

Impulse response estimation via flexible local projections

Haroon Mumtaz and Michele Piffer

Working Paper No. 938 April 2022

ISSN 1473-0278

Ù&@ [|Á ÀÒ&[} [{ æ• Áæ åÁÖā æ &^



Impulse response estimation via flexible local projections

-

Latest draft available at this [link](#)*

Haroon Mumtaz[†] Michele Piffer[‡]

April 15, 2022

Abstract

This paper introduces a flexible local projection that generalises the model by [Jordà \(2005\)](#) to a non-parametric setting using Bayesian Additive Regression Trees. Monte Carlo experiments show that our BART-LP model is able to capture non-linearities in the impulse responses. Our first application shows that the fiscal multiplier is stronger in recession than expansion only in response to contractionary fiscal shocks, but not in response to expansionary fiscal shocks. We then show that financial shocks generate effects on the economy that increase more than proportionately in the size of the shock when the shock is negative, but not when the shock is positive.

JEL classification: C14, C11, C32, E52.

Keywords: Non-linear models, non-parametric techniques, identification.

*We are thankful to Vito Polito for helpful comments and suggestions.

[†]Queen Mary, University of London, School of Economics and Finance, Mile End road, London E1 4LJ, UK. e-mail: h.mumtaz@qmul.ac.uk

[‡]King's Business School, King's College London, Bush House, 30 Aldwych, London WC2B 4BG, UK. e-mail: m.b.piffer@gmail.com

1 Introduction

Estimation of impulse responses (IRFs) via local projections (LP) by [Jordà \(2005\)](#) has become increasingly common in applied Macroeconometric analysis. A key feature of the local projection estimator is that it estimates IRFs of variable y_t to an innovation to variable x_t directly via linear regressions of the form $y_{t+h} = \beta_h x_t + d_h w_t + u_{t+h}$, where w_t denotes control variables. Given their flexibility, considerable attention has been given to investigate the properties of the LP estimator, see, for instance, [Stock and Watson \(2018\)](#) and [Plagborg-Møller and Wolf \(2021\)](#).

In their most popular specification, LP estimators impose a linearity between y_{t+h} and (x_t, w_t) . This limitation implies that linear LPs cannot be used to study non-linear effects of the shocks of interest, for instance non-linearities on the sign or size of the shock, or on the economic conditions when the shock occurs. Some extensions of the linear LP estimator have been proposed, but they all rely on the functional form introduced to model the non-linearity. [Jordà \(2005\)](#) proposes the use of quadratic and cubic terms. [Auerbach and Gorodnichenko \(2013a,b\)](#) and [Ramey and Zubairy \(2018\)](#) use a smooth transition function and a threshold function, respectively. [Ruisi \(2019\)](#) and [Lusompa \(2021\)](#) use a time varying extension of LP based on parametric state-space models, while [Inoue et al. \(2022\)](#) provides a more general framework for modelling structural shifts.

In this paper we propose a flexible non-linear extension of the LP estimator that does not require assumptions on the functional form of the LP regression equation. We propose a non-parametric LP estimator that uses the Bayesian Additive Regression Trees (BART) model to approximate the unknown function $m_h(z_t)$ in the more general equation $y_{t+h} = m_h(z_t) + u_{t+h}$, with $z_t = (x_t, w_t)$. Introduced by [Chipman et al. \(2010\)](#), BART uses regression trees as its building block. Regression trees split the space of explanatory variables z_t into sub-groups based on rules of the form $z_{t,j} < C$, where $j = 1, 2, \dots, K$. The function $m_h(z_t)$ is approximated as a sum of a large number of small trees. [Chipman et al. \(2010\)](#) show that BART is able to approximate highly

non-linear functions accurately.

We first illustrate how BART techniques can be applied in a non-linear LP estimator, and we refer to this new methodology as BART-LP. We show that BART-LP can handle autocorrelation in the error terms, a problem already discussed in the literature of linear LP estimators. We then document the performance of BART-LP using Monte Carlo analysis. We build our simulations on three models. First, we use the same SVAR-GARCH model employed in the simulation by [Jordà \(2005\)](#), where a structural shock generates non-linear effects that affect the variance of the shock. Second, we use a recursive Threshold VAR model in which shocks generate stronger effects in one of the two regimes of the model. Third, we use a sign-dependent moving average model in which the true monetary policy shock generates different effects depending on the sign of the shock. In all cases, the BART-LP is capable of recovering the true impulse responses, while the linear LP typically estimates a weighted average of the true underlying non-linear impulse responses. We focus the discussion on the IRF estimation, which is separate from identification of the structural shocks. As also in [Jordà \(2005\)](#), we do not investigate the topic of identification in our framework.

We then apply the new methodology to shed some light on two separate ongoing debates in the literature. We first apply our methodology to US fiscal spending shocks. While [Auerbach and Gorodnichenko \(2013b\)](#) has argued that US fiscal multipliers are stronger in recession than in boom, their evidence has been called into question by [Ramey and Zubairy \(2018\)](#). We use our BART-LP procedure to show that the answer to this question depends on the sign of the shock. We show that the multiplier in response to a positive shock does not change significantly depending on whether the shock hits the economy in an expansion or in a recession. However, in response to a negative shock the multiplier is stronger in a recession. Compared to the models used by [Auerbach and Gorodnichenko \(2013b\)](#) and [Ramey and Zubairy \(2018\)](#), our methodology can detect both non-linearities over the state of the economy and non-linearities over the sign and size of the shock. Our results hence help reconcile the

difference found by [Auerbach and Gorodnichenko \(2013b\)](#) and [Ramey and Zubairy \(2018\)](#), who used models that are more suitable to detect non-linearities over the state of the economy, but not over the sign of the shocks. Last, we revisit the analysis by [Forni et al. \(2021\)](#) on financial shocks. We confirm their result that negative financial shocks generate detrimental effects on output that increase more than proportionally in the size of the shock. These effects cannot otherwise be detected by a linear model, which is doomed to underestimate the effect of strong negative financial shocks.

This paper relates to the literature that studies how BART techniques can be used in Macroeconometrics. [Huber and Rossini \(2021\)](#) introduce a VAR model where the dynamics of the endogenous variables are modelled using BART. The authors model the impact of uncertainty shocks using their proposed model. [Huber et al. \(2020\)](#) extend the BART-VAR to a mixed frequency setting and evaluate the forecasting performance of the model. [Clark et al. \(2021\)](#) show that multivariate BART regression models perform well in terms of tail forecasting. To the best of our knowledge, our paper is the first one to use BART in an LP framework. The paper is also part of a broad literature that studies the advantages of IRF estimation using LP estimators, relative to constructing IRFs on vector autoregressive models. Several contributions document the performance of LP estimators, including [Gonçalves et al. \(2022\)](#), [Kilian and Kim \(2011\)](#), [Alloza et al. \(2019\)](#), [Breitung et al. \(2019\)](#), [Herbst and Johannsen \(2021\)](#), and [Bruns and Lütkepohl \(2022\)](#). While LP estimators are usually proposed in a frequentist setting, we follow [Miranda-Agrippino and Ricco \(2021\)](#) and take a Bayesian approach to LP, yet in a non-linear framework.

The paper is organised as follows. [Section 2](#) presents the empirical model. [Section 3](#) reports the results from the simulation exercise. [Section 4](#) shows the application to fiscal and financial shocks. [Section 5](#) concludes.

2 Flexible local projections

In this section we outline the methodology, which we refer to as the BART-LP model, or flexible local projections. We discuss how BART-LP approximates the unknown conditional expectation function of local projection models, discuss the prior, and outline the posterior sampler. We then discuss how to compute generalized impulse responses to structural shocks within our framework.

2.1 The BART approximation

We work with the equation

$$y_{t+h} = m_h(x_t, \mathbf{z}_t, \mathbf{w}_{t+h}^{(h)}) + \epsilon_{t+h}^{(h)}, \quad (1)$$

where y_{t+h} denotes the scalar variable of interest, $h = 0, 1, \dots, H$ is the impulse response horizon, and $\epsilon_{t+h}^{(h)}$ satisfies $E(\epsilon_{t+h}^{(h)} | x_t, \mathbf{z}_t, \mathbf{w}_{t+h}^{(h)}) = 0$. We aim to study how y_{t+h} responds to a change in the scalar variable x_t . The vector \mathbf{z}_t contains observable control variables, possibly including lagged values of y_t and contemporaneous and/or lagged values of other variables. The vector $\mathbf{w}_{t+h}^{(h)}$ contains additional control variables in the form of estimated residuals, as explained in [Section 2.3](#). While (x_t, \mathbf{z}_t) are the same for every regression model h , $\mathbf{w}_{t+h}^{(h)}$ can potentially change. The function $m_h(\cdot)$ captures the true unknown conditional expectation function. The residual $\epsilon_{t+h}^{(h)}$ is assumed to be normally distributed with variance σ_{t+h}^2 . As noted in [Jordà \(2005\)](#), $\epsilon_{t+h}^{(h)}$ is serially correlated for $h \geq 1$.

It is common in the literature to assume a functional form for $m_h(x_t, \mathbf{z}_t, \mathbf{w}_{t+h}^{(h)})$. The most popular applications of LPs use a function of the type

$$y_{t+h} = g(q_t) [\alpha_0^{(h)} x_t + \alpha_1'^{(h)} \mathbf{z}_t + \alpha_2'^{(h)} \mathbf{w}_{t+h}^{(h)}] + \quad (2)$$

$$+ (1 - g(q_t)) [\beta_0^{(h)} x_t + \beta_1'^{(h)} \mathbf{z}_t + \beta_2'^{(h)} \mathbf{w}_{t+h}^{(h)}] + \epsilon_{t+h}^{(h)}. \quad (3)$$

The special case of a linear model sets $g(q_t) = 1, \forall t$, and estimates the impulse response to a shock to x_t using the estimates for $\{\alpha_0^{(h)}\}_{h=0}^H$ (Jordà, 2005). Non-linear applications usually specify a transition variable q_t , assume a specific non-linear functional form for $g(\cdot)$, and compute non-linear impulse responses as a function of the estimates for $\{\alpha_0^{(h)}, \beta_0^{(h)}\}_{h=0}^H$. For example, the smooth transition formulation by Auerbach and Gorodnichenko (2013a,b) sets $g(\cdot)$ equal to the logistic function, while the threshold formulation by Ramey and Zubairy (2018) and Alpanda et al. (2021) sets $g(\cdot)$ equal to the indicator function.

Our paper differs from the existing literature by approximating the unknown conditional expectation function $m_h(x_t, \mathbf{z}_t, \mathbf{w}_{t+h}^{(h)})$ using Bayesian Additive Regression Trees (BART). It is assumed that

$$m_h(x_t, \mathbf{z}_t, \mathbf{w}_{t+h}^{(h)}) \approx f_h(x_t, \mathbf{z}_t, \mathbf{w}_{t+h}^{(h)} | \Gamma^{(h)}, \boldsymbol{\mu}^{(h)}) = \sum_{j=1}^J f_{h,j}(x_t, \mathbf{z}_t, \mathbf{w}_{t+h}^{(h)} | \Gamma_j^{(h)}, \boldsymbol{\mu}_j^{(h)}), \quad (4)$$

where $f_{h,j}(x_t, \mathbf{z}_t, \mathbf{w}_{t+h}^{(h)} | \Gamma_j^{(h)}, \boldsymbol{\mu}_j^{(h)})$ denotes a single regression tree j at horizon h and $f_h(x_t, \mathbf{z}_t, \mathbf{w}_{t+h}^{(h)} | \Gamma^{(h)}, \boldsymbol{\mu}^{(h)})$ denotes the sum of J regression trees. The parameters of the regression trees are the tree structures $\Gamma^{(h)} = [\Gamma_1^{(h)}, \dots, \Gamma_J^{(h)}]$ and the terminal nodes (or leaves) $\boldsymbol{\mu}^{(h)} = (\boldsymbol{\mu}_1^{(h)'}, \dots, \boldsymbol{\mu}_J^{(h)'})'$, with $\boldsymbol{\mu}_j^{(h)}$ of dimension $B_j \times 1$ and B_j the number of terminal nodes of tree j . As an illustration, each regression tree divides the space of each explanatory variable by using binary splitting rules. Denoting X_i as the i -th entry of the vector $(x_t, \mathbf{z}_t, \mathbf{w}_{t+h}^{(h)})$, these rules are defined as:

$$X_i \leq c, \quad (5)$$

$$X_i > c,$$

with c the threshold value. Observations are assigned according to these splitting rules, and the terminal nodes return the fitted value conditional on the split. The fitted value of the dependent variable, based on a single regression tree, is then given

by

$$f_{h,j}(x_t, \mathbf{z}_t, \mathbf{w}_{t+h}^{(h)} | \Gamma_j^{(h)}, \boldsymbol{\mu}_j^{(h)}) = \sum_{b=1}^{B_j} I(X_i) \mu_{j,b}^{(h)}, \quad (6)$$

where $I(\cdot)$ denotes an indicator function that equals 1 if X_i belongs to the set defined by the splitting rule implicit in $\Gamma_j^{(h)}$. Note that the complexity of each tree is determined by B_j , the number of terminal nodes. We refer the reader to the Online Appendix for an illustrative example of a regression tree, and to [Hill et al. \(2020\)](#) for a comprehensive review.

The model in equation (4) approximates $m_h(x_t, \mathbf{z}_t, \mathbf{w}_{t+h}^{(h)})$ using a sum J trees. Each tree in the sum is restricted to be small *a priori* to avoid overfitting, and thus explains a small proportion of y_{t+h} and is a ‘weak learner’. [Chipman et al. \(2010\)](#) show that a low value of J reduces predictive accuracy. As J increases, predictive performance initially improves, but this improvement tapers off, eventually. In practice, studies such as [Huber et al. \(2020\)](#) note that the difference in predictive accuracy is negligible for $J > 150$ and complex functions can be easily approximated using $J = 200$ or 250.

The BART approximation of the relationship between y_{t+h} and x_t has implications for the properties of the impulse responses. As the regression trees split the space of the covariates via rules of the type shown in equation (5), the estimated predictions on the right-hand side of equation (15) are dependent on their history. Similarly, the shock d to variable x_t can lead to predictions that proportionally differ if the size and sign of the shock leads to the covariate space where the relationship between y_{t+h} and x_t is substantially different from the ‘average’ impact. We now discuss estimation, and then return to a detailed discussion of non-linear impulse responses in the BART-LP model in [Section 2.4](#).

2.2 Estimation

The model in equation (1) can be estimated using the MCMC algorithm described in [Chipman et al. \(2010\)](#), which we summarize here for completeness.

2.2.1 Priors

The prior distributions proposed by [Chipman et al. \(2010\)](#) play a crucial role, as they are devised to reduce the possibility of overfitting. The joint prior for the parameters of the J trees of the model at each horizon h is factored as follows:

$$p((\Gamma_1^{(h)}, \boldsymbol{\mu}_1^{(h)}), (\Gamma_2^{(h)}, \boldsymbol{\mu}_2^{(h)}), \dots, (\Gamma_J^{(h)}, \boldsymbol{\mu}_J^{(h)})) = \prod_{j=1}^J p(\boldsymbol{\mu}_j^{(h)} | \Gamma_j^{(h)}) p(\Gamma_j^{(h)}), \quad (7)$$

where $p(\boldsymbol{\mu}_j^{(h)} | \Gamma_j^{(h)}) = \prod_{b=1}^{B_j} p(\mu_{b,j}^{(h)} | \Gamma_j^{(h)})$.

The prior for the tree structure $\Gamma_j^{(h)}$ depends on the probability that the node at depth $d = 0, 1, 2, \dots$ is not a terminal node. This prior probability is given by $\alpha(1+d)^{-\beta}$ where $\alpha \in (0, 1)$ and $\beta > 0$. Higher values of β and smaller values of α reduce this probability and impose a stronger belief that the tree has a simple (i.e. shorter) structure. We follow the recommendation by [Chipman et al. \(2010\)](#) and set $\alpha = 0.95$ and $\beta = 2$. The prior for the threshold value c implies that this parameter is assumed to be uniform over the range of the values taken by the variables. In the default setting, the choice of splitting variable is also assumed to be uniform across the regressors.

To define $p(\boldsymbol{\mu}_j^{(h)} | \Gamma_j^{(h)})$ [Chipman et al. \(2010\)](#) first transform the dependent variable so that it lies between -0.5 and 0.5 . As a consequence, $m_h(x_t, \mathbf{z}_t, \mathbf{w}_{t+h}^{(h)})$ is also expected to lie between these values. The prior $p(\boldsymbol{\mu}_j^{(h)} | \Gamma_j^{(h)})$ is assumed to be normal $N(0, S)$. The variance S is set as $\frac{1}{2\kappa(J^{0.5})}$, with κ set to 2, the value recommended by [Chipman et al. \(2010\)](#). Under this default prior, there is a 95% probability that the conditional mean of the dependent variable lies between -0.5 and 0.5 .

A conjugate inverse χ^2 prior is used for the variance σ_{t+h}^2 . The hyperparameters of the prior distribution are set by using an estimate $\hat{\sigma}_{t+s}^2$ of the variance obtained from a linear regression. If the true model is non-linear $\hat{\sigma}_{t+s}^2$ will be biased upwards. Under the default prior, the hyperparameters are chosen so that $\Pr(\sigma_{t+s}^{(h)} < \hat{\sigma}_{t+s}^{(h)}) = 0.9$.

The total number of trees J is fixed.

2.2.2 MCMC algorithm

The MCMC algorithm devised by [Chipman et al. \(2010\)](#) samples from the conditional posterior distributions of σ_{t+h}^2 and the parameters of the regression trees in each iteration.¹ Each iteration of the algorithm samples from the following conditional posteriors:

1. conditional on the trees, the error variance can be easily drawn from the inverse Gamma distribution;
2. the conditional posterior distribution of the tree structure is not known in closed form and a Metropolis-Hastings algorithm is used. Define $R_j^{(h)}$ as the residual:

$$R_j^{(h)} = y_{t+h} - \sum_{i \neq j} f(x_t, \mathbf{z}_t, \mathbf{w}_{t+h}^{(h)} | \Gamma_j^{(h)}, \boldsymbol{\mu}_j^{(h)}). \quad (8)$$

The j – th tree is proposed using the density $q(\Gamma_j^{new}, \Gamma_j^{old})$. [Chipman et al. \(2010\)](#) use a proposal density that incorporates 4 moves: (i) splitting the node into two new nodes (grow), (ii) transforming adjacent nodes to terminal node (prune), (iii) changing the decision rule of an interior node (change), (iv) swapping a decision rule between a node that is above and the node before it (swap). The probabilities associated with these moves are fixed at 0.25, 0.25, 0.4 and 0.1 respectively. The proposed tree structure Γ_j^{new} is accepted with probability

$$\alpha = \frac{q(\Gamma_j^{new}, \Gamma_j^{old})p(R_j^{(h)} | \Gamma_j^{new}, \sigma_{t+h}^2)p(\Gamma_j^{new})}{q(\Gamma_j^{old}, \Gamma_j^{new})p(R_j^{(h)} | \Gamma_j^{old}, \sigma_{t+h}^2)p(\Gamma_j^{old})}, \quad (9)$$

where $p(R_j | \Gamma_j, \sigma_{t+h}^2)$ is the conditional likelihood and $p(\Gamma_j)$ denotes the prior. This step is repeated for $j = 1, 2, \dots, J$ trees;

3. the conditional posterior distribution of the terminal node parameters is Gaus-

¹Intuitive descriptions of this MCMC algorithm can be found in [Clark et al. \(2021\)](#) and [Hill et al. \(2020\)](#).

sian with the parameters known in closed form. Therefore, the draw of $\boldsymbol{\mu}_j^{(h)}$ for $j = 1, 2, \dots, J$ can be carried out in a straightforward manner;

4. given a draw of the model parameters conditioning on $(x_t, \mathbf{z}_t, \mathbf{w}_{t+h}^{(h)})$, the predicted value can be computed as

$$E\left(y_{t+h} | x_t, \mathbf{z}_t, \mathbf{w}_{t+h}^{(h)}\right) = \sum_{j=1}^J f_h(x_t, \mathbf{z}_t, \mathbf{w}_{t+h}^{(h)} | \Gamma_j^{(h)}, \boldsymbol{\mu}_j^{(h)}) = \sum_{j=1}^J \sum_{b=1}^{B_j} I(x_t, \mathbf{z}_t, \mathbf{w}_{t+h}^{(h)}) \mu_{j,b}^{(h)}, \quad (10)$$

with $I(x_t, \mathbf{z}_t, \mathbf{w}_{t+h}^{(h)})$ an indicator function equal to 1 if $(x_t, \mathbf{z}_t, \mathbf{w}_{t+h}^{(h)})$ belongs to the splitting rule implied by $[\Gamma_1^{(h)}, \dots, \Gamma_J^{(h)}]$.

2.3 Autocorrelation

The residual term in LP models is known to be autocorrelated, a feature that must be taken into account in the estimation. In the case of linear local projections, it has been shown that the residual at horizon h follows a $MA(h-1)$ process, see for instance [Lusompa \(2021\)](#). [Lusompa \(2021\)](#) suggests a GLS procedure whereby the autocorrelation is eliminated by including leads of the LP residuals from horizon $h=0$ in the conditioning set.²

The non-parametric setting considered in this paper encompasses non-linear models. For the purpose of illustration, consider a simple non-parametric AR(1) model

$$y_{t+1} = v(y_t; A_1) + e_{t+1}. \quad (11)$$

Iterating the process forward 3 periods as an example gives

$$\begin{aligned} y_{t+2} &= v(v(y_t; A_1) + e_{t+1}; A_1) + e_{t+2}, \\ y_{t+3} &= v(v(v(y_t; A_1) + e_{t+1}; A_1) + e_{t+2}; A_1) + e_{t+3}. \end{aligned}$$

²[Lusompa \(2021\)](#) suggests an efficient strategy that transforms the dependent variable of the LP regressions and does not require one to explicitly include the horizon 0 residuals as regressors.

It is useful to compare this with a BART-LP for this horizon:

$$y_{t+3} = f_3(y_t|\Gamma_3, \boldsymbol{\mu}_3) + \epsilon_{t+3}^{(3)}. \quad (12)$$

The function $f_3(y_t|\Gamma_3, \boldsymbol{\mu}_3)$ approximates the non-linear relationship between y_t and its lead, but does not account for the dependence between the dependent variable and lagged shocks. Thus in this setting, the residual $\epsilon_{t+3}^{(3)}$ is a non-linear function of e_{t+1} and e_{t+2} , and has a non-linear autocorrelation structure.

In general, the Volterra expansion of any non-linear time-series shows its complex dependence on past shocks:

$$y_t = \sum_{i=0}^{\infty} \phi^i e_{t-i} + \sum_{i=0}^{\infty} \sum_{j=i}^{\infty} \zeta_{ij} e_{t-i} e_{t-j} + \sum_{i=0}^{\infty} \sum_{j=0}^{\infty} \sum_{k=0}^{\infty} \phi_{ij} e_{t-i} e_{t-j} e_{t-k} + \dots \quad (13)$$

To account for this autocorrelation, we propose to include an estimate of the shocks, $\boldsymbol{w}_{t+h}^{(h)} = (\hat{e}_{t+1}, \hat{e}_{t+2}, \dots, \hat{e}_{t+h-1})$, as additional covariates in the h -period BART-LP. Following Lusompa (2021) we construct $\boldsymbol{w}_{t+h}^{(h)}$ at every horizon $h \geq 1$ using the residuals of the period $h = 0$ flexible local projection $y_t = f_0(x_t, \boldsymbol{z}_t, \boldsymbol{w}_t^{(0)}|\Gamma^{(0)}, \boldsymbol{\mu}^{(0)}) + \epsilon_t^{(0)}$, with $\boldsymbol{w}_t^{(0)} = \mathbf{0}$. Then, the flexible local projection for period h is specified as

$$y_{t+h} = f_h(x_t, \boldsymbol{z}_t, \boldsymbol{w}_{t+h}^{(h)}|\Gamma^{(h)}, \boldsymbol{\mu}^{(h)}) + \epsilon_{t+h}^{(h)}. \quad (14)$$

The BART approximation of the true non-linear function $m_h(x_t, \boldsymbol{z}_t, \hat{w}_{t+h}^{(h)})$ proxies the non-linear dependence of y_{t+h} on $e_{t+1}, e_{t+2}, \dots, e_{t+h-1}$ and ameliorates the autocorrelation in $\epsilon_{t+h}^{(h)}$.

2.4 Generalized structural impulse responses

The computation of impulse responses to structural shocks typically relies on two conceptually different pillars: an estimation procedure for the impulse responses, and an identification scheme for the structural shock of interest. Our paper only aims to

advance the literature on the former, and does not investigate the topic of identification in a non-parametric framework.

Define \mathbf{y}_t a $k \times 1$ vector of variables of interest, with x_t one of the entries of \mathbf{y}_t . Define $\phi_{i,h}$ the impulse response of variable $y_{i,t}$ to a structural shock to variable x_t of size d , with $i = 1, \dots, k$. Following [Koop et al. \(1996\)](#), $\phi_{i,h}$ is given by

$$\phi_{i,h} = E(y_{i,t+h}|x_t = d; \mathbf{n}_t) - E(y_{i,t+h}|x_t, \mathbf{n}_t). \quad (15)$$

We compute generalized impulse responses numerically using the following algorithm:

1. estimate the model for $h = 0$,

$$y_{i,t} = f_0(x_t, \mathbf{z}_t, \mathbf{w}_t^{(0)} | \Gamma^{(0)}, \boldsymbol{\mu}^{(0)}) + \epsilon_{i,t}^{(0)}, \quad (16)$$

with $\mathbf{w}_t^{(0)} = \mathbf{0}$, and store D vectors of dimension $T \times 1$ containing the estimated residuals associated with D posterior draws, $\{\hat{\epsilon}_{i,t}^{(0),d}\}_{t=1}^T$, $d = 1, \dots, D$;

2. for the generic draw d of the residuals $\{\hat{\epsilon}_{i,t}^{(0),d}\}_{t=1}^T$ and for the generic horizon h , estimate the model

$$y_{i,t+h} = f_h(x_t, \mathbf{z}_t, \mathbf{w}_{t+h}^{(h)} | \Gamma^{(h)}, \boldsymbol{\mu}^{(h)}) + \epsilon_{i,t+h}^{(h)}, \quad (17)$$

with

$$\mathbf{w}_{t+h}^{(h)} = (\hat{\epsilon}_{i,t}^{(0),d}, \hat{\epsilon}_{i,t-1}^{(0),d}, \dots, \hat{\epsilon}_{i,t-h+1}^{(0),d}); \quad (18)$$

3. compute the predicted values $(\hat{y}_{i,h}^{0,d}, \hat{y}_{i,h}^{1,d})$ associated with one posterior draw from model (17) conditioning on

$$(x_t = \bar{x}, \quad \mathbf{z}_t = \bar{\mathbf{z}}, \quad \mathbf{w}_{t+h}^{(h)} = \bar{\mathbf{w}}^{(h)}), \quad (19)$$

$$(x_t = \bar{x} + d, \quad \mathbf{z}_t = \bar{\mathbf{z}}, \quad \mathbf{w}_{t+h}^{(h)} = \bar{\mathbf{w}}^{(h)}), \quad (20)$$

respectively, with $(\bar{x}, \bar{z}, \bar{w}^{(h)})$ defined below. Compute $\psi_{i,h}^d = \hat{y}_{i,h}^{1,d} - \hat{y}_{i,h}^{0,d}$ and store

$$(\hat{y}_{i,h}^{0,d}, \hat{y}_{i,h}^{1,d}, \psi_{i,h}^d); \quad (21)$$

4. repeat steps 2-3 for $h = 0, \dots, H$, store $\{(\hat{y}_{i,h}^{0,d}, \hat{y}_{i,h}^{1,d}, \psi_{i,h}^d)\}_{h=0}^H$ given the same posterior draw d for $\{\hat{\epsilon}_{i,t}^{(0),d}\}_{t=1}^T$;
5. repeat steps 2-4 for $d = 1, \dots, D$;
6. compute the average across posterior draws

$$\psi_{i,h} = \frac{1}{D} \sum_{d=1}^D \hat{y}_{i,h}^{1,d} - \frac{1}{D} \sum_{d=1}^D \hat{y}_{i,h}^{0,d}. \quad (22)$$

The exact implementation of the above procedure depends on how the structural shocks are identified. Different options are in principle available. One option is to follow [Plagborg-Møller and Wolf \(2021\)](#) and [Barnichon and Brownlees \(2019\)](#) and replicate a recursive identification scheme by using appropriate control variables. For example, if the aim is to estimate the response of GDP to an interest rate shock that is restricted to have a zero contemporaneous impact on GDP and CPI in a trivariate model, one can set x_t equal to the policy interest rate and add contemporaneous GDP and CPI into \mathbf{z}_t , which will also include L lags of the variables. Alternatively, one can set x_t equal to either a proxy for the structural shock of interest, or equal to the true realizations of the shocks, if available, and set $\mathbf{z}_t = (\mathbf{y}'_{t-1}, \dots, \mathbf{y}'_{t-L})'$. Either way, the generalized impulse responses are computed as $\phi_{i,h} = \psi_{i,h}$, with $\psi_{i,h}$ from equation (22).

An alternative approach is to follow [Jordà \(2005\)](#) more closely and use a separate impulse vector. One can estimate a preliminary SVAR on \mathbf{y}_t and use a preferred identification approach to estimate the impulse vector \mathbf{d} of variables \mathbf{y}_t to a shock to variable x_t of size d . The above algorithm can then be run by setting $x_t = 0$ and $\mathbf{z}_t = (\mathbf{y}'_{t-1}, \dots, \mathbf{y}'_{t-L})'$. Step 1 of the algorithm is still required to generate $\{\hat{\epsilon}_{i,t}^{(0),d}\}_{t=1}^T$,

$d = 1, \dots, D$. Steps 2 to 4 are then run for $h = 1, 2, \dots, H$, replacing equations (19)-(20) with

$$(x_t = 0, \mathbf{z}_t = \bar{\mathbf{z}}, \quad \mathbf{w}_{t+h}^{(h)} = \bar{\mathbf{w}}^{(h)}), \quad (23)$$

$$(x_t = 0, \mathbf{z}_t = \bar{\mathbf{z}} + \mathbf{d}, \quad \mathbf{w}_{t+h}^{(h)} = \bar{\mathbf{w}}^{(h)}). \quad (24)$$

The generalized impulse responses are then computed as $\phi_{i,h} = d_i$ for $h = 0$ with d_i the $i - th$ entry of \mathbf{d} , and as $\phi_{i,h} = \psi_{i,h}$ from (22) for $h = 1, \dots, H$.

Last, to be operational, the above procedure requires specifying the conditioning values $(\bar{x}, \bar{\mathbf{z}}, \bar{\mathbf{w}}^{(h)})$ as well as the shock d . The former can be set differently depending on whether the shock is simulated to hit the economy at any point in time or on a subset of periods. Our baseline specification sets $(\bar{x}, \bar{\mathbf{z}}, \bar{\mathbf{w}}^{(h)})$ equal to the sample average, which is computed within each regression model h . d , instead, can be set in accordance to the sign and size of the intended shock.

3 Monte Carlo simulation

We use Monte Carlo simulations to assess if BART-LP can recover the true non-linear patterns of a data generating process. We use three different data generating processes: a SVAR-GARCH model, a Threshold VAR model, and a sign-dependent Moving Average model.

3.1 SVAR-GARCH model

The first model we use for simulations is the model from Section III.B in [Jordà \(2005\)](#),

$$\begin{pmatrix} y_{1t} \\ y_{2t} \\ y_{3t} \end{pmatrix} = A \begin{pmatrix} y_{1t-1} \\ y_{2t-1} \\ y_{3t-1} \end{pmatrix} + Bh_t + \begin{pmatrix} \sqrt{h_t}\epsilon_{1t} \\ \epsilon_{2t} \\ \epsilon_{3t} \end{pmatrix}, \quad (25)$$


$$(\epsilon_{1t}, \epsilon_{2t}, \epsilon_{3t})' = \epsilon_t \sim N(\mathbf{0}, I_3), \quad (26)$$

$$h_t = 0.5 + 0.5h_{t-1} + 0.3\sqrt{h_t}\epsilon_{1t}, \quad (27)$$

$$B = \begin{pmatrix} -1.75 \\ -1.5 \\ 1.75 \end{pmatrix} \quad A = \begin{pmatrix} 0.5 & -0.25 & 0.25 \\ 0.75 & 0.25 & 0.25 \\ -0.25 & -0.25 & 0.75 \end{pmatrix}. \quad (28)$$

In this model, each shock i affects only variable i contemporaneously, while all shocks affect all variables after one period. Contrary to the second and the third shock, the first shock features time-varying variance.³

Figure 1: Monte Carlo simulation - SVAR-GARCH



We use a simulation exercise to study how well the BART-LP methodology recovers

³We code the simulation exercise following the exact code available in the replication files from [Jordà \(2005\)](#).

the true impulse responses associated with a shock to y_{1t} of size 1. We first use generalized impulse responses to simulate the true impulse responses associated with $\epsilon_{1t} = 1$. We then generate simulated data for 300 periods and discard the first 100, as in [Jordà \(2005\)](#). The remaining 200 observations are used to estimate the pointwise linear LP impulse response, as well as the impulse response computed from the BART-LP method, setting $d = 1$. We replicate the exercise 100 times and store the generated estimates. The estimated model uses 2 lags, and sets x_t equal to the true realizations of ϵ_{1t} . We set $J = 250$.

[Figure 1](#) reports the results of the analysis. The squared black line shows the true generalized impulse responses. The shock increases y_{1t} on impact by 1, and generates no contemporaneous response in (y_{2t}, y_{3t}) . It subsequently generates an oscillating pattern in (y_{1t}, y_{2t}) and a hump-shaped response of y_{3t} . The dashed line shows the pointwise mean impulse response from the linear LP model, with the mean computed on the pointwise estimates along the 100 iterations. The continuous line and the shaded area report the pointwise median and 95% band computed over the 100 median responses from the BART-LP method, where within each iteration, the median response from BART-LP is computed over 2,000 posterior draws. We see that BART-LP does remarkably well in replicating the response of all variables, capturing both the timing and the magnitude of the response correctly. By contrast, the response estimated with the linear model estimated an attenuated effect, underestimating the effect of the shock in the short horizon.

3.2 Threshold VAR model

The second model we use for simulations is

$$\mathbf{y}_t = [\Pi_1 \mathbf{y}_{t-1} + B_1 \boldsymbol{\epsilon}_t] \cdot \mathbf{I}(y_{3,t-1} \leq 0) + \quad (29a)$$

$$+ [\Pi_2 \mathbf{y}_{t-1} + B_2 \boldsymbol{\epsilon}_t] \cdot \mathbf{I}(y_{3,t-1} > 0), \quad (29b)$$

$$\boldsymbol{\epsilon}_t \sim N(\mathbf{0}, I), \quad (29c)$$

with $\mathbf{y}_t = (y_{1t}, y_{2t}, y_{3t})'$, $\boldsymbol{\epsilon}_t = (\epsilon_{1t}, \epsilon_{2t}, \epsilon_{3t})'$ and


$$\Pi_1 = \begin{pmatrix} 0.25 & 0.25 & -0.25 \\ -0.25 & 0.25 & -0.25 \\ 0.25 & 0.25 & 0.15 \end{pmatrix}, \quad B_1 = \begin{pmatrix} 0.10 & 0 & 0 \\ -0.20 & 0.15 & 0 \\ 0.10 & -0.10 & 1 \end{pmatrix}, \quad (30)$$

$$\Pi_2 = \begin{pmatrix} 0.50 & 1.25 & -1.75 \\ -0.25 & 0.50 & -1.25 \\ 0.25 & 0.25 & 0.15 \end{pmatrix}, \quad B_2 = \begin{pmatrix} 0.10 & 0 & 0 \\ -0.20 & 0.15 & 0 \\ 0.10 & -0.10 & 0.40 \end{pmatrix}. \quad (31)$$

The model is a recursive Threshold Vector Autoregressive model that jumps across two regimes depending on the endogenous evolution of the third variable. See, for instance, [Castelnuovo and Pellegrino \(2018\)](#).

We first compute the true generalized impulse responses to a positive one-standard-deviation shock to a y_{3t} in regime 1 by setting the initial condition of the impulse response to $\mathbf{y}_0 = \mathbf{0}$. We then study how well the linear-LP and the BART-LP estimators recover the true impulse responses. We generate a dataset of 300 observations, discard the first 100 and use the remaining 200 to estimate the impulse responses. We compute the point estimates of the linear LP, and compute the pointwise median over 2,000 posterior draws from the BART-LP estimates. We then repeat the exercise over 100 iterations. The estimated models use 4 lags and set (x_t, \mathbf{z}_t) to replicate the recursive ordering within LP models, as discussed in [Section 2.4](#). We set $J = 250$.

Figure 2: Monte Carlo simulation - TVAR



The shock hits the system when the system is in regime 1. Under linearity, the model would stay in regime 1, with an impulse response uniquely pinned down by (Π_1, B_1) . Instead, the non-linearity of the model implies an endogenous evolution across regimes, in accordance with the endogenous response of y_{3t} . Figure 2 shows the results of the exercise. The pink continuous line shows the true linear impulse response associated with regime 1. The squared black line captures the true generalized impulse responses. The two lines differ. The dashed line shows the estimated response from the linear-LP. The figure shows that this response lies between the true generalized and the true linear impulse responses. By contrast, the impulse response estimated via BART-LP better estimates the true generalized impulse response. The model correctly captures the evolution of all three variables. It detects that the endogenous evolution across regimes makes the response of the first two variables more pronounced compared to a linear model that remains in regime 1.

3.3 Sign-dependent Moving Average model

The third model we use for simulation is

$$\mathbf{y}_t = \sum_{l=0}^{20} \boldsymbol{\beta}_{gdp, l} \cdot \epsilon_{gdp, t-l} + \sum_{l=0}^{20} \boldsymbol{\beta}_{\pi, l} \cdot \epsilon_{\pi, t-l} + \quad (32)$$

$$+ \sum_{l=0}^{20} \left[\boldsymbol{\beta}_{ff, l}^+ \cdot \mathbf{I}(\epsilon_{ff, t-l} \geq 0) + \boldsymbol{\beta}_{ff, l}^- \cdot \mathbf{I}(\epsilon_{ff, t-l} < 0) \right] \cdot \epsilon_{ff, t-l}, \quad (33)$$

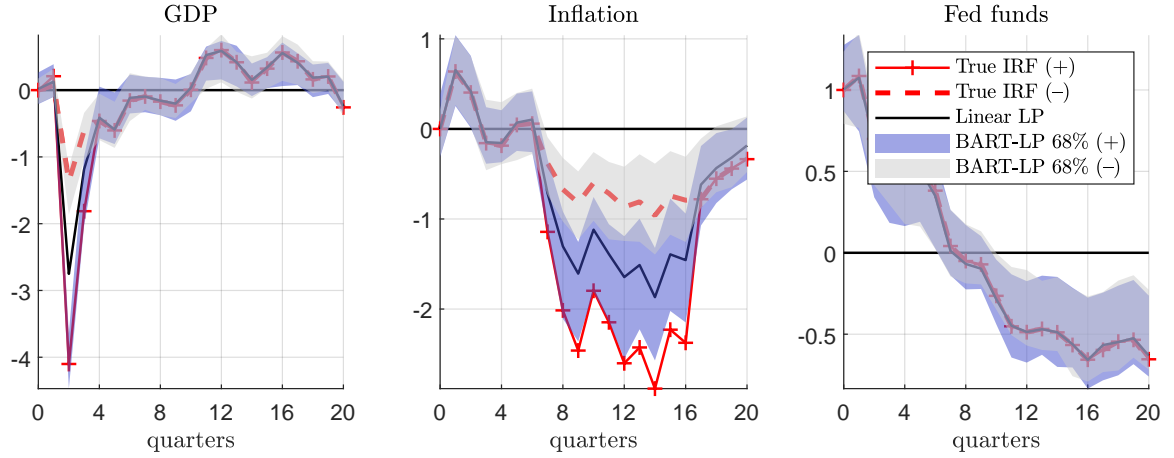
$$\boldsymbol{\epsilon}_t \sim N(\mathbf{0}, I), \quad (34)$$

with $\mathbf{y}_t = (gdp_t, \pi_t, ff_t)'$ a vector containing real GDP, inflation, and the federal funds rate, and $\boldsymbol{\epsilon} = (\epsilon_{gdp, t}, \epsilon_{\pi, t}, \epsilon_{t, ff})'$ a vector of structural shocks. The model is a moving average process of order 20 driven by three structural shocks, a GDP shock, an inflation shock, and a monetary policy shock. The monetary policy shock affects the variables differently at each horizon $t + h$, depending on whether the monetary policy shock at time t is positive or negative. The true impulse responses are captured by $\{\boldsymbol{\beta}_{gdp, l}\}_{l=0}^{20}$ for the GDP shock, by $\{\boldsymbol{\beta}_{\pi, l}\}_{l=0}^{20}$ for the inflation shock, and by $\{\boldsymbol{\beta}_{ff, l}^+, \boldsymbol{\beta}_{ff, l}^-\}_{l=0}^{20}$ for the positive and negative monetary policy shock.

We calibrate the model following an approach similar to [Barnichon and Brownlees \(2019\)](#). We first estimate a recursive VAR model on US real GDP, inflation and the federal funds rate to estimate linear impulse responses to a GDP shock, an inflation shock and a monetary policy shock. We then set $\{\boldsymbol{\beta}_{gdp, l}, \boldsymbol{\beta}_{\pi, l}, \boldsymbol{\beta}_{ff, l}^-\}_{l=0}^{20}$ equal to the estimated impulse responses to three shock: the GDP shock, the inflation shock, and the monetary policy shock. Last, we set $\boldsymbol{\beta}_{ff, l}^+ = \boldsymbol{\beta}_{ff, l}^-$ for every l , with two exceptions: (1) at horizons $l = 2, 3$ the first entry of $\boldsymbol{\beta}_{ff, l}^+$ equals 3 times the first entry $\boldsymbol{\beta}_{ff, l}^-$, and (2) at horizons $l = 7, 8, \dots, 20$ the second entry of $\boldsymbol{\beta}_{ff, l}^+$ equal 3 times the second entry of $\boldsymbol{\beta}_{ff, l}^-$. This implies a data generating process in which positive monetary policy shocks affect output and inflation 3 times more than negative shocks in the short term (for output) and in the medium term (for inflation) horizons.

We design the simulation exercise as follows. We generate an artificial dataset

Figure 3: Monte Carlo simulation - Sign-dependent Moving Average



of size $T = 400$, discard the first 100 observations and use the remaining 300 for estimation. We estimate the BART-LP model by controlling for the true shocks driving the data, together with up to 2 lags of the shocks and the endogenous variables. Last, we replicate the analysis 100 times and store the pointwise responses from the linear and the BART LPs, computing separately the generalized impulse responses associated with $d \pm 1$. We set $J = 250$.

Figure 3 shows the impulse responses. The (+) and (-) red lines show the true impulse responses associated with a positive and a negative monetary policy shock, respectively. The responses to a negative shock are reported with flipped sign to improve the comparison. By construction, the true response of the federal funds rate does not change in the sign of the shock, while the true response of GDP and inflation is 3 times stronger in response to a positive shock in periods in which the estimated linear responses imply strong responses of GDP and inflation. The black continuous line shows the pointwise median response associated with the linear LP. Note that it sits approximately halfway through the true positive and negative responses. The blue and grey shaded areas report the 68% pointwise posterior bands associated with the BART-LP model following a positive and a negative shock, respectively. The BART-LP methodology estimates the true response of GDP correctly both qualitatively and

quantitatively. It correctly detects that positive shocks have stronger short run effects on GDP than negative shocks, and also estimates the true values of the responses accurately. As for inflation, BART-LP correctly estimates the differential response qualitatively, detecting the stronger effect of positive shocks. While the response to a positive shock is somewhat underestimated, the response to a negative shock sits approximately in the middle of the estimated bands. All in the all, the model correctly detects that positive and negative shocks do not have the same effect.

Table 1: Monte Carlo Probability that the effect is stronger for a positive shock

| h | GDP | | | Inflation | | | Fed funds | | |
|-----|------------|-----------|------------|-----------|-----------|-----------|-----------|-----------|-----------|
| | $T = 100$ | $T = 200$ | $T = 300$ | $T = 100$ | $T = 200$ | $T = 300$ | $T = 100$ | $T = 200$ | $T = 300$ |
| 0 | 49 | 48 | 54 | 54 | 52 | 59 | 58 | 47 | 56 |
| 2 | 100 | 99 | 100 | 59 | 60 | 54 | 56 | 44 | 46 |
| 3 | 81 | 86 | 91 | 50 | 53 | 47 | 54 | 39 | 46 |
| 6 | 54 | 43 | 38 | 55 | 49 | 49 | 47 | 44 | 49 |
| 14 | 49 | 51 | 45 | 73 | 93 | 89 | 43 | 52 | 53 |
| 15 | 50 | 54 | 40 | 68 | 81 | 82 | 49 | 48 | 49 |

We further assess the ability of the model to detect differences in positive and negative monetary policy shocks by replicating the analysis over datasets of different lengths. We set the estimation sample period equal to $T = 100, 200, 300$ observations, keeping the initial discarded observations to 100. For each sample size, we replicate the analysis 100 times and compute the percentage of Monte Carlo iterations in which the effect is stronger after a positive rather than a negative shock in absolute value. [Table 1](#) reports the results for few illustrative horizons, indicating in bold the horizons and variables for which the data generating process features a non-linearity in the response to positive and negative shocks. On impact, the data generating process implies no non-linearity between positive and negative shocks. Indeed, the estimated impulse responses are well distributed across iterations, with around 50% of the iterations detecting a stronger effect of positive shocks and the remaining 50% detecting the opposite. At horizons 2-3, close to all iterations detect that the effects of a monetary shock are stronger on GDP, consistent with the data generating process, while correctly not finding evidence of non-linearities in the response of inflation and the policy rate. 6

horizons from the shocks the model correctly detects no non-linearity in any variable. At horizons 14-15 it correctly detects a stronger effect on inflation associated with positive compared to negative shocks, and no non-linearity in the remaining variables.

4 Empirical analysis

In this section, we apply the proposed model to two recent issues that have featured prominently in the empirical literature on non-linear macroeconomic dynamics.

4.1 Fiscal shocks during recessions and expansions

In a seminal contribution [Auerbach and Gorodnichenko \(2013a\)](#) use a smooth transition VAR model to show that the response of output to government spending shocks is larger during recessions. However, this evidence was disputed by [Ramey and Zubairy \(2018\)](#) who use a 120 year sample of quarterly data to estimate the response to military spending news shocks and show that there is no systematic difference between the spending multiplier in normal periods and those characterised by slack. Both papers use a non-linear LP model, but postulate a functional form for the transition, using either a logistic function or an indicator function.

We build on section VI-B of [Ramey and Zubairy \(2018\)](#) and re-visit the analysis of state-dependent fiscal multipliers using our flexible BART-LP model from equation (1). We set y_{t+h} equal to h -period ahead of (1) military spending news ($news_t$), (2) real per-capita government spending (g_t), and (3) real per-capita GDP (y_t).⁴ We first estimate a linear SVAR model in these three variables and estimated the impulse vector \mathbf{d} associated with the first shock. We then compute impulse responses as explained in [Section 2.4](#). We include 4 lags of the variables into the model. We set the number of

⁴The data set collated by [Ramey and Zubairy \(2018\)](#) is quarterly and runs from 1889 Q3 to 2015 Q3. The variables used in our analysis are downloaded from the website of the [Journal of Political Economy](#). Government spending and GDP are transformed via the procedure described in [Gordon and Krenn \(2010\)](#).

trees to 250 and use 2,000 posterior draws, with a burn-in of 1,000 draws.

We study non-linear effects of government spending shocks as follows. We define periods of recessions as periods where y_t is less than its 20th percentile, while expansions denote periods where y_t is above the 80th percentile. We then compute generalized impulse responses by generalizing over randomly drawn values of the conditioning values. Compared to [Auerbach and Gorodnichenko \(2013a\)](#) and [Ramey and Zubairy \(2018\)](#), our framework also allows for the computation of impulse responses that potentially differ for the size of the shock. For this reason, we compute the generalized impulse responses to either a contractionary or an expansionary fiscal spending shock, both studied either in a recession or an expansion.

Figure 4: Response to positive (top panel) and negative (bottom panel) military news shocks during recessions and expansions

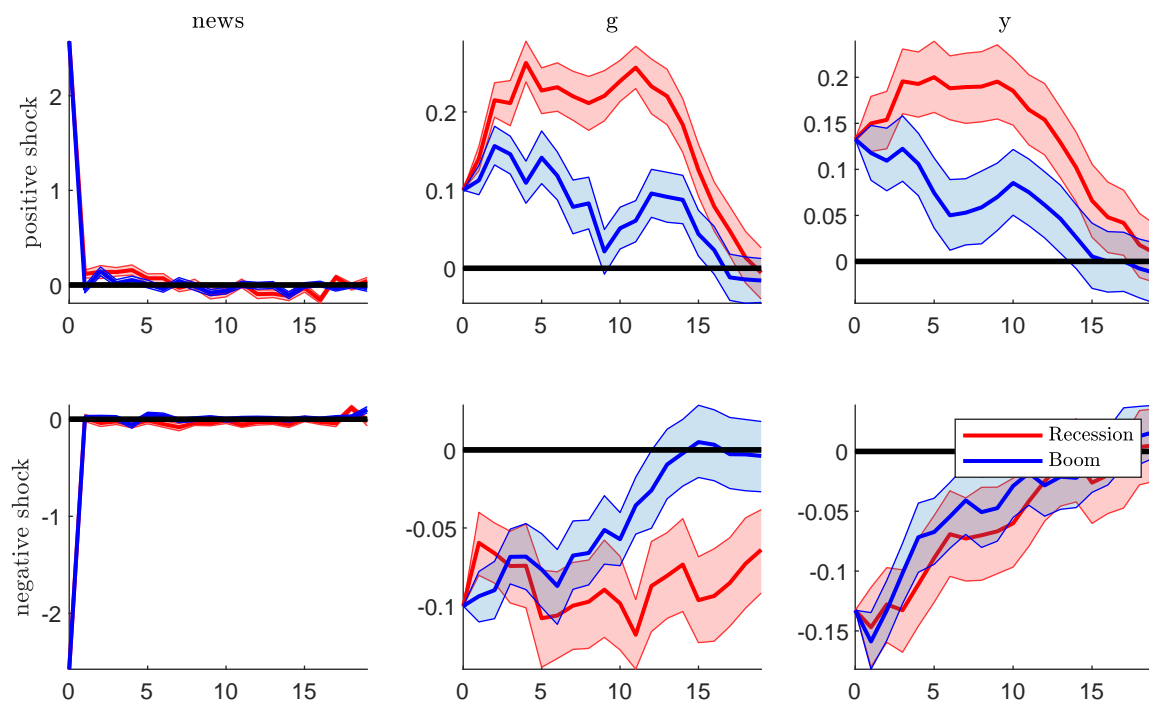
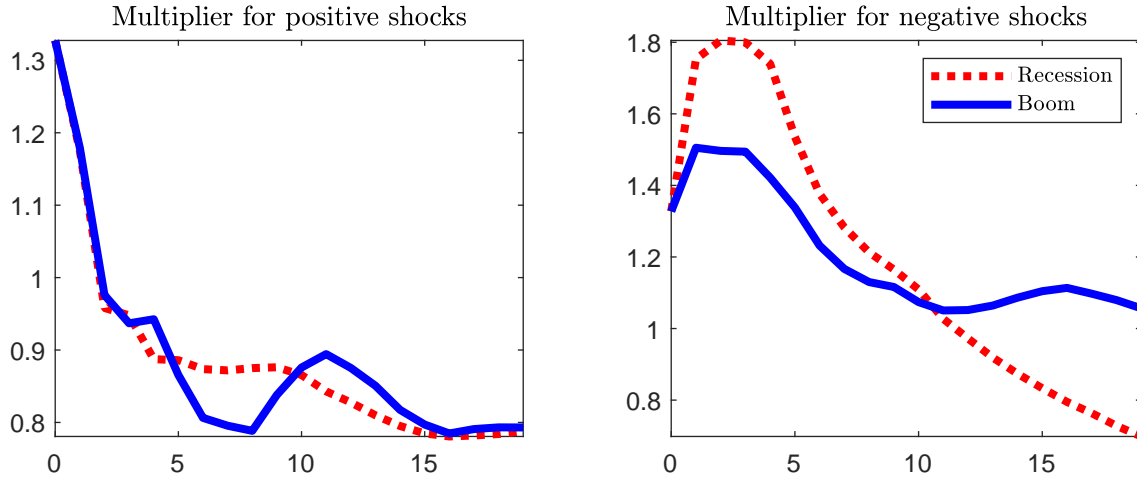


Figure 5: Cumulated multiplier



The top panel of [Figure 4](#) presents the response to an expansionary spending shock normalised to increase g by 0.1 units on impact. The response of g and y is estimated to be more persistent during recessions. However, the response of these two variables moves closely together in the two regimes. As a consequence, the spending multiplier (defined as the ratio of the cumulated response of y to the cumulated response of g) is similar across regimes. This can be seen from the left panel of [Figure 5](#), which displays the estimated multipliers and shows that evidence for a systematically larger multiplier during recessions is weak. However, it is interesting to note that contractionary spending shocks produce different dynamics. The bottom panel of [Figure 4](#) shows that the response of y to these shocks during expansions is less persistent. As a consequence the multiplier for spending cuts during these periods is systematically smaller than the estimate during recessions. This distinction helps reconcile the different results found by [Auerbach and Gorodnichenko \(2013a\)](#) and [Ramey and Zubairy \(2018\)](#) by documenting that state-specific non-linearities can be different depending on the sign of the shock.

4.2 Non-linear impact of Financial shocks


In a recent contribution [Forni et al. \(2021\)](#) use a non-linear vector moving average model to show that large negative financial shocks have a proportionally larger impact on the US economy. We revisit this question using the flexible local projection proposed in this paper. We use the model from equation (1) setting y_{t+h} equal to (1) growth of industrial production (IP_t), (2) CPI inflation (CPI_t), (3) unemployment rate (U_t), (4) Excess bond premium (EBP_t) of [Gilchrist and Zakrajsek \(2012\)](#), (5) stock returns ($STOCK_t$) and (6) federal funds rate (FFR_t). The data is monthly and the sample runs from 1973M1 to 2022 M2.⁵

We follow the identification approach by [Forni et al. \(2021\)](#) and identify the monetary policy shock as the shock that affects fast moving variables but not slow moving variables. We estimate a preliminary linear VAR model using the above ordering of the variables and estimate the impact vector associated with the entry of the excess bond premium in a recursive identification. We then add 6 lags of all variables. We simulate positive and negative shocks, considering either a shock that moves the excess bond premium by ± 50 basis points, or by ± 300 basis points.

The results of the analysis are shown in [Figure 6](#) and [Figure 7](#), which report the impulse responses associated with the small or the large exogenous variation of the excess bond premium, respectively. The response to expansionary shocks has been multiplied by -1, to improve the comparison. [Figure 6](#) shows that when the magnitude of the financial shock is relatively small, systematic evidence for differences in responses across the sign of the shock is largely absent. However, for large shocks there is clear evidence of sign non-linearity. [Figure 7](#) shows that shocks that increase EBP_t by 300 basis points are associated with large declines in output, inflation, stock returns and the interest rate while the unemployment rate rises. The responses to expansionary shocks of the same size are substantially smaller in magnitude. These results broadly

⁵The excess bond premium is downloaded from the federal reserve [website](#). Stock returns are calculated using the Standards and Poor total return index obtained from Global Financial database. The remaining variables are taken from the FRED database.

Figure 6: Response to contractionary and expansionary financial shocks that change the excess bond premium by 50 basis points. The response to expansionary shocks has been multiplied by -1 for the purpose of comparison.




confirm the findings reported by [Forni et al. \(2021\)](#).

5 Conclusions

Local projections are widely used in Macroeconometrics, as they provide a flexible tool to estimate impulse responses to structural shocks of interest. However, the most popular linear specification of local projections introduces the assumption of a linear relationship among variables within each horizon h considered. This paper introduces a flexible local projection that generalises the model of [Jordà \(2005\)](#) to a non-parametric setting by using Bayesian Additive Regression Trees (BART). Using

Figure 7: Response to contractionary and expansionary financial shocks that change the excess bond premium by 300 basis points. The response to expansionary shocks has been multiplied by -1 for the purpose of comparison



Monte Carlo experiments, we show that the model is able to capture impulse response non-linearities driven by state-dependence, sign-dependence or size-dependence.

We apply our methodology to US fiscal and financial shocks. We show that while it is true that the fiscal multiplier is stronger in recession, as advocated by [Auerbach and Gorodnichenko \(2013a\)](#), this holds true only in response to a contractionary fiscal shock. In response to an expansionary shock, we confirm the result by [Ramey and Zubairy \(2018\)](#), namely that the fiscal multiplier does not change significantly between recession and expansion. We then confirm the results by [Forni et al. \(2021\)](#) that financial shocks have non-linear effects on the economy. A financial shock that increases the cost of borrowing generates contractionary effects that increase more than propor-

tionately in the size of the shock. This suggests that strong negative financial shocks generate stronger effects that would be otherwise predicted using linear models.

References

- Alloza, M., Gonzalo, J. and Sanz, C. (2019), ‘Dynamic effects of persistent shocks’.
- Alpanda, S., Granziera, E. and Zubairy, S. (2021), ‘State dependence of monetary policy across business, credit and interest rate cycles’, *European Economic Review* **140**, 103936.
- Auerbach, A. J. and Gorodnichenko, Y. (2013a), Fiscal multipliers in recession and expansion, *in* A. Alesina and F. Giavazzi, eds, ‘Fiscal policy after the financial crisis’, University of Chicago Press, pp. 63–98.
- Auerbach, A. J. and Gorodnichenko, Y. (2013b), ‘Output spillovers from fiscal policy’, *American Economic Review* **103**(3), 141–46.
- Barnichon, R. and Brownlees, C. (2019), ‘Impulse Response Estimation by Smooth Local Projections’, *The Review of Economics and Statistics* **101**(3), 522–530.
- Breitung, J., Brüggemann, R. et al. (2019), ‘Projection estimators for structural impulse responses’, *University of Konstanz Department of Economics Working Paper Series* **5**.
- Bruns, M. and Lütkepohl, H. (2022), ‘Comparison of local projection estimators for proxy vector autoregressions’, *Journal of Economic Dynamics and Control* **134**, 104277.
- Castelnuovo, E. and Pellegrino, G. (2018), ‘Uncertainty-dependent effects of monetary policy shocks: A New-Keynesian interpretation’, *Journal of Economic Dynamics and Control* **93**, 277–296.
- 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. E., Huber, F., Koop, G., Marcellino, M. and Pfarrhofer, M. (2021), ‘Tail forecasting with multivariate Bayesian additive regression trees’.
- Forni, M., Gambetti, L., Maffei-Faccioli, N. and Sala, L. (2021), ‘Nonlinear transmission of financial shocks: Some new evidence’.
- Gilchrist, S. and Zakrajšek, E. (2012), ‘Credit spreads and business cycle fluctuations’, *American Economic Review* **102**(4), 1692–1720.
- Gonçalves, S., Herrera, A. M., Kilian, L. and Pesavento, E. (2022), ‘When do state-dependent local projections work?’.
- Gordon, R. and Krenn, R. (2010), ‘The end of the Great Depression 1939-41: Policy contributions and fiscal multipliers’, (16380).
- Herbst, E. and Johannsen, B. K. (2021), ‘Bias in local projections’.
- Hill, J., Linero, A. and Murray, J. (2020), ‘Bayesian additive regression trees: A review and look forward’, *Annual Review of Statistics and Its Application* **7**(1), 251–278.
- 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. (2021), ‘Inference in Bayesian additive vector autoregressive tree models’.
- Inoue, A., Rossi, B. and Wang, Y. (2022), ‘Local Projections in Unstable Environments’.
- Jordà, Ò. (2005), ‘Estimation and inference of impulse responses by local projections’, *American Economic Review* **95**(1), 161–182.
- Kilian, L. and Kim, Y. J. (2011), ‘How reliable are local projection estimators of impulse responses?’, *Review of Economics and Statistics* **93**(4), 1460–1466.

- Koop, G., Pesaran, M. H. and Potter, S. M. (1996), ‘Impulse response analysis in nonlinear multivariate models’, *Journal of Econometrics* **74**(1), 119–147.
- Lusompa, A. (2021), ‘Local Projections, Autocorrelation, and Efficiency’, (RWP 21-01).
- Miranda-Agrippino, S. and Ricco, G. (2021), ‘Bayesian Local Projections’.
- Plagborg-Møller, M. and Wolf, C. K. (2021), ‘Local Projections and VARs Estimate the Same Impulse Responses’, *Econometrica* **89**(2), 955–980.
- Ramey, V. A. and Zubairy, S. (2018), ‘Government spending multipliers in good times and in bad: Evidence from US historical data’, *Journal of Political Economy* **126**(2), 850–901.
- Ruisi, G. (2019), ‘Time-Varying Local Projections’, (891).
- Stock, J. H. and Watson, M. W. (2018), ‘Identification and estimation of dynamic causal effects in macroeconomics using external instruments’, *The Economic Journal* **128**(610), 917–948.

School of Economics and Finance



This working paper has been produced by the School of Economics and Finance at Queen Mary University of London

Copyright © 2022 Haroon Mumtaz and Michele Piffer

All rights reserved

School of Economics and Finance Queen Mary University of London

Mile End Road

London E1 4NS

Tel: +44 (0)20 7882 7356

Fax: +44 (0)20 8983 3580

Web: www.econ.qmul.ac.uk/research/workingpapers/