EUROPEAN CENTRAL BANK
EUROSYSTEM

Working Paper Series

Eoghan O'Neill, Sofia Velasco

Let the tree decide: FABART. A non-parametric factor model for nonlinear oil shock transmission

No 3265

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 combines a large macro-financial information set with a flexible nonparametric measurement equation, allowing nonlinear transmission to emerge from the data rather than being imposed through a pre-specified functional form. 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, financial variables, and equity prices. Employment responses are highly heterogeneous across U.S. states, with substantially stronger contractions in manufacturing-intensive regions, while energy-producing states display partially offsetting dynamics following adverse oil supply news shocks. Across shock magnitudes, nonlinearities arise mainly between very small and moderate oil-price movements: small shocks generate weak and imprecisely estimated responses, while moderate shocks already produce economically meaningful effects on industrial production and regional employment. Larger shocks do not systematically generate proportionally stronger responses across variables and shock signs.

Keywords: Dynamic factor model, Nonlinear Models, Non-parametric Techniques, Oil Price Shocks.

JEL Codes: C11, C32, E32, Q43

Non-technical summary

Oil price shocks have long been recognized as an important source of macroeconomic fluctuations, influencing inflation, economic activity and financial markets. Recent episodes of heightened geopolitical tensions and disruptions to global energy markets have renewed interest in understanding how these shocks propagate through the economy. At the same time, an extensive body of research suggests that the effects of oil shocks may be nonlinear: increases in oil prices often appear to have stronger economic consequences than comparable declines, while the effects of particularly large shocks may differ from those of more moderate disturbances. Despite this evidence, most empirical analyses continue to rely on linear models or impose specific functional forms to capture nonlinearities, potentially overlooking more complex patterns in the data.

This paper studies how the U.S. economy responds to oil supply news shocks using a flexible empirical framework that allows nonlinear relationships to emerge directly from the data. The analysis combines external identification of oil supply news shocks with a nonlinear factor model that summarizes information from a large set of macroeconomic and financial variables. Rather than imposing a particular type of asymmetry or threshold, the approach lets the data determine how the transmission of oil shocks varies across different shock sizes and directions.

The results show that adverse oil supply news shocks generate substantially larger macroeconomic effects than equally sized favourable shocks. Large positive shocks lead to persistent increases in oil prices, declines in industrial production and equity prices, and increases in inflation. By contrast, comparable negative shocks produce considerably weaker expansions in economic activity. These findings are consistent with the view that the transmission of oil shocks is asymmetric, particularly for real economic activity and financial markets.

The paper also examines whether larger shocks necessarily produce proportionally larger economic responses. The evidence suggests that this is generally not the case. Instead, the main source of nonlinearity arises because relatively small oil supply news shocks have only limited macroeconomic effects, while responses become more pronounced once shocks exceed moderate magnitudes. Beyond this point, however, the relationship is not proportional, indicating that the effects of increasingly large shocks eventually level off rather than continuing to grow at the same rate.

To investigate the regional consequences of oil shocks, the analysis further examines labour market outcomes across U.S. states. States with greater dependence on oil production experience smaller increases in unemployment and more resilient employment following adverse oil supply news shocks than states with lower exposure to the energy sector. These differences highlight the importance of regional industrial structure in shaping the transmission of energy shocks and suggest that aggregate responses mask substantial heterogeneity across local labour markets.

More broadly, the paper demonstrates the value of combining credible identification of structural shocks with flexible, data-driven empirical methods. Allowing nonlinearities to emerge from the data provides a richer characterization of macroeconomic transmission than conventional linear models while avoiding the need to specify the form of nonlinear relationships in advance. The findings contribute to a better understanding of how oil supply shocks affect economic activity, inflation and labour markets, offering insights that are relevant for policymakers assessing the macroeconomic consequences of future energy market disruptions.

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, financial, uncertainty, and state-level labor-market variables. The richer information set allows us to study heterogeneous transmission across regions differing in industrial structure, energy dependence, and labor-market composition, dimensions long considered central to the propagation of oil shocks (Davis and Haltiwanger, 2001; Herrera and Karaki, 2015; Hamilton and Owyang, 2012).

[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.

Our main findings are fourfold. First, expanding the information set changes the estimated transmission of oil supply news shocks, yielding sharper real contractions, stronger financial responses, and weaker inflationary effects. Second, adverse shocks generate substantially larger contractions than comparable favorable shocks. Third, the responses are highly heterogeneous across U.S. states, with manufacturing-intensive states contracting more strongly and oil-producing states displaying partially offsetting dynamics. Fourth, transmission also depends on shock size, with nonlinearities becoming economically important as shocks increase from small to moderate magnitudes.

Taken together, the results illustrate the value of combining richer information sets with flexible nonparametric measurement equations for structural analysis. They reinforce the importance of informational sufficiency (Mori and Peersman, 2024), complement recent evidence on nonlinear oil-shock transmission (Forni et al., 2025), and suggest that sign and size asymmetries are empirically distinct (Caravello and Martinez-Bruera, 2024). More generally, FABART addresses the specification concerns emphasized by Kilian and Vigfusson (2011a, 2017) by allowing nonlinear transmission to emerge from the data rather than imposing a predetermined functional form.

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, estimation 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 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

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

[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.

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:

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

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:

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

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 $\mu_{bs}^{(i)} \sim \mathcal{N}\left(0, \sigma_{\mu i}^2\right), \quad \text{with} \quad \sigma_{\mu i} = \frac{\max(X_{\cdot i}) - \min(X_{\cdot i})}{2\sqrt{\nu S}}$, 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: $\sigma_i^2 \sim \nu_i \xi_i / \chi^2_{\nu_i}$, 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]

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	$\mu_{bs} \sim \mathcal{N}(0, \sigma_{\mu i}^2)$	$\sigma_{\mu i} = \frac{\max(X_{\cdot i}) - \min(X_{\cdot i})}{2\sqrt{\nu S}}$
Error variances	$\sigma_i^2 \sim \nu_i \xi_i / \chi^2_{\nu_i}$	$\nu_i = 3,\ P(\sigma_i < \hat{\sigma}_i) = 0.9$
Number of trees	–	$S = 50$

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

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.[8]

[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.

[8] Similar results are obtained with a conjugate normal inverse Wishart priors the transition equation

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

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

The proposal distribution $q(\tau^j_{is}|\tau^*_{is})$ 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).

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.

[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).

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.
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.

[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.

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]

[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

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 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

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

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

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).

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

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

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 $j = 2, \dots, J$. 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).

[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.

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 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.

[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.

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.

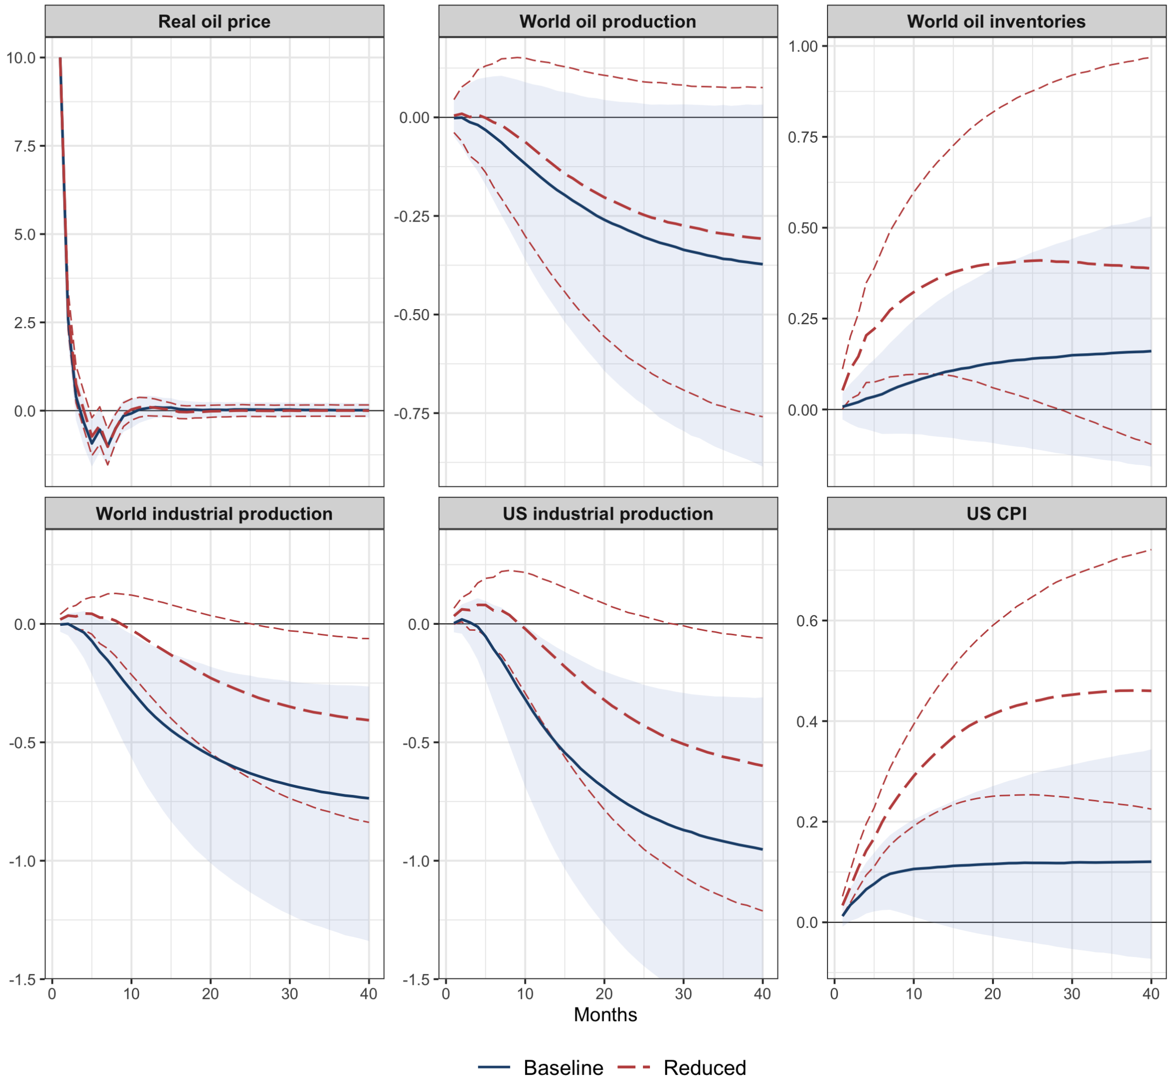

Note: GIRFs to an adverse oil supply news shock normalized to raise the real price of oil by 10% on impact. Solid lines correspond to the Baseline specification, which includes financial variables. Dashed lines correspond to the Reduced specification, which excludes financial variables. Shaded regions 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.

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 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

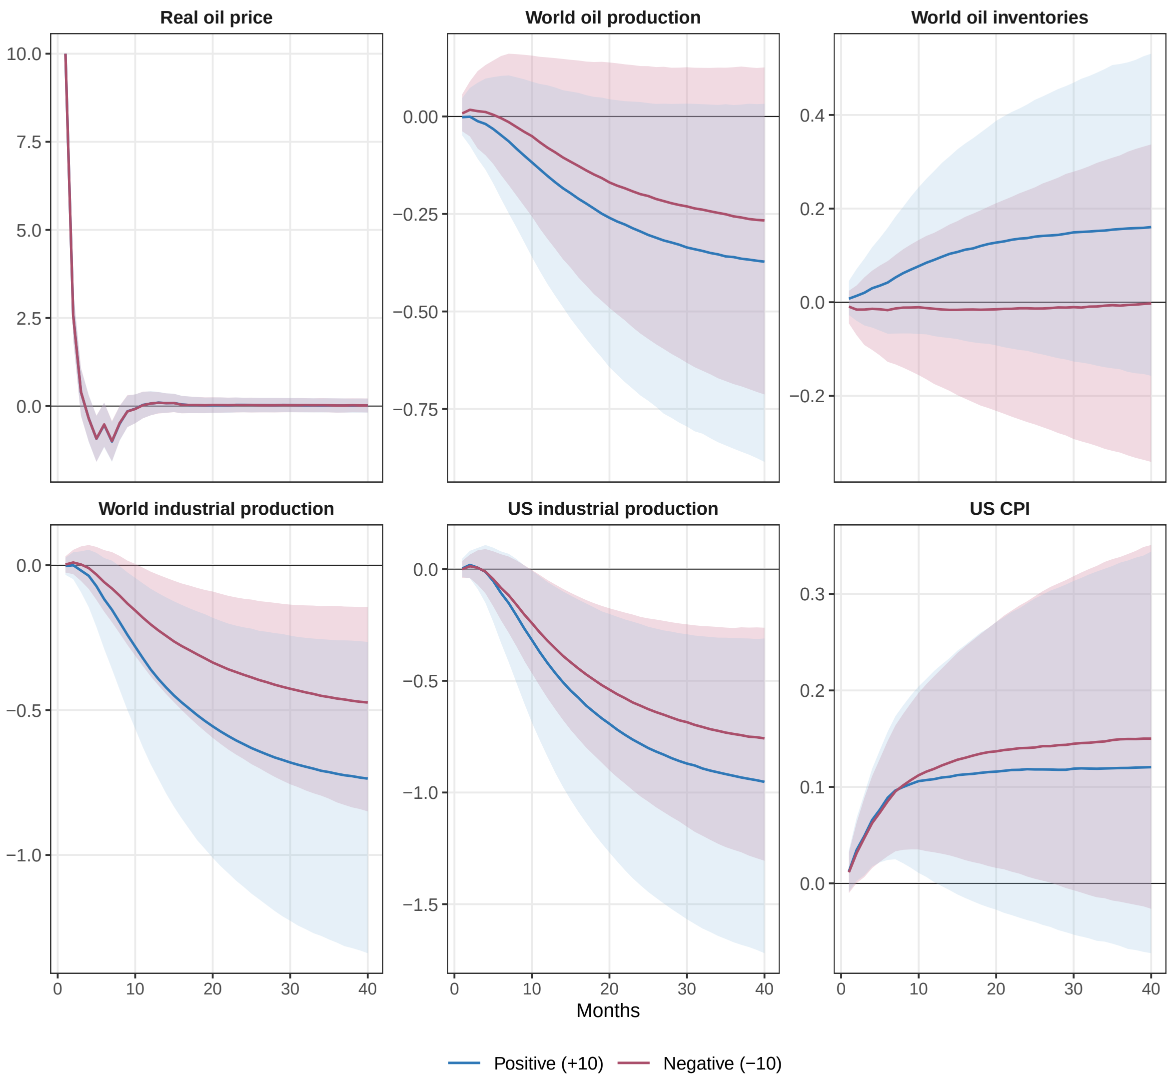

Note: GIRFs of oil supply news shocks normalized to move the real oil price by 10% on impact. Blue lines show adverse shocks that raise oil prices; red lines show favorable shocks that lower oil prices. Responses to favorable shocks are multiplied by minus one for visual comparability. Shaded areas represent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

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

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.

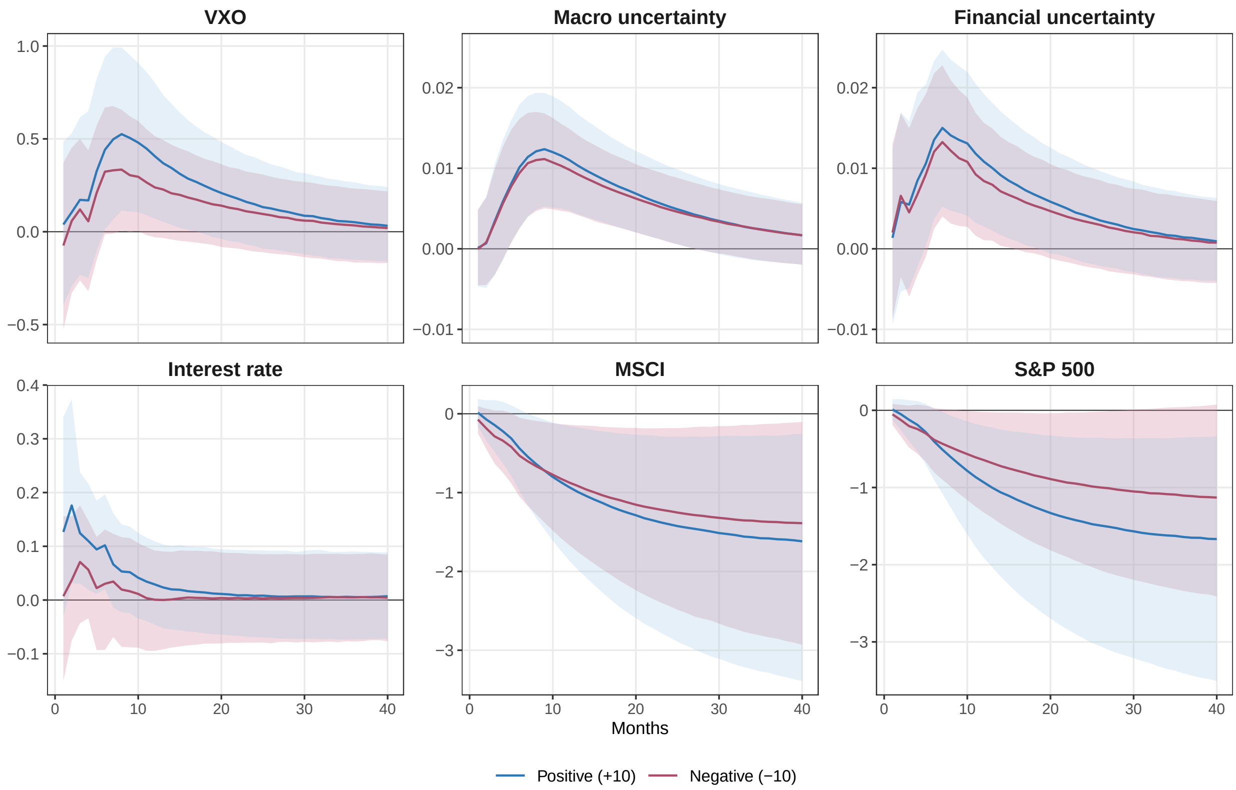

Note: GIRFs of financial and uncertainty variables following oil supply news shocks normalized to move the real oil price by 10% on impact. Blue lines show adverse shocks that raise oil prices; red lines show favorable shocks that lower oil prices. Responses to favorable shocks are multiplied by minus 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.

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 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$^{\text{th}}$ 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 26th, 57th, and 84th 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 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

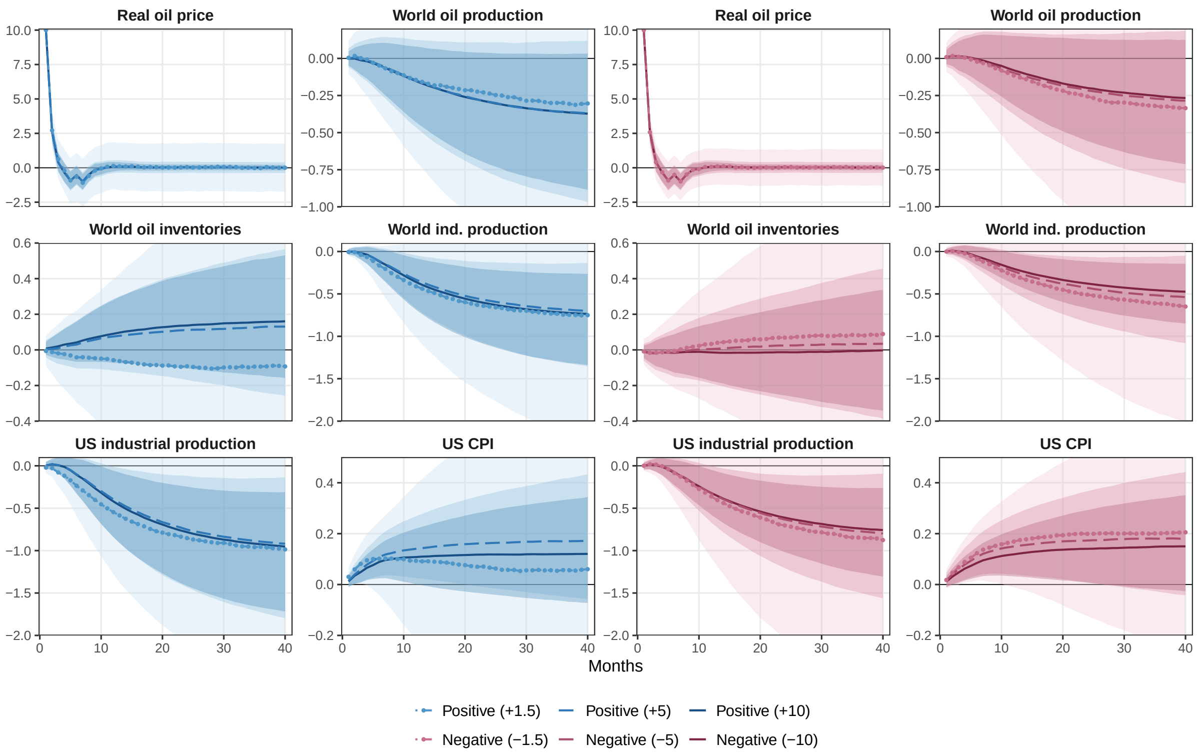

Note: GIRFs to adverse shocks that raise oil prices (blue) and favorable shocks that lower oil prices (red). The underlying shocks are associated with 1.5%, 5%, and 10% oil-price movements, and all responses are rescaled to a common 10% oil-price movement for comparability. Responses to favorable shocks 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.

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; Bjørnland 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.

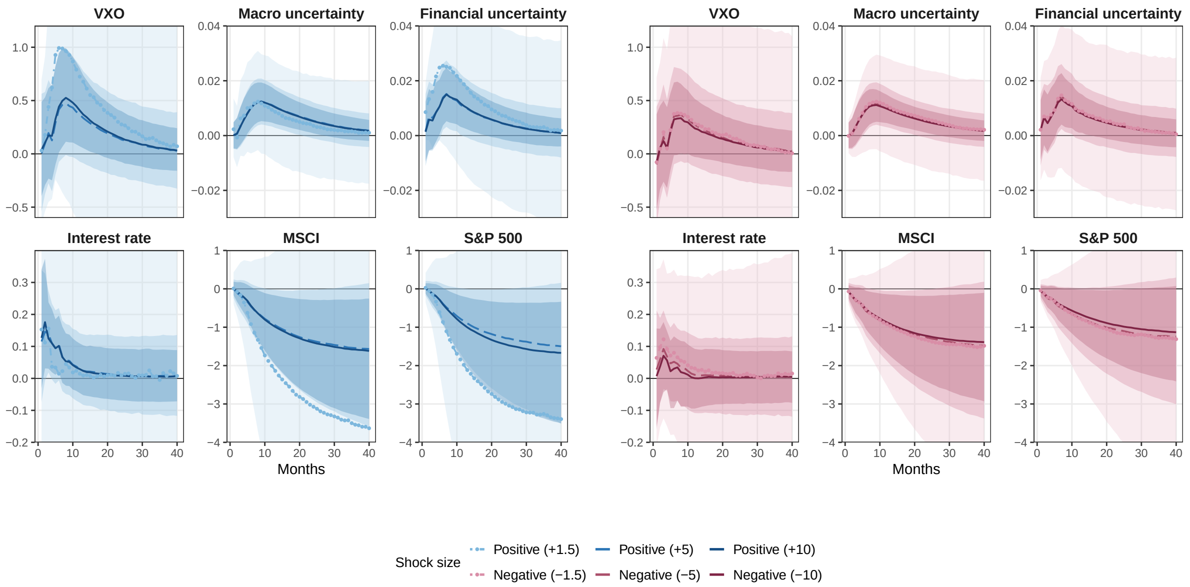

Note: GIRFs of financial and uncertainty variables to adverse shocks that raise oil prices (blue) and favorable shocks that lower oil prices (red). The underlying shocks are associated with 1.5%, 5%, and 10% oil-price movements, and all responses are rescaled to a common 10% oil-price movement for comparability. 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.

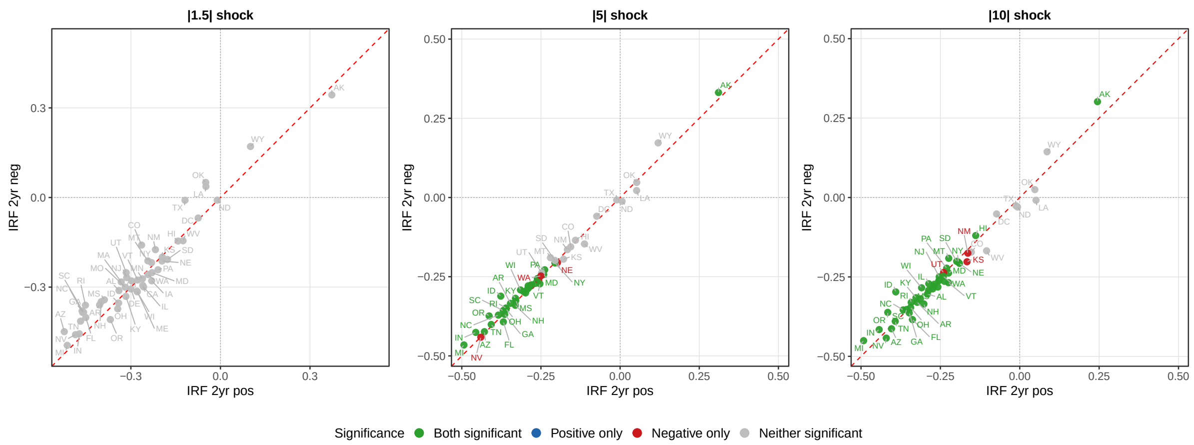

Note: This figure compares median two-year cumulative employment responses to adverse and favorable oil supply news shocks across U.S. states for shock sizes of 1.5%, 5%, and 10%. The x-axis reports responses to adverse shocks that raise oil prices. The y-axis reports responses to favorable shocks that 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 dependent. 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.

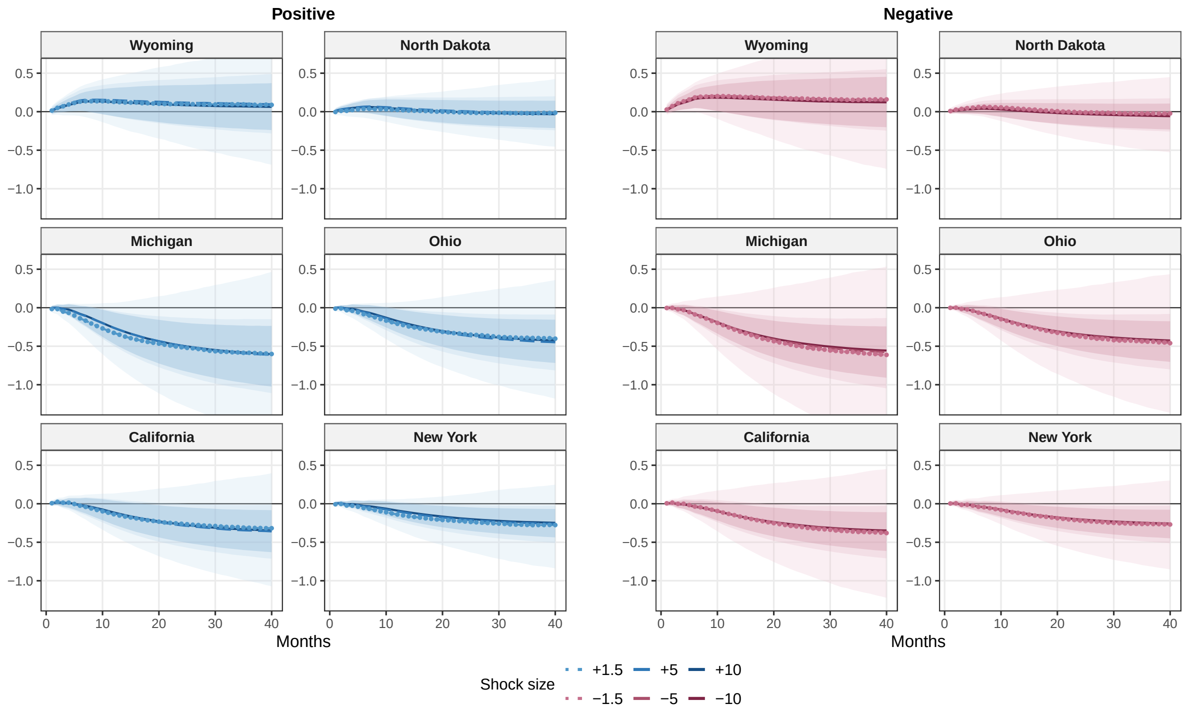

Note: Cumulative state-level employment GIRFs to oil supply news shocks calibrated to move the real oil price by 1.5%, 5%, and 10% on impact, for representative states. Blue lines show adverse shocks that raise oil prices; red lines show favorable shocks that lower oil prices. Responses to favorable 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 paper introduces the Factor Bayesian Additive Regression Tree (FABART), a nonparametric dynamic factor model for structural analysis in large information sets. FABART combines a linear vector autoregressive transition equation for the latent factors with outcome-specific Bayesian Additive Regression Tree (BART) functions in the measurement equations, allowing the relationship between latent factors and observables to be learned nonparametrically rather than imposed through a parametric functional form.

We study the transmission of externally identified oil supply news shocks using a large U.S. macro-financial and regional dataset. The results reveal pronounced sign asymmetries, with economically meaningful responses emerging only for sufficiently large 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. 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 magnitude across variables or between positive and negative shocks, indicating that nonlinear transmission is not limited to extreme oil-market disruptions.

Beyond the aggregate responses at the national level, the results reveal substantial regional heterogeneity in labor-market adjustment. Manufacturing-intensive states experience considerably stronger employment contractions following adverse oil supply news shocks, whereas energy-producing states display substantially weaker contractions and, in some cases, positive employment responses. These findings suggest that industrial composition and exposure to energy markets play an important role in shaping the regional transmission of oil supply news shocks.

The paper illustrates how a flexible nonparametric measurement equation can be combined with structural identification in a large-information dynamic factor model, allowing nonlinear relationships between latent factors and observed variables to be learned from the data rather than imposed through a parametric specification. Although the empirical application focuses on externally identified oil supply news shocks, the methodology is compatible with a broad class of structural identification strategies and can be applied to structural analysis and forecasting in large-information settings.

References

Andreini, P., Izzo, C., and Ricco, G. (2020). Deep dynamic factor models. *arXiv preprint arXiv:2007.11887*.

Bai, J. and Wang, P. (2015). Identification and Bayesian estimation of dynamic factor models. *Journal of Business & Economic Statistics*, 33(2):221–240.

Baumeister, C., Liu, P., and Mumtaz, H. (2013). Changes in the effects of monetary policy on disaggregate price dynamics. *Journal of Economic Dynamics and Control*, 37(3):543–560.

Bernanke, B. S., Boivin, J., and Eliasz, P. (2005). Measuring the effects of monetary policy: a factor-augmented vector autoregressive (FAVAR) approach. *The Quarterly Journal of Economics*, 120(1):387–422.

Bernanke, B. S., Gertler, M., and Watson, M. W. (1997). Systematic monetary policy and the effects of oil price shocks. *Brookings Papers on Economic Activity*, 1997(1):91–157.

Bjørnland, H. C., Chang, Y., and Cross, J. L. (2026). Oil and the stock market revisited: A mixed functional VAR approach. *Quantitative Economics*. Forthcoming.

Blanchard, O. J. and Riggi, M. (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):1032–1052.

Bobeica, E., Holton, S., Huber, F., and Martínez Hernández, C. (2025). Beware of large shocks! A non-parametric structural inflation model. Working Paper Series 3052, European Central Bank.

Caldara, D. and Herbst, E. (2019). Monetary policy, real activity, and credit spreads: Evidence from Bayesian proxy SVARs. *American Economic Journal: Macroeconomics*, 11(1):157–192.

Caravello, T. and Martinez-Bruera, P. (2024). Disentangling sign and size non-linearities. SSRN Working Paper 4704050.

Carter, C. K. and Kohn, R. (1994). On Gibbs sampling for state space models. *Biometrika*, 81(3):541–553.

Charnavoki, V. and Dolado, J. J. (2014). The effects of global shocks on small commodity-exporting economies: lessons from canada. *American Economic Journal: Macroeconomics*, 6(2):207–237.

Chen, L., Dolado, J. J., and Gonzalo, J. (2021). Quantile factor models. *Econometrica*, 89(2):875–910.

Chen, L., Dolado, J. J., Gonzalo, J., and Pan, H. (2026). Estimation of characteristics-based quantile factor models. *Journal of Econometrics*, 256:106267.

Chernis, T., Hauzenberger, N., Mumtaz, H., and Pfarrhofer, M. (2025). A Bayesian Gaussian process dynamic factor model. *arXiv preprint arXiv:2509.04928*.

Chipman, H. A., George, E. I., and McCulloch, R. E. (1998). Bayesian cart model search. *Journal of the American Statistical Association*, 93(443):935–948.

Chipman, H. A., George, E. I., and McCulloch, R. E. (2010). BART: Bayesian additive regression trees. *The Annals of Applied Statistics*, 4(1):266–298.

Clark, T., Huber, F., and Koop, G. (2025). A nonparametric approach to augmenting a Bayesian VAR with nonlinear factors. *arXiv preprint arXiv:2508.13972*.

Clark, T. E., Huber, F., Koop, G., Marcellino, M., and Pfarrhofer, M. (2024). Investigating growth-at-risk using a multicountry nonparametric quantile factor model. *Journal of Business & Economic Statistics*, 42(4):1302–1317.

Crawford, L., Flaxman, S. R., Runcie, D. E., and West, M. (2019). Variable prioritization in nonlinear black box methods: A genetic association case study. *The annals of applied statistics*, 13(2):958.

David, J. M., Giacomini, R., Jiao, X., and Wang, W. (2026). Causal state-dependent local projections. *arXiv preprint arXiv:2605.05404*.

Davis, S. J. and Haltiwanger, J. (2001). Sectoral job creation and destruction responses to oil price changes. *Journal of Monetary Economics*, 48(3):465–512.

Degasperi, R. (2026). Identification of expectational shocks in the oil market using OPEC announcements. Temi di Discussione 1516, Banca d'Italia.

Durbin, J. and Koopman, S. J. (2002). A simple and efficient simulation smoother for state space time series analysis. *Biometrika*, 89(3):603–616.

Forni, M., Franconi, A., Gambetti, L., and Sala, L. (2025). Asymmetric transmission of oil supply news. *Quantitative Economics*, 16(3):947–979.

Forni, M. and Gambetti, L. (2010). The dynamic effects of monetary policy: A structural factor model approach. *Journal of Monetary Economics*, 57(2):203–216.

Forni, M. and Gambetti, L. (2014). Sufficient information in structural VARs. *Journal of Monetary Economics*, 66:124–136.

Forni, M., Giannone, D., Lippi, M., and Reichlin, L. (2009). Opening the black box: Structural factor models with large cross sections. *Econometric Theory*, 25(5):1319–1347.

Gautier, E., Conflitti, C., Enderle, D., Fadejeva, L., Grimaud, A., Gutiérrez, E., Jouvanceau, V., Menz, J.-O., Paulus, A., Petroulas, P., Roldan-Blanco, P., and Wieland, E. (2026). Consumer price stickiness in the euro area during an inflation surge. Documentos de Trabajo 2608, Banco de España.

Gonçalves, S., Herrera, A. M., Kilian, L., and Pesavento, E. (2021). Impulse response analysis for structural dynamic models with nonlinear regressors. *Journal of Econometrics*, 225(1):107–130.

Goncalves, S., Herrera, A. M., Kilian, L., Pesavento, E., and Holban, I. K. (2026). Semiparametric local projections. *arXiv preprint arXiv:2606.13519*.

Hamilton, J. D. (1983). Oil and the macroeconomy since World War II. *Journal of Political Economy*, 91(2):228–248.

Hamilton, J. D. (2003). What is an oil shock? *Journal of Econometrics*, 113(2):363–398.

Hamilton, J. D. (2011). Nonlinearities and the macroeconomic effects of oil prices. *Macroeconomic Dynamics*, 15(S3):364–378.

Hamilton, J. D. and Herrera, A. M. (2004). Oil shocks and aggregate macroeconomic behavior: the role of monetary policy: a comment. *Journal of Money, Credit and Banking*, pages 265–286.

Hamilton, J. D. and Owyang, M. T. (2012). The propagation of regional recessions. *Review of Economics and Statistics*, 94(4):935–947.

Hauzenberger, N., Huber, F., and Klieber, K. (2023). Real-time inflation forecasting using non-linear dimension reduction techniques. *International Journal of Forecasting*, 39(2):901–921.

Herrera, A. M. and Karaki, M. B. (2015). The effects of oil price shocks on job reallocation. *Journal of Economic Dynamics and Control*, 61:95–113.

Herrera, A. M., Karaki, M. B., and Rangaraju, S. K. (2017). Where do jobs go when oil prices drop? *Energy Economics*, 64:469–482.

Huber, F., Koop, G., Onorante, L., Pfarrhofer, M., and Schreiner, J. (2020). Nowcasting in a pandemic using non-parametric mixed frequency VARs. *Journal of Econometrics*.

Huber, F. and Rossini, L. (2022). Inference in Bayesian additive vector autoregressive tree models. *The Annals of Applied Statistics*, 16(1):104–123.

Juvenal, L. and Petrella, I. (2015). Speculation in the oil market. *Journal of Applied Econometrics*, 30(4):621–649.

Känzig, D. R. (2021). The macroeconomic effects of oil supply news: Evidence from OPEC announcements. *American Economic Review*, 111(4):1092–1125.

Kapelner, A. and Bleich, J. (2016). bartMachine: Machine learning with Bayesian additive regression trees. *Journal of Statistical Software*, 70:1–40.

Kharroubi, E., Spigt, R., Igan, D., Takahashi, K., and Zakrajšek, E. (2023). Markups and the asymmetric pass-through of cost push shocks. BIS Working Papers 1150, Bank for International Settlements.

Kilian, L. (2008). The economic effects of energy price shocks. *Journal of Economic Literature*, 46(4):871–909.

Kilian, L. (2009). Not all oil price shocks are alike: Disentangling demand and supply shocks in the crude oil market. *American Economic Review*, 99(3):1053–1069.

Kilian, L. (2024). How to construct monthly VAR proxies based on daily surprises in futures markets. *Journal of Economic Dynamics and Control*, 168:104966.

Kilian, L. and Lewis, L. T. (2011). Does the fed respond to oil price shocks? *The Economic Journal*, 121(555):1047–1072.

Kilian, L. and Murphy, D. P. (2014). The role of inventories and speculative trading in the global market for crude oil. *Journal of Applied Econometrics*, 29(3):454–478.

Kilian, L. and Park, C. (2009). The impact of oil price shocks on the US stock market. *International Economic Review*, 50(4):1267–1287.

Kilian, L., Plante, M. D., and Richter, A. W. (2026). How times have changed: The impact of the 2026 Iran war on the US economy. CEPR Discussion Papers 21650, Centre for Economic Policy Research.

Kilian, L. and Vigfusson, R. J. (2011a). Are the responses of the US economy asymmetric in energy price increases and decreases? *Quantitative Economics*, 2(3):419–453.

Kilian, L. and Vigfusson, R. J. (2011b). Nonlinearities in the oil price–output relationship. *Macroeconomic Dynamics*, 15(S3):337–363.

Kilian, L. and Vigfusson, R. J. (2017). The role of oil price shocks in causing US recessions. *Journal of Money, Credit and Banking*, 49(8):1747–1776.

Kilian, L. and Zhou, X. (2023). *The Econometrics of Oil Market VAR Models*, volume 45 of *Advances in Econometrics*. Emerald Group Publishing Limited, Bingley, UK. Essays in Honor of Joon Y. Park: Econometric Methodology in Empirical Applications, pp. 65–95.

Koop, G., Pesaran, M. H., and Potter, S. M. (1996). Impulse response analysis in nonlinear multivariate models. *Journal of Econometrics*, 74(1):119–147.

Korobilis, D. and Schröder, M. (2024). Monitoring multi-country macroeconomic risk: A quantile factor-augmented vector autoregressive (qfavar) approach. *Journal of Econometrics*, page 105730.

Lindsten, F., Jordan, M. I., and Schon, T. B. (2014). Particle Gibbs with ancestor sampling. *Journal of Machine Learning Research*, 15:2145–2184.

Linero, A. R. (2018). Bayesian regression trees for high-dimensional prediction and variable selection. *Journal of the American Statistical Association*, 113(522):626–636.

Linero, A. R. and Yang, Y. (2018). Bayesian regression tree ensembles that adapt to smoothness and sparsity. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 80(5):1087–1110.

Maia, M., Murphy, K., and Parnell, A. C. (2024). GP-BART: A novel Bayesian additive regression trees approach using Gaussian processes. *Computational Statistics & Data Analysis*, 190:107858.

Marcellino, M. and Pfarrhofer, M. (2025). Nonparametric mixed frequency monitoring macro-at-risk. *Economics Letters*, page 112498.

Miescu, M., Mumtaz, H., and Theodoridis, K. (2024). Nonlinear dynamics of large oil supply news shocks. Cardiff Economics Working Papers E2024/18, Cardiff University.

Miranda-Agrippino, S. and Ricco, G. (2023). Identification with external instruments in structural VARs. *Journal of Monetary Economics*, 135:1–19.

Moran, G. E. (2019). *Bayesian approaches for modeling variation*. PhD thesis, University of Pennsylvania.

Mori, L. and Peersman, G. (2024). Estimating the macroeconomic effects of oil supply news. CESifo Working Papers 11532, CESifo.

Mork, K. A. (1989). Oil and the macroeconomy when prices go up and down: an extension of hamilton's results. *Journal of Political Economy*, 97(3):740–744.

Mumtaz, H. (2010). Evolving UK macroeconomic dynamics: a time-varying factor augmented VAR. Bank of England Working Papers 386, Bank of England.

Mumtaz, H. and Piffer, M. (2022). Impulse response estimation via flexible local projections. *arXiv preprint arXiv:2204.13150*.

Mumtaz, H. and Sunder-Plassmann, L. (2021). Nonlinear effects of government spending shocks in the usa: Evidence from state-level data. *Journal of Applied Econometrics*, 36(1):86–97.

Mumtaz, H., Sunder-Plassmann, L., and Theophilopoulou, A. (2018). The state-level impact of uncertainty shocks. *Journal of Money, Credit and Banking*, 50(8):1873–1913.

Mumtaz, H., Theophilopoulou, A., and Drossidis, T. (2024). The distributional effects of oil supply news shocks. *Economics Letters*, 240:111769.

Neri, S. (2025). Energy prices, inflation and the ECB's monetary policy during the 2021-22 energy crisis. Temi di Discussione 1481, Banca d'Italia.

Rotemberg, J. J. and Woodford, M. (1991). Markups and the business cycle. *NBER Macroeconomics Annual*, 6:63–129.

Stock, J. H. and Watson, M. W. (2005). Implications of dynamic factor models for VAR analysis. NBER Working Paper Series 11467, National Bureau of Economic Research.

Velasco, S. (2025). Let the tree decide: Fabart a non-parametric factor model. *arXiv preprint arXiv:2506.11551*.

Yalcin, I. and Amemiya, Y. (2001). Nonlinear factor analysis as a statistical method. *Statistical Science*, pages 275–294.

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

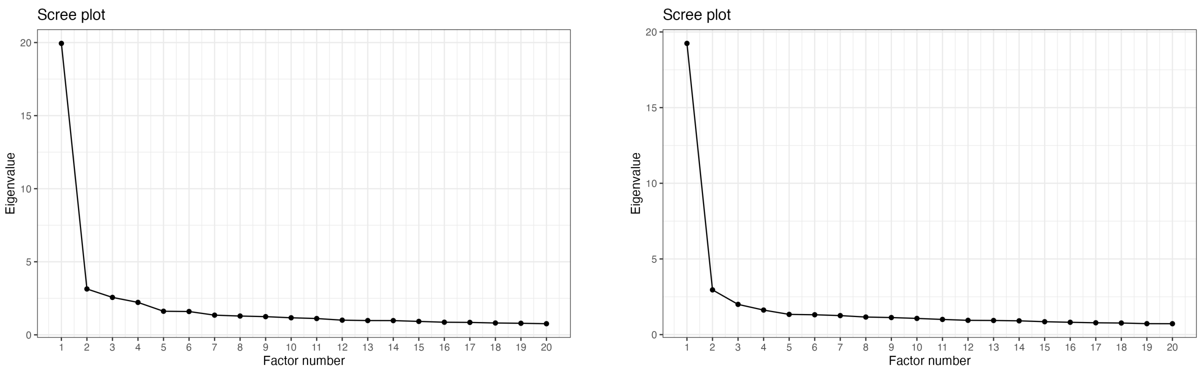

Baseline Reduced

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

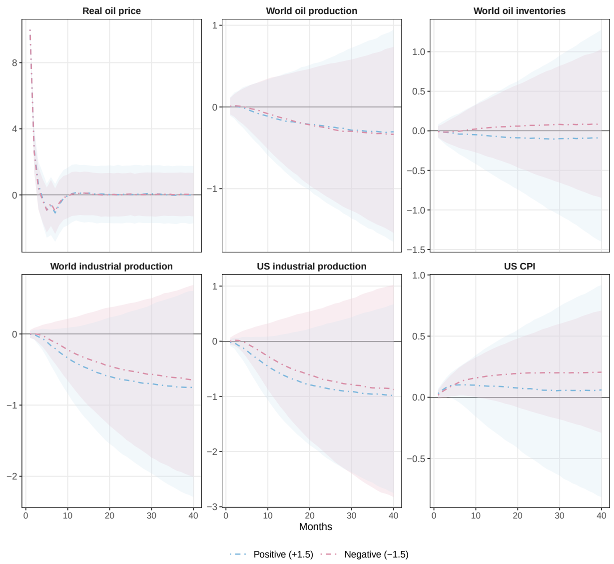

1.5% shock

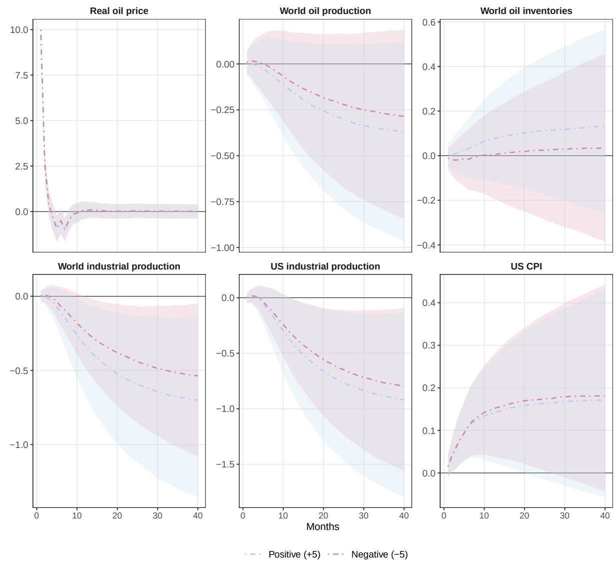

5% shock

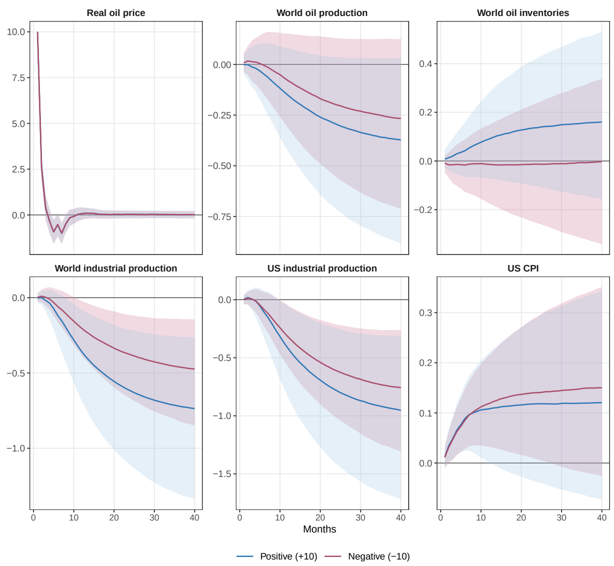

10% shock

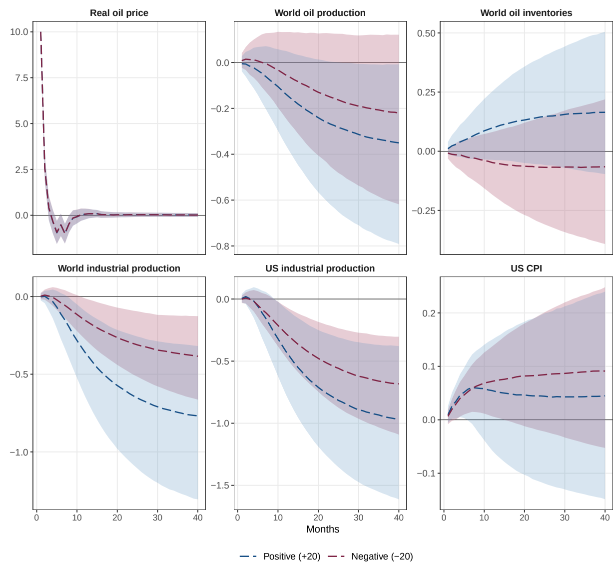

20% shock

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

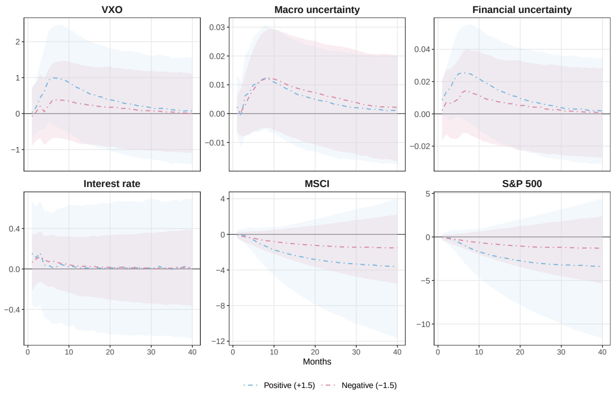

1.5% shock

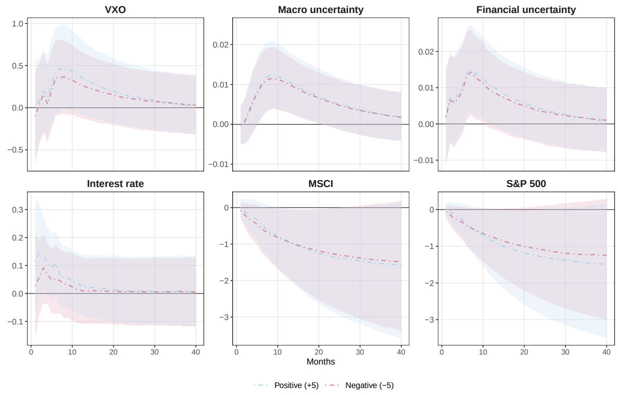

5% shock

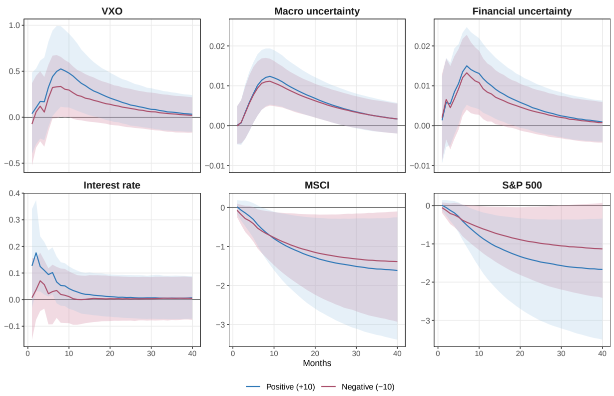

10% shock

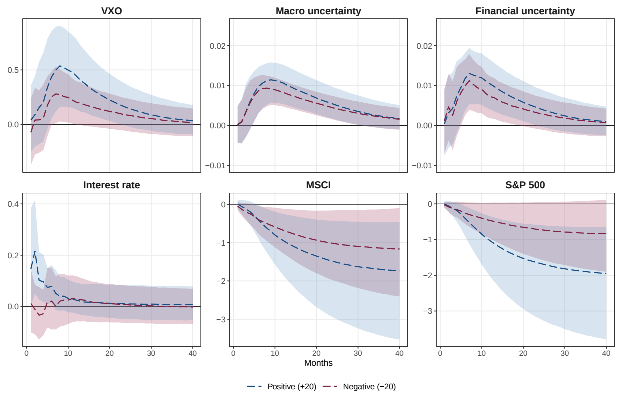

20% shock

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

Acknowledgements

The authors thank the Society for Nonlinear Dynamics & Econometrics for awarding this paper the SNDE Time Series Award. We are grateful without endorsement 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 XVIth Workshop in Time Series Econometrics (Zaragoza), the 33rd Symposium of the Society for Nonlinear Dynamics & Econometrics and the anonymous referee, for helpful comments and suggestions. An earlier version of this paper was circulated on arXiv as Velasco (2025), "Let the Tree Decide: FABART A Non-Parametric Factor Model". 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. Replication code will be made available on GitHub.
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).

Eoghan O'Neill
University College Dublin, Dublin, Ireland; email: eoghan.oneill1@ucd.ie

Sofia Velasco (corresponding author)
Banco de España, Madrid, España; email: sofia.velasco@bde.es

PDF ISBN 978-92-899-7994-8 ISSN 1725-2806 doi:10.2866/1207123 QB-01-26-193-EN-N