Online Appendix for
"Impulse response estimation via flexible local projections"

A An illustration of BART

We further support the discussion in Section 2 using a simple univariate illustration. Consider the following pointwise linear data generating process

$$\begin{aligned} y_t = & + 1.7\ x_t\ \mathrm{I}(x_t < -5) + \\ & - 0.1\ x_t\ \mathrm{I}(-5 \leq x_t < 5) + \\ & - 1.1\ x_t\ \mathrm{I}(x_t \geq 5) + \epsilon_t \end{aligned} \tag{A.1}$$

$$\epsilon_t \sim N(0, 1). \tag{A.2}$$

where the regressor is $x_t \sim N(0, 1)$ and (x_t, ϵ_t) are statistically independent. The true conditional expectation function $E(y|x)$ is pointwise linear with breaks at $x_t = -5$ and $x_t = 5$, but it can successfully be detected by BART. To show this, the blue dots in Figure A.1 report 200 observations generated from the model. The red dashed line shows the true conditional expectation function of the model. The black line shows the median conditional expectation estimated by BART using a standard prior parametrization as in Chipman et al. (2010), 1,000 posterior draws and 200 trees. As shown by the figure, the nonlinear nature of the conditional expectation function is detected successfully despite no assumption being made on the functional form in the relation between y_t and x_t.

Figure A.1: Illustration

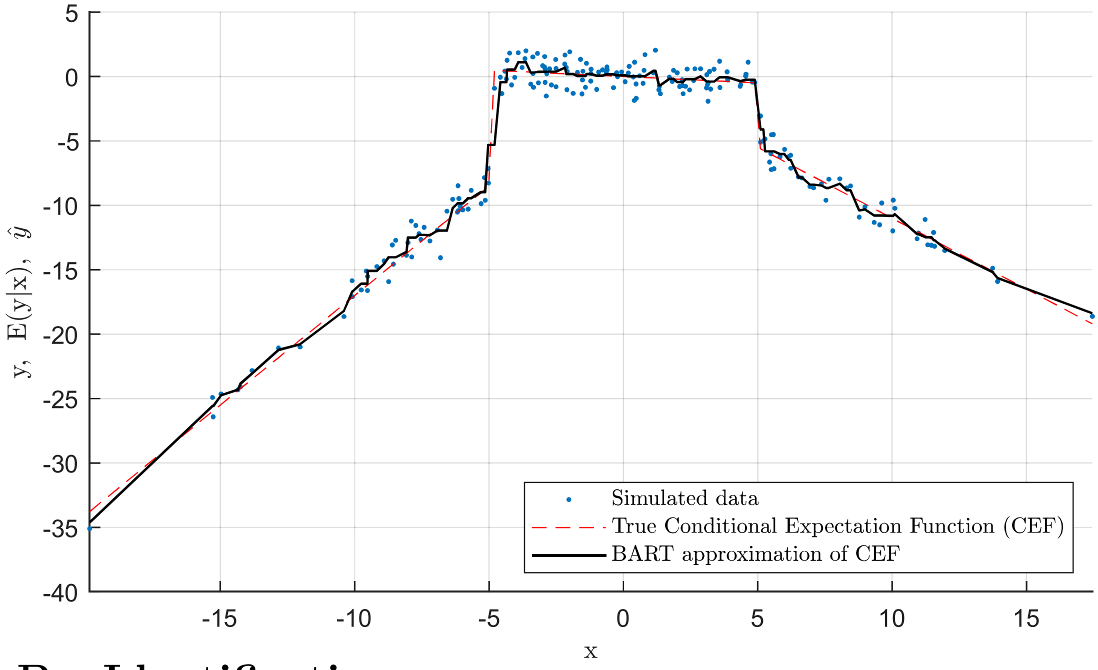

B Identification

This section discusses how to implement different identification strategies for the identification of the shocks in BART-LP.

One approach to identification in BART-LP is to follow Jordà (2005) and estimate a preliminary SVAR on $\boldsymbol{y}$, apply a identification strategy to the SVAR, and estimate the $k \times 1$ impulse vector $\boldsymbol{d}$. The contemporaneous impulse responses of the entire vector $\boldsymbol{y}_t$ can be set as $\boldsymbol{\phi}_0 = \boldsymbol{d}$, while future scalar impulse responses ϕ_h, $h = 1, .., H$, can be computed as follows:

1. set $\boldsymbol{z}_t$ equal to $\boldsymbol{y}_t$ as well as its first L lags;
2. set $\bar{\boldsymbol{z}}^0$ in accordance with the intended timing of the simulated impulse response;
3. set $\bar{\boldsymbol{z}}^1$ equal to $\bar{\boldsymbol{z}}^0$ except that the entry corresponding to $\boldsymbol{y}_t$ is augmented by $\boldsymbol{d} \cdot \bar{\epsilon}$.

One advantage of this procedure is that it makes the generated impulse response comparable to the approach by Jordà (2005). A second advantage is that it is compatible with a wide range of existing identifications strategies for linear SVAR models. A disadvantage is that it imposes a linear model at horizon 0, and allows for a non-linear response only after the first horizon. This is an inconvenient feature in a methodology that builds on non-parametric techniques. Note also that the identification assumes invertibility of the VAR model. Therefore, the researcher has to ensure that the VAR contains sufficient information and the reduced form residuals span the space of structural shocks (see Forni and Gambetti (2014)).

An alternative is to follow Barnichon and Brownlees (2019) and Plagborg-Møller and Wolf (2021) 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 $\boldsymbol{z}_t$. With BART-LP this requires to

1. set $\boldsymbol{z}_t$ equal to an appropriate subset of $\boldsymbol{y}_t$ as well as L lags of $\boldsymbol{y}_t$;

2. set $\bar{\boldsymbol{z}}^0$ in accordance with the intended timing of the simulated impulse response;

3. set $\bar{\boldsymbol{z}}^1$ equal to $\bar{\boldsymbol{z}}^0$ except that the entry corresponding to x_t is augmented by $\bar{\epsilon}$.

An advantage of this procedure is that it avoids a preliminary linear model to achieve identification of the shock of interest. In addition, identification and estimation are achieved jointly rather than in a two-step procedure. A disadvantage is that it does not retain the same flexibility on candidate identification strategies compared to when running a preliminary linear SVAR. In addition, it works less naturally if one wants to generate a shock such that a variable of interest responds by a desired amount.

A third option is to achieve identification using either the true series of shocks (an option available in simulations) or an instrument for the shock of interest. BART-LP can then be used as follows:

1. instead of adding x_t, set $\boldsymbol{z}_t$ equal to m_t (the observed shock to x_t or the instrument for the shock) as well as L lags of $\boldsymbol{y}_t$ (which includes x_t);

2. Alternatively x_t can be regressed on the instrument and the set of controls and the fitted value $\hat{x}_t$ can be used as the regressor of interest in the LP. As noted in the main text, the instrument should satisfy the relevance and exogeneity conditions set out in Stock and Watson (2018).

3. set $\bar{\boldsymbol{z}}^0$ in accordance with the intended timing of the simulated impulse response;

4. set $\bar{\boldsymbol{z}}^1$ equal to $\bar{\boldsymbol{z}}^0$ except that the entry corresponding to m_t is augmented by $\bar{\epsilon}$.

An advantage of this approach is that, in a simulation exercise, it allows isolating the computation of non-linear impulse responses from the issue of identification. A disadvantage is that instruments for the shocks of interest are not always available in applied work.

The above discussion simplifies the analysis by omitting term $\boldsymbol{w}_{t+h}^{(h)}$ from equation (1). Yet, $\boldsymbol{w}_{t+h}^{(h)}$ is needed to account for the autocorrelation structure in the residual, as discussed in Section 2.3. In light of this, the three options discussed above can be modified as follows:

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

$$y_t = m_0\big(\boldsymbol{z}_t\big) + \epsilon_t^{(0)}, \tag{B.3}$$

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

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

$$y_{t+h} = m_h\big(\boldsymbol{z}_t, \boldsymbol{w}_{t+h}^{(h)}\big) + \epsilon_{t+h}^{(h)}, \tag{B.4}$$

with

$$\boldsymbol{w}_{t+h}^{(h)} = (\hat{\epsilon}_t^{(0),d}, \hat{\epsilon}_{t-1}^{(0),d}, ..., \hat{\epsilon}_{t-h+1}^{(0),d}); \tag{B.5}$$

3. set the entry of $\bar{\boldsymbol{z}}^1, \bar{\boldsymbol{z}}^0$ associated with $\boldsymbol{w}_{t+h}^{(h)}$ equal to its estimated average.

An alternative approach is to follow the lag-augmentation strategy proposed by Montiel Olea and Plagborg-Møller (2021) and increase the number of lags in the model.

C Coverage of error bands

In this simulation exercise we assume that the data generating process is a VAR(3) model for 3 endogenous variables. The coefficients and variance-covariance of the error terms is set equal to the OLS estimates of a VAR(3) model using data on 3 variables employed in our empirical exercise: (1) CPI inflation, (2) IP growth, (3) Excess Bond Premium (EBP). We generate 500 observations after discarding an initial sample of 100 observations to account for starting values. Using this artificial data, we estimate two flexible LPs to estimate the response of inflation and IP to shocks to EBP.[1]. The experiment is repeated 500 times and we compute coverage probabilities using the estimated 90 percent highest posterior density intervals from the 2 LPs for IRFs of inflation and IP. Figure C.2 displays the estimated coverage rates for 15 periods. At short horizons, the coverage rates are close to 90%. There is a decline in the rates after 7 periods, but the decline appears to be limited to around 10 %.

D Impact of Monetary Policy shocks using a Smooth transition LP

We estimate the following non-linear LP:

$$y_{t+h} = \alpha_h + \sum_{p=1}^{P} B_{h,p,R_1}(y_{t-1} \cdot (1 - F(z_{t-1}))) + \sum_{p=1}^{P} B_{h,1,R_2}(y_{t-1} \cdot F(z_{t-1})) + \varepsilon_{t+h,h}, \quad \text{(D.6)}$$

with $h = 1, ..., H$. The vector $\boldsymbol{y}_t$ contains the variables: 1) the monetary policy shock proxy of Jarociński and Karadi (2020), (2) the one-year government bond yield, (3) real GDP growth, (4) GDP deflator inflation, (5) stock returns, (6) the excess bond premium. The data is monthly, with the pre-Covid sample running from 1979M7 to

[1]Both in the DGP and estimated models, the EBP shock is identified recursively with EBP ordered last).

Figure C.2: Coverage of Error bands

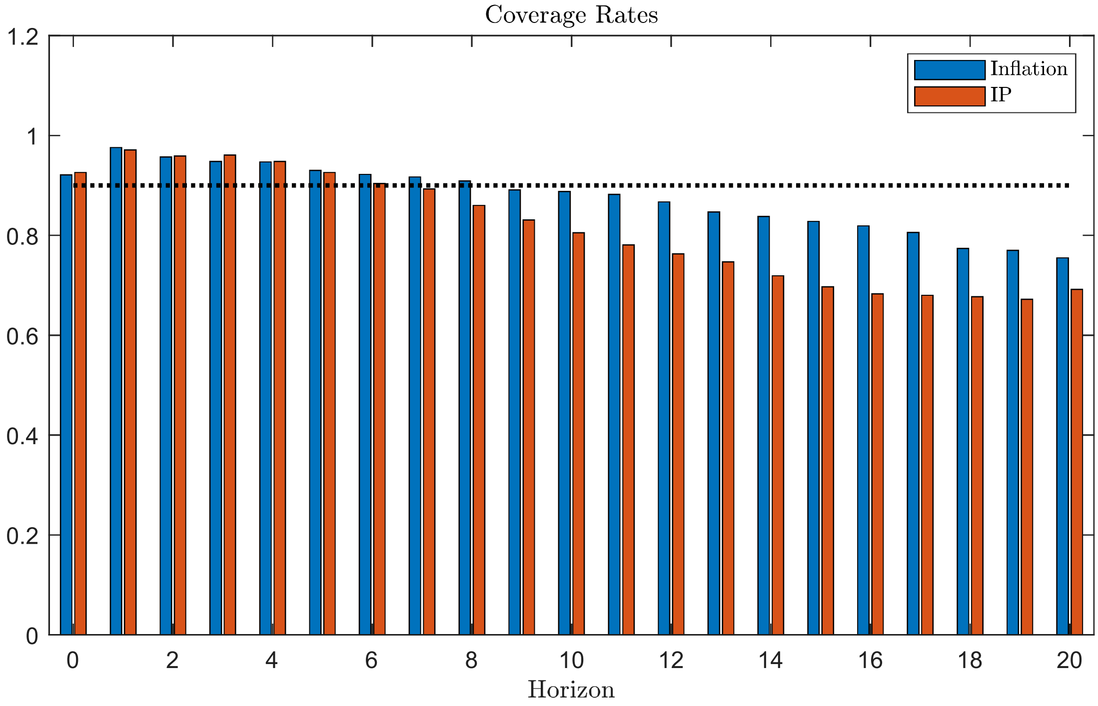

2020M2. We set the lag-length to 12. The transition function is defined as:

$$F(z_t) = \frac{e^{(-\gamma z_t)}}{1 + e^{(-\gamma z_t)}}, \quad \text{where } \gamma > 0. \tag{D.7}$$

Following Auerbach and Gorodnichenko (2013*a*) we set γ to 1.5 and the transition variable z_t is demeaned annual GDP growth. Note that $F(z_t)$ approaches 1 as z_t declines. As in the benchmark case, the policy shock is identified using a recursive linear VAR with the instrument ordered first. The models are estimated using OLS with robust standard errors (see the lpirfs R package). Figure D.3 displays the response to a contractionary monetary policy shock during recessions and expansions. Evidence that the impact of the shock on GDP growth is larger during expansions is not clear cut in this model. While GDP displays a stronger short-term decline during booms, the effect during recessions is more negative at medium horizons. Similarly, there is little evidence of a systematic difference in the response of stock returns. Moreover, the estimated IRFs are erratic, especially at longer horizons.

Figure D.3: Response to MP shock using a smooth transition LP. The shock increases the interest rate by 1 unit on impact

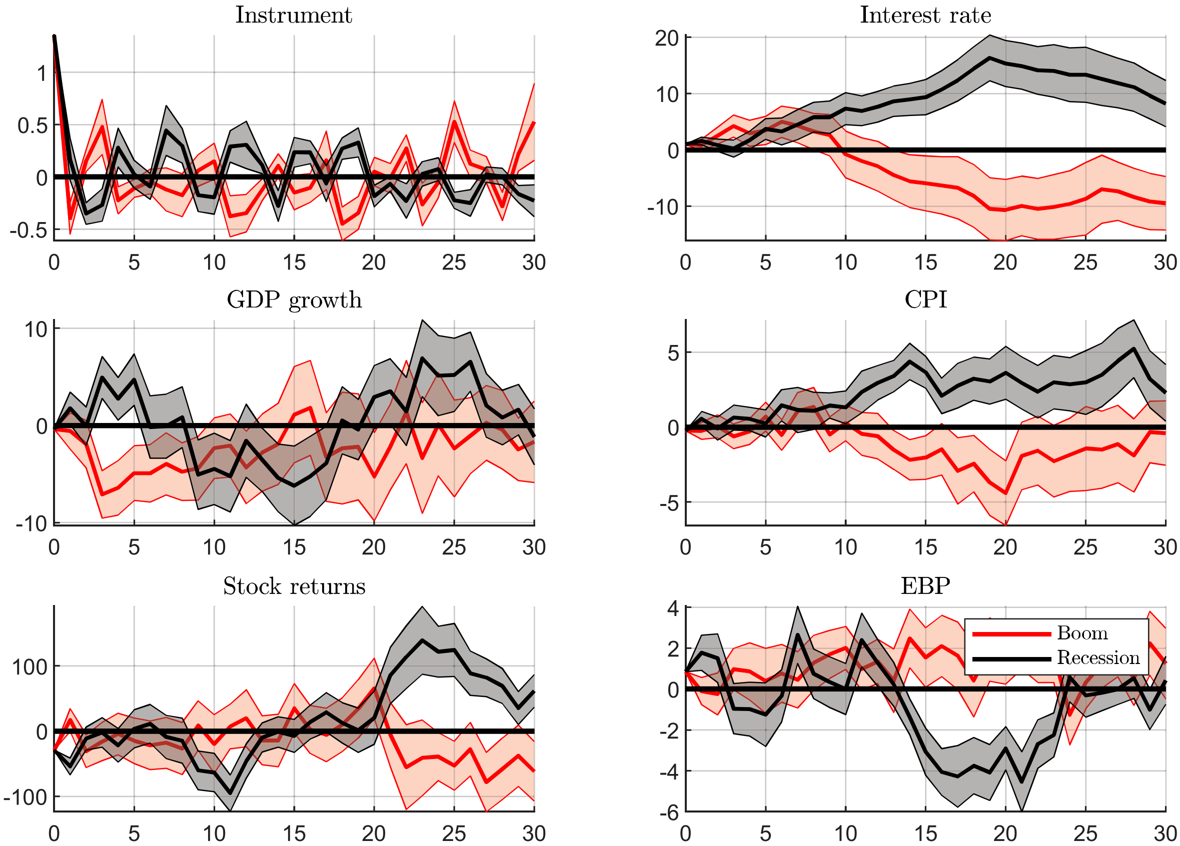

E Additional figures

Figure E.4: SVAR-GARCH - Robustness for $T = 100$

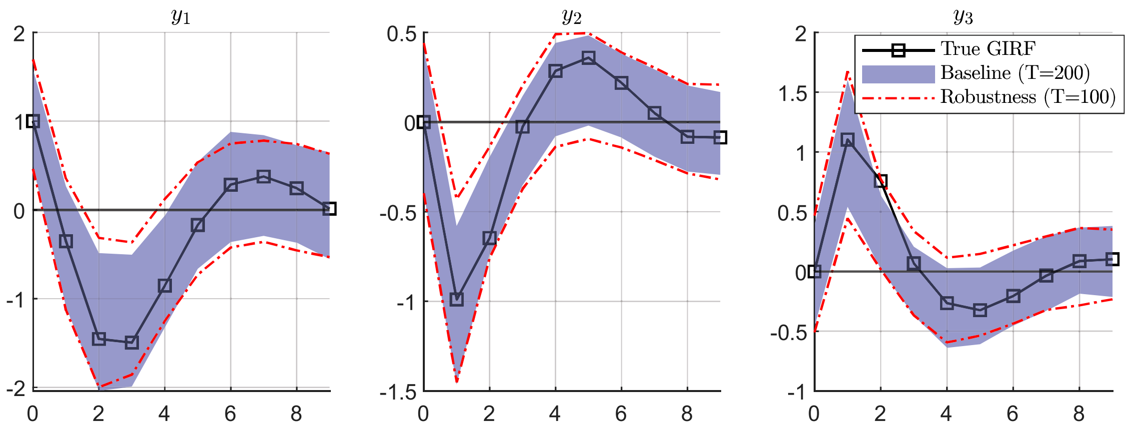

Figure E.5: Sign-dependent Moving Average - difference in the positive and negative response to a monetary policy shock - Robustness for $T = 100$

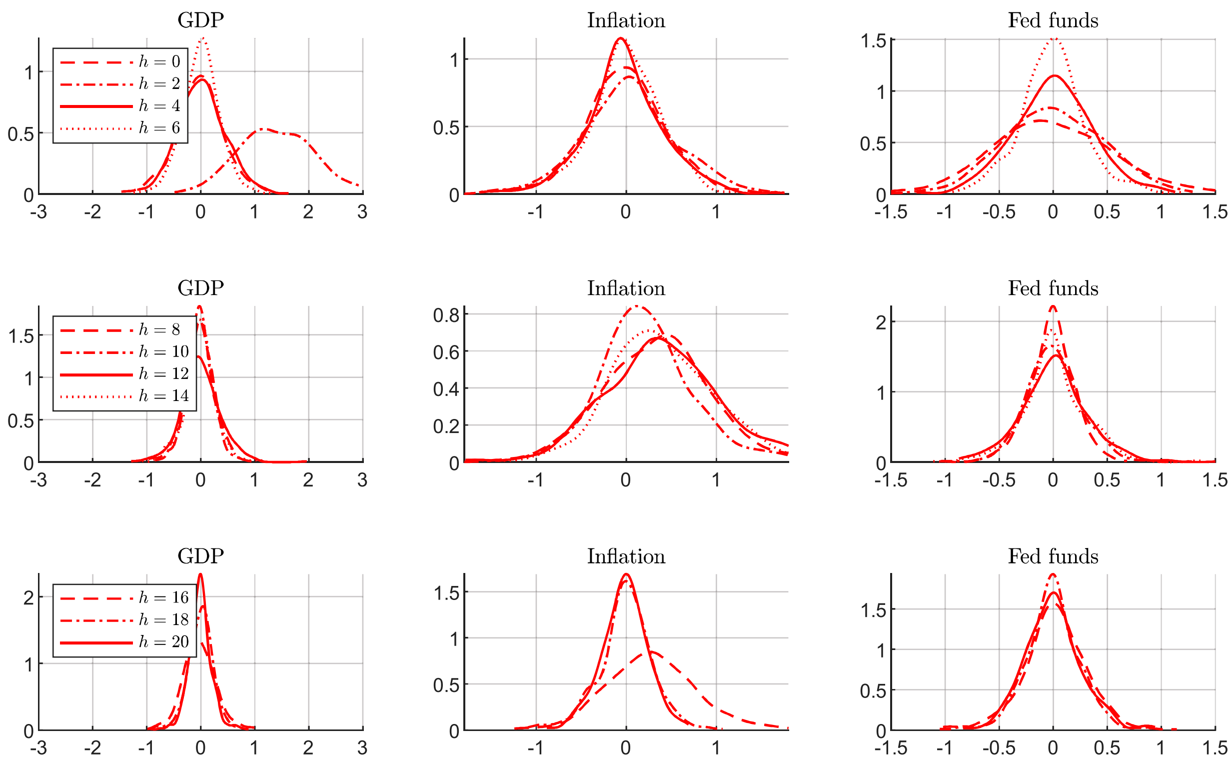

Figure E.6: Threshold VAR model - Robustness for $T = 100$

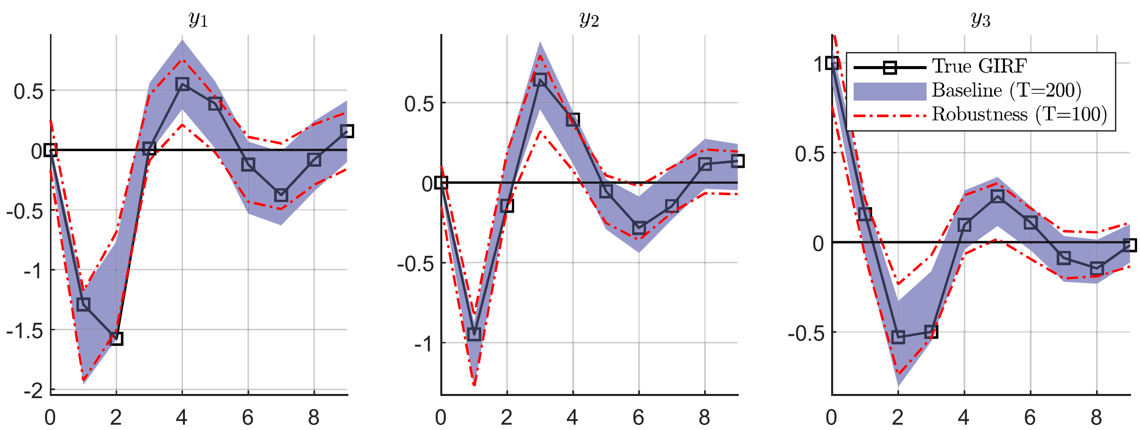

Figure E.7: Monte Carlo evidence of omitting residuals $\boldsymbol{w}_{t+h}^{(h)}$ (first part)

A) SVAR-GARCH

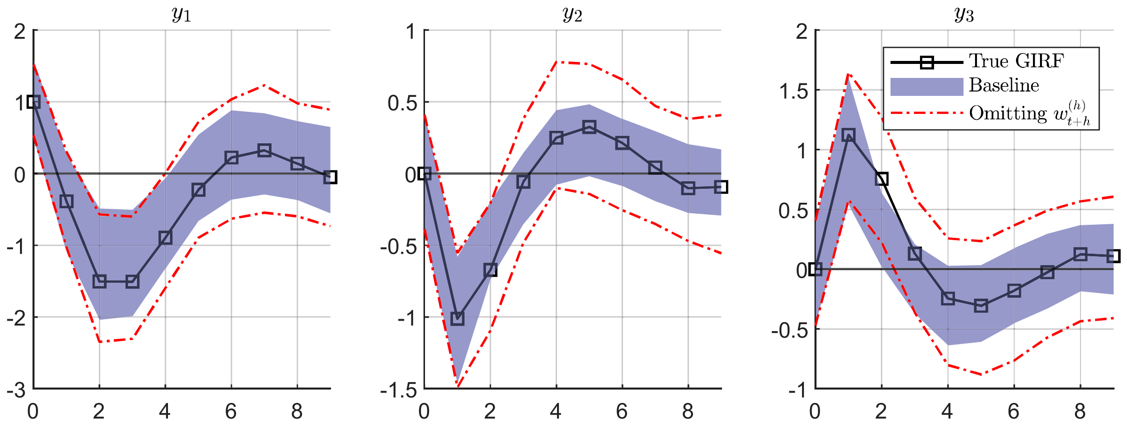

B) Sign-dependent MA (positive shock)

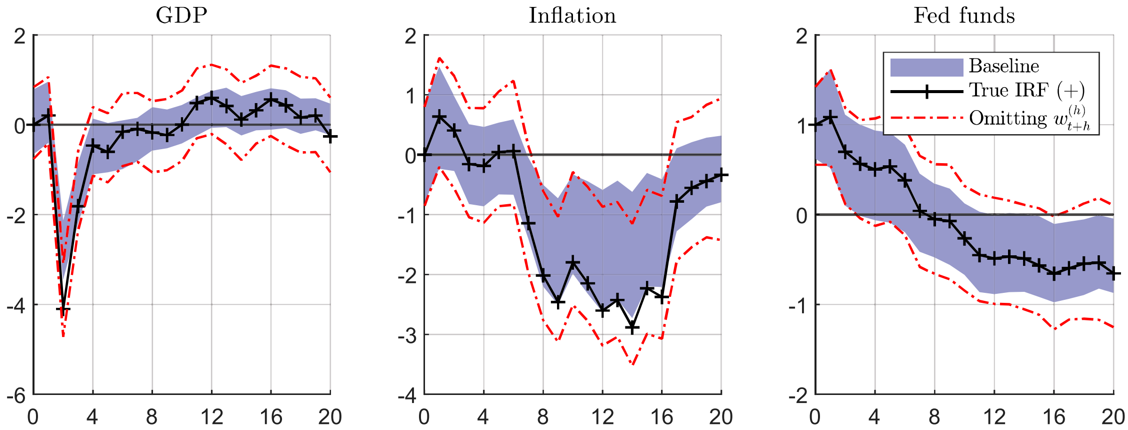

Figure E.8: Monte Carlo evidence of omitting residuals $\boldsymbol{w}_{t+h}^{(h)}$ (second part)

C) Sign-dependent MA (negative shock, flipped)

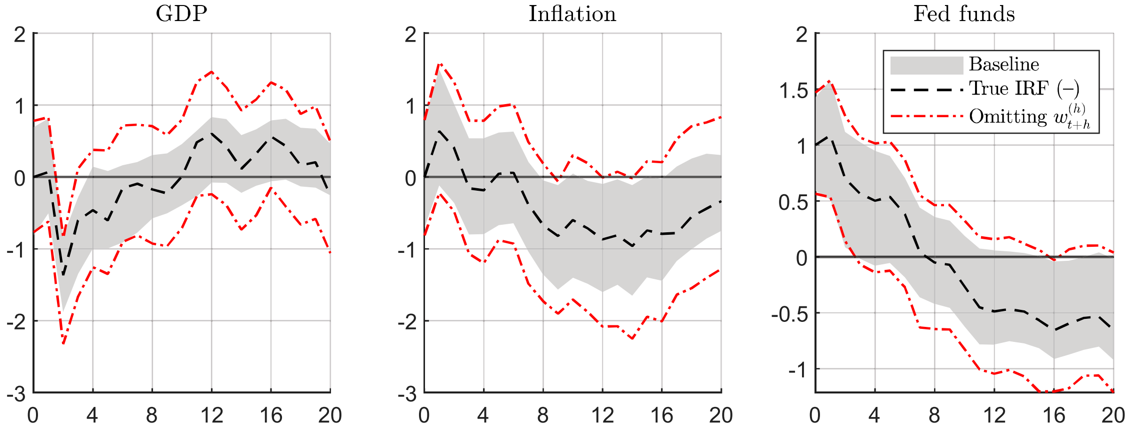

D) Threshold VAR

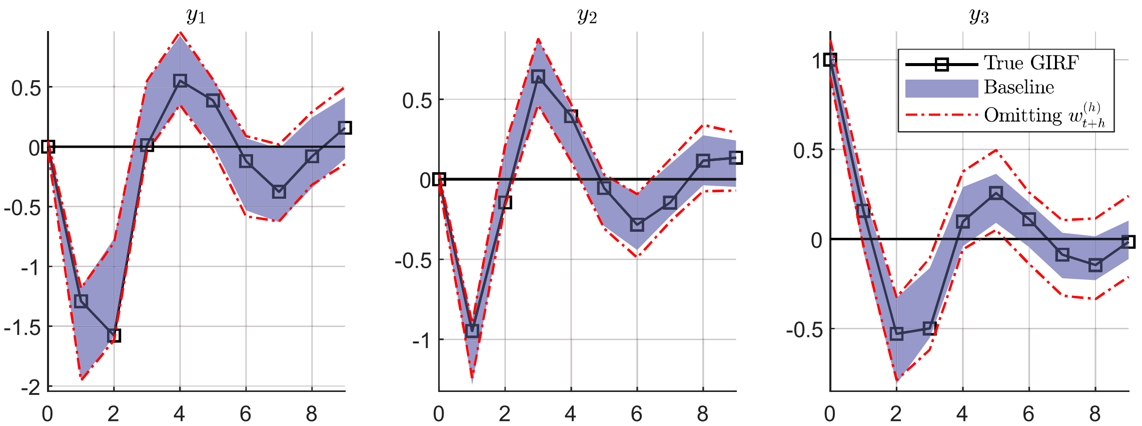

Figure E.9: Time varying effects of financial shocks on industrial production 12 months year after the shock

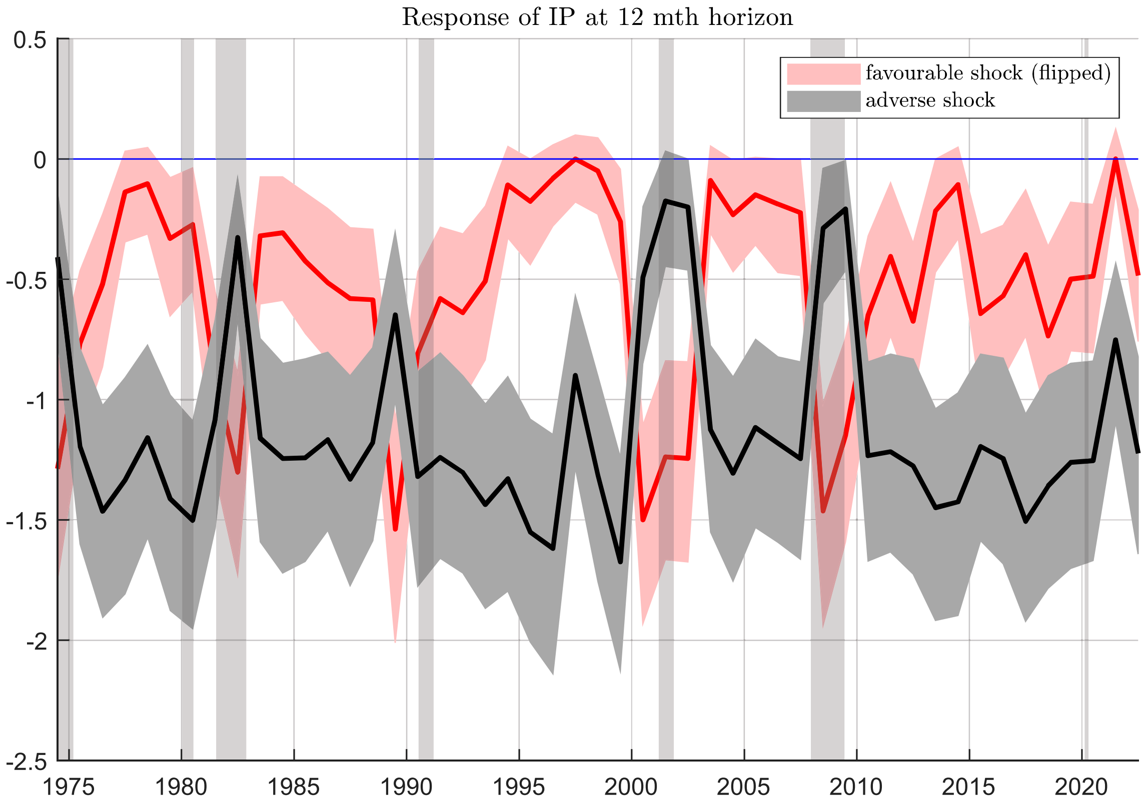

Effect of industrial production 12 months after a financial shock, setting the conditioning values for the generalized impulse responses equal to the value at time t over the full sample period. Grey bands report NBER recessions.