . \tag{4}$$

[4] In FAVAR applications identifying monetary policy shocks (e.g., Bernanke et al. (2005); Baumeister et al. (2013)), the Federal Funds Rate is often treated as an observed factor.

[5] Also, for identification, one can restrict $J - 1$ other measurement equations to be linear and load (with coefficient equal to one) only on each of the other $J - 1$ factors (Yalcin and Amemiya, 2001). i.e. an identity restriction is applied for identification. We also consider the application of this restriction to a linear approximation of the measurement equation for the purpose of sampling latent factors, without imposing the restriction in the sampling of fully nonlinear measurement equations.

Based on these recursive partitions, fitted values are obtained from a tree via the step function:

$$g(Y_t \mid \tau, \mu) = \sum_{r=1}^{b_\tau} \mu_r \cdot \mathbb{1}(Y_t \in R_r) \quad , \tag{5}$$

where R_r are the regions defined by the tree τ, and b_τ is the number of terminal nodes. Each μ_r is the parameter associated with terminal node r, drawn from a prior distribution detailed below. The indicator function $\mathbb{1}(Y_t \in R_r)$ equals one if Y_t lies within region R_r, and zero otherwise.

Priors

The Bayesian tree prior distributions introduced by Chipman et al. (2010) mitigate the risk of overfitting. The joint prior for the parameters governing the S trees for the i^{th} outcome variable decomposed as:

$$p((\tau_1^{(i)}, \mu_1^{(i)}), (\tau_2^{(i)}, \mu_2^{(i)}), \ldots, (\tau_S^{(i)}, \mu_S^{(i)})) = \prod_{s=1}^{S} p(\mu_s^{(i)} \mid \tau_s^{(i)}) p(\tau_s^{(i)}) \tag{6}$$

with conditional prior:

$$p(\mu_s^{(i)} \mid \tau_s^{(i)}) = \prod_{b=1}^{B_s} p(\mu_{bs}^{(i)} \mid \tau_s^{(i)}) \tag{7}$$

Each tree $\tau_s^{(i)}$ is defined by binary splitting rules. The tree prior $p(\tau_s^{(i)})$ penalizes tree depth and places uniform distributions over splitting variables and splitting points. [6] The prior probability that a node at depth d is non-terminal is $\alpha(1+d)^{-\beta}$, with $\alpha = 0.95$ and $\beta = 2$ following Chipman et al. (2010). Terminal node values $\mu_{bs}^{(i)}$ are independently drawn from a Gaussian prior centered at zero. Following Huber et al. (2020), the variance is calibrated as $\mu_{bs}^{(i)} \sim \mathcal{N}\left(0, \sigma_{\mu i}^2\right)$, with $\sigma_{\mu i} = \frac{\max(X_{\cdot i}) - \min(X_{\cdot i})}{2\sqrt{\nu S}}$, where S is the number of trees and ν reflects the prior confidence. For the error variance σ_i^2, we impose a conjugate scaled inverse-χ^2 prior: $\sigma_i^2 \sim \nu_i \xi_i / \chi_{\nu_i}^2$, with $\nu_i = 3$ and ξ_i defined such that $P(\sigma_i < \hat{\sigma}) = 0.9$, where $\hat{\sigma}_i$ is estimated from the data.[7]

The transition equation coefficients have a diffuse prior, and the covariance matrix of the transition equation errors has an inverse Wishart prior. The inverse Wishart prior on Ω has degrees of freedom equal to the number of factors plus 10, and scale equal to the

[6] We also considered a Dirichlet hyperprior on splitting rules as described by Linero (2018), although this is better suited to sparse settings than a dense factor model, and soft splitting rules introduced by Linero and Yang (2018).

[7] Chipman et al. (2010) use the sample variance of the residuals from a linear regression of the outcome on the predictors. We use the fifth percentile of the sample variances across all outcome variables, as suggested by Moran (2019). Alternatives include the sample variance of the residuals from a linear regression of the relevant outcome $X_{i,\cdot}$ on initial factor estimates from PCA, or the outcome-specific sample variance.

Table 1: Summary of prior hyperparameters for measurement equation.

Description	Parameters	Hyperparameters
Tree depth	$\alpha(1+d)^{-\beta}$	$\alpha = 0.95,\ \beta = 2$
Terminal node priors	$\mu_{bs} \sim \mathcal{N}(0, \sigma^2_{\mu i})$	$\sigma_{\mu i} = \frac{\max(X_{\cdot i}) - \min(X_{\cdot i})}{2\sqrt{\nu S}}$
Error variances	$\sigma_i^2 \sim \nu_i \xi_i / \chi^2_{\nu_i}$	$\nu_i = 3,\ P(\sigma_i < \hat{\sigma}_i) = 0.9$
Number of trees	–	$S = 50$

Notes: Prior distributions follow Chipman et al. (2010) and Huber et al. (2020).

sample covariance matrix of the residuals from OLS regression of principal components on their lags, multiplied by 9 to set the prior mean equal to the sample covariance of residuals.[8]

2.2 Estimation

2.2.1 Posterior simulation

The Bayesian backfitting Markov chain Monte Carlo (MCMC) algorithm of Chipman et al. (2010) is used to sample from the posterior distribution of the tree structures and terminal node parameters in the measurement equation. For each observable variable X_i, the tree structures τ_{is} are updated using the Metropolis–Hastings (MH) algorithm.

Integrating out the terminal node parameters μ_{is}, the marginal posterior for the s-th tree in equation i becomes:

$$p(\tau_{is} \mid R_{is}, \sigma_i) \propto p(\tau_{is}) \int p(R_{is} \mid \mu_{is}, \tau_{is}, \sigma_i) p(\mu_{is} \mid \tau_{is})\, d\mu_{is}, \tag{8}$$

where R_{is} denotes the partial residual for the s-th tree in equation i, given by:

$$R_{is} = X_{\cdot i} - \sum_{s' \neq s} g_{is'}(Y \mid \tau_{is'}, \mu_{is'}). \tag{9}$$

For each observed variable $X_{\cdot,i}$, Tree structures are sampled using the Metropolis–Hastings (MH) algorithm, as outlined in Chipman et al. (1998). A candidate tree τ^*_{is} is proposed

[8] Similar results are obtained with a conjugate normal inverse Wishart priors the transition equation in which the coefficients of factor lags have prior mean zero and variance equal to 10 divided by the sample variance of initial factor estimates from PCA. For robustness, we also consider a Minnesota prior on the coefficients, i.e. a normal prior with mean zero, and variance equal to 10 divided by the product of the sample variance of the initial factor estimate and the square of the order of the lag, for greater shrinkage on coefficients of higher order lags. We also implemented, but do not include results for, a transition equation specification with stochastic volatility.

from a probability distribution $q(\tau_{is}^{j}|\tau_{is}^{*})$, where the superscript j denotes the accepted draw at iteration j. The acceptance probability for setting $\tau_{is}^{j+1} = \tau_{is}^{*}$ is given by:

$$a(\tau_{is}^{j}, \tau_{is}^{*}) = \min\left(1, \frac{q(\tau_{is}^{*}|\tau_{is}^{j})p(R_{is} \mid \tau_{is}^{*}, \sigma_i)p(\tau_{is}^{*})}{q(\tau_{is}^{j}|\tau_{is}^{*})p(R_{is} \mid \tau_{is}^{j}, \sigma_i)p(\tau_{is}^{j})}\right). \tag{10}$$

The proposal distribution $q(\tau_{is}^{j}|\tau_{is}^{*})$ incorporates four types of moves:[9]

1. *Grow*: Expands a terminal node into two new child nodes (probability 0.25).
2. *Prune*: Merges two terminal nodes back into a single parent node (probability 0.25).
3. *Change*: Modifies the splitting rule at an internal node (probability 0.40).
4. *Swap*: Exchanges decision rules between a parent and child node (probability 0.10).

The algorithm for sampling from the full conditional posterior distribution is detailed below. The MCMC algorithm sequentially draws samples from the full conditional posterior distributions described in the preceding subsections. Given a set of initial values, the algorithm follows these steps in each iteration:[10]

1. Conditional on a draw of the unobserved components, sample the transition equation VAR coefficients and the reduced-form covariance matrix from their conditional posterior distributions.[11]
2. Sample the S trees for each of the N nonlinear measurement equations using the Metropolis–Hastings algorithm. The acceptance probability is determined by equation (10). After each tree structure is accepted or rejected, the leaf parameters are sampled from a Gaussian full conditional distribution.
3. Obtain the projected factor loadings through a linear approximation to the BART-implied fitted values (Crawford et al. (2019); Huber et al. (2020)).
4. Conditional on the draws obtained in Steps 1–3, apply the Carter and Kohn (1994) or Durbin and Koopman (2002) simulation smoother to sample the factors within the linearized state-space model described in 2.4.

[9] We also consider setting the swap proposal probability to 0, as recent papers suggest that this proposal has higher computational complexity and swap proposals might not be feasible for shallow trees (Maia et al., 2024; Kapelner and Bleich, 2016).

[10] Steps 2 and 5 are implemented using the R package `dbarts`.

[11] The transition equation is estimated using an inverse Wishart prior on the reduced-form covariance matrix Σ, with hyperparameters calibrated from a data-based linear VAR.

5. The out of sample factor values Y_{T+h} are sampled over the selected horizon and forecasts are sampled for the N observable variables $X_{i,t}$ conditional on the factor samples.

The algorithm generates a total of 30,000 draws, with the initial 15,000 discarded as burn-in.

2.3 Identification of Oil Supply News Shocks

This section describes the identification of oil supply news shocks and the computation of their dynamic effects within the FABART framework. We follow the external-instrument approach of Känzig (2021), which allows us to recover the structural shock using exogenous variation in oil futures markets.

Let $\mathbf{u}_t$ denote the reduced-form innovations from the VAR representation of the latent factors. These innovations are related to structural shocks $\boldsymbol{\varepsilon}_t$ through

$$\mathbf{u}_t = \mathbf{A}\boldsymbol{\varepsilon}_t, \tag{11}$$

where $\mathbf{A}$ is the contemporaneous impact matrix. In the absence of additional restrictions, $\mathbf{A}$ is not identified.

We identify an oil supply news shock using an external instrument m_t satisfying

$$m_t = \beta\varepsilon_{1t} + \sigma v_t, \quad v_t \sim \mathcal{N}(0,1), \tag{12}$$

together with the orthogonality condition

$$E(m_t\varepsilon_{jt}) = 0 \quad \text{for all } j \neq 1. \tag{13}$$

The first structural shock, ε_{1t}, is therefore interpreted as an oil supply news shock. In the empirical application, m_t is constructed from high-frequency changes in oil futures prices around OPEC announcements, following Känzig (2021).

Rather than augmenting the system with equation (12), identification is implemented by projecting the instrument onto the reduced-form residuals for each posterior draw of the VAR parameters, as in standard proxy-SVAR approaches (see Miranda-Agrippino and Ricco (2023)). This procedure recovers the impact vector associated with the oil supply news shock while preserving its external-instrument interpretation.

Given the nonlinear structure of the measurement equation, impulse responses are computed using the Generalized Impulse Response Function (GIRF) framework of Koop et al. (1996). For variable i at horizon h, the GIRF is defined as

$$\text{GIRF}_{ij}(h, \Psi_t, e_j, \overline{Y}) = E\left[X_{i,t+h} \mid \Psi_t, \overline{Y}, e_j\right] - E\left[X_{i,t+h} \mid \Psi_t, \overline{Y}\right], \tag{14}$$

where Ψ_t denotes the parameter draw, e_j the identified structural shock, and $\overline{Y}$ the reference state at which the response is evaluated.

In the baseline implementation, generalized impulse responses are evaluated at the unconditional mean of this transition state. Since the transition equation contains no intercept and the latent factors are assumed to be stationary, the unconditional mean of the transition state is centered at zero: $\overline{Y} = 0$.[12]

GIRFs are computed by simulating conditional paths from the estimated model. For each posterior draw, two paths are generated from the same reference state and under a common sequence of future innovations: one path includes the identified oil supply news shock at impact, while the other does not. The impulse response at each horizon is given by the difference between these two paths.

2.4 Linear approximation and state estimation

To sample the latent factors efficiently, we follow the strategy of Huber et al. (2020), who use the projection based approximation of Crawford et al. (2019), which allows for (approximate) sampling of latent variables in a nonlinear state space model using Forward Filter Backward Sampling.[13] The idea is to approximate the nonlinear measurement function by a local linear representation obtained from the fitted values of the nonlinear model. This approach was originally developed to interpret effect sizes in black box models, but it can be applied whenever the nonlinear function can be evaluated at the observed sample points.

Let $\widehat{\mathcal{F}}$ denote the $T \times N$ matrix of fitted values obtained from the BART measurement equation, with ith column $\widehat{\mathcal{F}}_i = (\widehat{f}_i(Y_1), \ldots, \widehat{f}_i(Y_T))'$. To avoid confusion with the latent

[12] More generally, if the transition equation includes a constant term C, the unconditional mean is given by $\overline{Y} = (I - \Phi)^{-1}C$, where Φ denotes the companion-form transition matrix, provided the process is stationary. Evaluating responses at this reference state abstracts from the influence of specific historical initial conditions and ensures that differences across responses reflect the imposed shock rather than the particular starting point of the system.

[13] We also considered replacing the linear projection with auxiliary penalized regressions as suggested by Marcellino and Pfarrhofer (2025). Furthermore, we have also implemented factor sampling without a linear approximation by using Particle Gibbs and Particle Gibbs with Ancestor Sampling (Lindsten et al., 2014).

factors, we use $\widehat{\mathcal{F}}$ for fitted BART values and reserve F_t^j for latent factors. Let $\mathbf{Y}$ denote the $T \times J$ matrix collecting the observed factor and latent factors over time, with rows $Y_t' = (Z_t, F_t^2, \ldots, F_t^J)$.

Following the projection argument in Huber et al. (2020), and adapting it to the FABART measurement equation, the projected linear coefficient for observable X_i is obtained as

$$\tilde{\Gamma}_i = \text{Proj}(\mathbf{Y}, \widehat{\mathcal{F}}_i) = \mathbf{Y}^\dagger \widehat{\mathcal{F}}_i, \tag{15}$$

where $\mathbf{Y}^\dagger$ denotes the Moore–Penrose pseudoinverse of $\mathbf{Y}$. If $\mathbf{Y}$ has full column rank, this is equivalent to $\tilde{\Gamma}_i = (\mathbf{Y}'\mathbf{Y})^{-1}\mathbf{Y}'\widehat{\mathcal{F}}_i$. Thus, $\tilde{\Gamma}_i$ is the linear projection of the BART fitted values for X_i onto the observed and latent factors.

This gives the approximate measurement relationship $X_{i,t} \approx \tilde{\Gamma}_i' Y_t + \epsilon_{i,t}$, where $\epsilon_{i,t}$ captures the residual component not explained by the projected nonlinear fit. Stacking across the N observables gives

$$X_t \approx \tilde{\Gamma} Y_t + \epsilon_t, \quad \epsilon_t \sim \mathcal{N}(0, \Sigma), \tag{16}$$

where the ith row of $\tilde{\Gamma}$ is $\tilde{\Gamma}_i'$.

The purpose of this step is computational. Conditional on the projected loadings, the nonlinear measurement equation is replaced by a Gaussian linear approximation, which allows the latent factors to be sampled using standard simulation smoothing methods. The nonlinear BART measurement equation is still used to update the fitted values and the projected loadings at each iteration, so the approximation is recomputed as the sampler cycles through the posterior draws.

In the full state space representation, the observed factor Z_t is included as the first element of the state vector and its loading is normalized to one. The resulting linearized measurement system can be written as

$$\begin{pmatrix} Z_t \\ X_{1,t} \\ \vdots \\ X_{N,t} \end{pmatrix} = \begin{pmatrix} 1 & 0 & \ldots & 0 \\ \Psi_1 & \Lambda_{12} & \ldots & \Lambda_{1J} \\ \Psi_2 & \Lambda_{22} & \ldots & \Lambda_{2J} \\ \vdots & \vdots & \ddots & \vdots \\ \Psi_N & \Lambda_{N2} & \ldots & \Lambda_{NJ} \end{pmatrix} \begin{pmatrix} Z_t \\ F_t^2 \\ \vdots \\ F_t^J \end{pmatrix} + \begin{pmatrix} 0 \\ \epsilon_{1,t} \\ \vdots \\ \epsilon_{N,t} \end{pmatrix}. \tag{17}$$

Here, Ψ_i denotes the loading of observable $X_{i,t}$ on the observed factor Z_t, while Λ_{ij} denotes the loading of $X_{i,t}$ on latent factor F_t^j for $j = 2, \ldots, J$. These coefficients are obtained from the projected BART fitted values. In the case with one observed factor, the loading matrix has dimension $(N+1) \times J$.

As in standard factor models, the latent factors and loadings are not separately identified without parameter restrictions, since rotations of the factor space can generate observationally equivalent representations. Following the factor model identification literature, we impose parameter restrictions on the linearized measurement matrix to fix the scale, sign, and ordering of the factors. In particular, the square submatrix of the first J rows and columns of the factor loading matrix in the approximated measurement equation is restricted to be an identity matrix.[14] Based on the resulting Gaussian linear approximation, the latent factors are drawn using the simulation smoother of Carter and Kohn (1994), as implemented in Mumtaz (2010).

3 Results

This section studies how externally identified oil supply news shocks transmit to real activity, financial conditions, and state labor markets. We first assess how the information set affects identification and transmission. We then examine nonlinearities across shock signs and magnitudes. Finally, we use state-level employment to study how aggregate responses are distributed across regions with different production structures.

3.1 Oil Market and U.S. Macroeconomic Activity

The dataset consists of the oil-market and macroeconomic variables from Känzig (2021), state-level labor market indicators, and forward-looking financial variables.

Our first specification follows the baseline set of variables in Känzig (2021), which combines the oil-market block, consisting of real oil prices, global industrial production, world oil production, and oil inventories, with key macroeconomic aggregates, including U.S. industrial production and consumer prices. We augment this information set with state-level total nonfarm employment series.[15] We refer to this as the *Reduced* specification. In this specification, we retain $J = 4$ factors, as suggested by the scree plot in Annex Figure 8. State-level employment is central to the analysis because it provides a high-frequency regional indicator of business-cycle propagation and because oil shocks transmit through sectoral reallocation and adjustment frictions that operate

[14] See Bai and Wang (2015) for a discussion of identification restrictions for linear DFMs. A further option is to apply this restriction to the actual measurement equations instead of the approximation, i.e. restrict J measurement equations so that the conditional means are equal to J distinct factors instead of being nonlinear sum-of-tree functions. This is required for identification of the factors (Yalcin and Amemiya, 2001), and we include this option our code.

[15] Series are obtained from the Bureau of Labor Statistics and accessed via FRED. Variable names are reported in the Appendix. Each series is seasonally adjusted using standard X-13ARIMA-SEATS procedures and transformed into log differences to isolate cyclical dynamics.

unevenly across regions. Differences in industrial composition imply that common oil shocks generate heterogeneous employment responses across states, while costly reallocation of labor and capital can delay adjustment and amplify job destruction relative to job creation (Davis and Haltiwanger, 2001; Herrera and Karaki, 2015; Herrera et al., 2017). From this perspective, cross-state variation is not merely dispersion, but reflects how a common aggregate shock is transmitted across a system of regionally differentiated economies (Hamilton and Owyang, 2012).

All variables are observed at monthly frequency. The sample spans 1974:01 to 2019:12. The endpoint excludes the COVID-19 period, whose unprecedented disruptions to oil-market and macroeconomic dynamics could dominate the estimated nonlinear relationships, while maintaining comparability with much of the recent literature on oil supply news shocks.

Consistent with standard practice in the dynamic factor model literature (e.g., Stock and Watson, 2005; Bernanke et al., 2005), the series are transformed to induce stationarity prior to factor extraction; the transformations are reported in Appendix Table 3. The external instrument based on high-frequency surprises around OPEC announcements is only available from 1983 onward, so identification relies on the corresponding subsample, even though the underlying monthly dataset begins in 1974:01. The VAR dynamics are estimated with six monthly lags, which provides a parsimonious lag structure for the high-dimensional factor-augmented nonlinear specification while preserving the monthly dynamics needed to trace oil-shock transmission. Our sample period is also consistent with recent studies of oil supply news shocks and their macroeconomic effects (Känzig, 2021; Miescu et al., 2024; Caravello and Martinez-Bruera, 2024; Forni et al., 2025; Mumtaz et al., 2024; Mori and Peersman, 2024). MCMC diagnostics are included in the Online Appendix.

We enrich the information set with forward-looking financial variables to better capture the transmission of oil supply news shocks and to improve structural identification. Financial variables have long been incorporated into empirical models of oil shocks to study their transmission to asset prices and financial markets, with Kilian and Park (2009) providing an early example by augmenting a structural oil-market VAR with U.S. stock returns. More recently, Mori and Peersman (2024) show that financial variables also help address omitted-information concerns when identifying oil supply news shocks. We also incorporate state-level employment series to provide a monthly measure of regional business-cycle conditions. The measure most closely aligned with GDP at the state level is gross state product, but GSP is available only annually and with a substantial publication lag, making it unsuitable for tracing monthly business-cycle dynamics. Following the regional-business-cycle perspective in Hamilton and Owyang (2012), we therefore use

state employment to study whether aggregate responses mask heterogeneous regional adjustment across states with different industrial structures and exposure to energy markets.

We therefore augment the information set with forward-looking financial and uncertainty variables, including VXO, the one-year Treasury yield, the S&P 500, Jurado–Ludvigson–Ng uncertainty measures, and the MSCI index, while retaining the full set of state-level employment indicators discussed above. Altogether, the Baseline specification includes 62 variables; the complete list of series and transformations is reported in Table 3. We refer to this specification as the *Baseline* specification. Consistent with the larger information set, we retain $J = 7$ factors in this specification, in line with the scree plot evidence reported in Annex Figure 8. This richer information set mitigates omitted-information concerns and helps separate oil supply news from broader revisions in expectations, demand conditions, and financial factors.

Figure 1 compares the responses obtained under the Reduced and Baseline specifications. The blue line and shaded area correspond to the baseline specification, which includes financial variables, whereas the dashed red lines correspond to the reduced specification, which excludes financial variables. The Reduced specification delivers responses broadly consistent with the benchmark evidence: adverse oil supply news shocks raise oil prices, reduce real activity, and increase global oil inventories. However, as emphasized by Mori and Peersman (2024), the joint increase in oil prices and inventories may also reflect anticipated demand expansions, implying potential contamination from omitted information.

Including financial variables changes the transmission pattern in systematic ways. In the Reduced specification, inventory accumulation and consumer prices display larger and more persistent responses. Once financial information is included, these responses are attenuated, while real activity responds more strongly and more rapidly. In our setting, the inventory response is not only smaller than in the Reduced specification, but also becomes statistically insignificant in the Baseline specification. This feature should be distinguished from Mori and Peersman (2024), where the inclusion of financial variables changes the transmission pattern but does not imply the same inventory response in our FABART specification. In the Baseline specification, declines in industrial production materialize earlier, are larger in magnitude, and are more precisely estimated.

This pattern is consistent with the identification concerns emphasized by Mori and Peersman (2024). When financial variables are omitted, the identified shock partly reflects broader revisions in expectations, amplifying forward-looking responses and blurring real-side dynamics. Including financial information reduces this contamination and yields a

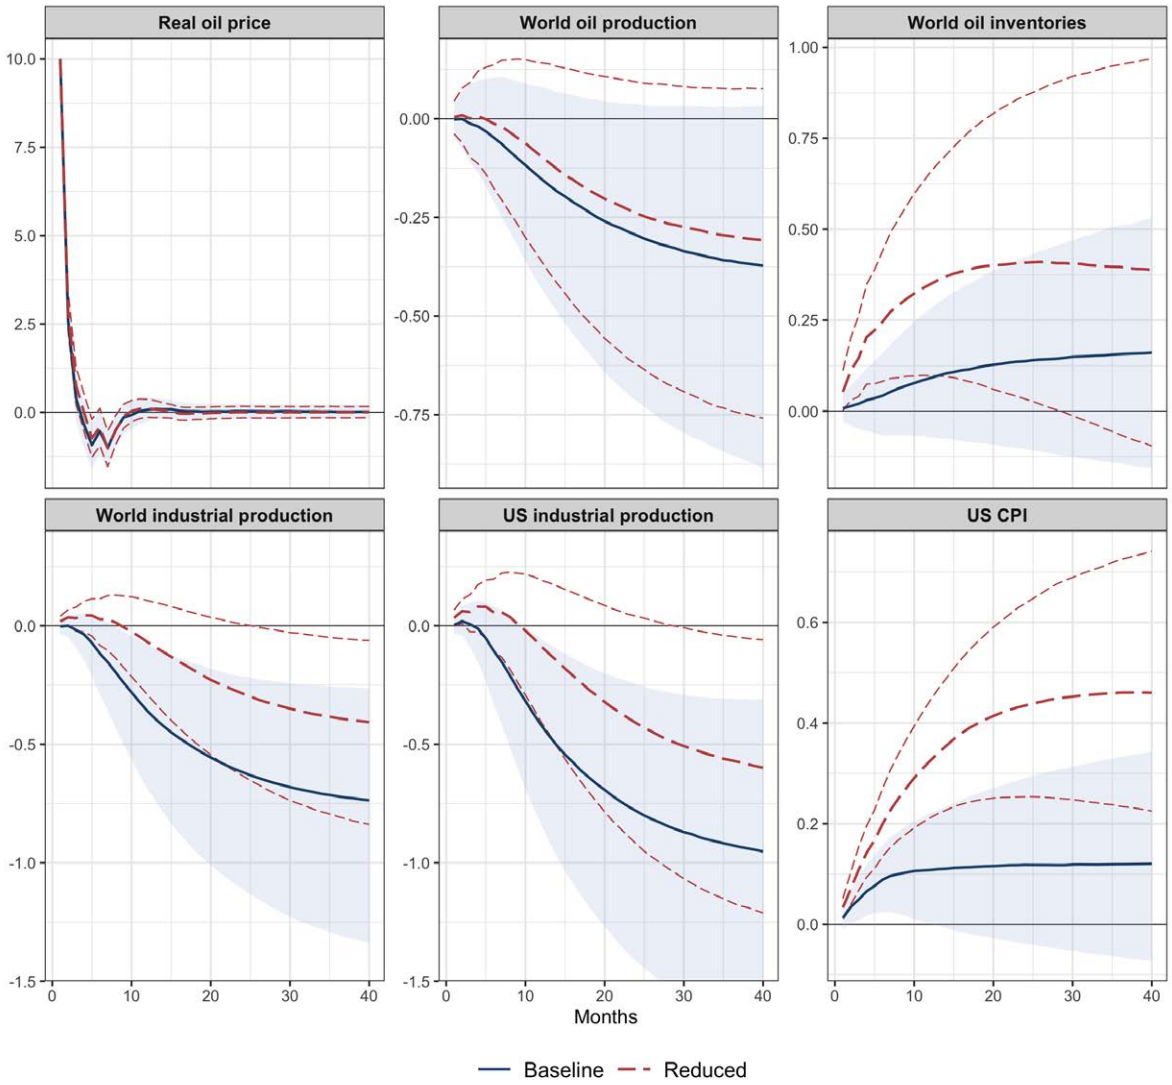

Note: GIRFs to an adverse oil supply news shock normalized to raise the real price of oil by 10% on impact. Solid lines correspond to the Baseline specification, which includes financial variables. Dashed lines correspond to the Reduced specification, which excludes financial variables. Shaded regions represent 68% credible intervals constructed from 16^{th} and 84^{th} percentiles.

Figure 1: Impact of Oil Supply News Shocks: The Role of Financial Variables.

sharper contraction in activity, consistent with an externally identified deterioration in expected oil supply.

3.2 Asymmetries in the transmission of oil supply news shocks

The previous results establish the Baseline specification as our preferred information set for studying oil supply news shocks. We now ask whether the transmission of these shocks is symmetric and proportional. These questions are central because the mechanisms emphasized in the oil-shock literature, including sectoral reallocation, adjustment frictions, uncertainty, and labor-market rigidities, may generate responses that differ across adverse and favorable shocks and across moderate and large disturbances. At the same time, evidence on asymmetry is sensitive to identification, specification, and the way impulse responses are constructed, as emphasized by Kilian and Vigfusson (2011a) and Kilian and Vigfusson (2017). We therefore study nonlinear transmission using externally identified oil supply news shocks within the FABART framework, allowing asymmetries and state dependence to emerge from the data rather than imposing them through prespecified transformations of oil prices or shocks.

We distinguish two dimensions of nonlinearity. The first is *sign asymmetry*: whether adverse oil supply news shocks that raise oil prices have effects that mirror those of favorable shocks that lower oil prices. The second is *size asymmetry*: whether larger shocks propagate proportionally relative to smaller shocks. This distinction follows Caravello and Martinez-Bruera (2024), who show that sign and size nonlinearities are distinct empirical objects and find stronger evidence for size nonlinearities than for sign-dependent responses. It also relates to Miescu et al. (2024), who document that large oil supply news shocks have stronger effects on real activity, labor-market variables, and risk measures, while prices and monetary policy display weaker size dependence. Our objective is to assess whether these margins appear separately or interact in the aggregate responses.

Sign asymmetry We begin by comparing positive and negative oil supply news shocks after normalizing responses to a common 10% oil-price impact. Positive shocks correspond to adverse oil supply news that raises oil prices, while negative shocks correspond to favorable oil supply news that lowers oil prices. This normalization allows us to isolate sign asymmetries while keeping the initial oil-price movement comparable across shocks. Throughout the analysis, we jointly study macroeconomic, financial, and uncertainty variables in order to assess whether nonlinear real effects are accompanied by asymmetric movements in risk and financial conditions. The size dimension is studied separately below.

Figure 2 extends the previous exercise by comparing adverse and favorable oil supply news shocks within the Baseline specification. The adverse shock corresponds to the Baseline response in Figure 1: it raises the real price of oil on impact, is associated with a gradual decline in world oil production, reduces industrial production, and increases consumer prices. The production response should be interpreted cautiously, however, since it is imprecisely estimated in the Baseline model and the credible interval includes zero over most horizons. The more robust features of the response are the sharp oil-price increase, the contraction in real activity, and the increase in consumer prices. This pattern broadly resembles the benchmark oil-supply-news evidence in Känzig (2021), but with the qualifications already discussed above. In particular, the inclusion of financial variables attenuates the response of world inventories relative to the Reduced specification, and inventories are not precisely estimated in the Baseline model. In the sign comparison, inventories also remain tightly concentrated around zero after favorable shocks. The relevant object in Figure 2 is therefore not a new validation of the adverse shock, but the comparison between the adverse response and the response to favorable oil supply news. This comparison shows that the Baseline oil supply news shock generates sizable real-side contractionary effects when oil prices rise, while the expansionary effects of comparable oil-price decreases are more limited.

The responses are not mirror images across signs. Adverse oil supply news shocks generate larger and more persistent contractions in real activity than the expansions associated with favorable shocks. This result speaks directly to the asymmetry debate initiated by Mork (1989), who extended Hamilton (1983) by showing that oil price increases had stronger predictive content for real activity than oil price decreases. Subsequent work has emphasized both the economic importance of nonlinear responses to large oil-price movements (Hamilton, 2003, 2011) and the importance of correctly specified nonlinear models and impulse-response-based inference when assessing asymmetries (Kilian and Vigfusson, 2011a,b). Our framework is designed to allow nonlinear impulse responses to emerge flexibly from the data within a structurally identified setting.

Price responses display a different form of asymmetry. CPI responses are initially similar across adverse and favorable oil supply news shocks, but they diverge over the medium horizon: the inflationary response following adverse shocks gradually levels off, whereas the sign-adjusted disinflationary response following favorable shocks continues to increase. This real–price pattern is related to Forni et al. (2025), who find that shocks raising oil prices generate stronger real effects, while shocks lowering oil prices produce relatively larger price responses. The underlying transmission mechanism, however, is different in our setting. In Forni et al. (2025), uncertainty rises after shocks of both signs and amplifies the contractionary effects of adverse shocks through a common uncertainty channel. By

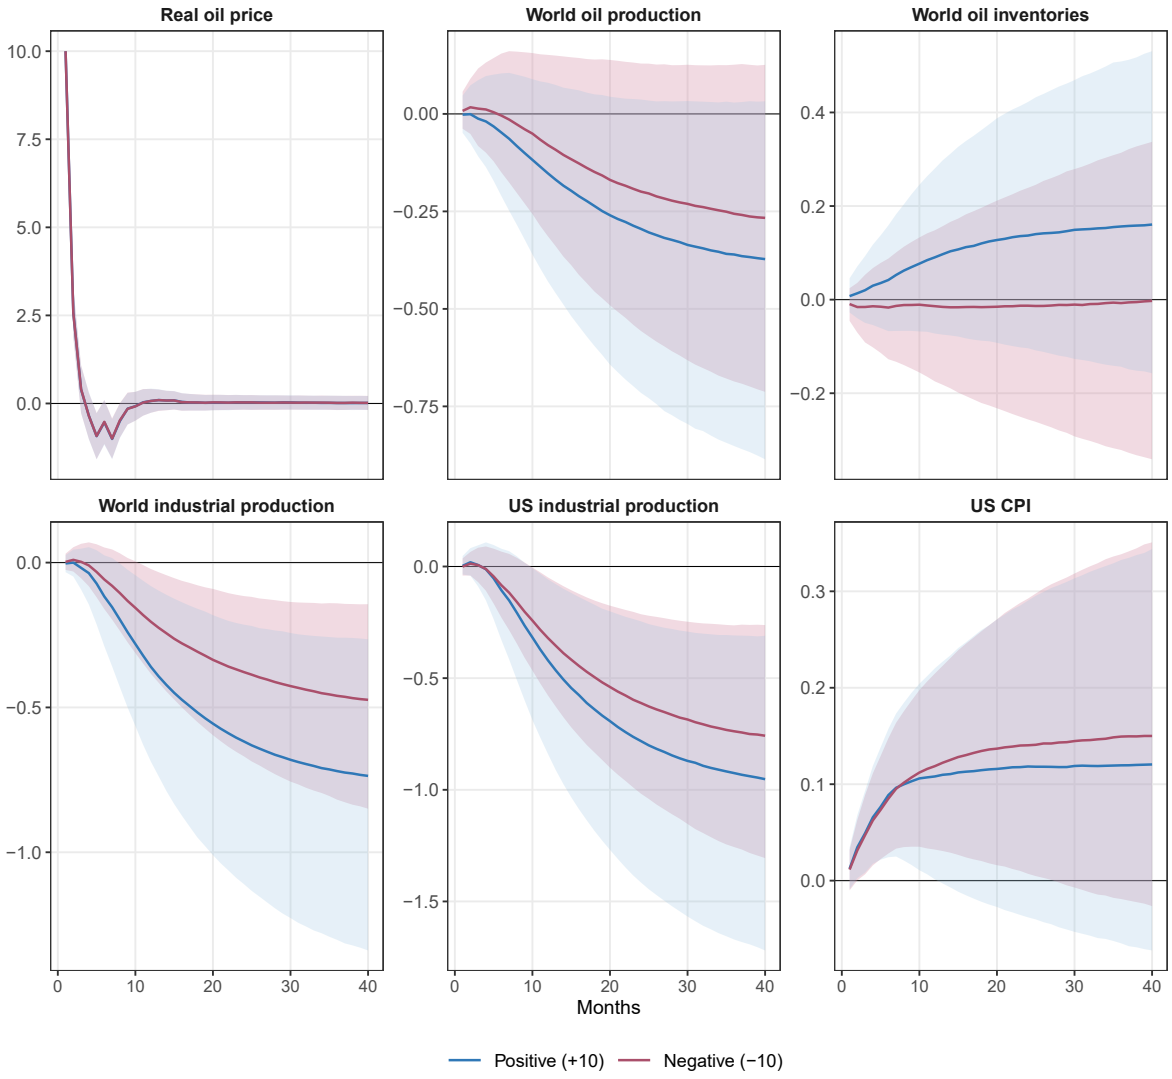

Note: GIRFs of oil supply news shocks normalized to move the real oil price by 10% on impact. Blue lines show adverse shocks that raise oil prices; red lines show favorable shocks that lower oil prices. Responses to favorable shocks are multiplied by minus one for visual comparability. Shaded areas represent 68% credible intervals constructed from 16^{th} and 84^{th} percentiles.

Figure 2: Sign asymmetry in the transmission of oil supply news shocks.

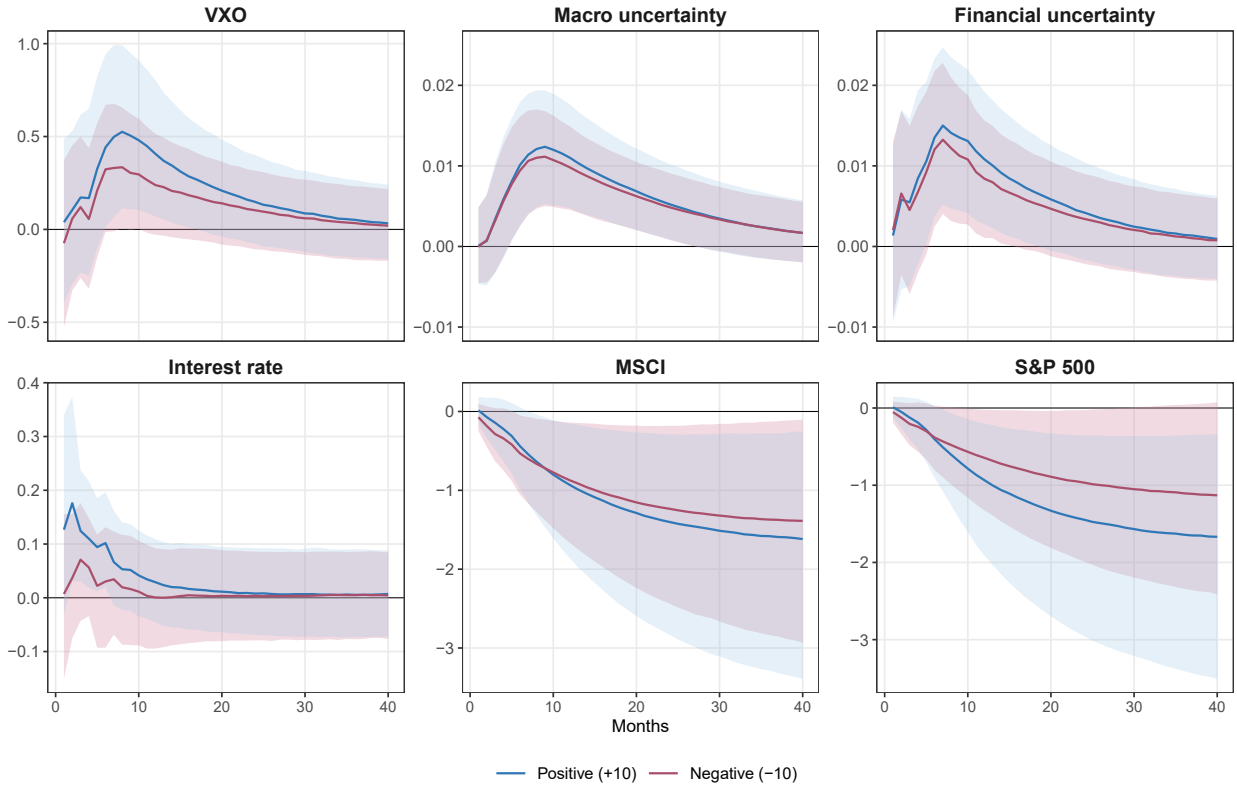

Note: GIRFs of financial and uncertainty variables following oil supply news shocks normalized to move the real oil price by 10% on impact. Blue lines show adverse shocks that raise oil prices; red lines show favorable shocks that lower oil prices. Responses to favorable shocks are multiplied by minus one for visual comparability. Shaded areas represent 68% credible intervals constructed from 16^{th} and 84^{th} percentiles.

Figure 3: Sign asymmetry in the transmission of oil supply news shocks to financial variables.

contrast, Figure 3 shows that, once the sign adjustment is undone, favorable shocks are associated with declines in VXO, macro uncertainty, and financial uncertainty, together with improving financial conditions and equity-market responses. Thus, our results do not point to an increase in uncertainty after favorable shocks, but rather to an easing of financial conditions.

One possible interpretation of the medium-run inflation asymmetry is that firms adjust pass-through and markups differently following oil-price increases and decreases. In particular, firms may absorb part of the input-cost change through margins rather than passing input-cost movements one-for-one into consumer prices. This interpretation is related to recent evidence showing that markups shape the transmission of cost-push shocks to inflation (Kharroubi et al., 2023). More broadly, it is also consistent with models of imperfect competition in which oil-related cost shocks generate endogenous adjustments in markups and pass-through (Rotemberg and Woodford, 1991; Blanchard and Riggi, 2013). At the same time, these mechanisms should be viewed as suggestive interpretations of the reduced-form evidence rather than as directly identified channels within the present framework.

Short-term interest rates also respond asymmetrically across shock signs. Following adverse oil supply news shocks, the one-year Treasury yield increases on impact before gradually declining back toward baseline. This response is consistent with higher expected short rates or compensation for inflation risk after energy-driven price pressures. It is also consistent with the idea that expected short rates may react when oil-price shocks feed into consumer prices (Bernanke et al., 1997; Neri, 2025). At the same time, the magnitude and interpretation of this response depend on the underlying source of oil-price movements and on model specification (Hamilton and Herrera, 2004; Kilian and Lewis, 2011). While favorable shocks generate declines in yields, these responses are considerably smaller and less persistent in our estimates. Since our specification includes the one-year Treasury yield rather than the federal funds rate, these responses should be interpreted more broadly as capturing movements in expected short rates, inflation compensation, and financial conditions, rather than as a direct estimate of the monetary-policy reaction function.

Size asymmetry The baseline sign comparison considers shocks that generate a 10% movement in the real price of oil. This normalization follows the benchmark calibration in Känzig (2021) and makes the results directly comparable to the asymmetry analysis of Forni et al. (2025). A shock of this magnitude is nevertheless economically meaningful. Over the 1974:01–2019:12 sample, a 10% absolute monthly oil-price movement corresponds to the 84^{th} percentile of the historical distribution of absolute monthly oil-price

changes.

To examine whether transmission also depends on shock size, we additionally consider smaller realizations of the same externally identified oil supply news shock in the framework of Känzig (2021). In particular, we study shocks associated with oil-price movements of 1.5%, 5%, and 10%. These correspond to the 26^{th}, 57^{th}, and 84^{th} percentiles, respectively, of the historical distribution of absolute monthly oil-price movements over the 1974:01–2019:12 sample.

Importantly, the exercise does not rely on identifying different structural shocks or imposing separate nonlinear regimes across shock sizes. Instead, we consider different realizations of the same identified oil supply news shock and examine whether the nonlinear transmission mechanism embedded in the FABART framework generates responses that vary with the magnitude of the underlying disturbance.

Figure 4 compares responses associated with 1.5%, 5%, and 10% oil-price movements, with all responses rescaled to a common 10% oil-price movement for comparability. Under linear proportional transmission, the normalized responses would coincide. Instead, the normalized responses differ across shock magnitudes. The 1.5% calibration generally produces weak responses, whereas the differences between the 5% and 10% calibrations vary across variables. For some variables, responses are already largely saturated at the 5% calibration, while for others the 10% calibration generates additional amplification. Appendix Figure 9 reports the corresponding comparison including the 20% calibration, which is best interpreted as an extreme upper-tail scenario rather than part of the main empirical design.

For real activity, very small shocks generate weak and generally insignificant responses, whereas shocks associated with a 5% oil-price movement already produce persistent declines in both world and U.S. industrial production. By contrast, increasing the oil-price movement from 5% to 10% leads to comparatively little additional amplification in the responses, suggesting that moderate oil-price movements are sufficient to generate economically meaningful real-side effects. Following adverse shocks, normalized industrial production responses for the 5% and 10% calibrations remain broadly similar. After favorable shocks, larger oil-price declines do not generate proportionally larger expansions, and the normalized responses in world industrial production and U.S. industrial production become proportionally smaller as the underlying shock becomes larger. Overall, the evidence suggests that the main real-side nonlinearity arises between very small and moderate shocks rather than through monotonic amplification as shock size increases.

CPI responses display a different pattern. The 1.5% shocks generate weak and generally

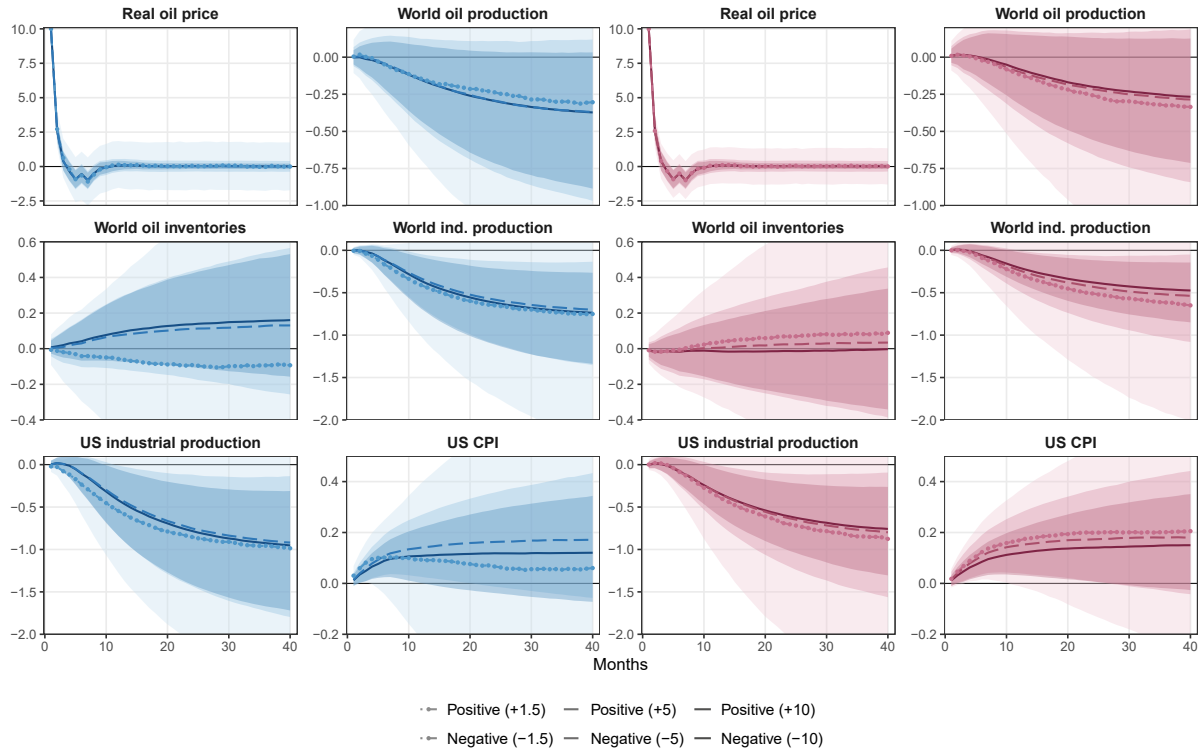

Note: GIRFs to adverse shocks that raise oil prices (blue) and favorable shocks that lower oil prices (red). The underlying shocks are associated with 1.5%, 5%, and 10% oil-price movements, and all responses are rescaled to a common 10% oil-price movement for comparability. Responses to favorable shocks are multiplied by minus one. Shaded areas denote 16–84% credible intervals.

Figure 4: Nonlinear responses to oil supply news shocks: sign and size effects.

insignificant inflation responses, making it difficult to draw strong conclusions about pass-through at very small magnitudes. Once shocks reach the 5% range, however, inflation responses become stronger and more persistent. As shown more clearly in Appendix Figure 9, the CPI responses associated with the 5% shocks remain statistically significant for more than two years, exceeding the persistence observed for both the smaller and the larger shock calibrations. By contrast, the normalized inflation response for the 10% shock is weaker relative to the underlying oil-price movement, pointing to less-than-proportional pass-through at larger magnitudes.

One possible interpretation is that moderate oil-price movements correspond more closely to the type of recurrent cost shocks that firms adjust to more directly in practice, whereas very large shocks may trigger broader general-equilibrium adjustments, precautionary behavior, or offsetting demand effects that partly dampen pass-through into consumer prices. This interpretation is consistent with recent evidence emphasizing that the transmission of energy-cost shocks varies with the size and persistence of the disturbance (Bobeica et al., 2025; Gautier et al., 2026). At the same time, our evidence should be interpreted as reduced-form evidence on nonlinear pass-through rather than as direct evidence on firms' pricing decisions.

Taken together, the evidence suggests that the relationship between shock size and

macroeconomic transmission is nonlinear but not monotonic. These findings reinforce the distinction between sign and size nonlinearities emphasized by Caravello and Martinez-Bruera (2024), while avoiding a specification that mechanically attributes nonlinear responses to either dimension *ex ante*. The main departures from proportional transmission arise between very small and moderate shocks rather than through monotonic amplification as shock size increases. Larger shocks matter, but their effects do not scale proportionally across variables or shock signs.

The financial responses also display size dependence, although the evidence is less clear-cut than for real activity and prices. The 1.5% shocks are estimated imprecisely, with credible intervals that are wide and generally include zero, suggesting that the effects of very small oil supply news shocks cannot be estimated precisely. For the 5% calibration, uncertainty measures display a pattern broadly similar to the 10% calibration. After adverse shocks, the volatility index increases and remains statistically significant for about one year, while macroeconomic and financial uncertainty also rise. After favorable shocks, the sign-adjusted uncertainty responses are weaker and often imprecisely estimated; in particular, the response of the volatility index is not statistically significant. Overall, the uncertainty responses suggest that adverse oil supply news shocks are more systematically associated with deteriorating financial and risk conditions than favorable shocks are with improvements in those conditions, although the evidence for favorable shocks remains substantially less precise. Appendix Figure 10 reports the corresponding financial responses including the 20% calibration.

Among the financial variables, equity prices display the clearest evidence of size-dependent transmission. After adverse shocks, the decline in equity prices is stronger for the 10% calibration than for the 5% calibration, suggesting that larger adverse oil supply news shocks transmit more strongly to financial markets. This pattern is consistent with the broader literature documenting asymmetric stock-market responses to oil shocks and emphasizing that adverse oil-price movements are associated with stronger financial and macroeconomic effects than favorable shocks (Kilian and Lewis, 2011; Bjornland et al., 2026). It is also consistent with the view that adverse oil-price movements propagate more strongly through financial and uncertainty channels. By contrast, after favorable shocks the equity responses are statistically significant only at very short horizons. Although some panels suggest a reversed ordering across shock sizes, the rapid widening of credible intervals implies that this pattern should not be interpreted too strongly. Overall, the financial evidence points to size-dependent transmission mainly for adverse shocks, while the evidence for favorable shocks remains weaker and largely confined to the short run.

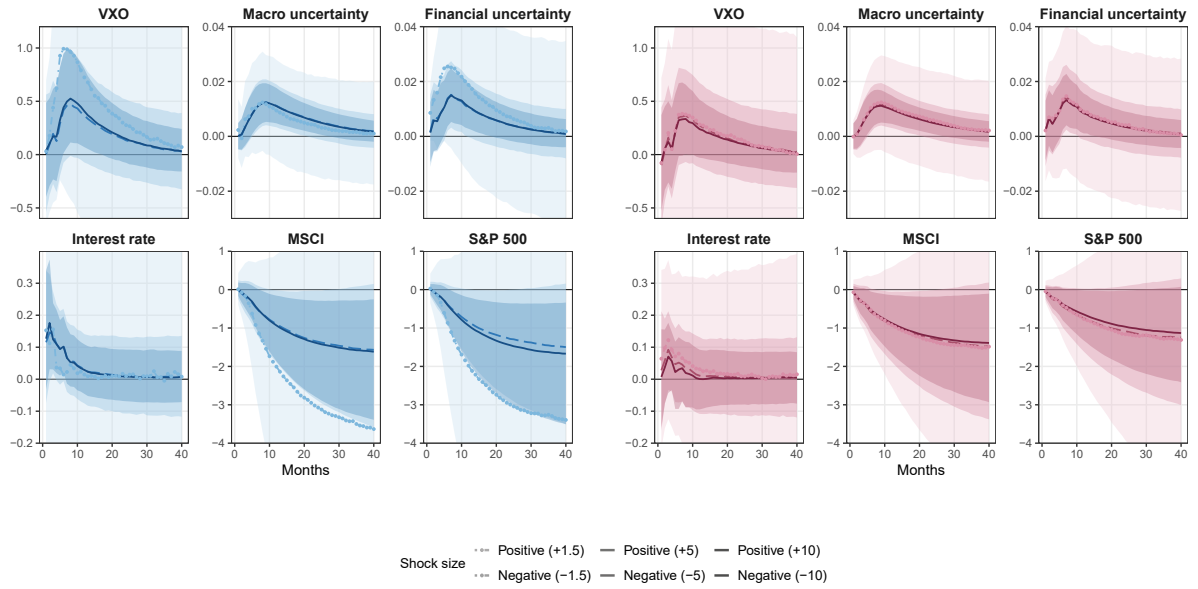

Note: GIRFs of financial and uncertainty variables to adverse shocks that raise oil prices (blue) and favorable shocks that lower oil prices (red). The underlying shocks are associated with 1.5%, 5%, and 10% oil-price movements, and all responses are rescaled to a common 10% oil-price movement for comparability. Responses to favorable shocks are multiplied by minus one. Shaded areas denote 16–84% credible intervals.

Figure 5: Nonlinear financial and uncertainty responses to oil supply news shocks: sign and size effects.

3.3 U.S. State-level information and regional transmission

We next turn to state-level employment responses and include regional labor-market variables in the information set because labor markets are one of the main channels through which oil supply news shocks propagate to the broader economy. Oil supply news shocks can affect employment through sectoral reallocation, input-cost exposure, changes in demand for durable and energy-intensive goods, and local uncertainty effects, channels that are likely to vary substantially across states with different industrial structures and exposure to the energy sector (Davis and Haltiwanger, 2001; Hamilton and Owyang, 2012). From the perspective of the asymmetry debate, regional labor markets are particularly informative because labor adjustment, sectoral reallocation, and job destruction need not operate symmetrically after adverse and favorable shocks. In particular, adverse oil-price movements may trigger contractions in energy-intensive sectors and durable manufacturing, whereas favorable shocks need not generate equally strong employment recoveries. The aim is therefore not to identify separate regional shocks, but to study how a common oil supply news shock is absorbed and propagated across heterogeneous state labor markets.

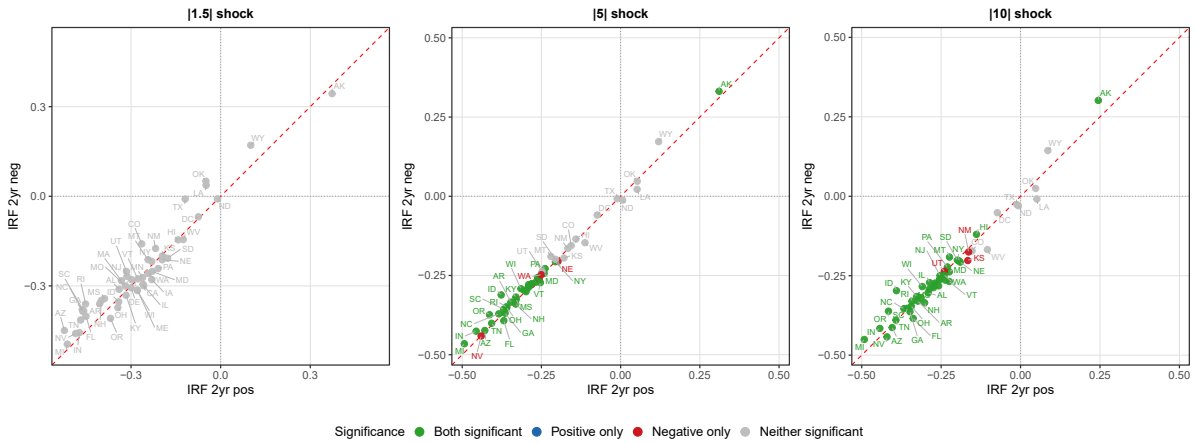

Note: This figure compares median two-year cumulative employment responses to adverse and favorable oil supply news shocks across U.S. states for shock sizes of 1.5%, 5%, and 10%. The x-axis reports responses to adverse shocks that raise oil prices. The y-axis reports responses to favorable shocks that lower oil prices, multiplied by minus one for visual comparability. The red dashed line indicates equality between the adverse response and the sign-adjusted favorable response. Dots are colored according to whether the response is statistically significant for both shocks, only for the adverse shock, only for the favorable shock, or for neither shock.

Figure 6: Cross-state heterogeneity in employment responses to oil supply news shocks

Following Hamilton and Owyang (2012), we use state employment as a high-frequency indicator of regional business-cycle conditions. This choice is partly practical. Gross state product is the closest state-level analogue of GDP, but it is available only annually and with a substantial publication lag. Monthly employment therefore provides a more suitable measure for tracing regional propagation.

Figure 6 provides a first characterization of the cross-state transmission of oil supply news shocks by comparing two-year cumulative employment responses across states for the 1.5%, 5%, and 10% oil-price calibrations. The figure reveals substantial heterogeneity in how state labor markets absorb oil supply news shocks, with the strongest contractions concentrated in manufacturing-intensive states and comparatively weaker responses in more service-oriented economies. The main difference across shock sizes is not simply the magnitude of the median response, but whether employment responses become statistically distinguishable from zero across states. For the 1.5% calibration, most observations cluster tightly around the 45-degree line and only a limited number of states display statistically significant responses. This suggests that relatively small oil-price movements generate weak and imprecisely estimated regional employment effects. Once the shock size reaches 5%, however, a substantially larger set of states exhibits significant employment responses and the cross-state dispersion becomes more pronounced. The 10% calibration further increases the dispersion of responses, although the evidence points more strongly to heterogeneous regional transmission than to a uniform amplification affecting all states proportionally. Overall, the figure suggests that the main nonlinearity arises between very small and moderate oil-price movements: moderate shocks already generate economically meaningful regional employment adjustments, while the additional amplification from 5% to 10% is comparatively limited and highly state de-

pendent. The figure also highlights important heterogeneity across types of states. Most states display larger employment contractions following adverse oil supply news shocks than employment gains following favorable shocks, although the main message of the figure is cross-state heterogeneity rather than a uniform sign asymmetry. This pattern is particularly visible for manufacturing-oriented states located in the lower-left quadrant, where adverse shocks generate persistent employment losses. The transmission is therefore strongest precisely in sectors where labor adjustment costs, durable-goods demand, and sectoral reallocation frictions are likely to be most relevant, consistent with the mechanisms emphasized by Davis and Haltiwanger (2001). At the same time, a small number of oil-producing states, most notably Alaska, display positive employment responses following adverse shocks that raise oil prices. These patterns are consistent with the regional propagation mechanisms emphasized by Hamilton and Owyang (2012) and Davis and Haltiwanger (2001): states differ substantially in industrial composition, energy-sector exposure, and labor-market structure, so a common oil supply news shock need not propagate uniformly across regions.

Motivated by this evidence, Figure 7 reports cumulative employment GIRFs for a set of representative states selected to illustrate these transmission channels more directly. Following the regional propagation framework of Hamilton and Owyang (2012), the selected states capture distinct forms of exposure to oil supply news shocks and different industrial structures. Wyoming and North Dakota represent energy-producing states with substantial exposure to the oil sector. Michigan and Ohio capture manufacturing-intensive states exposed to energy-sensitive durable production, labor reallocation, and adjustment-cost effects. California and New York instead represent more service-oriented economies with lower direct dependence on energy-intensive production. The figure also reinforces the message from the scatter plots: the 1.5% shocks are generally imprecise, while employment responses become more clearly distinguishable from zero starting around the 5% calibration.

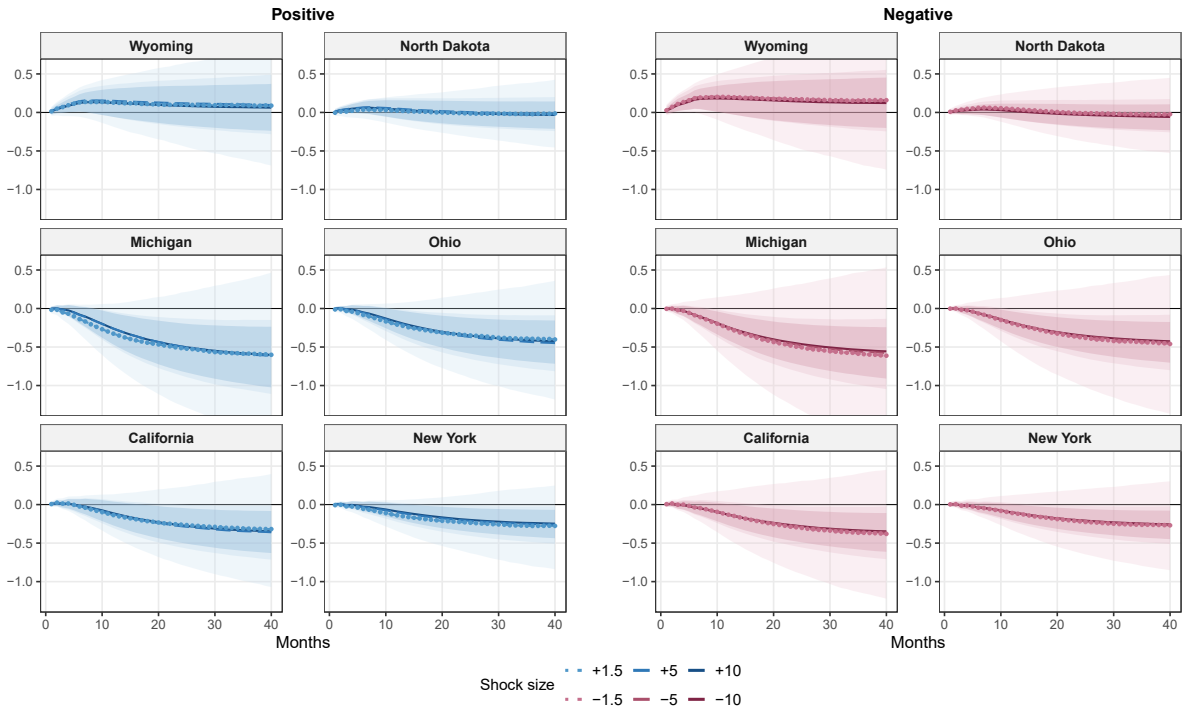

Note: Cumulative state-level employment GIRFs to oil supply news shocks calibrated to move the real oil price by 1.5%, 5%, and 10% on impact, for representative states. Blue lines show adverse shocks that raise oil prices; red lines show favorable shocks that lower oil prices. Responses to favorable shocks are multiplied by minus one for visual comparability. Shaded areas represent 16–84% credible intervals.

Figure 7: Sign asymmetry in state-level employment responses to oil supply news shocks.

The state-level GIRFs reinforce the evidence of heterogeneous transmission. Energy-producing states such as Wyoming display positive employment responses following adverse oil supply news shocks, in contrast to the contractionary responses observed in most other states. This pattern is consistent with the offsetting role of local energy-sector activity after oil-price increases. At the same time, the persistence and statistical significance of these responses differ across states. Wyoming displays positive and statistically significant employment responses for roughly the first year following adverse shocks and somewhat more persistent sign-adjusted responses after favorable shocks that lower oil prices. North Dakota exhibits a qualitatively similar pattern, although its responses are considerably less precisely estimated, especially following adverse shocks. More generally, the responses associated with the 1.5% calibration are rarely statistically distinguishable from zero, while the 5% and 10% calibrations produce substantially clearer regional employment effects.

Manufacturing-intensive states display the opposite pattern. Michigan and Ohio exhibit some of the strongest contractionary employment responses following adverse oil supply news shocks, while the corresponding employment gains after favorable shocks are more limited. These responses are consistent with the sectoral reallocation and adjustment-cost mechanisms emphasized by Davis and Haltiwanger (2001), whereby higher energy costs

disproportionately affect energy-intensive production and durable manufacturing sectors. The evidence is also consistent with the regional propagation mechanisms emphasized by Hamilton and Owyang (2012), where the effects of oil shocks differ systematically according to local industrial composition and energy dependence. While service-oriented states such as California and New York also display negative employment responses following adverse oil supply news shocks, these responses are comparatively weaker and less persistent than in manufacturing-intensive states. This pattern is consistent with their lower direct exposure to energy-intensive production and durable-goods sectors, where higher energy costs and sectoral reallocation effects are likely to be more pronounced. Across representative states, there is little evidence that size dependence differs systematically across shock signs. Instead, larger calibrations mainly make the employment responses more clearly distinguishable from zero and reveal stronger cross-state heterogeneity, rather than generating a uniform amplification pattern across all states.

Taken together, Figures 6 and 7 show that the aggregate asymmetries documented above mask substantial regional differences in labor-market adjustment. The transmission of oil supply news shocks depends importantly on state industrial composition and exposure to the energy sector, rather than following a uniform national pattern. More broadly, the state-level evidence suggests that labor markets are not simply an additional margin of adjustment, but one of the main channels through which nonlinear oil-shock transmission materializes in the real economy. The results are therefore closely aligned with the view that oil shocks propagate partly through costly labor reallocation and heterogeneous regional adjustment dynamics.

Because the scatter plots reveal substantial cross-state heterogeneity, we next analyze adverse and favorable shocks separately using cross-sectional regressions. Specifically, we examine whether variation in industrial composition, particularly the relative importance of manufacturing and mining sectors, can account for part of the observed cross-sectional heterogeneity, following the approach of Mumtaz et al. (2018) and related evidence on heterogeneous regional responses to oil shocks (Hamilton and Owyang, 2012; Herrera et al., 2017).

The cross-sectional analysis focuses on three predetermined state characteristics. Manufacturing and construction are state-level GDP industry shares constructed as in Mumtaz et al. (2018). Manufacturing captures exposure to energy-intensive production and adjustment costs, while construction captures a locally cyclical sector linked to regional demand and housing conditions. Oil & gas is constructed following Mumtaz and Sunder-Plassmann (2021) and measures the share of state value added accounted for by oil and gas extraction. These variables are not intended to provide a complete model of state-level transmission. They provide a parsimonious way to summarize the main regional margins

emphasized by the literature: production structure, direct energy-sector exposure, and local cyclical sensitivity.

Table 2: State Characteristics and Employment Responses to Oil Shocks: First-Stage Posterior Uncertainty

	1.5%		5%		10%	
	+	−	+	−	+	−
Manufacturing	-1.084	1.029	-1.029	1.006	-0.984	0.962
2.5^{th}percentile	−3.526	−1.702	−2.330	−0.152	−2.030	0.052
97.5^{th}percentile	1.269	3.762	0.195	2.244	−0.028	1.930
Oil & gas	2.329	-2.513	2.457	-2.390	2.246	-2.177
2.5^{th}percentile	−3.880	−8.460	−0.417	−5.480	−0.041	−4.546
97.5^{th}percentile	8.712	3.323	5.630	0.306	4.700	−0.027
Construction	-5.065	4.120	-4.240	4.057	-4.286	4.001
2.5^{th}percentile	−24.153	−16.376	−13.934	−4.475	−12.704	−2.461
97.5^{th}percentile	13.042	25.115	4.795	13.221	3.215	11.274
Median adjusted R^2	0.295	0.345	0.496	0.484	0.513	0.500

Dependent variable: posterior draws of the two-year cumulative state employment response. Intervals are quantiles of draw-specific cross-sectional regression coefficients and account for first-stage posterior uncertainty in the estimated impulse responses. Point estimates are means of draw-specific cross-sectional regression coefficients. Regressions are across 50 U.S. states. Negative-shock responses are multiplied by minus one. Interval probability: 95%.

The results in Table 2 point to systematic regional heterogeneity in the labor-market transmission of oil supply news shocks. The dependent variable in these regressions is the estimated state-level employment response obtained from the impulse response analysis. Consequently, the exercise is best viewed as a descriptive summary of the posterior samples that shows which state characteristics are associated with stronger or weaker employment responses, rather than as a causal model of the impact of state characteristics on employment responses. The intervals account for uncertainty from the posterior draws of the dependent variable.[16] Although a large proportion of intervals include zero, a smaller proportion of 68% intervals include zero (see Online Appendix).

Several patterns emerge. States with larger manufacturing sectors exhibit stronger employment responses to oil supply news shocks, consistent with the sectoral reallocation and energy-cost channels emphasized by Davis and Haltiwanger (2001) and Hamilton and Owyang (2012). By contrast, states with larger oil and gas sectors display more muted employment responses, suggesting that local energy production partly insulates regional labor markets from oil-market disturbances. Construction intensity is also associated with stronger employment responses, indicating that cyclical local industries play an important role in the regional propagation of oil supply news shocks.

[16]We do not account for second-stage uncertainty, that is, uncertainty associated with the error term in the regression of a particular posterior sample on state characteristics. These descriptive regressions are not assumed to be true models of impulse responses, and do not provide inference for some true coefficient parameters pertaining to a larger superpopulation of states. However, we include intervals that account for second stage uncertainty in the online Supplemental Appendix.

Importantly, these relationships are broadly similar for adverse and favorable shocks. This finding is consistent with Figure 6, where most states lie close to the 45-degree line, particularly for the 5% and 10% calibrations. Sectoral composition therefore appears to be an important correlate of the magnitude of state-level employment responses, but provides limited evidence that regional asymmetries are driven by different transmission mechanisms across shock signs.

Overall, the state-level evidence highlights the importance of labor-market adjustment and sectoral exposure in shaping the regional transmission of oil supply news shocks. More broadly, the results suggest that cross-state heterogeneity is a key feature of oil-shock transmission, with differences in industrial structure helping explain why employment responses vary substantially across states.

4 Conclusion

This article studies how oil supply news shocks affect real activity, financial conditions, and regional labor markets, and whether these effects depend on the sign and magnitude of the shock. To address this question, we combine external identification based on the oil supply news instrument of Känzig (2021) with a factor-augmented nonlinear framework that incorporates a large macro-financial information set and allows nonlinear transmission to emerge flexibly from the data. The proposed Factor Bayesian Additive Regression Tree (FABART) model combines latent-factor dynamics with a nonparametric measurement equation, allowing nonlinear transmission to emerge from the data rather than being imposed through a particular functional specification.

The results provide evidence of important nonlinearities in the transmission of oil supply news shocks. Shocks that raise oil prices generate persistent declines in world and U.S. industrial production together with increases in inflation, financial stress, and uncertainty measures, whereas shocks that lower oil prices generate comparatively weaker expansionary effects. The evidence on shock size is more nuanced. Very small shocks produce weak and often statistically insignificant responses across most variables, while moderate oil-price movements already generate economically meaningful effects on industrial production, financial conditions, and regional employment. At the same time, the responses do not scale proportionally with shock size across all variables and shock signs, indicating that nonlinear transmission is not limited to extreme oil-market disruptions.

The state-level results further reveal substantial regional heterogeneity in labor-market adjustment. Manufacturing-intensive states experience considerably stronger employment contractions following adverse oil supply news shocks, while energy-producing states

display substantially weaker contractions and, in some cases, positive employment responses. More broadly, the results suggest that industrial composition and exposure to energy markets play an important role in shaping the regional propagation of oil supply news shocks.

More generally, the results highlight that conclusions about nonlinear oil-shock transmission depend critically on both credible structural identification and the way nonlinear dynamics are modeled. Low-dimensional nonlinear oil-market models may confound omitted macro-financial information and restrictive functional-form assumptions with genuine economic asymmetries. Embedding externally identified oil supply news shocks within a large-information nonlinear framework helps distinguish robust transmission patterns from artifacts of informational insufficiency or restrictive functional-form assumptions. More broadly, the results suggest that part of the disagreement in the literature may reflect differences in identification strategies, information sets, and nonlinear specifications rather than fundamentally different economic mechanisms. In this sense, the FABART framework provides a flexible setting for revisiting the asymmetry debate while allowing nonlinear transmission to emerge from the data under credible structural identification.

References

Amemiya, Yasuo, and Ilker Yalcin. (2001). "Nonlinear factor analysis as a statistical method". *Statistical Science*, 16(3), pp. 275-294. https://doi.org/10.1214/ss/1009213729

Andreini, Paolo, Cosimo Izzo and Giovanni Ricco. (2020). "Deep dynamic factor models". *arXiv preprint*, 2007.11887. https://arxiv.org/abs/2007.11887

Bai, Jushan, and Peng Wang. (2015). "Identification and Bayesian estimation of dynamic factor models". *Journal of Business & Economic Statistics*, 33(2), pp. 221-240. https://doi.org/10.1080/07350015.2014.941467

Baumeister, Christiane, Philip Liu and Haroon Mumtaz. (2013). "Changes in the effects of monetary policy on disaggregate price dynamics". *Journal of Economic Dynamics and Control*, 37(3), pp. 543-560. https://doi.org/10.1016/j.jedc.2012.09.009

Bernanke, Ben S., Jean Boivin and Piotr Eliasz. (2005). "Measuring the effects of monetary policy: a factor-augmented vector autoregressive (FAVAR) approach". *The Quarterly Journal of Economics*, 120(1), pp. 387-422. https://doi.org/10.1162/0033553053327452

Bernanke, Ben S., Mark Gertler and Mark W. Watson. (1997). "Systematic monetary policy and the effects of oil price shocks". *Brookings Papers on Economic Activity*, 1997(1), pp. 91-157. https://doi.org/10.2307/2534702

Bjørnland, Hilde C., Yoosoon Chang and Jamie L. Cross. (2026). "Oil and the stock market revisited: A mixed functional VAR approach". *Quantitative Economics*. Forthcoming.

Blanchard, Olivier J, and Marianna Riggi. (2013). "Why are the 2000s so different from the 1970s? A structural interpretation of changes in the macroeconomic effects of oil prices". *Journal of the European Economic Association*, 11(5), pp. 1032-1052. https://doi.org/10.1111/jeea.12029

Bobeica, Elena, Sarah Holton, Florian Huber and Catalina Martínez Hernández. (2025). "Beware of large shocks! A non-parametric structural inflation model". Working Paper Series, 3052, European Central Bank. https://doi.org/10.2866/4689671

Caldara, Dario, and Edward Herbst. (2019). "Monetary policy, real activity, and credit spreads: Evidence from Bayesian proxy SVARs". *American Economic Journal: Macroeconomics*, 11(1), pp. 157-192. https://doi.org/10.1257/mac.20170294

Caravello, Tomás, and Pedro Martinez-Bruera. (2024). *Disentangling sign and size non-linearities*. SSRN Working Paper, 4704050. https://doi.org/10.2139/ssrn.4704050

Carter, Chris K., and Robert Kohn. (1994). "On Gibbs sampling for state space models". *Biometrika*, 81(3), pp. 541-553. https://doi.org/10.1093/biomet/81.3.541

Charnavoki, Valery, and Juan J. Dolado. (2014). "The effects of global shocks on small commodity-exporting economies: Lessons from Canada". *American Economic Journal: Macroeconomics*, 6(2), pp. 207-237. https://doi.org/10.1257/mac.6.2.207

Chen, Liang, Juan J. Dolado and Jesús Gonzalo. (2021). "Quantile factor models". *Econometrica*, 89(2), pp. 875-910. https://doi.org/10.3982/ecta15746

Chen, Liang, Juan J. Dolado, Jesús Gonzalo and Haozi Pan. (2026). "Estimation of characteristics-based quantile factor models". *Journal of Econometrics*, 256-A(106267). https://doi.org/10.1016/j.jeconom.2026.106267

Chernis, Tony, Niko Hauzenberger, Haroon Mumtaz and Michael Pfarrhofer. (2025). "A Bayesian Gaussian process dynamic factor model". *arXiv preprint*, 2509.04928. https://arxiv.org/abs/2509.04928

Chipman, Hugh A., Edward I. George and Robert E. McCulloch. (1998). "Bayesian CART model search". *Journal of the American Statistical Association*, 93(443), pp. 935-948. https://doi.org/10.1080/01621459.1998.10473750

Chipman, Hugh A., Edward I. George and Robert E. McCulloch. (2010). "BART: Bayesian additive regression trees". *The Annals of Applied Statistics*, 4(1), pp. 266-298. https://doi.org/10.1214/09-aoas285

Clark, Todd, Florian Huber and Gary Koop. (2025). "A flexible approach to augmenting a Bayesian VAR with nonlinear factors". *arXiv preprint*, 2508.13972. https://arxiv.org/abs/2508.13972

Clark, Todd E., Florian Huber, Gary Koop, Massimiliano Marcellino and Michael Pfarrhofer. (2024). "Investigating growth-at-risk using a multicountry nonparametric quantile factor model". *Journal of Business & Economic Statistics*, 42(4), pp. 1302-1317. https://doi.org/10.1080/07350015.2024.2310020

Crawford, Lorin, Seth R. Flaxman, Daniel E. Runcie and Mike West. (2019). "Variable prioritization in nonlinear black box methods: A genetic association case study". *TheAnnals of Applied Statistics*, 13(2), pp. 958-989. https://doi.org/10.1214/18-aoas1222

David, Joel M., Raffaella Giacomini, Xiyu Jiao and Weining Wang. (2026). "Causal state-dependent local projections". *arXiv preprint*, 2605.05404. https://arxiv.org/abs/2605.05404

Davis, Steven J., and John Haltiwanger. (2001). "Sectoral job creation and destruction responses to oil price changes". *Journal of Monetary Economics*, 48(3), pp. 465-512. https://doi.org/10.1016/s0304-3932(01)00086-1

Degasperi, Riccardo. (2026). "Identification of expectational shocks in the oil market using OPEC announcements". Temi di Discussione (Working Papers), 1516, Banca d'Italia. https://www.bancaditalia.it/pubblicazioni/temi-discussione/2026/2026-1516/index.html

Drossidis, Theo, Haroon Mumtaz and Angeliki Theophilopoulou. (2024). "The distributional effects of oil supply news shocks". *Economics Letters*, 240(111769). https://doi.org/10.1016/j.econlet.2024.111769

Durbin, James, and Siem Jan Koopman. (2002). "A simple and efficient simulation smoother for state space time series analysis". *Biometrika*, 89(3), pp. 603-616. https://doi.org/10.1093/biomet/89.3.603

Forni, Mario, Alessandro Franconi, Luca Gambetti and Luca Sala. (2025). "Asymmetric transmission of oil supply news". *Quantitative Economics*, 16(3), pp. 947-979. https://doi.org/10.3982/qe2548

Forni, Mario, and Luca Gambetti. (2010). "The dynamic effects of monetary policy: A structural factor model approach". *Journal of Monetary Economics*, 57(2), pp. 203-216. https://doi.org/10.1016/j.jmoneco.2009.11.009

Forni, Mario, and Luca Gambetti. (2014). "Sufficient information in structural VARs". *Journal of Monetary Economics*, 66, pp. 124-136. https://doi.org/10.1016/j.jmoneco.2014.04.005

Forni, Mario, Domenico Giannone, Marco Lippi and Lucrezia Reichlin. (2009). "Opening the black box: Structural factor models with large cross sections". *Econometric Theory*, 25(5), pp. 1319-1347. https://doi.org/10.1017/s026646660809052x

Gautier, Erwan, Cristina Conflitti, Daniel Enderle, Ludmila Fadejeva, Alex Grimaud, Eduardo Gutiérrez, Valentin Jouvanceau, Jan-Oliver Menz, Alari Paulus, Pavlos Petroulas, Pau Roldan-Blanco and Elisabeth Wieland. (2026). "Consumer price stickiness in the euro area during an inflation surge". Documentos de Trabajo, 2608, Banco de España. https://repositorio.bde.es/handle/123456789/42545

Gonçalves, Sílvia, Ana María Herrera, Lutz Kilian and Elena Pesavento. (2021). "Impulse response analysis for structural dynamic models with nonlinear regressors". *Journal of Econometrics*, 225(1), pp. 107-130. https://doi.org/10.1016/j.jeconom.2021.06.009

Gonçalves, Silvia, Ana Maria Herrera, Lutz Kilian, Elena Pesavento and Ione Kelanemer Holban. (2026). "Semiparametric local projections". *arXiv preprint*, 2606.13519. https://arxiv.org/abs/2606.13519

Hamilton, James D. (1983). "Oil and the macroeconomy since World War II". *Journal of Political Economy*, 91(2), pp. 228-248. https://doi.org/10.1086/261140

Hamilton, James D. (2003). "What is an oil shock?". *Journal of Econometrics*, 113(2), pp. 363-398. https://doi.org/10.1016/s0304-4076(02)00207-5

Hamilton, James D. (2011). "Nonlinearities and the macroeconomic effects of oil prices". *Macroeconomic Dynamics*, 15(S3), pp. 364-378. https://doi.org/10.1017/s1365100511000307

Hamilton, James D., and Ana María Herrera. (2004). "Oil shocks and aggregate macroeconomic behavior: the role of monetary policy: A comment". *Journal of Money, Credit and Banking*, 36(2), pp. 265-286. https://muse.jhu.edu/pub/30/article/53949

Hamilton, James D., and Michael T. Owyang. (2012). "The propagation of regional recessions". *The Review of Economics and Statistics*, 94(4), pp. 935-947. https://doi.org/10.1162/REST_a_00197

Hauzenberger, Niko, Florian Huber and Karin Klieber. (2023). "Real-time inflation forecasting using non-linear dimension reduction techniques". *International Journal of Forecasting*, 39(2), pp. 901-921. https://doi.org/10.1016/j.ijforecast.2022.03.002

Herrera, Ana María, and Mohamad B. Karaki. (2015). "The effects of oil price shocks on job reallocation". *Journal of Economic Dynamics and Control*, 61, pp. 95-113. https://doi.org/10.1016/j.jedc.2015.08.006

Herrera, Ana María, Mohamad B Karaki and Sandeep Kumar Rangaraju. (2017). "Where do jobs go when oil prices drop?". *Energy Economics*, 64, pp. 469-482. https://doi.org/10.1016/j.eneco.2016.02.006

Huber, Florian, Gary Koop, Luca Onorante, Michael Pfarrhofer and Josef Schreiner. (2023). "Nowcasting in a pandemic using non-parametric mixed frequency VARs". *Journal of Econometrics* 232(1), pp.52-69. https://doi.org/10.1016/j.jeconom.2020.11.006

Huber, Florian, and Luca Rossini. (2022). "Inference in Bayesian additive vector autoregressive tree models". *The Annals of Applied Statistics*, 16(1), pp. 104-123. https://doi.org/10.1214/21-aoas1488

Juvenal, Luciana, and Ivan Petrella. (2015). "Speculation in the oil market". *Journal of Applied Econometrics*, 30(4), pp. 621-649. https://doi.org/10.1002/jae.2388

Känzig, Diego R. (2021). "The macroeconomic effects of oil supply news: Evidence from OPEC announcements". *American Economic Review*, 111(4), pp. 1092-1125. https://doi.org/10.1257/aer.20190964

Kapelner, Adam, and Justin Bleich. (2016). "bartMachine: Machine learning with Bayesian additive regression trees". *Journal of Statistical Software*, 70(4), pp. 1-40. https://doi.org/10.18637/jss.v070.i04

Kharroubi, Enisse, Renée Spigt, Deniz Igan, Koji Takahashi and Egon Zakrajšek. (2023). "Markups and the asymmetric pass-through of cost push shocks". BIS Working Papers, 1150, Bank for International Settlements. https://www.bis.org/publ/work1150.htm

Kilian, Lutz. (2008). "The economic effects of energy price shocks". *Journal of Economic Literature*, 46(4), pp. 871-909. https://doi.org/10.1257/jel.46.4.871

Kilian, Lutz. (2009). "Not all oil price shocks are alike: Disentangling demand and supply shocks in the crude oil market". *American Economic Review*, 99(3), pp. 1053-1069. https://doi.org/10.1257/aer.99.3.1053

Kilian, Lutz. (2024). "How to construct monthly VAR proxies based on daily surprises in futures markets". *Journal of Economic Dynamics and Control*, 168(104966). https://doi.org/10.1016/j.jedc.2024.104966

Kilian, Lutz, and Logan T. Lewis. (2011). "Does the Fed respond to oil price shocks?". *The Economic Journal*, 121(555), pp. 1047-1072. https://doi.org/10.1111/j.1468-0297.2011.02437.x

Kilian, Lutz, and Daniel P. Murphy. (2014). "The role of inventories and speculative trading in the global market for crude oil". *Journal of Applied Econometrics*, 29(3), pp. 454-478. https://doi.org/10.1002/jae.2322

Kilian, Lutz, and Cheolbeom Park. (2009). "The impact of oil price shocks on the U.S. stock market". *International Economic Review*, 50(4), pp. 1267-1287. https://doi.org/10.1111/j.1468-2354.2009.00568.x

Kilian, Lutz, Michael D. Plante and Alexander W. Richter. (2026). "How times have changed: The impact of the 2026 Iran war on the U.S. economy". CEPR Discussion Papers, 21650, Centre for Economic Policy Research. https://cepr.org/publications/dp21650

Kilian, Lutz, and Robert J. Vigfusson. (2011a). "Are the responses of the U.S. economy asymmetric in energy price increases and decreases?". *Quantitative Economics*, 2(3), pp. 419-453. https://doi.org/10.3982/qe99

Kilian, Lutz, and Robert J. Vigfusson. (2011b). "Nonlinearities in the oil price–output relationship". *Macroeconomic Dynamics*, 15(S3), pp. 337-363. https://doi.org/10.1017/s1365100511000186

Kilian, Lutz, and Robert J. Vigfusson. (2017). "The role of oil price shocks in causing U.S. recessions". *Journal of Money, Credit and Banking*, 49(8), pp. 1747-1776. https://doi.org/10.1111/jmcb.12430

Kilian, Lutz, and Xiaoqing Zhou. (2023). "Chapter 3: The Econometrics of Oil Market VAR Models". In Yoosoon Chang, Sokbae Lee and J. Isaac Miller (eds.), *Advances in Econometrics. Essays in Honor of Joon Y. Park: Econometric Methodology in Empirical Applications.* Emerald Group Publishing Limited, 45, pp. 65-95. https://doi.org/10.1108/s0731-90532023000045b003

Koop, Gary, M. Hashem Pesaran and Simon M. Potter. (1996). "Impulse response analysis in nonlinear multivariate models". *Journal of Econometrics*, 74(1), pp. 119-147. https://doi.org/10.1016/0304-4076(95)01753-4

Korobilis, Dimitris, and Maximilian Schröder. (2024). "Monitoring multi-country macroeconomic risk: A quantile factor-augmented vector autoregressive (QFAVAR) approach". *Journal of Econometrics*, 249C(105730). https://doi.org/10.1016/j.jeconom.2024.105730

Lindsten, Fredrik, Michael I. Jordan and Thomas B. Schön. (2014). "Particle Gibbs with ancestor sampling". *Journal of Machine Learning Research*, 15(63), pp. 2145-2184. https://jmlr.org/papers/v15/lindsten14a.html

Linero, Antonio R. (2018). "Bayesian regression trees for high-dimensional prediction and variable selection". *Journal of the American Statistical Association*, 113(522), pp. 626-636. https://doi.org/10.1080/01621459.2016.1264957

Linero, Antonio R, and Yun Yang. (2018). "Bayesian regression tree ensembles that adapt to smoothness and sparsity". *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 80(5), pp. 1087-1110. https://doi.org/10.1111/rssb.12293

Maia, Mateus, Keefe Murphy and Andrew C. Parnell. (2024). "GP-BART: A novel Bayesian additive regression trees approach using Gaussian processes". *Computational Statistics & Data Analysis*, 190(107858). https://doi.org/10.1016/j.csda.2023.107858

Marcellino, Massimiliano, and Michael Pfarrhofer. (2025). "Nonparametric mixed frequency monitoring macro-at-risk". *Economics Letters*, 255(112498). https://doi.org/10.1016/j.econlet.2025.112498

Miescu, Mirela, Haroon Mumtaz and Konstantinos Theodoridis. (2024). "Nonlinear dynamics of large oil supply news shocks". Cardiff Economics Working Papers, E2024/18, Cardiff University. https://econpapers.repec.org/RePEc:cdf:wpaper:2024/18

Miranda-Agrippino, Silvia, and Giovanni Ricco. (2023). "Identification with external instruments in structural VARs". *Journal of Monetary Economics*, 135, pp. 1-19. https://doi.org/10.1016/j.jmoneco.2023.01.006

Moran, Gemma Elyse. (2019). *Bayesian approaches for modeling variation* [Ph. D. thesis]. University of Pennsylvania. https://repository.upenn.edu/server/api/core/bitstreams/a7df6139-e4d6-451d-aac6-94555378cb51/content

Mori, Lorenzo, and Gert Peersman. (2024). "Estimating the macroeconomic effects of oil supply news". CESifo Working Papers, 11532, CESifo. https://www.econstor.eu/bitstream/10419/308428/1/cesifo1_wp11532.pdf

Mork, Knut Anton. (1989). "Oil and the macroeconomy when prices go up and down: an extension of Hamilton's results". *Journal of Political Economy*, 97(3), pp. 740-744. https://doi.org/10.1086/261625

Mumtaz, Haroon. (2010). "Evolving UK macroeconomic dynamics: a time-varying factor augmented VAR". Working Papers, 386, Bank of England. https://www.bankofengland.co.uk/-/media/boe/files/working-paper/2010/evolving-uk-macroeconomic-dynamics-a-time-varying-factor-augmented-var.pdf

Mumtaz, Haroon, and Michele Piffer. (2022). "Impulse response estimation via flexible local projections". *arXiv preprint*, 2204.13150. https://arxiv.org/abs/2204.13150

Mumtaz, Haroon, and Laura Sunder-Plassmann. (2021). "Nonlinear effects of govern-ment spending shocks in the USA: Evidence from state-level data". *Journal of Applied Econometrics*, 36(1), pp. 86-97. https://doi.org/10.1002/jae.2800

Mumtaz, Haroon, Laura Sunder-Plassmann and Angeliki Theophilopoulou. (2018). "The state-level impact of uncertainty shocks". *Journal of Money, Credit and Banking*, 50(8), pp. 1879-1899. https://doi.org/10.1111/jmcb.12509

Neri, Stefano. (2025). "Energy prices, inflation and the ECB's monetary policy during the 2021-22 energy crisis". Temi di Discussione, 1481, Banca d'Italia. https://www.bancaditalia.it/pubblicazioni/temi-discussione/2025/2025-1481/index.html?com.dotmarketing.htmlpage.language=1

Rotemberg, Julio J., and Michael Woodford. (1991). "Markups and the business cycle". *NBER Macroeconomics Annual*, 6, pp. 63-129. https://doi.org/10.1086/654159

Stock, James H., and Mark W. Watson. (2005). "Implications of dynamic factor models for VAR analysis". NBER Working Paper Series, 11467, National Bureau of Economic Research. https://doi.org/10.3386/w11467

Velasco, Sofia. (2025). "Let the tree decide: FABART a non-parametric factor model". *arXiv preprint*, 2506.11551. https://arxiv.org/abs/2506.11551

Appendix

A Data Descriptions

Main Data Set

Table 3 lists the variables included in the baseline specification, together with their FRED mnemonics and transformation codes. Transformation codes follow the FRED-MD convention: 1 – no transformation (levels); 2 – first difference; 3 – first difference of logarithm. All variables cover the sample span of $1974:01 - 2019:12$.

In addition to the time-series variables, the state-level analysis uses a set of cross-sectional variables that capture persistent differences in industrial composition and regional economic conditions. These variables include sectoral shares (manufacturing, mining, and construction), measures of local slack (home vacancy rates), and regional indicators. Consistent with Mumtaz et al. (2018) and Mumtaz and Sunder-Plassmann (2021), they are constructed as time averages to proxy structural heterogeneity across states. Variables 13 to 62 are state-specific Total Nonform Employment variables sourced from the Bureau of Labor Statistics via FRED, with the transformation code number 3 (first difference of logarithms). The names of these state-level variables are included in the online appendix. The other variables are detailed in Table 4.

Table 3: Variable Descriptions and Transformations

No.	Mnemonic	T-code	Description	Source
Oil Market Variables				
1	POIL	3	Real oil price (WTI deflated by CPI)	EIA, FRED
2	WIP	3	World Industrial Production	IEA
3	OILPROD	3	World Oil Production	EIA
4	OILSTOCKS	3	OECD Petroleum Inventories	EIA
U.S. Macroeconomic Aggregates				
5	INDPRO	3	U.S. Industrial Production Index	FRED
6	CPI	3	U.S. Consumer Price Index	FRED
Financial Variables (Baseline only)				
7	VXOCLS	1	VXO Volatility Index	FRED
8	Macro_JLN	1	Macroeconomic Uncertainty (Jurado–Ludvigson–Ng)	JLN website
9	Financial_JLN	1	Financial Uncertainty (Jurado–Ludvigson–Ng)	JLN website
10	GS1	1	1-Year Treasury Constant Maturity Rate	FRED
11	SP500	3	S&P 500 Index	FRED
12	MSCI	3	MSCI World Index	Bloomberg

Notes: EIA = Energy Information Administration; FRED = Federal Reserve Economic Data (St. Louis Fed); IEA = International Energy Agency; JLN = Jurado–Ludvigson–Ng; BLS = Bureau of Labor Statistics. Transformation codes follow the FRED-MD convention: 1 – no transformation (levels); 2 – first difference; 3 – first difference of logarithm.

Table 4: Data for cross-section regressions

Variable	Construction	Source
Manufacturing (`man`)	Time average (1963–2014) of state value added in manufacturing	BEA (State GDP by industry)
Oil & gas (`oil`)	Time average (1963–2014) of state value added in oil & gas extraction (NAICS 211)	BEA (State GDP by industry)
Construction (`con`)	Time average (1963–2014) of state value added in construction	BEA (State GDP by industry)
Regional dummies (`regdum`)	Time-invariant categorical variables for U.S. Census regions (Northeast, Midwest, South, West)	U.S. Census Bureau

Notes: Cross-sectional variables are constructed as time averages of state-level GDP shares by industry. Manufacturing and construction follow Mumtaz et al. (2018). The oil variable is constructed following Mumtaz and Sunder-Plassmann (2021) using the BEA oil and gas extraction sector (NAICS 211).

A.1 Supplementary results: Empirical analysis

This subsection provides additional results from the empirical application.

Selection of factors

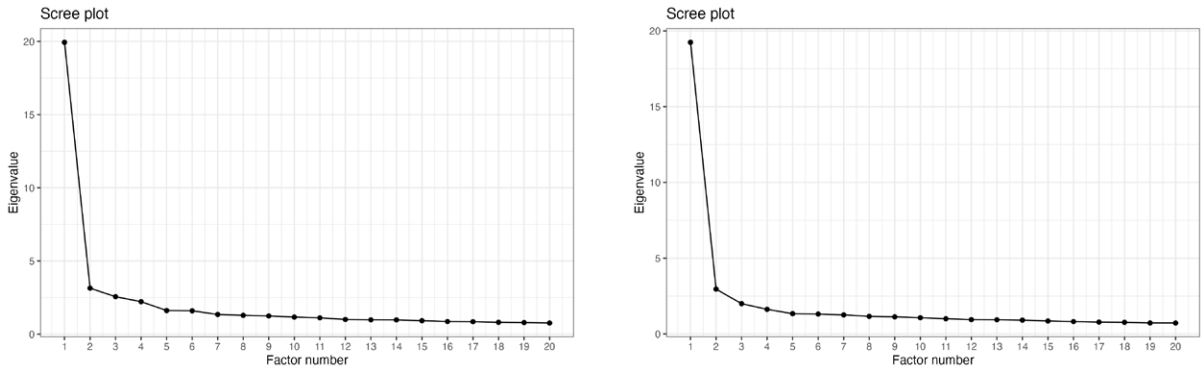

Baseline Reduced

Figure 8: Scree plots for factor selection: Baseline (left) and Reduced (right) specifications.

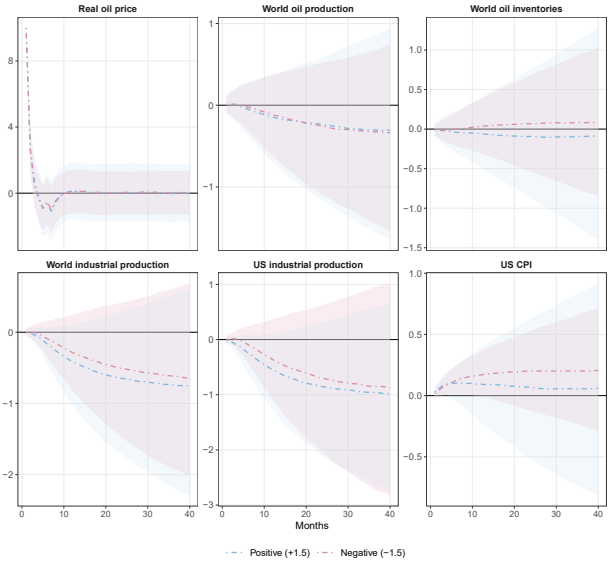

1.5% shock

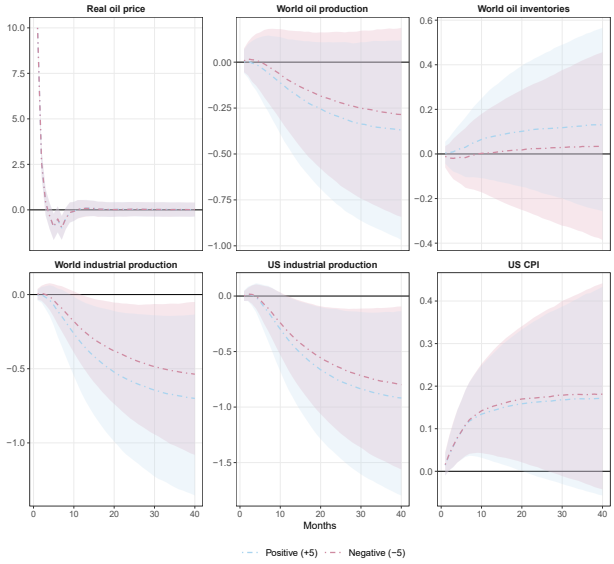

5% shock

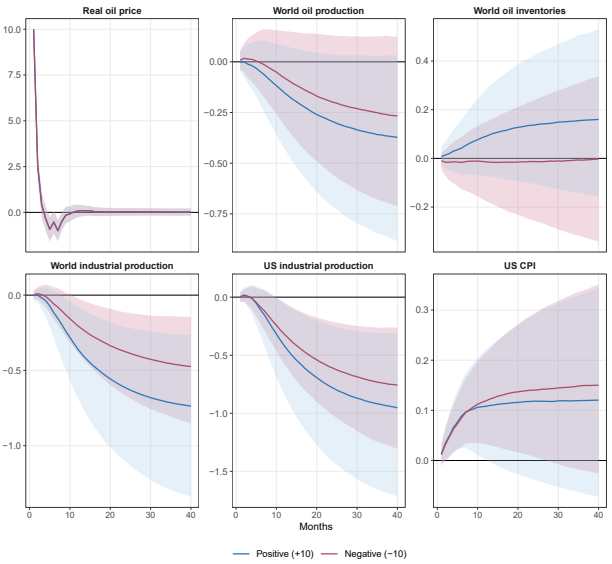

10% shock

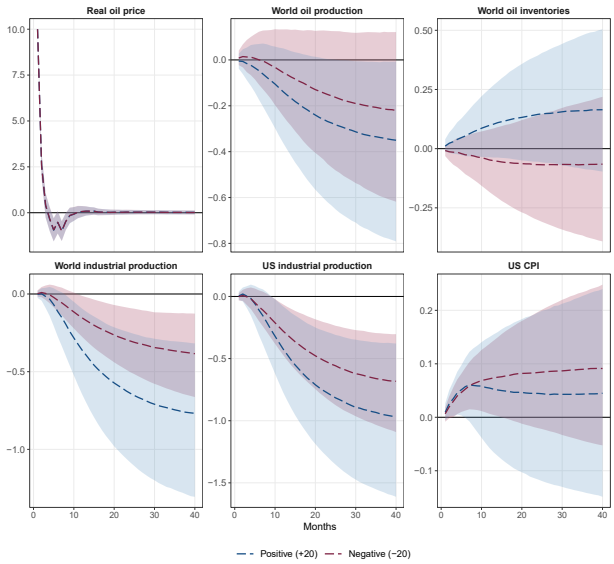

20% shock

Figure 9: Baseline responses to oil supply news shocks by size.

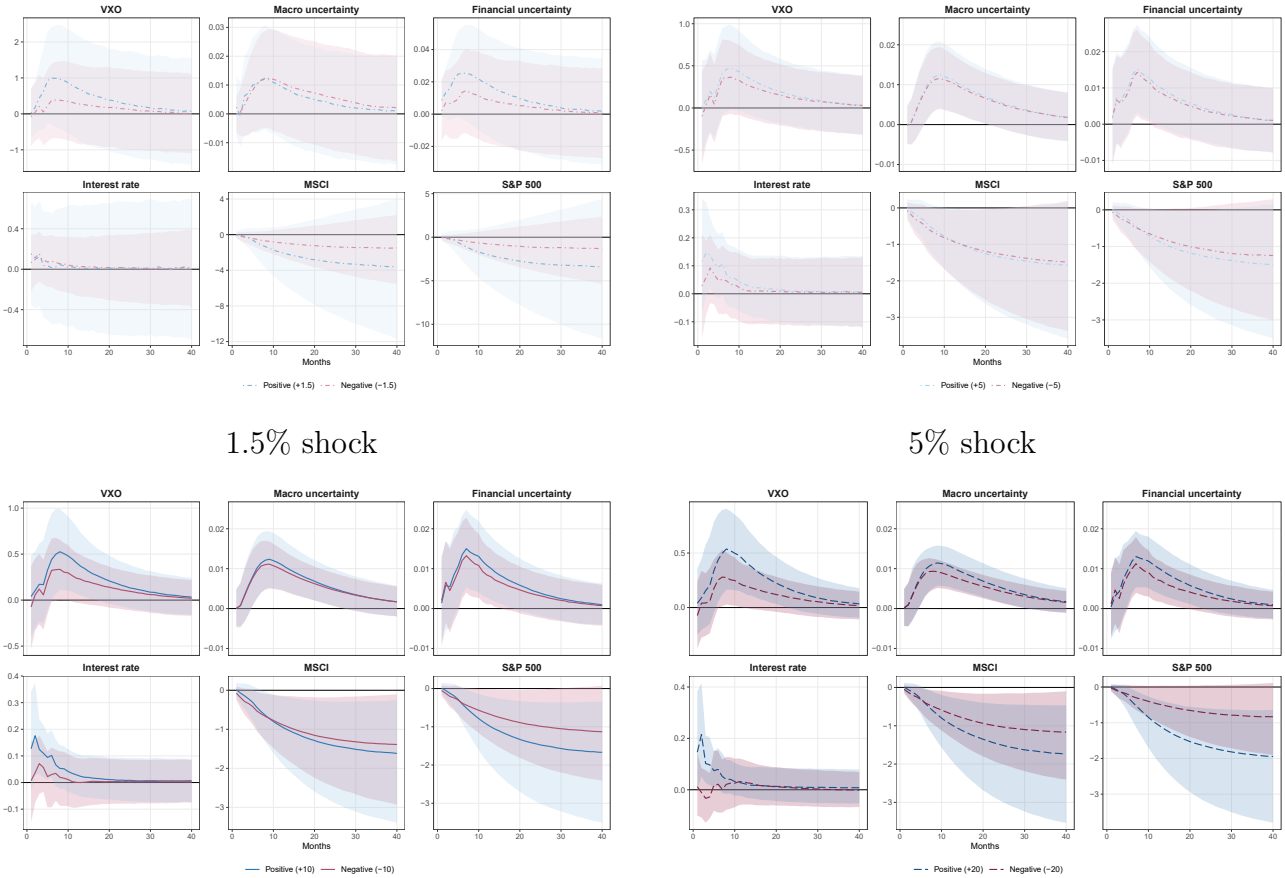

1.5% shock

5% shock

10% shock

20% shock

Figure 10: Financial-variable responses to oil supply news shocks by size.

FABART Online Supplementary Appendix

1 Extra Results: Cross Sectional Regressions of posterior draws of state-level employment responses on state characteristics

Table 1 presents the results for the cross-sectional regressions of posterior draws of state-level employment responses on state characteristics, as in the main text but with 68% intervals instead of 95% intervals. Table 2 includes 95% intervals that account for second stage uncertainty, i.e. uncertainty in each cross-sectional regression coefficient estimate, by approximating the second stage posterior (conditional on the draw of the dependent variable) with the estimated OLS coefficients and their sample covariance matrix. For each cross-sectional regression, a sample is made from the approximate conditional distribution of the regression coefficient. Quantiles of these samples are used to construct the intervals.

Table 1: State Characteristics and Employment Responses to Oil Shocks: First-Stage Posterior Uncertainty, 68% intervals

	1.5%		5%		10%	
	+	−	+	−	+	−
Manufacturing	-1.084	1.029	-1.029	1.006	-0.984	0.962
16^{th} percentile	−2.238	−0.281	−1.638	0.423	−1.484	0.504
84^{th} percentile	0.039	2.324	−0.419	1.581	−0.489	1.416
Oil & gas	2.329	-2.513	2.457	-2.390	2.246	-2.177
16^{th} percentile	−0.614	−5.387	0.988	−3.778	1.077	−3.273
84^{th} percentile	5.232	0.295	3.943	−0.957	3.430	−1.071
Construction	-5.065	4.120	-4.240	4.057	-4.286	4.001
16^{th} percentile	−13.635	−5.847	−8.720	−0.126	−8.142	0.635
84^{th} percentile	3.529	13.696	0.238	8.274	−0.460	7.319
Median adjusted R^2	0.295	0.345	0.496	0.484	0.513	0.500

Dependent variable: posterior draws of the two-year cumulative state employment response. Intervals are quantiles of draw-specific cross-sectional regression coefficients and account for first-stage posterior uncertainty in the estimated impulse responses. They account for first-stage posterior uncertainty in the estimated impulse responses. Point estimates are means of draw-specific cross-sectional regression coefficients. Regressions are across 50 U.S. states. Negative-shock responses are multiplied by minus one. Interval probability: 68%.

Table 2: State Characteristics and Employment Responses to Oil Shocks: First- and Second-Stage Uncertainty

	1.5%		5%		10%	
	+	−	+	−	+	−
Manufacturing	-1.084	1.029	-1.029	1.006	-0.984	0.962
2.5^{th}percentile	−3.839	−1.864	−2.543	−0.228	−2.263	−0.029
97.5^{th}percentile	1.427	4.062	0.242	2.478	0.016	2.167
Oil & gas	2.329	-2.513	2.457	-2.390	2.246	-2.177
2.5^{th}percentile	−4.274	−9.556	−0.524	−6.127	−0.148	−5.200
97.5^{th}percentile	9.763	3.928	6.204	0.546	5.106	0.191
Construction	-5.065	4.120	-4.240	4.057	-4.286	4.001
2.5^{th}percentile	−28.568	−18.548	−15.990	−5.531	−14.662	−3.612
97.5^{th}percentile	14.811	28.319	5.636	15.397	3.877	13.472
Median adjusted R^2	0.295	0.345	0.496	0.484	0.513	0.500

Dependent variable: posterior draws of the two-year cumulative state employment response. For each posterior IRF draw, a cross-sectional regression is estimated and one coefficient vector is sampled from a normal approximation using the draw-specific coefficient estimate and HC1 covariance matrix. Intervals are quantiles of these sampled coefficients and therefore incorporate both first-stage IRF uncertainty and conventional second-stage regression uncertainty. Regressions are across 50 U.S. states. Negative-shock responses are multiplied by minus one. Interval probability: 95%.

2 MCMC Diagnostics for Data Application

The tables below contain MCMC diagnostic results for the real data application. The number of post burn-in iterations is 10,000. [1]

Tables 3 to 6 contain MCMC diagnostics results for all outcome equation variances and for a subset of other variables flagged as having a low Effective Sample Size (ESS)

[1]Microsoft Copilot wrote the script that implemented the tests.

ratio (below 0.05),[2] Geweke (1992) test statistic p-values below 0.05, or a (single chain) stable GR test (Vats and Knudson, 2021) Potential Scale Reduction Factor (PSRF) below 1.05. Therefore, variables not included in these MCMC diagnostic results have better MCMC diagnostic results. The set of variables to which MCMC diagnostics tests were applied includes: measurement equation error variances, GIRFs for oil, factors $F_{2,t}$ to $F_{5,t}$, and 8 outcome variables at horizons $k = 1, 6, 12, 24, 40$, baseline steady-state expectations $E\left[X_{i,t+k} \mid \Psi_t, \overline{Y}\right]$ for all variables, expectations with shock equal to 10 $E\left[X_{i,t+k} \mid \Psi_t, \overline{Y}, e_j\right]$ for all outcomes, across-time means of factors $\frac{1}{T}\sum_{t=1}^{T} \boldsymbol{F}_{j,t}$ for $j = 2, 3, 4, 5$, and factors at time T $F_{j,T}$ for $j = 2, 3, 4, 5$.

Although the Geweke tests reject the null hypothesis for some selected variables, it is important to note that these are a subset of a much larger set of variables to which the tests were applied, and it can be seen in the tables that the other diagnostics tests mostly suggest MCMC convergence.

Table 7 contains the Multivariate Potential Scale Reduction Factor (MPSRF) obtained from the Stable Gelman-Rubin test of Vats and Knudson (2021). The blocks to which the test is applied are: measurement equation error variances, GRIFs for OIL at horizons $k = 1, 6, 12, 24, 40$, GIRFs for factors $F_{2,t}$ to $F_{5,t}$ at horizons $k = 1, 6, 12, 24, 40$, GIRFs for 5 outcome variables at horizons $k = 1, 6, 12, 24, 40$, steady state expectations and shock equal to 10 combinations for 5 outcome variables at horizons $k = 1, 6, 12, 24, 40$, and a group containing both across-time means of factors $\frac{1}{T}\sum_{t=1}^{T} \boldsymbol{F}_{j,t}$ for $j = 2, 3, 4, 5$, and factors at time T $F_{j,T}$ for $j = 2, 3, 4, 5$. The results for all blocks suggest that the chain has converged.

[2]ESS and other MCMC diagnostics except the Stable GR test are implemented under the default settings using the `R` package `coda` (Plummer et al., 2006).

	Mean	SD	ESS	ESS ratio	ACF lag 1	ACF lag 10	Geweke z	Geweke p	HW pass	Stable GR PSRF
σ_1	0.92	0.03	6060.91	0.61	0.12	0.12	1.19	0.23	TRUE	1.00
σ_2	0.66	0.03	669.95	0.07	0.38	0.38	-4.57	0.00	TRUE	1.00
σ_3	0.71	0.03	4627.28	0.46	0.21	0.21	-1.42	0.16	TRUE	1.00
σ_4	0.69	0.02	4128.31	0.41	0.17	0.17	1.13	0.26	TRUE	1.00
σ_5	0.62	0.02	2091.16	0.21	0.27	0.27	-6.43	0.00	TRUE	1.00
σ_6	0.62	0.02	2536.10	0.25	0.32	0.32	-1.54	0.12	TRUE	1.00
σ_7	0.74	0.03	1790.96	0.18	0.32	0.32	6.05	0.00	FALSE	1.00
σ_8	0.94	0.03	6243.75	0.62	0.13	0.13	2.00	0.05	TRUE	1.00
σ_9	0.92	0.03	6837.94	0.68	0.11	0.11	0.86	0.39	TRUE	1.00
σ_{10}	0.71	0.03	3143.81	0.31	0.23	0.23	-1.67	0.09	TRUE	1.00
σ_{11}	0.68	0.02	3987.62	0.40	0.18	0.18	-0.90	0.37	TRUE	1.00
σ_{12}	0.91	0.03	7743.17	0.77	0.09	0.09	0.64	0.52	TRUE	1.00
σ_{13}	0.80	0.03	5430.89	0.54	0.16	0.16	0.16	0.87	TRUE	1.00
σ_{14}	0.78	0.03	5465.57	0.55	0.16	0.16	2.33	0.02	TRUE	1.00
σ_{15}	0.76	0.03	3823.16	0.38	0.17	0.17	1.50	0.13	TRUE	1.00
σ_{16}	0.63	0.02	4522.97	0.45	0.21	0.21	-0.49	0.62	TRUE	1.00
σ_{17}	0.84	0.03	6574.39	0.66	0.11	0.11	2.68	0.01	TRUE	1.00
σ_{18}	0.86	0.03	3844.49	0.38	0.16	0.16	1.61	0.11	TRUE	1.00
σ_{19}	0.91	0.03	3104.27	0.31	0.15	0.15	1.51	0.13	TRUE	1.00
σ_{20}	0.81	0.03	3526.50	0.35	0.16	0.16	1.20	0.23	TRUE	1.00
σ_{21}	0.83	0.03	5775.01	0.58	0.14	0.14	0.38	0.70	TRUE	1.00
σ_{22}	0.84	0.03	4018.93	0.40	0.18	0.18	2.10	0.04	TRUE	1.00
σ_{23}	0.76	0.03	4426.65	0.44	0.15	0.15	0.94	0.35	TRUE	1.00
σ_{24}	0.71	0.02	6658.32	0.67	0.13	0.13	-2.64	0.01	TRUE	1.00
σ_{25}	0.78	0.03	4692.02	0.47	0.18	0.18	-0.52	0.61	TRUE	1.00
σ_{26}	0.75	0.03	3659.61	0.37	0.17	0.17	-4.11	0.00	TRUE	1.00
σ_{27}	0.88	0.03	6056.88	0.61	0.15	0.15	0.26	0.80	TRUE	1.00
σ_{28}	0.66	0.02	4529.88	0.45	0.18	0.18	0.62	0.54	TRUE	1.00
σ_{29}	0.90	0.03	5037.85	0.50	0.16	0.16	0.93	0.35	TRUE	1.00
σ_{30}	0.84	0.03	6372.70	0.64	0.13	0.13	-0.51	0.61	TRUE	1.00

Table 3: Selected MCMC diagnostics for FABART output. ESS = Effective Sample Size, ACF = Autocorrelation function, Geweke = Geweke (1992) test with 10% from the begining of the post-burn-in chain and 50% fromt he end of the chain (the z-value and p-value are included), HW = Heidelberger and Welch (1983) test for stationarity of the MCMC chain (TRUE if failed to reject null of stationarity), Stable GR PSRF = Stable Gelman-Rubin (Gelman and Rubin, 1992) Potential Scale Reduction Factor of the stable GR test for one chain (Vats and Knudson, 2021).

	Mean	SD	ESS	ESS ratio	ACF lag 1	ACF lag 10	Geweke z	Geweke p	HW pass	Stable GR PSRF
σ_{31}	0.74	0.03	3417.71	0.34	0.19	0.19	-2.40	0.02	TRUE	1.00
σ_{32}	0.73	0.03	4630.60	0.46	0.21	0.21	1.80	0.07	TRUE	1.00
σ_{33}	0.78	0.03	3747.06	0.37	0.20	0.20	1.38	0.17	TRUE	1.00
σ_{34}	0.75	0.03	4552.87	0.46	0.14	0.14	-0.66	0.51	TRUE	1.00
σ_{35}	0.75	0.03	4627.75	0.46	0.21	0.21	0.78	0.44	TRUE	1.00
σ_{36}	0.59	0.02	3533.88	0.35	0.17	0.17	-0.43	0.66	TRUE	1.00
σ_{37}	0.84	0.03	5311.89	0.53	0.15	0.15	1.44	0.15	TRUE	1.00
σ_{38}	0.69	0.02	5518.59	0.55	0.15	0.15	1.68	0.09	TRUE	1.00
σ_{39}	0.76	0.03	4206.52	0.42	0.14	0.14	-0.02	0.98	TRUE	1.00
σ_{40}	0.84	0.03	5702.60	0.57	0.14	0.14	-0.85	0.40	TRUE	1.00
σ_{41}	0.76	0.03	4400.94	0.44	0.15	0.15	1.52	0.13	TRUE	1.00
σ_{42}	0.77	0.03	4063.34	0.41	0.22	0.22	-3.56	0.00	TRUE	1.00
σ_{43}	0.65	0.02	4085.02	0.41	0.17	0.17	-2.25	0.02	TRUE	1.00
σ_{44}	0.61	0.02	1853.02	0.19	0.29	0.29	-1.18	0.24	TRUE	1.00
σ_{45}	0.68	0.02	3185.88	0.32	0.19	0.19	2.30	0.02	TRUE	1.00
σ_{46}	0.86	0.03	6169.97	0.62	0.13	0.13	-1.60	0.11	TRUE	1.00
σ_{47}	0.75	0.03	6060.84	0.61	0.13	0.13	1.29	0.20	TRUE	1.00
σ_{48}	0.67	0.02	4935.01	0.49	0.15	0.15	0.22	0.83	TRUE	1.00
σ_{49}	0.99	0.03	5119.80	0.51	0.14	0.14	2.81	0.00	TRUE	1.00
σ_{50}	0.88	0.03	3949.75	0.40	0.21	0.21	1.00	0.32	TRUE	1.00
σ_{51}	0.87	0.03	2677.39	0.27	0.14	0.14	-5.74	0.00	TRUE	1.00
σ_{52}	1.01	0.03	6842.87	0.68	0.07	0.07	-0.94	0.35	TRUE	1.00
σ_{53}	0.94	0.03	7186.90	0.72	0.09	0.09	-0.30	0.76	TRUE	1.00
σ_{54}	0.79	0.03	795.94	0.08	0.25	0.25	-4.34	0.00	FALSE	1.00
σ_{55}	0.79	0.03	2377.25	0.24	0.25	0.25	0.28	0.78	TRUE	1.00
σ_{56}	0.65	0.03	1701.15	0.17	0.34	0.34	-0.12	0.90	TRUE	1.00
σ_{57}	0.17	0.02	121.62	0.01	0.88	0.88	2.33	0.02	TRUE	1.00
σ_{58}	0.43	0.03	366.29	0.04	0.62	0.62	0.29	0.77	FALSE	1.00
σ_{59}	0.37	0.02	491.96	0.05	0.62	0.62	-1.01	0.31	TRUE	1.00
σ_{60}	0.97	0.03	5177.02	0.52	0.12	0.12	-0.80	0.42	TRUE	1.00
σ_{61}	0.98	0.03	5817.69	0.58	0.09	0.09	0.26	0.79	TRUE	1.00

Table 4: Selected MCMC diagnostics for FABART output. ESS = Effective Sample Size, ACF = Autocorrelation function, Geweke = Geweke (1992) test with 10% from the begining of the post-burn-in chain and 50% fromt he end of the chain (the z-value and p-value are included), HW = Heidelberger and Welch (1983) test for stationarity of the MCMC chain (TRUE if failed to reject null of stationarity), Stable GR PSRF = Stable Gelman-Rubin (Gelman and Rubin, 1992) Potential Scale Reduction Factor of the stable GR test for one chain (Vats and Knudson, 2021).

	Mean	SD	ESS	ESS ratio	ACF lag 1	ACF lag 10	Geweke z	Geweke p	HW pass	Stable GR PSRF
$GIRF$, OIL, $k = 1$	10.00	0.14	10000	1.00	-0.01	-0.01	-2.04	0.04	TRUE	1.00
$GIRF$, F_2, $k = 1$	-0.07	1.15	7279.27	0.73	0.05	0.05	-2.56	0.01	TRUE	1.00
$GIRF$, F_2, $k = 6$	-0.87	0.86	6696.73	0.67	0.07	0.07	2.10	0.04	TRUE	1.00
$GIRF$, F_3, $k = 6$	-0.72	0.75	5578.33	0.56	0.11	0.11	-4.66	0.00	TRUE	1.00
$GIRF$, F_2, $k = 12$	-0.69	0.48	8055.20	0.81	0.05	0.05	-2.68	0.01	TRUE	1.00
$GIRF$, F_4, $k = 12$	-0.88	0.70	8118.25	0.81	0.03	0.03	-3.08	0.00	TRUE	1.00
$GIRF$, F_4, $k = 24$	-0.44	0.55	9013.53	0.90	0.00	0.00	-2.39	0.02	TRUE	1.00
$GIRF$, $aknan$, $k = 6$	0.21	0.20	3801.44	0.38	0.13	0.13	2.68	0.01	FALSE	1.00
$GIRF$, $aknan$, $k = 12$	0.11	0.13	2393.81	0.24	0.15	0.15	2.19	0.03	FALSE	1.00
$GIRF$, $aknan$, $k = 24$	0.06	0.09	4128.96	0.41	0.07	0.07	2.42	0.02	TRUE	1.00
$GIRF$, $alnan$, $k = 1$	0.10	0.60	8423.91	0.84	0.03	0.03	-2.07	0.04	TRUE	1.00
$GIRF$, $alnan$, $k = 12$	-0.49	0.32	5244.30	0.52	0.07	0.07	-2.97	0.00	TRUE	1.00
$GIRF$, $arnan$, $k = 12$	-0.46	0.29	4274.70	0.43	0.09	0.09	-3.05	0.00	FALSE	1.00
$GIRF$, $aznan$, $k = 12$	-0.50	0.31	4672.79	0.47	0.08	0.08	-3.00	0.00	TRUE	1.00
$GIRF$, $aznan$, $k = 24$	-0.24	0.23	6094.15	0.61	0.02	0.02	-2.49	0.01	TRUE	1.00
$GIRF$, $canan$, $k = 12$	-0.56	0.36	4892.91	0.49	0.07	0.07	-2.22	0.03	TRUE	1.00
$GIRF$, $conan$, $k = 12$	-0.37	0.33	6458.16	0.65	0.04	0.04	-3.38	0.00	TRUE	1.00
$GIRF$, $ctnan$, $k = 12$	-0.41	0.26	4838.12	0.48	0.07	0.07	-3.77	0.00	TRUE	1.00
$GIRF$, $ctnan$, $k = 24$	-0.20	0.19	7454.54	0.75	0.01	0.01	-3.72	0.00	TRUE	1.00
$GIRF$, $dcnan$, $k = 6$	-0.17	0.20	5010.46	0.50	0.10	0.10	-2.82	0.00	TRUE	1.00
$GIRF$, $dcnan$, $k = 12$	-0.13	0.13	5173.43	0.52	0.07	0.07	-2.56	0.01	TRUE	1.00

Table 5: Selected MCMC diagnostics for FABART output. ESS = Effective Sample Size, ACF = Autocorrelation function, Geweke = Geweke (1992) test with 10% from the begining of the post-burn-in chain and 50% fromt he end of the chain (the z-value and p-value are included), HW = Heidelberger and Welch (1983) test for stationarity of the MCMC chain (TRUE if failed to reject null of stationarity), Stable GR PSRF = Stable Gelman-Rubin (Gelman and Rubin, 1992) Potential Scale Reduction Factor of the stable GR test for one chain (Vats and Knudson, 2021).

	Mean	SD	ESS	ESS ratio	ACF lag 1	ACF lag 10	Geweke z	Geweke p	HW pass	Stable GR PSRF
$E\left[X_{i,t+k} \mid \Psi_t, \overline{Y}\right]$, *alnan*, $k = 1$	-0.33	0.27	2267.89	0.23	0.16	0.16	-3.67	0.00	TRUE	1.00
$E\left[X_{i,t+k} \mid \Psi_t, \overline{Y}\right]$, canan, $k = 6$	-0.69	0.25	5978.80	0.60	0.09	0.09	2.42	0.02	TRUE	1.00
$E\left[X_{i,t+k} \mid \Psi_t, \overline{Y}\right]$, canan, $k = 12$	-0.51	0.23	4189.06	0.42	0.10	0.10	2.48	0.01	TRUE	1.00
$E\left[X_{i,t+k} \mid \Psi_t, \overline{Y}\right]$, canan, $k = 24$	-0.32	0.22	3419.04	0.34	0.13	0.13	2.58	0.01	TRUE	1.00
$E\left[X_{i,t+k} \mid \Psi_t, \overline{Y}\right]$, canan, $k = 40$	-0.27	0.20	1847.87	0.18	0.17	0.17	2.10	0.04	TRUE	1.00
$E\left[X_{i,t+k} \mid \Psi_t, \overline{Y}\right]$, dcnan, $k = 6$	-0.03	0.14	6238.94	0.62	0.08	0.08	2.76	0.01	FALSE	1.00
$E\left[X_{i,t+k} \mid \Psi_t, \overline{Y}\right]$, dcnan, $k = 12$	-0.03	0.12	4365.87	0.44	0.09	0.09	3.75	0.00	TRUE	1.00
$E\left[X_{i,t+k} \mid \Psi_t, \overline{Y}\right]$, dcnan, $k = 24$	-0.04	0.12	4665.90	0.47	0.11	0.11	3.19	0.00	TRUE	1.00
$E\left[X_{i,t+k} \mid \Psi_t, \overline{Y}\right]$, dcnan, $k = 40$	-0.06	0.12	4751.38	0.48	0.13	0.13	2.85	0.00	TRUE	1.00
$E\left[X_{i,t+k} \mid \Psi_t, \overline{Y}, e_j\right]$, aknan, $k = 6$	0.40	0.22	2704.49	0.27	0.17	0.17	2.40	0.02	TRUE	1.00
$E\left[X_{i,t+k} \mid \Psi_t, \overline{Y}, e_j\right]$, aknan, $k = 12$	0.24	0.17	2094.71	0.21	0.18	0.18	3.00	0.00	TRUE	1.00
$E\left[X_{i,t+k} \mid \Psi_t, \overline{Y}, e_j\right]$, alnan, $k = 1$	-0.23	0.63	7132.77	0.71	0.06	0.06	-4.11	0.00	FALSE	1.00
$E\left[X_{i,t+k} \mid \Psi_t, \overline{Y}, e_j\right]$, alnan, $k = 6$	-0.98	0.49	6739.41	0.67	0.06	0.06	-2.51	0.01	TRUE	1.00
$E\left[X_{i,t+k} \mid \Psi_t, \overline{Y}, e_j\right]$, alnan, $k = 12$	-0.87	0.35	7074.03	0.71	0.06	0.06	-3.38	0.00	TRUE	1.00
$E\left[X_{i,t+k} \mid \Psi_t, \overline{Y}, e_j\right]$, arnan, $k = 12$	-0.81	0.32	5151.87	0.52	0.08	0.08	-2.88	0.00	TRUE	1.00
$E\left[X_{i,t+k} \mid \Psi_t, \overline{Y}, e_j\right]$, aznan, $k = 12$	-1.00	0.35	3803.31	0.38	0.10	0.10	-3.12	0.00	TRUE	1.00
$E\left[X_{i,t+k} \mid \Psi_t, \overline{Y}, e_j\right]$, aznan, $k = 24$	-0.53	0.28	4965.96	0.50	0.05	0.05	-3.13	0.00	TRUE	1.00
$E\left[X_{i,t+k} \mid \Psi_t, \overline{Y}, e_j\right]$, canan, $k = 40$	-0.36	0.25	5120.26	0.51	0.08	0.08	2.82	0.00	TRUE	1.00
$E\left[X_{i,t+k} \mid \Psi_t, \overline{Y}, e_j\right]$, conan, $k = 12$	-0.81	0.37	5591.88	0.56	0.08	0.08	-2.73	0.01	TRUE	1.00
$E\left[X_{i,t+k} \mid \Psi_t, \overline{Y}, e_j\right]$, ctnan, $k = 12$	-0.71	0.29	5071.05	0.51	0.07	0.07	-2.77	0.01	TRUE	1.00
$E\left[X_{i,t+k} \mid \Psi_t, \overline{Y}, e_j\right]$, dcnan, $k = 24$	-0.09	0.15	7779.62	0.78	0.07	0.07	2.19	0.03	TRUE	1.00
$E\left[X_{i,t+k} \mid \Psi_t, \overline{Y}, e_j\right]$, dcnan, $k = 40$	-0.08	0.14	6201.30	0.62	0.11	0.11	2.69	0.01	TRUE	1.00
$\frac{1}{T}\sum_{t=1}^{T} F_{1,t}$	-0.01	0.01	6456.30	0.65	0.02	0.02	-2.36	0.02	TRUE	1.00
$F_{1,T}$	0.69	0.13	411.36	0.04	0.28	0.28	-2.51	0.01	TRUE	1.00
$\frac{1}{T}\sum_{t=1}^{T} F_{2,t}$	-0.00	0.01	10000	1.00	0.01	0.01	-3.61	0.00	TRUE	1.00
$F_{2,T}$	0.06	0.31	1053.33	0.11	0.13	0.13	-2.06	0.04	TRUE	1.00

Table 6: Selected MCMC diagnostics for FABART output. ESS = Effective Sample Size, ACF = Autocorrelation function, Geweke = Geweke (1992) test with 10% from the begining of the post-burn-in chain and 50% fromt he end of the chain (the z-value and p-value are included), HW = Heidelberger and Welch (1983) test for stationarity of the MCMC chain (TRUE if failed to reject null of stationarity), Stable GR PSRF = Stable Gelman-Rubin (Gelman and Rubin, 1992) Potential Scale Reduction Factor of the stable GR test for one chain (Vats and Knudson, 2021).

Block	Number of Variables	Stable GR MPSRF
Σ	50.00	1.00
GIRF OIL	5.00	1.00
Group GIRF $\boldsymbol{F}$	20.00	1.00
Group GIRF X	25.00	1.00
Group $E\left[X_{i,t+k} \mid \Psi_t, \overline{Y}\right]$	25.00	1.00
Group $E\left[X_{i,t+k} \mid \Psi_t, \overline{Y}, e_j\right]$	25.00	1.00
Group $\boldsymbol{F}$	8.00	1.00

Table 7: Multivariate stable Gelman–Rubin diagnostics for FABART output blocks. "Number of Variables" refers to the number of variables in the block to which the Multivariate Stable GR test is applied. MPSRD = Multivariate Potential Scale Reduction Factor. See Vats and Knudson (2021) for a description of the test.

3 Monte Carlo Study

The objective of this appendix is to assess the FABART model's ability to capture nonlinear dynamics in economic data. We develop a simulation experiment that generates data with nonlinear effects of latent factors. The central question is whether FABART can successfully recover these nonlinear relationships without imposing restrictive parametric assumptions. [3]

3.1 DGP Description

Model performance is evaluated within a factor model framework by comparing forecasts to true outcomes across 18 data-generating processes (DGPs):

- A-C: Two linear and one approximately linear measurement equation DGPs, with random draws of parameters across 100 simulated datasets. There is one factor variable.
- D-F: The same DGPs as A-C, except with fixed parameters across 100 simulated datasets. There is one factor variable.
- G-I: Three DGPs with nonlinear measurement equations and single factors.
- J-R: Nine DGPs with nonlinear measurement equations and three factors.

The factor evolves according to an autoregressive process of order p:

$$F_t = c + b_1 F_{t-1} + b_2 F_{t-2} + ... + b_p F_{t-p} + e_t, \quad e_t \sim \mathcal{N}(0, \sigma^2) \tag{1}$$

where e_t is normally distributed with variance σ^2.

The setup follows a standard factor model in which the N observed variables X_t is driven by J latent factors, $F_t \in \mathbb{R}^J$. The baseline (linear) specification, DGP A, is given by:

$$X_t = BF_t + V_t, \quad V_t \sim \mathcal{N}(0, R), \tag{2}$$

where X_t is a $N \times T$ vector of observed variables, B is a $N \times J$ factor loading matrix, and V_t follows a normal distribution with mean zero and diagonal covariance matrix R $(N \times N)$.

Throughout this simulation study, FABART is implemented with 2000 burn-in and 2000 post-burn-in MCMC iterations, non-sparse splitting rules, diffuse transition equation coefficient prior, inverse Wishart prior for the transition equation error covariance matrix, and an identity restriction on the factor loadings in the linear approximation to the measurement equation, but not restriction on the true measurement equation. The benchmark linear dynamic factor model (LDFM) is estimated in a multiple step approach by first obtaining principal components, and fitting separately principal components regressions for the measurement equations and a linear VAR for the transition equation.The number of lags specified for FABART and the LDFM is the true number of lags for the relavant DGP. A linear autoregessive model (AR1), with separate autoregressions for each outcome variable, is also included as a benchmark. Results are averaged across all outcomes/factors, samples from the DGP, and horizons, unless otherwise stated.

[3] Extensive use was made of Google Gemini and OpenAI GPT LLMs in the design, implementation, and summarization of this Monte Carlo study.

For all DGPs, the number of time periods in the training data is 304 (350 minus 40 hold-out observations and removing the first 6 simulated time periods), and the first p observations (the number of lags) are dropped in the estimation of FABART, and only used for setting priors and initial values of the factors. The forecasting exercise is repeated 100 times for each DGP.

The DGPs are defined as follows:

(A) Linear DGP, $J = 1$, $N = 20$.
Transition: $F_t = 0.6F_{t-1} - 0.3F_{t-2} + 0.2F_{t-3} + e_t, \quad e_t \sim \mathcal{N}(0, 1)$.
Measurement: $X_t = BF_t + V_t, \quad V_t \sim \mathcal{N}(0, R)$, $B = \left[9, 8, 7, 6, \ldots, 1, 0, 9, 8, \ldots, 1, 0\right]^T$, $R = 0.1 \times \mathbf{I}_{20}$.

(B) Linear (in factors nonlinear in coefficients) DGP, $J = 1$, $N = 20$.
Transition: $F_t = 0.6F_{t-1} - 0.3F_{t-2} + 0.2F_{t-3} + e_t, \quad e_t \sim \mathcal{N}(0, 1)$.
Measurement: $X_t = B^2F_t + V_t, \quad V_t \sim \mathcal{N}(0, R)$, $B = \frac{1}{10}\left[9, 8, 7, 6, \ldots, 1, 0, 9, 8, \ldots, 1, 0\right]^T$, $R = 0.1 \times \mathbf{I}_{20}$.

(C) Nonlinear DGP, $J = 1$, $N = 20$.
Transition: $F_t = 0.6F_{t-1} - 0.3F_{t-2} + 0.2F_{t-3} + e_t, \quad e_t \sim \mathcal{N}(0, 1)$.
Measurement: $X_t = \tanh(BF_t) + V_t, \quad V_t \sim \mathcal{N}(0, R)$, $B = \frac{1}{10}\left[9, 8, 7, 6, \ldots, 1, 0, 9, 8, \ldots, 1, 0\right]^T$, $R = 0.1 \times \mathbf{I}_{20}$.

(D) Nonlinear DGP, $J = 1$, $N = 20$.
Transition: $F_t = 0.6F_{t-1} - 0.3F_{t-2} + 0.2F_{t-3} + e_t, \quad e_t \sim \mathcal{N}(0, 1)$.
Measurement: $X_t = BF_t + \boldsymbol{\gamma}^T \left(F_t - \frac{1}{T}\sum_{s=1}^{T} F_s\right) + V_t, \quad V_t \sim \mathcal{N}(0, R)$,
$B = \frac{1}{10}\left[2, 3, 4, 5, \ldots, 9, 10, 5, 2, 3 \ldots, 10, 5\right]^T$, $R = 0.25 \times \mathbf{I}_{20}$,
$\boldsymbol{\gamma} = [0.25 \times \mathbf{1}_5^T, 0 \times \mathbf{1}_5^T, -0.20 \times \mathbf{1}_5^T, 0.15 \times \mathbf{1}_5^T]^T$.

(E) Nonlinear DGP, $J = 1$, $N = 20$.
Transition: $F_t = 0.6F_{t-1} - 0.3F_{t-2} + 0.2F_{t-3} + e_t, \quad e_t \sim \mathcal{N}(0, 1)$.
Measurement: $X_t = BF_t + \boldsymbol{\gamma}^T F_t \times \mathbb{I}\{F_t > 0\} + V_t, \quad V_t \sim \mathcal{N}(0, R)$,
$B = \frac{1}{10}\left[2, 3, 4, 5, \ldots, 9, 10, 5, 2, 3 \ldots, 10, 5\right]^T$, $R = 0.25 \times \mathbf{I}_{20}$,
$\boldsymbol{\gamma} = [0.3 \times \mathbf{1}_7^T, 0 \times \mathbf{1}_6^T, -0.25 \times \mathbf{1}_7^T]^T$.

(F) Nonlinear DGP, $J = 1$, $N = 20$.
Transition: $F_t = 0.6F_{t-1} - 0.3F_{t-2} + 0.2F_{t-3} + e_t, \quad e_t \sim \mathcal{N}(0, 1)$.
Measurement: $X_t = BF_t + \boldsymbol{\gamma}^T \tanh(0.8F_t) + V_t, \quad V_t \sim \mathcal{N}(0, R)$,
$B = \frac{1}{10}\left[2, 3, 4, 5, \ldots, 9, 10, 5, 2, 3 \ldots, 10, 5\right]^T$, $R = 0.25 \times \mathbf{I}_{20}$,
$\boldsymbol{\gamma} = [0.6 \times \mathbf{1}_6^T, 0 \times \mathbf{1}_8^T, -0.35 \times \mathbf{1}_6^T]^T$.

(G) Nonlinear DGP, $J = 3$, $N = 12$.
Transition: $\boldsymbol{F}_t = \boldsymbol{\Phi}\boldsymbol{F}_{t-1} + e_t, \quad e_t \sim \mathcal{N}(\mathbf{0}_3, \boldsymbol{I}_3)$.

Measurement: $\boldsymbol{\Lambda F}_t + \boldsymbol{\Gamma q}_t + \boldsymbol{\eta}_t \quad , \quad \boldsymbol{\eta}_t \sim \mathcal{N}(\mathbf{0}_{12}, \boldsymbol{R})$

$$\boldsymbol{\Phi} = \begin{bmatrix} 0.85 & 0.00 & 0.00 \\ 0.00 & 0.80 & 0.00 \\ 0.00 & 0.00 & 0.75 \end{bmatrix} \quad , \quad \Lambda = \begin{bmatrix} 1 & 0 & 0 \\ 1 & 0 & 0 \\ 1 & 0 & 0 \\ 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 1 & 0 \\ 0 & 1 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \\ 0 & 0 & 1 \\ 0 & 0 & 1 \\ 0 & 0 & 1 \end{bmatrix} \quad , \quad \Gamma = \begin{bmatrix} 0 & 0 & 0 & 0.3 & 0 & 0 \\ 0 & 0 & 0 & 0.3 & 0 & 0 \\ 0 & 0 & 0 & 0.3 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0.3 \\ 0 & 0 & 0 & 0 & 0 & 0.3 \\ 0 & 0 & 0 & 0 & 0 & 0.3 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0.3 & 0 \\ 0 & 0 & 0 & 0 & 0.3 & 0 \\ 0 & 0 & 0 & 0 & 0.3 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \end{bmatrix}$$

$$\boldsymbol{R} = \begin{bmatrix} 2\boldsymbol{I}_4 & \mathbf{0}_{4\times4} & \mathbf{0}_{4\times4} \\ \mathbf{0}_{4\times4} & 3\boldsymbol{I}_4 & \mathbf{0}_{4\times4} \\ \mathbf{0}_{4\times4} & \mathbf{0}_{4\times4} & 4\boldsymbol{I}_4 \end{bmatrix}, \boldsymbol{q}_t = \begin{bmatrix} F_{1t}^2 \\ F_{2t}^2 \\ F_{3t}^2 \\ F_{1t}F_{2t} \\ F_{1t}F_{3t} \\ F_{2t}F_{3t} \end{bmatrix}$$

(H) Nonlinear DGP, $J = 3$, $N = 12$.
Transition: $\boldsymbol{F}_t = \boldsymbol{\Phi F}_{t-1} + e_t, \quad e_t \sim \mathcal{N}(\mathbf{0}_3, \boldsymbol{I}_3)$.
Measurement: $\boldsymbol{\Lambda F}_t + \boldsymbol{\Gamma q}_t + \boldsymbol{\eta}_t \quad , \quad \boldsymbol{\eta}_t \sim \mathcal{N}(\mathbf{0}_{12}, \boldsymbol{R})$

$$\boldsymbol{\Phi} = \begin{bmatrix} 0.85 & 0.00 & 0.00 \\ 0.00 & 0.80 & 0.00 \\ 0.00 & 0.00 & 0.75 \end{bmatrix} \quad , \quad \Lambda = \begin{bmatrix} 1 & 0 & 0 \\ 1 & 0 & 0 \\ 1 & 0 & 0 \\ 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 1 & 0 \\ 0 & 1 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \\ 0 & 0 & 1 \\ 0 & 0 & 1 \\ 0 & 0 & 1 \end{bmatrix} \quad , \quad \Gamma = \begin{bmatrix} 0.35 & 0 & 0 \\ 0.35 & 0 & 0 \\ 0.35 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 0.35 \\ 0 & 0 & 0.35 \\ 0 & 0 & 0.35 \\ 0 & 0 & 0 \\ 0 & 0.35 & 0 \\ 0 & 0.35 & 0 \\ 0 & 0.35 & 0 \\ 0 & 0 & 0 \end{bmatrix}$$

$$\boldsymbol{R} = \begin{bmatrix} 2\boldsymbol{I}_4 & \mathbf{0}_{4\times4} & \mathbf{0}_{4\times4} \\ \mathbf{0}_{4\times4} & 3\boldsymbol{I}_4 & \mathbf{0}_{4\times4} \\ \mathbf{0}_{4\times4} & \mathbf{0}_{4\times4} & 4\boldsymbol{I}_4 \end{bmatrix}, \boldsymbol{q}_t = \begin{bmatrix} F_{1t} \times \mathbb{I}\{F_{2t} > 0\} \\ F_{2t} \times \mathbb{I}\{F_{3t} > 0\} \\ F_{3t} \times \mathbb{I}\{F_{1t} > 0\} \end{bmatrix}$$

(I) Nonlinear DGP, $J = 3$, $N = 12$.
Transition: $\boldsymbol{F}_t = \boldsymbol{\Phi F}_{t-1} + e_t, \quad e_t \sim \mathcal{N}(\mathbf{0}_3, \boldsymbol{I}_3)$.

Measurement: $\boldsymbol{\Lambda F}_t + \boldsymbol{\Gamma q}_t + \boldsymbol{\eta}_t \quad , \quad \boldsymbol{\eta}_t \sim \mathcal{N}(\mathbf{0}_{12}, \boldsymbol{R})$

$$\boldsymbol{\Phi} = \begin{bmatrix} 0.85 & 0.00 & 0.00 \\ 0.00 & 0.80 & 0.00 \\ 0.00 & 0.00 & 0.75 \end{bmatrix} \quad , \quad \Lambda = \begin{bmatrix} 1 & 0 & 0 \\ 1 & 0 & 0 \\ 1 & 0 & 0 \\ 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 1 & 0 \\ 0 & 1 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \\ 0 & 0 & 1 \\ 0 & 0 & 1 \\ 0 & 0 & 1 \end{bmatrix} \quad , \quad \Gamma = \begin{bmatrix} 0.2 & 0 & 0 & 0.25 & 0 & 0 \\ 0.2 & 0 & 0 & 0.25 & 0 & 0 \\ 0.2 & 0 & 0 & 0.25 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0.2 & 0 & 0 & 0 & 0.25 \\ 0 & 0.2 & 0 & 0 & 0 & 0.25 \\ 0 & 0.2 & 0 & 0 & 0 & 0.25 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0.2 & 0 & 0.25 & 0 \\ 0 & 0 & 0.2 & 0 & 0.25 & 0 \\ 0 & 0 & 0.2 & 0 & 0.25 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \end{bmatrix}$$

$$\boldsymbol{R} = \begin{bmatrix} 2\boldsymbol{I}_4 & \mathbf{0}_{4\times4} & \mathbf{0}_{4\times4} \\ \mathbf{0}_{4\times4} & 3\boldsymbol{I}_4 & \mathbf{0}_{4\times4} \\ \mathbf{0}_{4\times4} & \mathbf{0}_{4\times4} & 4\boldsymbol{I}_4 \end{bmatrix}, \boldsymbol{q}_t = \begin{bmatrix} \tanh(0.8 \times F_{1t}) \\ \tanh(0.8 \times F_{2t}) \\ \tanh(0.8 \times F_{3t}) \\ F_{1t}F_{2t} \\ F_{1t}F_{3t} \\ F_{2t}F_{3t} \end{bmatrix}$$

(J) Nonlinear DGP, $J = 3$, $N = 12$.
Transition: $\boldsymbol{F}_t = \boldsymbol{\Phi F}_{t-1} + e_t, \quad e_t \sim \mathcal{N}(\mathbf{0}_3, \boldsymbol{I}_3)$.
Measurement: $\boldsymbol{\Lambda F}_t + \boldsymbol{\Gamma q}_t + \boldsymbol{\eta}_t \quad , \quad \boldsymbol{\eta}_t \sim \mathcal{N}(\mathbf{0}_{12}, \boldsymbol{R})$

$$\boldsymbol{\Phi} = \begin{bmatrix} 0.85 & 0.00 & 0.00 \\ 0.00 & 0.80 & 0.00 \\ 0.00 & 0.00 & 0.75 \end{bmatrix} \quad , \quad \Lambda = \begin{bmatrix} 1 & 0 & 0 \\ 0.45 & 0 & 0 \\ 0.45 & 0 & 0 \\ 0.45 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0.45 & 0 \\ 0 & 0.45 & 0 \\ 0 & 0.45 & 0 \\ 0 & 0 & 1 \\ 0 & 0 & 0.45 \\ 0 & 0 & 0.45 \\ 0 & 0 & 0.45 \end{bmatrix} \quad , \quad \Gamma = \begin{bmatrix} 0 & 0 & 0 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0.8 & 0 & 0 \\ 1 & 0 & 0 & 0.8 & 0 & 0 \\ 1 & 0 & 0 & 0.8 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0.8 & 0 \\ 0 & 1 & 0 & 0 & 0.8 & 0 \\ 0 & 1 & 0 & 0 & 0.8 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0.8 \\ 0 & 0 & 1 & 0 & 0 & 0.8 \\ 0 & 0 & 1 & 0 & 0 & 0.8 \end{bmatrix}$$

$$\boldsymbol{R} = \begin{bmatrix} \boldsymbol{I}_4 & \mathbf{0}_{4\times4} & \mathbf{0}_{4\times4} \\ \mathbf{0}_{4\times4} & 1.25\boldsymbol{I}_4 & \mathbf{0}_{4\times4} \\ \mathbf{0}_{4\times4} & \mathbf{0}_{4\times4} & 1.5\boldsymbol{I}_4 \end{bmatrix}, \boldsymbol{q}_t = \begin{bmatrix} \max(F_{1t}, 0) \\ \max(F_{2t}, 0) \\ \max(F_{3t}, 0) \\ F_{1t} \times \mathbb{I}\{F_{2t} > 0\} \\ F_{2t} \times \mathbb{I}\{F_{3t} > 0\} \\ F_{3t} \times \mathbb{I}\{F_{1t} > 0\} \end{bmatrix}$$

(K) Nonlinear DGP, $J = 3$, $N = 12$.
Transition: $\boldsymbol{F}_t = \boldsymbol{\Phi F}_{t-1} + e_t, \quad e_t \sim \mathcal{N}(\mathbf{0}_3, \boldsymbol{I}_3)$.

Measurement: $\boldsymbol{\Lambda F}_t + \boldsymbol{\Gamma q}_t + \boldsymbol{\eta}_t \quad , \quad \boldsymbol{\eta}_t \sim \mathcal{N}(\mathbf{0}_{12}, \boldsymbol{R})$

$$\boldsymbol{\Phi} = \begin{bmatrix} 0.85 & 0.00 & 0.00 \\ 0.00 & 0.80 & 0.00 \\ 0.00 & 0.00 & 0.75 \end{bmatrix} \quad , \quad \Lambda = \begin{bmatrix} 1 & 0 & 0 \\ 0.45 & 0 & 0 \\ 0.45 & 0 & 0 \\ 0.45 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0.45 & 0 \\ 0 & 0.45 & 0 \\ 0 & 0.45 & 0 \\ 0 & 0 & 1 \\ 0 & 0 & 0.45 \\ 0 & 0 & 0.45 \\ 0 & 0 & 0.45 \end{bmatrix} \quad , \quad \Gamma = \begin{bmatrix} 0 & 0 & 0 & 0 & 0 & 0 \\ 0.8 & 0 & 0 & 1 & 0 & 0 \\ 0.8 & 0 & 0 & 1 & 0 & 0 \\ 0.8 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0.8 & 0 & 0 & 0 & 1 \\ 0 & 0.8 & 0 & 0 & 0 & 1 \\ 0 & 0.8 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0.8 & 0 & 1 & 0 \\ 0 & 0 & 0.8 & 0 & 1 & 0 \\ 0 & 0 & 0.8 & 0 & 1 & 0 \end{bmatrix}$$

$$\boldsymbol{R} = \begin{bmatrix} \boldsymbol{I}_4 & \mathbf{0}_{4\times4} & \mathbf{0}_{4\times4} \\ \mathbf{0}_{4\times4} & 1.25\boldsymbol{I}_4 & \mathbf{0}_{4\times4} \\ \mathbf{0}_{4\times4} & \mathbf{0}_{4\times4} & 1.5\boldsymbol{I}_4 \end{bmatrix}, \boldsymbol{q}_t = \begin{bmatrix} F_{1t}^2 \\ F_{2t}^2 \\ F_{3t}^2 \\ F_{1t}F_{2t} \\ F_{1t}F_{3t} \\ F_{2t}F_{3t} \end{bmatrix}$$

(L) Nonlinear DGP, $J = 3$, $N = 12$.
Transition: $\boldsymbol{F}_t = \boldsymbol{\Phi F}_{t-1} + e_t, \quad e_t \sim \mathcal{N}(\mathbf{0}_3, \boldsymbol{I}_3)$.
Measurement: $\boldsymbol{\Lambda F}_t + \boldsymbol{\Gamma q}_t + \boldsymbol{\eta}_t \quad , \quad \boldsymbol{\eta}_t \sim \mathcal{N}(\mathbf{0}_{12}, \boldsymbol{R})$

$$\boldsymbol{\Phi} = \begin{bmatrix} 0.85 & 0.00 & 0.00 \\ 0.00 & 0.80 & 0.00 \\ 0.00 & 0.00 & 0.75 \end{bmatrix} \quad , \quad \Lambda = \begin{bmatrix} 1 & 0 & 0 \\ 0.45 & 0 & 0 \\ 0.45 & 0 & 0 \\ 0.45 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0.45 & 0 \\ 0 & 0.45 & 0 \\ 0 & 0.45 & 0 \\ 0 & 0 & 1 \\ 0 & 0 & 0.45 \\ 0 & 0 & 0.45 \\ 0 & 0 & 0.45 \end{bmatrix} \quad , \quad \Gamma = \begin{bmatrix} 0 & 0 & 0 & 0 & 0 & 0 \\ 1.2 & 0 & 0 & 0.9 & 0 & 0 \\ 1.2 & 0 & 0 & 0.9 & 0 & 0 \\ 1.2 & 0 & 0 & 0.9 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1.2 & 0 & 0 & 0 & 0.9 \\ 0 & 1.2 & 0 & 0 & 0 & 0.9 \\ 0 & 1.2 & 0 & 0 & 0 & 0.9 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1.2 & 0 & 0.9 & 0 \\ 0 & 0 & 1.2 & 0 & 0.9 & 0 \\ 0 & 0 & 1.2 & 0 & 0.9 & 0 \end{bmatrix}$$

$$\boldsymbol{R} = \begin{bmatrix} 0.75\boldsymbol{I}_4 & \mathbf{0}_{4\times4} & \mathbf{0}_{4\times4} \\ \mathbf{0}_{4\times4} & 1.0\boldsymbol{I}_4 & \mathbf{0}_{4\times4} \\ \mathbf{0}_{4\times4} & \mathbf{0}_{4\times4} & 1.25\boldsymbol{I}_4 \end{bmatrix}, \boldsymbol{q}_t = \begin{bmatrix} \tanh(1.8 \times F_{1t}) \\ \tanh(1.8 \times F_{2t}) \\ \tanh(1.8 \times F_{3t}) \\ \tanh(1.2 \times F_{1t}F_{2t}) \\ \tanh(1.2 \times F_{1t}F_{3t}) \\ \tanh(1.2 \times F_{2t}F_{3t}) \end{bmatrix}$$

(M) Nonlinear DGP, $J = 3$, $N = 12$.
Transition: $\boldsymbol{F}_t = \boldsymbol{\Phi F}_{t-1} + e_t, \quad e_t \sim \mathcal{N}(\mathbf{0}_3, 0.04 \times \boldsymbol{I}_3)$.

Measurement: $\boldsymbol{\Lambda F}_t + \boldsymbol{\Gamma q}_t + \boldsymbol{\eta}_t \quad , \quad \boldsymbol{\eta}_t \sim \mathcal{N}(\mathbf{0}_{12}, \boldsymbol{R})$

$$\boldsymbol{\Phi} = \begin{bmatrix} 0.98 & 0.00 & 0.00 \\ 0.00 & 0.97 & 0.00 \\ 0.00 & 0.00 & 0.96 \end{bmatrix} \quad , \quad \Lambda = \begin{bmatrix} 1 & 0 & 0 \\ 0.45 & 0 & 0 \\ 0.45 & 0 & 0 \\ 0.45 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0.45 & 0 \\ 0 & 0.45 & 0 \\ 0 & 0.45 & 0 \\ 0 & 0 & 1 \\ 0 & 0 & 0.45 \\ 0 & 0 & 0.45 \\ 0 & 0 & 0.45 \end{bmatrix} \quad , \quad \Gamma = \begin{bmatrix} 0 & 0 & 0 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0.8 & 0 & 0 \\ 1 & 0 & 0 & 0.8 & 0 & 0 \\ 1 & 0 & 0 & 0.8 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0.8 & 0 \\ 0 & 1 & 0 & 0 & 0.8 & 0 \\ 0 & 1 & 0 & 0 & 0.8 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0.8 \\ 0 & 0 & 1 & 0 & 0 & 0.8 \\ 0 & 0 & 1 & 0 & 0 & 0.8 \end{bmatrix}$$

$$\boldsymbol{R} = \begin{bmatrix} 1.0\boldsymbol{I}_4 & \mathbf{0}_{4\times4} & \mathbf{0}_{4\times4} \\ \mathbf{0}_{4\times4} & 1.25\boldsymbol{I}_4 & \mathbf{0}_{4\times4} \\ \mathbf{0}_{4\times4} & \mathbf{0}_{4\times4} & 1.5\boldsymbol{I}_4 \end{bmatrix}, \begin{bmatrix} \max(F_{1t}, 0) \\ \max(F_{2t}, 0) \\ \max(F_{3t}, 0) \\ F_{1t} \times \mathbb{I}\{F_{2t} > 0\} \\ F_{2t} \times \mathbb{I}\{F_{3t} > 0\} \\ F_{3t} \times \mathbb{I}\{F_{1t} > 0\} \end{bmatrix}$$

(N) Nonlinear DGP, $J = 3$, $N = 12$.
Transition: $\boldsymbol{F}_t = \boldsymbol{\Phi F}_{t-1} + e_t, \quad e_t \sim \mathcal{N}(\mathbf{0}_3, 0.04 \times \boldsymbol{I}_3)$.
Measurement: $\boldsymbol{\Lambda F}_t + \boldsymbol{\Gamma q}_t + \boldsymbol{\eta}_t \quad , \quad \boldsymbol{\eta}_t \sim \mathcal{N}(\mathbf{0}_{12}, \boldsymbol{R})$

$$\boldsymbol{\Phi} = \begin{bmatrix} 0.98 & 0.00 & 0.00 \\ 0.00 & 0.97 & 0.00 \\ 0.00 & 0.00 & 0.96 \end{bmatrix} \quad , \quad \Lambda = \begin{bmatrix} 1 & 0 & 0 \\ 0.45 & 0 & 0 \\ 0.45 & 0 & 0 \\ 0.45 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0.45 & 0 \\ 0 & 0.45 & 0 \\ 0 & 0.45 & 0 \\ 0 & 0 & 1 \\ 0 & 0 & 0.45 \\ 0 & 0 & 0.45 \\ 0 & 0 & 0.45 \end{bmatrix} \quad , \quad \Gamma = \begin{bmatrix} 0 & 0 & 0 & 0 & 0 & 0 \\ 0.8 & 0 & 0 & 1 & 0 & 0 \\ 0.8 & 0 & 0 & 1 & 0 & 0 \\ 0.8 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0.8 & 0 & 0 & 0 & 1 \\ 0 & 0.8 & 0 & 0 & 0 & 1 \\ 0 & 0.8 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0.8 & 0 & 1 & 0 \\ 0 & 0 & 0.8 & 0 & 1 & 0 \\ 0 & 0 & 0.8 & 0 & 1 & 0 \end{bmatrix}$$

$$\boldsymbol{R} = \begin{bmatrix} \boldsymbol{I}_4 & \mathbf{0}_{4\times4} & \mathbf{0}_{4\times4} \\ \mathbf{0}_{4\times4} & 1.25\boldsymbol{I}_4 & \mathbf{0}_{4\times4} \\ \mathbf{0}_{4\times4} & \mathbf{0}_{4\times4} & 1.5\boldsymbol{I}_4 \end{bmatrix}, \boldsymbol{q}_t = \begin{bmatrix} F_{1t}^2 \\ F_{2t}^2 \\ F_{3t}^2 \\ F_{1t}F_{2t} \\ F_{1t}F_{3t} \\ F_{2t}F_{3t} \end{bmatrix}$$

(O) Nonlinear DGP, $J = 3$, $N = 12$.
Transition: $\boldsymbol{F}_t = \boldsymbol{\Phi F}_{t-1} + e_t, \quad e_t \sim \mathcal{N}(\mathbf{0}_3, 0.04 \times \boldsymbol{I}_3)$.

Measurement: $\boldsymbol{\Lambda F}_t + \boldsymbol{\Gamma q}_t + \boldsymbol{\eta}_t \quad , \quad \boldsymbol{\eta}_t \sim \mathcal{N}(\mathbf{0}_{12}, \boldsymbol{R})$

$$\boldsymbol{\Phi} = \begin{bmatrix} 0.98 & 0.00 & 0.00 \\ 0.00 & 0.97 & 0.00 \\ 0.00 & 0.00 & 0.96 \end{bmatrix} \quad , \quad \Lambda = \begin{bmatrix} 1 & 0 & 0 \\ 0.45 & 0 & 0 \\ 0.45 & 0 & 0 \\ 0.45 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0.45 & 0 \\ 0 & 0.45 & 0 \\ 0 & 0.45 & 0 \\ 0 & 0 & 1 \\ 0 & 0 & 0.45 \\ 0 & 0 & 0.45 \\ 0 & 0 & 0.45 \end{bmatrix} \quad , \quad \Gamma = \begin{bmatrix} 0 & 0 & 0 & 0 & 0 & 0 \\ 1.2 & 0 & 0 & 0.9 & 0 & 0 \\ 1.2 & 0 & 0 & 0.9 & 0 & 0 \\ 1.2 & 0 & 0 & 0.9 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1.2 & 0 & 0 & 0 & 0.9 \\ 0 & 1.2 & 0 & 0 & 0 & 0.9 \\ 0 & 1.2 & 0 & 0 & 0 & 0.9 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1.2 & 0 & 0.9 & 0 \\ 0 & 0 & 1.2 & 0 & 0.9 & 0 \\ 0 & 0 & 1.2 & 0 & 0.9 & 0 \end{bmatrix}$$

$$\boldsymbol{R} = \begin{bmatrix} 0.75\boldsymbol{I}_4 & \mathbf{0}_{4\times4} & \mathbf{0}_{4\times4} \\ \mathbf{0}_{4\times4} & 1.0\boldsymbol{I}_4 & \mathbf{0}_{4\times4} \\ \mathbf{0}_{4\times4} & \mathbf{0}_{4\times4} & 1.25\boldsymbol{I}_4 \end{bmatrix}, \boldsymbol{q}_t = \begin{bmatrix} \tanh(1.8 \times F_{1t}) \\ \tanh(1.8 \times F_{2t}) \\ \tanh(1.8 \times F_{3t}) \\ \tanh(1.2 \times F_{1t}F_{2t}) \\ \tanh(1.2 \times F_{1t}F_{3t}) \\ \tanh(1.2 \times F_{2t}F_{3t}) \end{bmatrix}$$

(P) Nonlinear DGP, $J = 3$, $N = 12$.
Transition: $\boldsymbol{F}_t = \boldsymbol{\Phi F}_{t-1} + e_t, \quad e_t \sim \mathcal{N}(\mathbf{0}_3, 0.08 \times \boldsymbol{I}_3)$.
Measurement: $\boldsymbol{\Lambda F}_t + \boldsymbol{\Gamma q}_t + \boldsymbol{\eta}_t \quad , \quad \boldsymbol{\eta}_t \sim \mathcal{N}(\mathbf{0}_{12}, \boldsymbol{R})$

$$\boldsymbol{\Phi} = \begin{bmatrix} 0.94 & 0.00 & 0.00 \\ 0.00 & 0.92 & 0.00 \\ 0.00 & 0.00 & 0.90 \end{bmatrix} \quad , \quad \Lambda = \begin{bmatrix} 1 & 0 & 0 \\ 0.25 & 0 & 0 \\ 0.25 & 0 & 0 \\ 0.25 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0.25 & 0 \\ 0 & 0.25 & 0 \\ 0 & 0.25 & 0 \\ 0 & 0 & 1 \\ 0 & 0 & 0.25 \\ 0 & 0 & 0.25 \\ 0 & 0 & 0.25 \end{bmatrix} \quad , \quad \Gamma = \begin{bmatrix} 0 & 0 & 0 & 0 & 0 & 0 \\ 1.4 & 0 & 0 & 1.2 & 0 & 0 \\ 1.4 & 0 & 0 & 1.2 & 0 & 0 \\ 1.4 & 0 & 0 & 1.2 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1.4 & 0 & 0 & 0 & 1.2 \\ 0 & 1.4 & 0 & 0 & 0 & 1.2 \\ 0 & 1.4 & 0 & 0 & 0 & 1.2 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1.4 & 0 & 1.2 & 0 \\ 0 & 0 & 1.4 & 0 & 1.2 & 0 \\ 0 & 0 & 1.4 & 0 & 1.2 & 0 \end{bmatrix}$$

$$\boldsymbol{R} = \begin{bmatrix} 0.35\boldsymbol{I}_4 & \mathbf{0}_{4\times4} & \mathbf{0}_{4\times4} \\ \mathbf{0}_{4\times4} & 0.45\boldsymbol{I}_4 & \mathbf{0}_{4\times4} \\ \mathbf{0}_{4\times4} & \mathbf{0}_{4\times4} & 0.55\boldsymbol{I}_4 \end{bmatrix}, \boldsymbol{q}_t = \begin{bmatrix} F_{1t}^2 \\ F_{2t}^2 \\ F_{3t}^2 \\ F_{1t}F_{2t} \\ F_{1t}F_{3t} \\ F_{2t}F_{3t} \end{bmatrix}$$

(Q) Nonlinear DGP, $J = 3$, $N = 12$.
Transition: $\boldsymbol{F}_t = \boldsymbol{\Phi F}_{t-1} + e_t, \quad e_t \sim \mathcal{N}(\mathbf{0}_3, 0.08 \times \boldsymbol{I}_3)$.

Measurement: $\boldsymbol{\Lambda F}_t + \boldsymbol{\Gamma q}_t + \boldsymbol{\eta}_t \quad , \quad \boldsymbol{\eta}_t \sim \mathcal{N}(\mathbf{0}_{12}, \boldsymbol{R})$

$$\boldsymbol{\Phi} = \begin{bmatrix} 0.94 & 0.00 & 0.00 \\ 0.00 & 0.92 & 0.00 \\ 0.00 & 0.00 & 0.90 \end{bmatrix} \quad , \quad \Lambda = \begin{bmatrix} 1 & 0 & 0 \\ 0.25 & 0 & 0 \\ 0.25 & 0 & 0 \\ 0.25 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0.25 & 0 \\ 0 & 0.25 & 0 \\ 0 & 0.25 & 0 \\ 0 & 0 & 1 \\ 0 & 0 & 0.25 \\ 0 & 0 & 0.25 \\ 0 & 0 & 0.25 \end{bmatrix} \quad , \quad \Gamma = \begin{bmatrix} 0 & 0 & 0 & 0 & 0 & 0 \\ 1.5 & 0 & 0 & 0.9 & 0 & 0 \\ 1.5 & 0 & 0 & 0.9 & 0 & 0 \\ 1.5 & 0 & 0 & 0.9 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1.5 & 0 & 0 & 0 & 0.9 \\ 0 & 1.5 & 0 & 0 & 0 & 0.9 \\ 0 & 1.5 & 0 & 0 & 0 & 0.9 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1.5 & 0 & 0.9 & 0 \\ 0 & 0 & 1.5 & 0 & 0.9 & 0 \\ 0 & 0 & 1.5 & 0 & 0.9 & 0 \end{bmatrix}$$

$$\boldsymbol{R} = \begin{bmatrix} 0.35\boldsymbol{I}_4 & \mathbf{0}_{4\times4} & \mathbf{0}_{4\times4} \\ \mathbf{0}_{4\times4} & 0.45\boldsymbol{I}_4 & \mathbf{0}_{4\times4} \\ \mathbf{0}_{4\times4} & \mathbf{0}_{4\times4} & 0.55\boldsymbol{I}_4 \end{bmatrix}, \boldsymbol{q}_t = \begin{bmatrix} |F_{1t}| \\ |F_{2t}| \\ |F_{3t}| \\ |F_{1t} + F_{2t}| \\ |F_{1t} + F_{3t}| \\ |F_{2t} + F_{3t}| \end{bmatrix}$$

(R) Nonlinear DGP, $J = 3$, $N = 12$.
Transition: $\boldsymbol{F}_t = \boldsymbol{\Phi F}_{t-1} + e_t, \quad e_t \sim \mathcal{N}(\mathbf{0}_3, 0.08 \times \boldsymbol{I}_3)$.
Measurement: $\boldsymbol{\Lambda F}_t + \boldsymbol{\Gamma q}_t + \boldsymbol{\eta}_t \quad , \quad \boldsymbol{\eta}_t \sim \mathcal{N}(\mathbf{0}_{12}, \boldsymbol{R})$

$$\boldsymbol{\Phi} = \begin{bmatrix} 0.94 & 0.00 & 0.00 \\ 0.00 & 0.92 & 0.00 \\ 0.00 & 0.00 & 0.90 \end{bmatrix} \quad , \quad \Lambda = \begin{bmatrix} 1 & 0 & 0 \\ 0.25 & 0 & 0 \\ 0.25 & 0 & 0 \\ 0.25 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0.25 & 0 \\ 0 & 0.25 & 0 \\ 0 & 0.25 & 0 \\ 0 & 0 & 1 \\ 0 & 0 & 0.25 \\ 0 & 0 & 0.25 \\ 0 & 0 & 0.25 \end{bmatrix} \quad , \quad \Gamma = \begin{bmatrix} 0 & 0 & 0 & 0 & 0 & 0 \\ 2 & 0 & 0 & 1 & 0 & 0 \\ 2 & 0 & 0 & 1 & 0 & 0 \\ 2 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 2 & 0 & 0 & 0 & 1 \\ 0 & 2 & 0 & 0 & 0 & 1 \\ 0 & 2 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 2 & 0 & 1 & 0 \\ 0 & 0 & 2 & 0 & 1 & 0 \\ 0 & 0 & 2 & 0 & 1 & 0 \end{bmatrix}$$

$$\boldsymbol{R} = \begin{bmatrix} 0.35\boldsymbol{I}_4 & \mathbf{0}_{4\times4} & \mathbf{0}_{4\times4} \\ \mathbf{0}_{4\times4} & 0.45\boldsymbol{I}_4 & \mathbf{0}_{4\times4} \\ \mathbf{0}_{4\times4} & \mathbf{0}_{4\times4} & 0.55\boldsymbol{I}_4 \end{bmatrix}, \boldsymbol{q}_t = \begin{bmatrix} \frac{1}{1+\exp(-2.5F_{1t})} \\ \frac{1}{1+\exp(-2.5F_{2t})} \\ \frac{1}{1+\exp(-2.5F_{3t})} \\ \frac{1}{1+\exp(-2.5F_{1t}F_{2t})} \\ \frac{1}{1+\exp(-2.5F_{1t}F_{3t})} \\ \frac{1}{1+\exp(-2.5F_{2t}F_{3t})} \end{bmatrix}$$

3.2 Prior settings

The prior settings differ to those for the real data application. However, we observe that these settings often have little impact on the results.

Following Huber et al. (2020), the variance is calibrated as $\mu_{bs}^{(i)} \sim \mathcal{N}\left(0, \sigma_{\mu i}^2\right)$, with $\sigma_{\mu i} = \frac{\max(X_{\cdot i}) - \min(X_{\cdot i})}{2\sqrt{\nu S}}$ where S is the number of trees and ν reflects the prior confidence. For the error variance σ_i^2, we impose a conjugate scaled inverse-χ^2 prior: $\sigma_i^2 \sim \nu_i \xi_i / \chi_{\nu_i}^2$, with $\nu_i = 3$ and ξ_i defined such that $P(\sigma_i < \hat{\sigma}_i) = 0.9$ where $\hat{\sigma}_i$ is estimated from a preliminary regression.[4]

Table 8: Summary of prior hyperparameters for measurement equation.

Description	Parameters	Hyperparameters
Tree depth	$\alpha(1+d)^{-\beta}$	$\alpha = 0.95$, $\beta = 2$
Terminal node priors	$\mu_{bs} \sim \mathcal{N}(0, \sigma_{\mu i}^2)$	$\sigma_{\mu i} = \frac{\max(X_{\cdot i}) - \min(X_{\cdot i})}{2\sqrt{\nu S}}$
Error variances	$\sigma_i^2 \sim \nu_i \xi_i / \chi_{\nu_i}^2$	$\nu_i = 3$, $P(\sigma_i < \hat{\sigma}_i) = 0.9$
Number of trees	–	$S = 100$

Notes: Prior distributions follow Chipman et al. (2010) and Huber et al. (2020).

The transition equation coefficients have a diffuse prior, and the covariance matrix of the transition equation errors has an inverse Wishart prior. The inverse Wishart prior on Ω has degrees of freedom equal to the number of factors plus 10, and scale equal to the sample covariance matrix of the residuals from OLS regression of principal components on their lags, multiplied by 9 to set the prior mean equal to the sample covariance of residuals.[5]

3.3 Results, outcome forecasts

[4] Chipman et al. (2010) use the sample variance of the residuals from a linear regression of the outcome on the predictors. We use the sample variance of the residuals from a linear regression of the relevant outcome $X_{i,\cdot}$ on initial factor estimates from PCA. Alternatives include the sample variance of the relevant outcome and the fifth percentile of the sample variances across all outcome variables, as suggested by Moran (2019).

[5] Similar results are obtained with a conjugate normal inverse Wishart priors the transition equation in which the coefficients of factor lags have prior mean zero and variance equal to 10 divided by the sample variance of initial factor estimates from PCA. For robustness, we also consider a Minnesota prior on the coefficients, i.e. a normal prior with mean zero, and variance equal to 10 divided by the product of the sample variance of the initial factor estimate and the square of the order of the lag, for greater shrinkage on coefficients of higher order lags. We also implemented, but do not include results for, a transition equation specification with stochastic volatility.

Table 9: RMSE relative to RW by DGP at horizon 1

DGP	AR1	FABART	LDFM	RW
DGP A	0.8190	0.7850	0.7820	1.0000
DGP B	0.7850	0.7630	0.7660	1.0000
DGP C	0.8650	0.9050	0.8300	1.0000
DGP D	0.8230	0.7900	0.7940	1.0000
DGP E	0.8210	0.7850	0.7840	1.0000
DGP F	0.8210	0.7840	0.7830	1.0000
DGP G	0.8640	0.8100	0.8040	1.0000
DGP H	0.8600	0.8050	0.8000	1.0000
DGP I	0.8660	0.8050	0.8040	1.0000
DGP J	0.8730	0.8290	0.8380	1.0000
DGP K	0.9220	0.9520	0.9360	1.0000
DGP L	0.9170	0.8820	0.8810	1.0000
DGP M	0.8790	0.7800	0.8150	1.0000
DGP N	0.8210	0.7550	0.7940	1.0000
DGP O	0.8840	0.7700	0.8010	1.0000
DGP P	0.9600	0.9320	0.9390	1.0000
DGP Q	0.9250	0.9150	0.9180	1.0000
DGP R	0.9040	0.8300	0.8360	1.0000

Table 10: RMSE relative to RW by DGP at horizon 3

DGP	AR1	FABART	LDFM	RW
DGP A	0.7600	0.7560	0.7550	1.0000
DGP B	0.7460	0.7430	0.7470	1.0000
DGP C	0.7550	0.7700	0.7550	1.0000
DGP D	0.7690	0.7680	0.7700	1.0000
DGP E	0.7900	0.7870	0.7900	1.0000
DGP F	0.7910	0.7900	0.7920	1.0000
DGP G	0.8220	0.7920	0.7970	1.0000
DGP H	0.8380	0.7980	0.8050	1.0000
DGP I	0.8280	0.7930	0.7990	1.0000
DGP J	0.8570	0.8450	0.8450	1.0000
DGP K	0.8810	0.8990	0.8900	1.0000
DGP L	0.8550	0.8460	0.8430	1.0000
DGP M	0.8930	0.7720	0.8180	1.0000
DGP N	0.8980	0.8190	0.8600	1.0000
DGP O	0.9150	0.7620	0.8210	1.0000
DGP P	0.8990	0.8820	0.8800	1.0000
DGP Q	0.8870	0.8710	0.8770	1.0000
DGP R	0.9190	0.8380	0.8590	1.0000

Table 11: RMSE relative to RW by DGP at horizon 6

DGP	AR1	FABART	LDFM	RW
DGP A	0.7130	0.7150	0.7140	1.0000
DGP B	0.7060	0.7080	0.7090	1.0000
DGP C	0.7340	0.7200	0.7350	1.0000
DGP D	0.6560	0.6620	0.6590	1.0000
DGP E	0.6760	0.6800	0.6790	1.0000
DGP F	0.6760	0.6800	0.6780	1.0000
DGP G	0.7520	0.7480	0.7470	1.0000
DGP H	0.7810	0.7670	0.7730	1.0000
DGP I	0.7590	0.7520	0.7520	1.0000
DGP J	0.7640	0.7690	0.7670	1.0000
DGP K	0.7620	0.7620	0.7620	1.0000
DGP L	0.7760	0.7810	0.7780	1.0000
DGP M	0.9250	0.8190	0.8790	1.0000
DGP N	0.8690	0.8050	0.8450	1.0000
DGP O	0.9650	0.8080	0.8930	1.0000
DGP P	0.8150	0.8310	0.8110	1.0000
DGP Q	0.8160	0.8160	0.8100	1.0000
DGP R	0.8720	0.8270	0.8410	1.0000

Table 12: RMSE relative to RW by DGP at horizon 12

DGP	AR1	FABART	LDFM	RW
DGP A	0.7530	0.7520	0.7530	1.0000
DGP B	0.7370	0.7360	0.7400	1.0000
DGP C	0.7900	0.7710	0.7900	1.0000
DGP D	0.7900	0.7880	0.7900	1.0000
DGP E	0.7850	0.7850	0.7850	1.0000
DGP F	0.7860	0.7880	0.7870	1.0000
DGP G	0.7270	0.7290	0.7280	1.0000
DGP H	0.7360	0.7370	0.7360	1.0000
DGP I	0.7280	0.7300	0.7280	1.0000
DGP J	0.7290	0.7310	0.7280	1.0000
DGP K	0.7540	0.7620	0.7550	1.0000
DGP L	0.7390	0.7400	0.7390	1.0000
DGP M	0.8720	0.8090	0.8580	1.0000
DGP N	0.8110	0.7760	0.8030	1.0000
DGP O	0.9200	0.8080	0.8880	1.0000
DGP P	0.8190	0.8230	0.8180	1.0000
DGP Q	0.7580	0.7620	0.7580	1.0000
DGP R	0.8300	0.8160	0.8230	1.0000

Table 13: RMSE relative to RW by DGP at horizon 24

DGP	AR1	FABART	LDFM	RW
DGP A	0.6970	0.7000	0.6970	1.0000
DGP B	0.7090	0.7080	0.7080	1.0000
DGP C	0.6790	0.6710	0.6790	1.0000
DGP D	0.7500	0.7500	0.7500	1.0000
DGP E	0.7560	0.7530	0.7560	1.0000
DGP F	0.7580	0.7570	0.7580	1.0000
DGP G	0.6980	0.6990	0.6980	1.0000
DGP H	0.7060	0.7080	0.7070	1.0000
DGP I	0.6980	0.6990	0.6990	1.0000
DGP J	0.7400	0.7420	0.7400	1.0000
DGP K	0.7540	0.7580	0.7540	1.0000
DGP L	0.7200	0.7190	0.7200	1.0000
DGP M	0.8230	0.8040	0.8220	1.0000
DGP N	0.7700	0.7700	0.7700	1.0000
DGP O	0.8380	0.7990	0.8340	1.0000
DGP P	0.7220	0.7280	0.7220	1.0000
DGP Q	0.7190	0.7190	0.7190	1.0000
DGP R	0.7610	0.7610	0.7600	1.0000

Table 14: RMSE relative to RW for X1 by DGP at horizon 1

DGP	AR1	FABART	LDFM	RW
DGP A	0.8960	0.8640	0.8610	1.0000
DGP B	0.8930	0.8730	0.8850	1.0000
DGP C	0.8860	0.8730	0.8690	1.0000
DGP D	0.8730	0.8250	0.8400	1.0000
DGP E	0.9010	0.8450	0.8490	1.0000
DGP F	0.8860	0.8310	0.8280	1.0000
DGP G	0.9320	0.8590	0.8450	1.0000
DGP H	0.9130	0.8400	0.8310	1.0000
DGP I	0.9270	0.8470	0.8390	1.0000
DGP J	0.8580	0.7550	0.7750	1.0000
DGP K	0.8640	1.0020	0.9790	1.0000
DGP L	0.9300	0.8350	0.8460	1.0000
DGP M	0.9450	0.7870	0.8460	1.0000
DGP N	0.7980	0.7470	0.8210	1.0000
DGP O	0.8040	0.7360	0.7500	1.0000
DGP P	0.8770	0.9790	1.0220	1.0000
DGP Q	0.8950	0.9920	0.9930	1.0000
DGP R	0.9420	0.8210	0.8510	1.0000

Table 15: RMSE relative to RW for X3 by DGP at horizon 1

DGP	AR1	FABART	LDFM	RW
DGP A	0.9110	0.8660	0.8550	1.0000
DGP B	0.8950	0.8510	0.8460	1.0000
DGP C	0.8970	0.9030	0.8370	1.0000
DGP D	0.8680	0.8290	0.8330	1.0000
DGP E	0.8660	0.8160	0.8170	1.0000
DGP F	0.8580	0.8010	0.8010	1.0000
DGP G	0.8480	0.7540	0.7350	1.0000
DGP H	0.8380	0.7310	0.7200	1.0000
DGP I	0.8450	0.7430	0.7280	1.0000
DGP J	0.8900	0.8890	0.8770	1.0000
DGP K	0.9270	0.9420	0.9120	1.0000
DGP L	0.8770	0.8840	0.8790	1.0000
DGP M	0.9210	0.7710	0.8010	1.0000
DGP N	0.8190	0.7220	0.7460	1.0000
DGP O	0.9250	0.7870	0.8210	1.0000
DGP P	1.0870	0.9620	0.9720	1.0000
DGP Q	0.9470	0.9280	0.9360	1.0000
DGP R	0.9390	0.9380	0.9260	1.0000

Table 16: FABART posterior predictive interval coverage and length for observed outcomes

DGP	Coverage All	Length All	Coverage H12	Length H12	Coverage H24	Length H24
DGP A	0.9758	4.5613	0.9815	4.5731	0.9760	4.5715
DGP B	0.9683	4.3344	0.9705	4.3416	0.9690	4.3347
DGP C	0.9665	4.0481	0.9645	4.0534	0.9705	4.0546
DGP D	0.9532	4.1355	0.9445	4.1442	0.9340	4.1464
DGP E	0.9557	4.1069	0.9415	4.1224	0.9490	4.1182
DGP F	0.9570	4.1234	0.9470	4.1374	0.9480	4.1329
DGP G	0.9498	4.0834	0.9317	4.1155	0.9492	4.1303
DGP H	0.9517	4.0432	0.9375	4.0742	0.9508	4.0862
DGP I	0.9501	4.0585	0.9342	4.0835	0.9500	4.1024
DGP J	0.9515	4.0130	0.9567	4.0590	0.9433	4.0800
DGP K	0.9534	3.9346	0.9650	3.9649	0.9383	3.9649
DGP L	0.9445	3.7563	0.9408	3.7896	0.9492	3.7982
DGP M	0.9435	3.9258	0.9625	3.8759	0.9417	4.0152
DGP N	0.9435	4.0068	0.9492	3.9790	0.9492	4.0835
DGP O	0.9450	3.7833	0.9592	3.7257	0.9542	3.8932
DGP P	0.9441	3.9798	0.9367	4.0103	0.9442	4.0571
DGP Q	0.9426	3.9651	0.9317	3.9969	0.9458	4.0292
DGP R	0.9389	3.9590	0.9367	3.9552	0.9408	4.0690

Table 17: FABART posterior predictive interval coverage by horizon

DGP	H1	H3	H6	H12	H24
DGP A	0.9800	0.9680	0.9750	0.9815	0.9760
DGP B	0.9700	0.9655	0.9675	0.9705	0.9690
DGP C	0.9690	0.9640	0.9630	0.9645	0.9705
DGP D	0.9455	0.9375	0.9695	0.9445	0.9340
DGP E	0.9460	0.9435	0.9785	0.9415	0.9490
DGP F	0.9455	0.9430	0.9770	0.9470	0.9480
DGP G	0.9633	0.9525	0.9633	0.9317	0.9492
DGP H	0.9633	0.9542	0.9467	0.9375	0.9508
DGP I	0.9608	0.9533	0.9583	0.9342	0.9500
DGP J	0.9767	0.9550	0.9633	0.9567	0.9433
DGP K	0.9775	0.9650	0.9650	0.9650	0.9383
DGP L	0.9500	0.9533	0.9342	0.9408	0.9492
DGP M	0.9583	0.9583	0.9542	0.9625	0.9417
DGP N	0.9592	0.9500	0.9450	0.9492	0.9492
DGP O	0.9617	0.9667	0.9567	0.9592	0.9542
DGP P	0.9450	0.9550	0.9325	0.9367	0.9442
DGP Q	0.9425	0.9475	0.9333	0.9317	0.9458
DGP R	0.9542	0.9617	0.9342	0.9367	0.9408

3.4 Factor Estimation

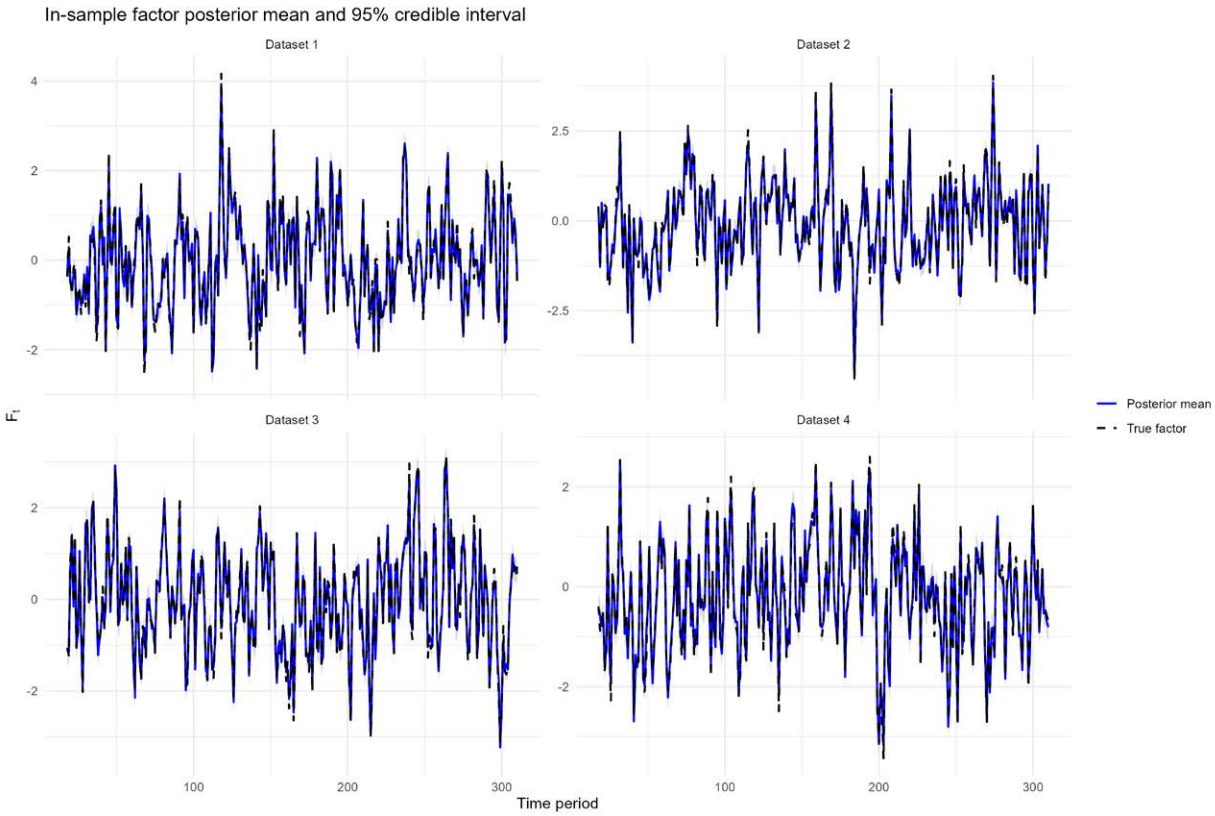

Note: In-sample posterior mean estimates (in blue) of the latent factor for one replication of each of the Monte Carlo studies for DGP A to DGP D. The true factor values are represented by the dashed black line. 95% credible intervals are in grey.

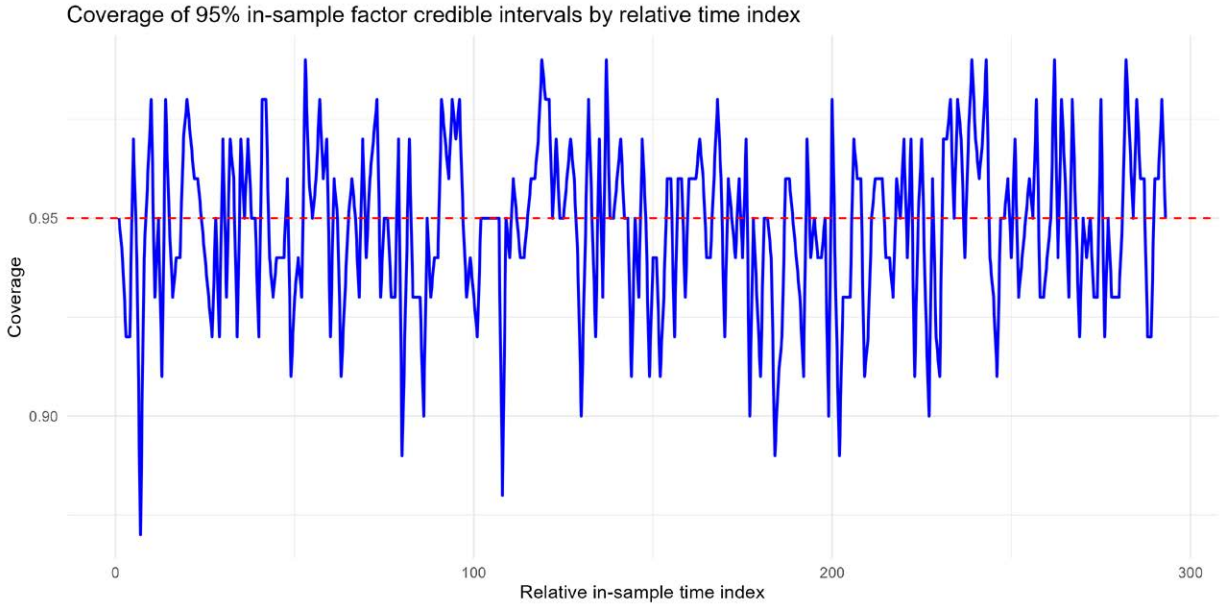

Note: In-sample posterior credible interval coverage by time index, averaged over 100 replications of DGP A.

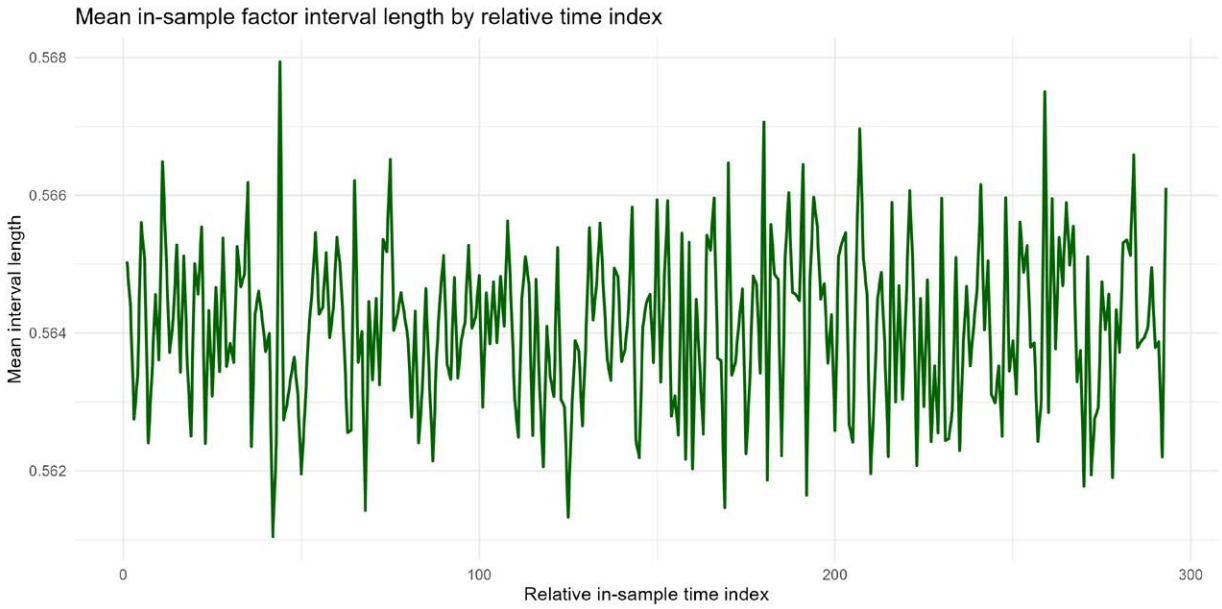

Note: In-sample posterior credible interval length by time index, averaged over 100 replications of DGP A.

3.5 In-sample factor estimation

Table 18: In-sample factor recovery summary by DGP

DGP	RMSE	MAE	Corr	Coverage	Length
DGP A	0.1441	0.1150	0.9924	0.9479	0.5649
DGP B	0.2047	0.1634	0.9849	0.9423	0.7806
DGP C	0.2988	0.2127	0.9673	0.8562	0.7400
DGP D	0.2513	0.1867	0.9774	0.8510	0.6579
DGP E	0.1750	0.1395	0.9891	0.9416	0.6666
DGP F	0.1770	0.1411	0.9889	0.9411	0.6731
DGP G	1.1881	0.9295	0.8546	0.3685	1.0625
DGP H	1.1411	0.9055	0.8913	0.3568	1.0354
DGP I	1.1567	0.9041	0.8791	0.3722	1.0405
DGP J	1.0124	0.7848	0.9258	0.3576	0.8450
DGP K	1.6081	1.2423	0.3279	0.2128	0.8491
DGP L	1.0701	0.7833	0.8968	0.4152	0.8314
DGP M	0.4983	0.3920	0.9127	0.5829	0.7667
DGP N	0.6342	0.5061	0.7134	0.5084	0.8606
DGP O	0.4929	0.3799	0.9126	0.5710	0.6964
DGP P	0.6890	0.5393	0.4708	0.4458	0.8155
DGP Q	0.6624	0.5330	0.4639	0.4058	0.7663
DGP R	0.4382	0.3281	0.8333	0.6918	0.8011

Table 19: In-sample factor recovery Coverage by DGP

DGP	F1	Average	F2	F3
DGP A	0.9479	0.9479	-	-
DGP B	0.9423	0.9423	-	-
DGP C	0.8562	0.8562	-	-
DGP D	0.8510	0.8510	-	-
DGP E	0.9416	0.9416	-	-
DGP F	0.9411	0.9411	-	-
DGP G	0.3569	0.3685	0.3794	0.3692
DGP H	0.3467	0.3568	0.3775	0.3461
DGP I	0.3549	0.3722	0.3885	0.3731
DGP J	0.2899	0.3576	0.3709	0.4119
DGP K	0.1712	0.2128	0.2060	0.2611
DGP L	0.3720	0.4152	0.4380	0.4355
DGP M	0.4614	0.5829	0.6325	0.6547
DGP N	0.4107	0.5084	0.5571	0.5573
DGP O	0.4438	0.5710	0.6301	0.6389
DGP P	0.3849	0.4458	0.4507	0.5016
DGP Q	0.3729	0.4058	0.4127	0.4318
DGP R	0.6666	0.6918	0.7350	0.6738

3.6 Comparison of factor estimates to linear benchmark

Table 20: In-sample factor recovery: FABART (posterior mean) versus linear DFM

	RMSE		MAE		Correlation	
DGP	FABART	LDFM	FABART	LDFM	FABART	LDFM
DGP A	0.1441	0.1626	0.1150	0.1298	0.9924	0.9902
DGP B	0.2047	0.1199	0.1634	0.0945	0.9849	0.9947
DGP C	0.2988	0.1927	0.2127	0.1377	0.9673	0.9864
DGP D	0.2513	0.2661	0.1867	0.1983	0.9774	0.9744
DGP E	0.1750	0.1913	0.1395	0.1524	0.9891	0.9867
DGP F	0.1770	0.1906	0.1411	0.1518	0.9889	0.9868
DGP G	1.1881	0.8862	0.9295	0.6805	0.8546	0.8376
DGP H	1.1411	0.7521	0.9055	0.5989	0.8913	0.8787
DGP I	1.1567	0.8138	0.9041	0.6305	0.8791	0.8640
DGP J	1.0124	0.6012	0.7848	0.4637	0.9258	0.9291
DGP K	1.6081	1.5035	1.2423	1.1582	0.3279	0.3873
DGP L	1.0701	0.7091	0.7833	0.5538	0.8968	0.9008
DGP M	0.4983	0.4067	0.3920	0.3184	0.9127	0.8233
DGP N	0.6342	0.5968	0.5061	0.4615	0.7134	0.5938
DGP O	0.4929	0.3772	0.3799	0.2915	0.9126	0.8574
DGP P	0.6890	0.5956	0.5393	0.4584	0.4708	0.4815
DGP Q	0.6624	0.5971	0.5330	0.4601	0.4639	0.4811
DGP R	0.4382	0.3927	0.3281	0.2876	0.8333	0.8180

References

Chipman, Hugh A., Edward I. George and Robert E. McCulloch. (2010). "BART: Bayesian additive regression trees". *The Annals of Applied Statistics*, 4(1), pp. 266-298. https://doi.org/10.1214/09-AOAS285

Gelman, Andrew, and Donald B. Rubin. (1992). "Inference from iterative simulation using multiple sequences". *Statistical Science*, 7(4), pp. 457-472. https://doi.org/10.1214/ss/1177011136

Geweke, John. (1992). "Evaluating the accuracy of sampling-based approaches to the calculating of posterior moments". In Bernardo, Berger, Dawid and Smith (eds.), *Bayesian Statistics 4: Proceedings of the Fourth Valencia International Meeting, Dedicated to the memory of Morris H. DeGroot, 1931 -1989*. Oxford Science Publications, pp. 696-700. https://doi.org/10.1093/oso/9780198522669.003.0010

Heidelberger, Philip, and Peter D. Welch. (1983). "Simulation run length control in the presence of an initial transient". *Operations Research*, 31(6), pp. 1109-1144. https://doi.org/10.1287/opre.31.6.1109

Huber, Florian, Gary Koop, Luca Onorante, Michael Pfarrhofer, and Josef Schreiner. (2023). "Nowcasting in a pandemic using non-parametric mixed frequency VARs". *Journal of Econometrics*, 232(1), pp. 52-69. https://doi.org/10.1016/j.jeconom.2020.11.006

Moran, Gemma E. (2019). *Bayesian approaches for modeling variation* [Ph. D. Tesis].

Plummer, Martin, Nicky Best, Kate Cowles, Karen Vines. (2006). "CODA: convergence diagnosis and output analysis for MCMC". *R news*, 6(1), pp. 7-11. https://journal.r-project.org/articles/RN-2006-002/

Vats, Dootika, and Christina Knudson. (2021). "Revisiting the Gelman-Rubin diagnostic". *Statistical Science*, 36(4), pp. 518-529. https://doi.org/10.1214/20-STS812

BANCO DE ESPAÑA PUBLICATIONS

WORKING PAPERS

2520 ALICIA AGUILAR: Beyond fragmentation: unraveling the drivers of yield divergence in the euro area.
2521 RUBÉN DOMÍNGUEZ-DÍAZ and DONGHAI ZHANG: The macroeconomic effects of unemployment insurance extensions: A policy rule-based identification approach.
2522 IRMA ALONSO-ALVAREZ, MARINA DIAKONOVA and JAVIER J. PÉREZ: Rethinking GPR: The sources of geopolitical risk.
2523 ALBERTO MARTÍN, SERGIO MAYORDOMO and VICTORIA VANASCO: Banks vs. Firms: Who Benefits from Credit Guarantees?
2524 SUMIT AGARWAL, SERGIO MAYORDOMO, MARÍA RODRÍGUEZ-MORENO and EMANUELE TARANTINO: Household Heterogeneity and the Lending Channel of Monetary Policy.
2525 DIEGO BONELLI, BERARDINO PALAZZO, and RAM YAMARTHY: Good inflation, bad inflation: implications for risky asset prices.
2526 STÉPHANE BONHOMME and ANGELA DENIS: Fixed Effects and Beyond. Bias Reduction, Groups, Shrinkage and Factors in Panel Data.
2527 ÁLVARO FERNÁNDEZ-GALLARDO and IVÁN PAYÁ: Public debt burden and crisis severity.
2528 GALO NUÑO: Three Theories of Natural Rate Dynamics.
2529 GALO NUÑO, PHILIPP RENNER and SIMON SCHEIDEGGER: Monetary policy with persistent supply shocks.
2530 MIGUEL ACOSTA-HENAO, MARÍA ALEJANDRA AMADO, MONTSERRAT MARTÍ and DAVID PÉREZ-REYNA: Heterogeneous UIPDs across Firms: Spillovers from U.S. Monetary Policy Shocks.
2531 LUIS HERRERA and JESÚS VÁZQUEZ: Learning from news.
2532 MORTEZA GHOMI, JOCHEN MANKART, RIGAS OIKONOMOU and ROMANOS PRIFTIS: Debt maturity and government spending multipliers.
2533 MARINA DIAKONOVA, CORINNA GHIRELLI and JAVIER J. PÉREZ: Political polarization in Europe.
2534 NICOLÁS FORTEZA and SERGIO PUENTE: Measuring non-workers' labor market attachment with machine learning.
2535 GERGELY GANICS and LLUC PUIG CODINA: Simple Tests for the Correct Specification of Conditional Predictive Densities.
2536 HENRIQUE S. BASSO and OMAR RACHEDI: Robot adoption and inflation dynamics.
2537 PABLO GARCIA, PASCAL JACQUINOT, ČRT LENARČIČ, KOSTAS MAVROMATIS, NIKI PAPADOPOULOU and EDGAR SILGADO-GÓMEZ: Green transition in the Euro area: domestic and global factors.
2538 MARÍA ALEJANDRA AMADO, CARLOS BURGA and JOSÉ E. GUTIÉRREZ: Cross-border spillovers of bank regulations: Evidence of a trade channel.
2539 ALEJANDRO CASADO and DAVID MARTÍNEZ-MIERA: Banks' specialization and private information.
2540 CHRISTIAN E. CASTRO, ÁNGEL ESTRADA GARCÍA and GONZALO FERNÁNDEZ DIONIS: Diversifying sovereign risk in the Euro area: empirical analysis of different policy proposals.
2541 RAFAEL GUNTIN and FEDERICO KOCHEN: The Origins of Top Firms.
2542 ÁLVARO FERNÁNDEZ-GALLARDO: Natural disasters, economic activity, and property insurance: evidence from weekly U.S. state-level data.
2543 JOSÉ ELÍAS GALLEGOS, ESTEBAN GARCÍA-MIRALLES, IVÁN KATARYNIUK and SUSANA PÁRRAGA RODRÍGUEZ: Fiscal Announcements and Households' Beliefs: Evidence from the Euro Area.
2544 LUIS HERRERA, MARA PIROVANO and VALERIO SCALONE: From risk to buffer: Calibrating the positive neutral CCyB rate.
2545 ESTEBAN GARCÍA-MIRALLES et al.: Fiscal drag in theory and in practice: A European perspective.
2546 TATSURO SENGA and IACOPO VAROTTO: Investment Irreversibility in a Granular World.
2547 OLYMPIA BOVER, NEZIH GUNER, YULIYA KULIKOVA, ALESSANDRO RUGGIERI and CARLOS SANZ: Family-friendly policies and fertility: What firms have to do with it?
2548 ADINA-ELENA FUDULACHE and MARIA DEL CARMEN CASTILLO LOZOYA: Demand drivers of central bank liquidity: A time-to-exit TLTRO analysis.
2549 ERIK ANDRES-ESCAYOLA, LUIS MOLINA, JAVIER J. PÉREZ and ELENA VIDAL: How economic policy uncertainty spreads across borders: the case of Latin America.
2550 MATTHIAS BURGERT, MATTHIEU DARRACQ PARIÈS, LUIGI DURAND, MARIO GONZÁLEZ, ROMANOS PRIFTIS, OKE RÖHE, MATTHIAS ROTTNER, EDGAR SILGADO-GÓMEZ, NIKOLAI STÄHLER and JANOS VARGA: Macroeconomic effects of carbon-intensive energy price changes: A model comparison.
2601 IACOPO VAROTTO: Blocking the Blockers? Diversity Matters.
2602 CARLOS CAÑIZARES MARTÍNEZ, ADRIANA LOJSCHOVÁ and ALICIA AGUILAR: Non-linear effects of monetary

policy shocks on housing: Evidence from a CESEE country.

2603 DIEGO BONELLI: Inflation risk and yield spread changes.

2604 MARÍA ALEJANDRA AMADO: Macroprudential FX Regulations and Small Firms: Unintended Consequences for Credit Growth.

2605 FERNANDO ÁVALOS, BORIS HOFMANN and JOSE M. SERENA: Monetary policy and private equity acquisitions: tracing the links.

2606 RICARD GREBOL, MARGARITA MACHELETT, JAN STUHLER and ERNESTO VILLANUEVA: Assortative Mating, Inequality, and Rising Educational Mobility in Spain.

2607 PABLO AGUILAR, RUBÉN DOMÍNGUEZ-DÍAZ, JOSÉ-ELÍAS GALLEGOS and JAVIER QUINTANA: The Transmission of Foreign Shocks in a Networked Economy.

2608 ERWAN GAUTIER, CRISTINA CONFLITTI, DANIEL ENDERLE, LUDMILA FADEJEVA, ALEX GRIMAUD, EDUARDO GUTIÉRREZ, VALENTIN JOUVANCEAU, JAN-OLIVER MENZ, ALARI PAULUS, PAVLOS PETROULAS, PAU ROLDAN-BLANCO and ELISABETH WIELAND: Consumer price stickiness in the euro area during an inflation surge.

2609 MORTEZA GHOMI and SAMUEL HURTADO: RAUI: Uncertainty Indicators Built With Artificial Intelligence.

2610 MORTEZA GHOMI and EVI PAPPA: Stimulating avenues: EIB loans and returns to Public Investment.

2611 JUAN S. MORA-SANGUINETTI, CRISTINA PEÑASCO and ROK SPRUK: The impact of "Green Regulation" on firms' innovation.

2612 ÁLVARO FERNÁNDEZ-GALLARDO and EVI PAPPA: Natural disasters and fiscal shelters.

2613 NICOLÁS BONINO-GAYOSO and MÓNICA CORREA-LÓPEZ: Unexpecting the expected in real-time inflation forecasting: The inflation expectations channel?

2614 ANDRÉS AZQUETA-GAVALDÓN, MARINA DIAKONOVA, CORINNA GHIRELLI and JAVIER J. PÉREZ: Diverging signals from economic uncertainty measures: Uncovering coherence through news narratives.

2615 JOSÉ E. GUTIÉRREZ, ENRIC MARTORELL and MARIYA MELNYCHUK: The nexus between the deposit and risk-taking channels of monetary policy.

2616 MARTA GARCÍA-RODRÍGUEZ and CLEMENTE PINILLA-TORREMOCHA: The role of confidence measures in European unemployment dynamics.

2617 DANIEL DEJUAN-BITRIA: Processing forward-looking loan loss provisions: evidence from the adoption of the CECL model.

2618 RYAN BANERJEE, FRANCISCO GONZÁLEZ, JOSÉ E. GUTIÉRREZ and JOSE-MARÍA SERENA: Credit Supply in the Wake of Distressed Bank Acquisitions.

2619 ALEJANDRO CASADO, MARCO GIOMETTI, JOSÉ E. GUTIÉRREZ, DAVID MARTÍNEZ-MIERA, ALEXANDRA MATYUNINA and TAMMARO TERRACCIANO: Digitalization and credit markets: Evidence from e-Invoicing.

2620 JORGE ABAD, SAKI BIGIO, SALOMÓN GARCÍA-VILLEGAS, JOËL MARBET and GALO NUÑO: The Heterogeneous Bank Lending Channel of Monetary Policy.

2621 NICOLÁS FORTEZA and JOSÉ M. LABEAGA: Pollution, Density and Low Emission Zones: European Evidence.

2622 ANDREA DE POLIS, ÁLVARO FERNÁNDEZ-GALLARDO and JOSÉ NICOLÁS ROSAS: Oil supply shocks and inflation tail risks.

2623 LUIS HERRERA, CATERINA MENDICINO, KALIN NIKOLOV and VALERIO SCALONE: Monetary Tightening, Financial Stability and the Role of Macroprudential Policy.

2624 MARINA GÓMEZ-GARCÍA: The impact of air pollution on the housing market: The case of Madrid.

2625 ANA AGUILAR, RAFAEL GUERRA, CAROLA MÜLLER and ALEXANDRE TOMBINI: The macro-financial impact of economic policy uncertainty in Latin America.

2626 ANDREA DE POLIS, LEONARDO MELOSI and IVAN PETRELLA: The Taming of the Skew: Asymmetric Inflation Risk and Monetary Policy.

2627 EOGHAN O'NEILL and SOFIA VELASCO: Let the Tree Decide: FABART. A Non-parametric factor model for nonlinear oil shock transmission.