Impulse response estimation via flexible local projections

-

PRELIMINARY AND INCOMPLETE

Latest draft available at this link

Haroon Mumtaz* Michele Piffer†

February 15, 2022

Abstract

This paper introduces a flexible local projection that generalises the model of Jordà (2005) to a nonparameteric setting by using Bayesian Additive Regression Trees. Using Monte Carlo experiments, we show that the model is able to capture non-linearities in the impulse responses. Our empirical analysis uncovers three facts about US monetary policy shocks. First, that contractionary shocks have stronger effects on CPI and GDP than expansionary shocks. Second, that the effects on CPI progressively weaken out as the size of the shock increases. Third, that contractionary shocks generate stronger and faster responses if the shock hits the economy during expansionary periods.

JEL classification: C14, C11, C32, E52.

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

*Queen Mary, University of London, School of Economics and Finance, Mile End road, London E1 4LJ. 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 $y_{t+h} = B_h x_t + d_h w_t + u_{t+h}$, 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. 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 (2013) 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 function 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 $j = 1, 2, .., K$. 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 how to combine BART-LP with identification of the shocks of interest using the recursive approach, sign restrictions or external instruments. We document the performance of BART-LP using Monte Carlo analysis. We then apply BART-LP to study the effects of monetary policy shocks in the US. We document three facts about monetary policy shocks in the US. First, CPI and GDP decrease more strongly after a monetary contraction than they increase after a monetary expansion. Second, the response of GDP scales up proportionally as the size of the monetary shock increases, but the effect on CPI weakens out, suggesting a progressively weaker effect on CPI. Third, that the decrease in GDP after a monetary contraction is faster if the shock hits the economy during periods of high expansions, while CPI decreases more quickly.

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 the proposed model. Huber et al. (2020) extend the BART-VAR to a mixed frequency setting and evaluate the forecasting perfomance 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 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 et al. (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 and provides details on identification and estimation. Section 3 shows the application to monetary shocks. Section 4 concludes.

2 Flexible local projections

In this section we outline the methodology, which we refer to as the BART-LP model. We first discuss how BART approximates the unknown linear function, discuss the prior, and outline the posterior sampler. We then discuss how to achieve identification of the structural shocks of interest within this model.

2.1 The BART approximation

As noted above, the model is defined as:

$$y_{t+h} = m_h(x_t, w_t) + u_{t+h}, \tag{1}$$

where $\underbrace{y_t}_{1\times 1}$ denotes the variable of interest and $h = 0, 1, ..., H$ is the impulse response horizon to a change in x_t. The vector $\underbrace{w_t}_{(K-1)\times 1}$ consists of control variables that can include further lags of y_t. The residual u_{t+h} is assumed to be normally distributed with a variance σ^2_{t+h}. As noted in Jordà (2005), u_{t+h} is serially correlated for $h > 1$.

The impulse response at horizon h can be calculated as:

$$IRF_h = E\left(y_{t+h}|x_t = d; w_t\right) - E\left(y_{t+h}|x_t, w_t\right), \tag{2}$$

where d denotes the innovation to x_t. As we describe below, the structural shock can be identified by using appropriate control variables to implement restrictions or by using sign restrictions or instruments for x_t.

The function $m_h(x_t, w_t)$ is approximated by using BART:

$$m_h(x_t, w_t) \approx \sum_{j=1}^{J} f(x_t, w_t|\Gamma_j, \mu_j), \tag{3}$$

where $f(x_t, w_t|\Gamma_j, \mu_j)$ denotes a single regression tree. The parameters of the regression tree are the tree structure Γ_j and terminal nodes (or leaves) μ_j and thresholds c. Each regression tree divides the space of each explanatory variable by using binary splitting rules. Denoting X_i as the *i-th* column of $X_t = (x_t, w_t)$, these rules are defined as:

$$\begin{aligned} X_i &\leq c, \\ X_i &> c. \end{aligned} \tag{4}$$

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

$$f(x_t, w_t|\Gamma, \mu) = \sum_{b=1}^{B} I(X_i, c)\mu_b, \tag{5}$$

where $I(.)$ denotes an indicator function that equals 1 if X_i belongs to the set defined by the splitting rule. Note that the complexity of the tree is determined by B, the number of terminal nodes.

The model in equation 3 approximates $m_h(x_t, w_t)$ by 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 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 beyond a certain level. In practice, studies such as Huber et al. (2020) note that 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 IRF_h. As the regression trees split the space of the covariates via rules of the type shown in equation 4, the estimated predictions on the RHS of equation 2 are dependent on their history. Similarly, the shock d can lead to predictions that proportionally differ if the the size and sign of the shock leads to use of the covariate space where the relationship between y_{t+h} and x_t is substantially different from the 'average' impact.

2.2 Estimation

The model in equation (1) can be estimated using the MCMC algorithm described in Chipman et al. (2010).

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 prior for the tree parameters is factored as follows:

$$p((\Gamma_1, \mu_1), (\Gamma_2, \mu_2), ..., (\Gamma_J, \mu_J)) = \prod_{j=1}^{J} p(\mu_j|\Gamma_j)p(\Gamma_j), \tag{6}$$

where $p(\mu_j|\Gamma_j) = \prod_{b=1}^{B} p(\mu_{b,j}|\Gamma_j)$.

The prior for the tree structure Γ_j depends on the probability that the node at depth $d = 0, 1, 2, ..$ is not a terminal node. This 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 structure. As recommended by Chipman et al. (2010) we set $\alpha = 0.95$ and $\beta = 2$. The prior for 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 $p(\mu_j|\Gamma_j)$ Chipman et al. (2010) first tranform the dependent variable so that it lies between -0.5 and 0.5. As a consequence, $m_h(x_t, w_t)$ is also expected to lie between these values. The prior $p(\mu_j|\Gamma_j)$ is assumed to be normal $N(0, S)$. The variance S is set as $\frac{1}{2\kappa(J^{0.5})}$ where J denotes the total number of trees and the default value of κ recommended in Chipman et al. (2010) is 2. 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^2_{t+h}. The hyperparameters of the prior distribution are set by using $\hat{\sigma}^2_{t+s}$ an estimate of the variance obtained from a linear regression. If the true model is non-linear $\hat{\sigma}^2_{t+s}$ will be biased upwards. Under the default prior, the hyperparameters are chosen so that $\Pr(\sigma_{t+s} < \hat{\sigma}_{t+s}) = 0.9$.

2.2.2 MCMC algorithm

The MCMC algorithm devised by Chipman et al. (2010) samples from the conditional posterior distributions of σ^2_{t+h} 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 as the residual:

$$R_j = y - \sum_{i \neq j} f(x_t, w_t|\Gamma_j, \mu_j). \tag{7}$$

   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

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

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

$$\alpha = \frac{q(\Gamma_j^{new}, \Gamma_j^{old}) p(R_j|\Gamma_j^{new}, \sigma_{t+h}^2) p(\Gamma_j^{new})}{q(\Gamma_j^{old}, \Gamma_j^{new}) p(R_j|\Gamma_j^{old}, \sigma_{t+h}^2) p(\Gamma_j^{old})}, \tag{8}$$

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 $j = 1, 2, .., J$ trees;

3. the conditional posterior distribution of the terminal node parameters is Gaussian with the parameters known in closed form. Therefore, the draw of μ_j for $j = 1, 2, .., J$ can be carried out in a straightforward manner;

4. given a draw of the model parameters, the IRF is calculated by using equation (2). This simply requires the calculation of the predicted value $E(y_{t+h}|x_t, w_t) = \sum_{j=1}^{J} f(x_t, w_t|\Gamma_j, \mu_j)$ where x_t, w_t are fixed at the required histories for the base prediction and a shock is added to x_t for the prediction conditional on the shock.

2.2.3 Autocorrelation

In the case of linear local projections, it has been shown that the residual at horizon h follows a $MA(h-1)$ process (see Lusompa, 2021). Lusompa (2021) suggests a GLS procedure whereby the autocorrelation is eliminated by including leads of the residuals from the horizon 0 LP in the conditioning set.[2]

The non-parametric setting considered in this paper encompasses non-linear mod-

[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.

els. For the purpose of illustration consider a simple non-parametric AR(1) model:

$$y_{t+1} = f(y_t; A_1) + e_{t+1}. \tag{9}$$

Iterating the AR forward 3 periods as an example:

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

It is useful to compare this with a flexible local projection for this horizon:

$$y_{t+3} = m_3(y_t; A_1) + u_{t+3}. \tag{10}$$

The function $m_3(y_t; A_1)$ 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 u_{t+3} is a non-linear function of e_{t+1} and e_{t+2} and has a non-linear autocorrelation structure.

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

$$y_t = \sum_{i=0}^{\infty} \phi^i e_{t-i} + \sum_{i=0}^{\infty} \sum_{i=j}^{\infty} \zeta_{ij} e_{t-i} e_{t-j} + \sum_{i=0}^{\infty} \sum_{i=0}^{\infty} \sum_{k=0}^{\infty} \phi_{ij} e_{t-i} e_{t-j} e_{t-k} + ... \tag{11}$$

To account for this autocorrelation, we propose to include an estimate of the shocks $\hat{e}_{h-1} = (\hat{e}_{t+1}, \hat{e}_{t+2,}..., \hat{e}_{t+h-1})$ as additional covariates in the h period flexible local projection. Following Lusompa (2021) we use the residuals of period 0 flexible local projection $y_t = m_0(x_t, w_t) + u_t$ to build the vector $\hat{e}_{h-1}$. Then, the flexible local projection for period h is specified as:

$$y_{t+h} = m_h(x_t, w_t, \hat{e}_{h-1}) + u_{t+h}. \tag{12}$$

The BART approximation of $m_h(x_t, w_t, \hat{e}_{h-1})$ proxies the non-linear dependence of y_{t+h} on $e_{t+1}, e_{t+2,}..., e_{t+h-1}$ and ameliorates the autocorrelation in u_{t+h}.

2.3 Identification

In order to interpret the responses from an economic perspective, it is crucial to impose identifying restrictions. As discussed in Plagborg-Møller and Wolf (2021) and Barnichon and Brownlees (2019), a recursive identification scheme can be easily implemented 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 impact on GDP and CPI at horizon 0, then $y_{t+h} = GDP_{t+h}, x_t = R_t$ and $w_t = GDP_t, CPI_t, GDP_{t-1}, CPI_{t-1}, R_{t-1}, ..., GDP_{t-p}, CPI_{t-p}, R_{t-p}$. The inclusion of GDP_t in the set of control variables implies that the response at $h = 0$ is restricted to be zero.

An instrumental variable approach can also be used to identify the structural impulse response. Suppose that an instrument or proxy z_t is available that satisfies exogeneity and relevance. In other words, $cov[z, \varepsilon_-] = 0$ and $cov[z, \varepsilon] = \alpha, \alpha > 0$ where ε denotes the shock of interest, while ε_- denotes other shocks. Following studies such as Auerbach and Gorodnichenko (2013) and Tenreyro and Thwaites (2016), a simple approach involves setting $x_t = z_t$ with the control variables w_t including lags of y_t, z_t and other relevant covariates. Note that the timing of the controls implies that z_t is ordered first in a recursive setting.[3]

An alternative approach consists of using sign restrictions. Within BART, sign restrictions are conveniently introduced by adding an accept/reject step in the algorithm.

[3]In the context of linear VAR models, Plagborg-Møller and Wolf (2021) show that this recursive approach yields valid relative impulse responses.

3 Empirical application

We use the proposed model to investigate non-linearities in the transmission of monetary policy in the US and estimate the response of GDP and CPI to a shock to the Federal Funds rate (FFR). The model is defined as follows:

$$y_{t+h} = m_h(x_t, w_t, \hat{e}_{h-1}) + u_{t+h}, \tag{13}$$

where $y_{t+h} = \Delta GDP_{t+h}$ or $y_{t+h} = \Delta CPI_{t+h}$ where Δ denotes the quarterly growth rate. x_t is set equal to the level of FFR. The shock is identified via timing assumptions: the interest rate can only affect GDP growth and inflation after a one period lag. As discussed above, these timing restrictions can be incorporated by using the contemporaneous values of GDP growth and CPI inflation as controls. In addition to these, we also include 2 lags of these variables and the FFR in w_t. To control for autocorrelation, we include the vector residuals $\hat{e}_{h-1} = (\hat{e}_{t+1}, \hat{e}_{t+2}, ..., \hat{e}_{t+h-1})$ where $\hat{e}_t$ is obtained from the period 0 LP.

We use default prior set-up of Chipman et al. (2010), setting the number of trees to 250. The total number of iterations is set to 2000, with a burn-in of 1000. The data is obtained from the FRED database. As our focus is on conventional monetary policy, the sample runs from 1966Q1 to 2007Q4.

3.1 Sign and size non-linearity

We first investigate the possibility of non-linearity associated with the sign of the shock by computing the responses to a 1 unit increase and decrease in FFR. As the IRFs depend on the history of the covariates, we compute the responses for a range of histories and report the average. In particular, when approximating the expectations in equation (2), we set the value of the covariates used to construct the predictions $E(y_{t+h}|x_t; w_t)$ on the basis of quantiles of ΔGDP. We consider the 10th to 90th quantiles with increments of 10.

Figure 1: Sign non-linearity

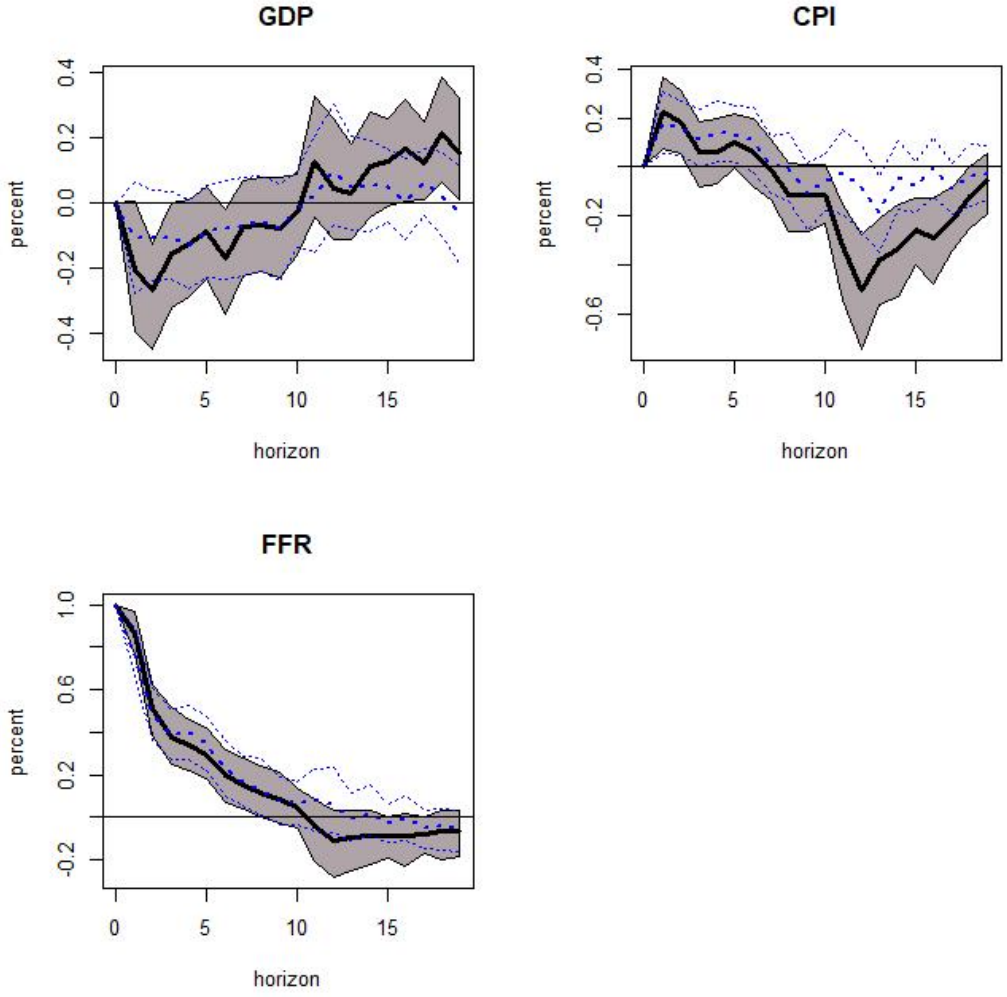

Note: Response to shocks that change FFR by 1 unit. Black lines are response to an increase in FFR with the shaded area showing the 68% error band. Blue dotted lines are the point estimated and 68% error band of the response an fall in FFR. The sign of these has been reversed for comparison.

Figure 1 shows the response to positive (black, continuous line) and negative (blue, dashed line) shocks to FFR of one unit. Note that the sign of responses to negative shock has been reversed ex-post for the purpose of comparison. It is clear from the response of GDP growth that shocks that increase in FFR have a larger effect. At the 2-3 quarter horizon, the response of GDP growth to contractionary shocks is twice as large the response to expansionary shocks. The response of CPI inflation displays a

Figure 2: Size non-linearity

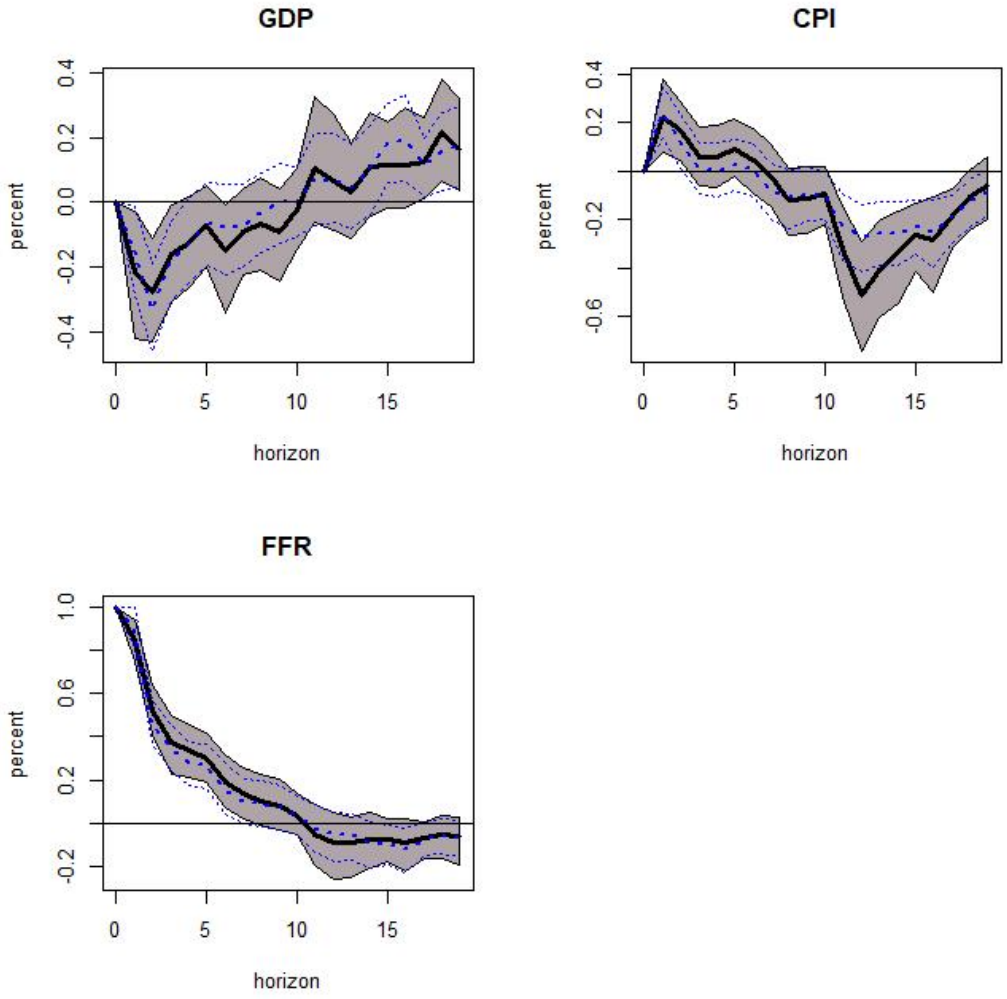

Note: Response to shocks that change FFR by 1 and 2 units. Black lines are response to an increase in FFR by 1 unit with the shaded area showing the 68% error band. Blue dotted lines are the point estimated and 68% error band of the response to a shock that increases the FFR by 2 units. These responses have been divided by 2 for comaprison.

price puzzle, but becomes negative after about two years. It is interesting to note that this negative response is substantially larger for contractionary shocks.

Figure 2 considers the possibility of non-linearity associated with the size of the shock. It shows the response to contractionary shocks that increase the FFR by 1 unit (black, continuous line) and 2 units (blue, dashed line), respectively. The responses to the larger shocks are divided by 2 for comparison. The estimates indicate that there

Figure 3: State non-linearity

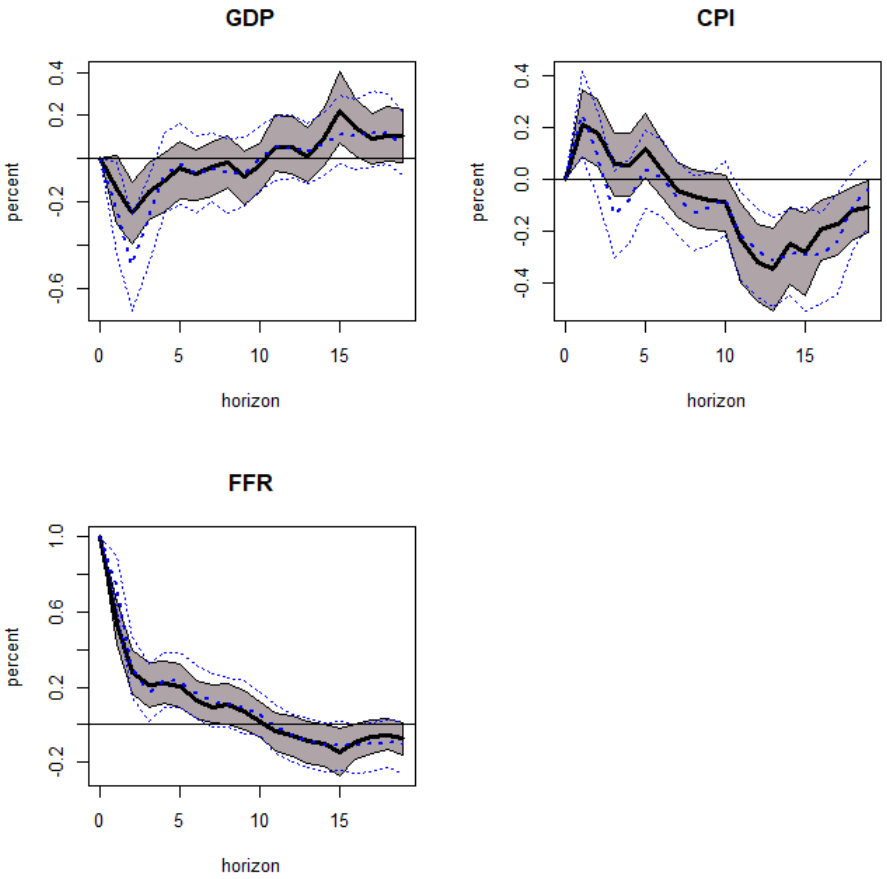

Note: Response in recessions (black lines and shaded area) and booms (blue lines)

is no evidence to suggest that changing the size of the shock affects the magnitude of the response of GDP. However, the response of CPI weakens out as the shock becomes larger.

3.2 State dependence

To check if the magnitude of the responses is different across business cycle regimes, we estimate the IRFs using different histories. In particular, we condition on two histories: 1) x_t^L and 2) x_t^H where x_t^L corresponds to the value of covariates when GDP growth is below the 10th percentile and x_t^H captures the value of the regressors when

GDP growth is above the 90th percentile. We intepret these responses as estimates in recessions and booms, respectively.

Figure 3 presents the estimated response to contractionary monetary policy shocks in the two regimes. The estimated response of GDP growth during booms is substantially larger – 3 quarters after the shock, the median response in recessions is 3 times smaller the response when GDP is high. Similarly, the response of CPI inflation declines much faster when the estimates are conditioned on high GDP, with the median response negative within 4 quarters. This supports the general conclusion reached by Tenreyro and Thwaites (2016) that monetary policy is more effective when GDP growth is high.

4 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 introduce the assumption of a linear relationship among variables within each horizon considered. This paper introduces a flexible local projection that generalises the model of Jordà (2005) to a non-parameteric setting by using BART. Using Monte Carlo experiments, we show that the model is able to capture non-linearity in the IRF due to the size and sign of the shock or because differences in initial conditions originating from different states of the business cycle.

We apply our methodology to US monetary policy shocks. We first find that CPI and GDP contract more after a monetary expansion than they expand after a monetary contraction. We then find that the effect of a monetary contraction on CPI progressively weakens out as the central bank generates a stronger monetary shock. Last, we document that during periods of economic expansions, GDP and CPI respond more and faster to a monetary contraction compared to periods of economic contractions. All in all, our results suggest the presence of non-linearities in the transmission

of US monetary policy.

References

Alloza, M., Gonzalo, J. and Sanz, C. (2019), 'Dynamic effects of persistent shocks'.

Auerbach, A. J. and Gorodnichenko, Y. (2013), '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.

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'.

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.

Lusompa, A. (2021), 'Local Projections, Autocorrelation, and Efficiency', (RWP 21-01).

Miranda-Agrippino, S., Ricco, G. et al. (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.

Tenreyro, S. and Thwaites, G. (2016), 'Pushing on a string: US monetary policy is less powerful in recessions', *American Economic Journal: Macroeconomics* **8**(4), 43–74.