

Impulse response analysis for structural dynamic models with nonlinear regressors*

Sílvia Gonçalves[†], Ana María Herrera[‡], Lutz Kilian[§] and Elena Pesavento[¶]

July 1, 2021

Abstract

We study the construction of nonlinear impulse responses in linear structural dynamic models that include nonlinearly transformed regressors. We derive the closed-form solution for the population impulse responses to a given shock and propose a control function approach to estimating these responses without taking a stand on how the remainder of the model is identified. Our plug-in estimator dispenses with the need for simulations and, unlike conventional local projection (LP) estimators, is consistent. A modified LP estimator is shown to be consistent in special cases, but less accurate in finite samples than the plug-in estimator.

JEL codes : C22, C32, C51

Key Words: structural model, censored regressor, nonlinear transformation, nonlinear responses, partial identification, control function, block recursive model, Monte Carlo integration, local projection.

*Acknowledgments: The views expressed in this paper are our own and should not be interpreted as reflecting the views of the Federal Reserve Bank of Dallas or any other member of the Federal Reserve System. We thank Jörg Breitung, Atsushi Inoue, Helmut Lutkepohl, Mikkel Plagborg-Møller, the editor, the associate editor, and three anonymous referees for helpful comments.

[†]McGill University, Department of Economics, 855 Sherbrooke St. W., Montréal, Québec, H3A 2T7, Canada. E-mail: silvia.goncalves@mcgill.ca.

[‡]University of Kentucky, Department of Economics, 550 South Limestone, Lexington, KY 40506-0034, USA. E-mail: amherrera@uky.edu.

[§]Federal Reserve Bank of Dallas, Research Department, 2200 N. Pearl St., Dallas, TX 75201, USA. E-mail: lkilian2019@gmail.com.

[¶]Emory University, Economics Department, 1602 Fishburne Dr. Atlanta, GA 30322, USA. E-mail: epe-save@emory.edu.

1 Introduction

We study the construction of nonlinear impulse responses generated by linear models that include regressors that are censored or otherwise nonlinearly transformed. Such models have played an important role in recent years in capturing asymmetries, thresholds and other nonlinearities in the responses of macroeconomic variables to exogenous shocks.¹

For any scalar variable x_t , let $f(x_t)$ denote a nonlinear transformation of x_t . For example, we may define $f(x_t) = \max(0, x_t)$, in which case $f(x_t)$ corresponds to a censored version of x_t . Our analysis covers a range of censored regressors that have been employed in the empirical literature as well as powers of regressors. We follow a large existing literature in postulating that $(x_t \ y_t)'$ is a multivariate structural dynamic process such that, in general, x_t may linearly depend on its own lags and lags of y_t , whereas y_t depends on current and lagged values of x_t and $f(x_t)$ as well as its own lags. Thus, x_t is assumed to be predetermined with respect to y_t . In many applications, x_t reduces to a directly observed exogenous shock. For example, x_t may be an exogenous policy shock and y_t may consist of macroeconomic outcome variables such as inflation and output growth. We are interested in estimating the response of the elements of y_{t+h} , $h = 1, \dots, H$, to a one-time shock to the innovation in x_t of size δ . This response is nonlinear in general, even though the model is linear in the parameters.

Traditionally, such structural models have been estimated using equation-by-equation ordinary least squares (OLS) under the assumption that the model is fully recursive. The impulse response functions have been evaluated numerically by Monte Carlo integration (MCI) (Kilian and Lütkepohl 2017). There are two drawbacks of this approach. One is that the construction of impulse response estimators by MCI is computationally demanding. The other is that the assumption of a fully recursive structural model is rarely economically plausible, except in the bivariate setting. Without this assumption it is not possible to directly estimate the structural model by OLS, as required for the implementation of the MCI approach.

In this paper, we show how to estimate the population responses under the weaker assumption of a block recursive data generating process (DGP) using a control function approach (see e.g. Wooldridge 2010). This approach takes advantage of the fact that the identification of the shock to x_t is typically uncontroversial in applied work, while allowing us to remain agnostic about how the remainder of the model is identified. We propose a simple plug-in estimator that dispenses with the need for simulations. We formally prove the consistency of this plug-in estimator. The proof covers situations in which x_t is a directly observed i.i.d. shock, an exogenous serially correlated variable, or a predetermined endogenous

¹Examples include Alsalman and Herrera (2015), Barnichon, Matthes and Ziegenbein (2020), Hamilton (2011), Herrera, Lagalo and Wada (2011, 2015), Herrera and Karaki (2015), Hussain and Malik (2016), Hwa, Kapinos and Ramirez (2018), Kilian and Vigfusson (2011a,b, 2017), Tenreyro and Thwaites (2016) and Venditti (2013).

variable.

In some cases, the nonlinear responses of interest in this paper have instead been estimated using a variation of the linear local projection (LP) estimator of Jordà (2005, 2009), as popularized by Ramey (2016) (e.g., Tenreyro and Thwaites 2016, Hwa, Kapinos and Ramirez 2018, Barnichon, Matthes and Ziegenbein 2020). The asymptotic validity of the LP estimator under a wide range of conditions has been discussed in Plagborg-Møller and Wolf (2021), but there has been no work justifying the use of LP estimators for models involving nonlinearly transformed regressors. We find that, unlike in vector autoregressive (VAR) models, the asymptotic equivalence between estimators based on the structural model and LP estimators breaks down in the presence of nonlinearly transformed regressors. Conventional linear LP estimators of the impulse response function are inconsistent in our setting, even when the structural model is recursive, because they ignore the nonlinearity of the responses.

We propose a modified LP estimator that remains asymptotically valid when x_t is a directly observed i.i.d. shock. This modified LP estimator, however, is not valid under the weaker assumptions that x_t is a serially correlated exogenous process or that x_t is only predetermined with respect to y_t . Simulation evidence shows that the modified LP estimator is less accurate in finite samples than the plug-in and MCI estimators. It tends to have much higher variance and hence higher mean-squared error (MSE). The relative performance of the modified LP estimator and the plug-in estimator tends to be the same when the regression model is dynamically misspecified.

In fully recursive models, there is little to choose between the MCI estimator and the plug-in estimator in terms of accuracy, but the plug-in estimator typically reduces the computational cost by at least 98%, which is a significant improvement when conducting bootstrap inference. For example, when estimating the variance of the impulse response estimator by bootstrap, as required for asymptotic tests of the symmetry of response functions, even in bivariate models with only one lag, the computation time drops from one hour to under one minute. When allowing for more lags, the speed gains become even more pronounced. More importantly, whereas the MCI approach is feasible only in the special case of fully recursive models, the more computationally efficient plug-in estimator can be used without taking a stand on the identification of the remaining shocks.

Consistent with the many empirical applications, our analysis mainly focuses on the estimation of the unconditional impulse response function. When x_t is a directly observed i.i.d. shock, the unconditional response function coincides with the response function conditional on the history of the data. As we illustrate by example, more generally, analytical solutions for the conditional response function require additional assumptions about the distribution of the error term of x_t .

The remainder of the paper is organized as follows. In Section 2, we introduce the model and we

discuss the definition of the unconditional and conditional impulse response functions in structural dynamic models with nonlinearly transformed regressors. We also review examples from the empirical literature of how researchers have transformed regressors and why. In Section 3, we derive an exact solution for the unconditional population response function in the general block recursive model. Section 4 introduces a plug-in estimator for this response function and establishes its consistency. Section 5 examines the ability of conventional LP estimators to recover the responses of interest. We also propose a modified LP estimator and establish its consistency in special cases of the general model. In Section 6, we use simulations to assess the relative accuracy of the LP estimator, the MCI estimator and the plug-in estimator. We also provide evidence on the computational advantages of the plug-in estimator and assess the consequences of dynamic model misspecification. Section 7 contains an empirical illustration. The concluding remarks are in Section 8. Details of the proofs can be found in the appendix.

2 Framework

2.1 The model

Let $z_t \equiv (x_t, y_t)'$ denote an $n \times 1$ vector of strictly stationary time series, where y_t is $n_1 \times 1$ with $n_1 = n - 1$. A widely used structural DGP that allows for the inclusion of nonlinear regressors is

$$B_0 z_t = b + B(L) z_{t-1} + C(L) f(x_t) + \varepsilon_t, \quad (1)$$

where $b = (b_1, b_2)'$ and $\varepsilon_t = (\varepsilon_{1t}, \varepsilon_{2t})'$ are partitioned accordingly. We let

$$B(L) = B_1 + B_2 L + \dots + B_p L^{p-1} \text{ and } C(L) = C_0 + C_1 L + \dots + C_p L^p,$$

where we assume that the order of $B(L)$ and $C(L)$ is $p-1$ and p , respectively.² For convenience, we collect all structural parameters in the vector $\theta = (b', \text{vec}(B_0)', \text{vec}(B_1)', \dots, \text{vec}(B_p)', \text{vec}(C_0)', \dots, \text{vec}(C_p)')'$.

We partition $B(L)$ and $C(L)$ as

$$B(L) = \begin{pmatrix} B_{11}(L) & B_{12}(L) \\ B_{21}(L) & B_{22}(L) \end{pmatrix} \text{ and } C(L) = \begin{pmatrix} C_{11}(L) \\ C_{21}(L) \end{pmatrix},$$

respectively, where \mathcal{A}_{ij} denotes the (i, j) block of any partitioned matrix \mathcal{A} .

We postulate that

$$B_0 = \begin{pmatrix} 1 & 0 \\ -B_{0,21} & B_{0,22} \end{pmatrix}, \quad (2)$$

²This is without loss of generality because we can set p equal to the maximum lag order and zero out the elements that exceed the true lag order.

where $B_{0,21}$ is $n_1 \times 1$ and $B_{0,22}$ is $n_1 \times n_1$. We assume that $B_{0,22}$ is nonsingular such that its diagonal elements are equal to 1, which is a standard normalization. Under these assumptions, x_t is predetermined with respect to y_t , which is a reasonable assumption in typical applications (e.g., Romer and Romer 2004, 2010; Kilian and Vega 2011). Note that we do not constrain $B_{0,22}$ to be lower triangular, allowing B_0 to be block recursive.

We also assume that the first element of $C(L)$ is zero. Under this assumption, the first equation does not include the nonlinear function $f(x_t)$ as a regressor. Realizations of x_t may depend on lagged values of $f(x_t)$ when $B_{12}(L) \neq 0$. Note that, in line with the existing literature, we also rule out the possibility that x_t depends on nonlinear functions of the remaining variables in the system. With these restrictions, we can rewrite (1) as

$$\begin{cases} x_t = b_1 + B_{11}(L)x_{t-1} + B_{12}(L)y_{t-1} + \varepsilon_{1t} \\ B_{0,22}y_t = b_2 + B_{0,21}x_t + B_{21}(L)x_{t-1} + B_{22}(L)y_{t-1} + C_{21}(L)f(x_t) + \varepsilon_{2t}. \end{cases} \quad (3)$$

When there are no nonlinearities, $C_{21}(L) = 0$, in which case (3) is a block recursive linear VAR model for $z_t = (x_t, y_t)'$. When $C_{21}(L) \neq 0$, y_t depends on $f(x_t)$ and lags of $f(x_t)$. We will discuss two economically interesting examples of $f(\cdot)$ at the end of this section.

Furthermore, we impose the following set of assumptions.

Assumption 1. $\{\varepsilon_{1t}\}$ and $\{\varepsilon_{2t}\}$ are mutually independent time series such that

$$\varepsilon_t \equiv \begin{pmatrix} \varepsilon_{1t} \\ \varepsilon_{2t} \end{pmatrix} \sim \text{i.i.d.} \left(\begin{pmatrix} 0 \\ 0 \end{pmatrix}, \begin{pmatrix} \sigma_1^2 & 0 \\ 0 & \Sigma_{22} \end{pmatrix} \right), \text{ where } \Sigma_{22} \text{ is diagonal.}$$

Assumption 2. The roots of the equation $\det(D(L)) = 0$ are outside the unit circle, where $D(L) \equiv I_n - B_0^{-1}B(L)L$, and the process z_t is strictly stationary and ergodic.

Assumption 3. $\sup_t E|x_t| < \infty$ and $\sup_t E|f(x_t)| < \infty$.

Under Assumption 1, the structural errors ε_{1t} and ε_{2t} are i.i.d. over time and follow mutually independent processes. Assumption 2 contains two parts. First, we assume that the roots of the determinantal equation associated with the matrix polynomial $D(L) = I_n - B_0^{-1}B(L)L$ are all outside the unit circle. This assumption ensures the absolute summability of the inverse filter $\Psi(L) \equiv D(L)^{-1}$, which will be used below to obtain the impulse response function. Although this condition is sufficient for the stationarity and ergodicity of z_t when $C_{21}(L) = 0$, it is not when $C_{21}(L) \neq 0$. We therefore impose stationarity and ergodicity of z_t as a high level assumption. For the bivariate case, when x_t is exogenous (i.e. $B_{12}(L) = 0$), we can show that y_t in model (3) is a special case of the nonlinear bivariate ARX model studied by Masry and Tjøstheim (1997), where the only additive nonlinear

term is itself an additive function of current and lagged values of $f(x_t)$. In this case, more primitive conditions on ε_t and $f(x_t)$ can be provided for the stationarity and ergodicity of z_t (see Section 5). When both $C_{21}(L) \neq 0$ and $B_{12}(L) \neq 0$, providing more primitive conditions under which stationarity and ergodicity hold is more challenging. This question is outside the scope of this paper.

Model (3) encompasses several special cases that feature prominently in the empirical literature. For example, x_t may be a directly observable shock such that $x_t = \varepsilon_{1t}$, as postulated in a number of recent studies on fiscal policy and monetary policy shocks (e.g., Romer and Romer 2004, 2010, Ramey 2011, 2016, Tenreyro and Thwaites 2016, Ramey and Zubairy 2018), OPEC oil supply shocks (e.g., Hamilton 2003, Kilian 2008, Bastianin and Manera 2018), technology shocks (e.g., Basu, Fernald and Kimball 2006), news shocks (e.g., Ramey 2011, Kilian and Vega 2011, Kilian and Hicks 2013), financial market shocks (e.g., Barnichon, Matthes and Ziegenbein 2020) and shocks to market expectations (e.g., Kuttner 2001, Cochrane and Piazzesi 2002, Piazzesi and Swanson 2008; Baumeister and Kilian 2016). Alternatively, we may allow for x_t to be an exogenous serially correlated variable. This specification accommodates the concern raised in Alloza, Gonzalo and Sanz (2020) that some measures of exogenous shocks used in the literature actually are serially correlated.³ Finally, we allow for x_t to be endogenously determined with respect to y_t , but predetermined, as in Kilian and Vigfusson’s (2011a,b) analysis of the relationship between the real price of oil and U.S. real GDP growth, for example.

Our goal is to estimate the impulse response function of selected elements of y_{t+h} at horizon $h = 0, 1, \dots, H$ to a shock of magnitude δ in ε_{1t} . For notational convenience, we suppress the subscript denoting this element. Note that even though model (3) is linear in the parameters, it is nonlinear in the variables. Hence, the impulse responses are inherently nonlinear. We discuss the identification and estimation of this response function in Section 4.

2.2 Nonlinear impulse responses

Nonlinear impulse response functions can be defined in many different ways (e.g., Gallant, Rossi and Tauchen 1993, Koop, Pesaran and Potter 1996, Potter 2000, Gouriéroux and Jasiak 2005, Kilian and Vigfusson 2011a). A natural starting point is to compare two sample paths for the variable of interest, one where ε_{1t} is subject to a shock at time t and another one where no such shock is present. The difference between the values of the outcome variable over time under these two scenarios can be interpreted as a measure of the impulse response function.

³This does not mean that all narrative shock measures are serially correlated. For example, the Romer and Romer (2004) monetary policy shock is arguably not serially correlated, and the Romer and Romer (2010) tax shock series is clearly serially uncorrelated. Likewise, the Hussain and Malik (2016) tax shock measure and the Kilian (2008) OPEC oil supply shock measure are serially uncorrelated as are VAR-based shock measures (e.g., Barnichon et al. 2020).

More specifically, we trace out the effect of a shock in ε_{1t} at time t on the future values of the outcome variables y_{t+h} , for $h = 0, 1, \dots, H$ by comparing two sample paths for y_{t+h} . One is the baseline path, which we denote by $\{y_{t+h}\}$. This path is implied by the sequence of structural shocks

$$\varepsilon^\infty = \{\dots, \varepsilon_{1t-1}, \varepsilon_{1t}, \varepsilon_{1t+1}, \dots, \varepsilon_{2t-1}, \varepsilon_{2t}, \varepsilon_{2t+1}, \dots, \}.$$

The other sample path is $\{y_{t+h}(\delta)\}$, which is the path implied by the sequence of shocks

$$\varepsilon^\infty(\delta) = \{\dots, \varepsilon_{1t-1}, \varepsilon_{1t} + \delta, \varepsilon_{1t+1}, \dots, \varepsilon_{2t-1}, \varepsilon_{2t}, \varepsilon_{2t+1}, \dots, \}.$$

The only difference between ε^∞ and $\varepsilon^\infty(\delta)$ occurs at time t , when $\varepsilon_{1t}(\delta) = \varepsilon_{1t} + \delta$. All other shocks are the same. Thus, this thought experiment involves perturbing by δ the structural innovation ε_{1t} that is driving the variable x_t in model (1). This shock translates into a contemporaneous change in x_t of the same magnitude, but a one-time shock to ε_{1t} may imply a persistent change in x_t over time.

Our definition of the nonlinear impulse response function is as follows.

Definition 1 (Unconditional IRF) *The unconditional nonlinear impulse response function of y_{t+h} to a shock of size δ in ε_{1t} is given by $IRF_{h,\delta} = E(y_{t+h}(\delta) - y_{t+h})$, for $h = 0, 1, 2, \dots, H$.*

Several remarks are in order. First, as expected from the literature on nonlinear impulse response functions, the presence of nonlinearities implies that the dynamic response of y_{t+h} to a shock in ε_{1t} depends on the entire sample path of the process. In particular, it depends on current and future values of the shocks in the model. Our approach in this paper is to integrate out this randomness, by considering the expected value of the difference between $\{y_{t+h}(\delta) : h = 0, 1, \dots, H\}$ and $\{y_{t+h} : h = 0, 1, \dots, H\}$.

Second, unlike Gallant, Rossi and Tauchen (1993), Koop, Pesaran and Potter (1996), Potter (2000), Gouriéroux and Jasiak (2005), our main results do not condition on the history of the process up to time $t - 1$, denoted Ω_{t-1} . Alternatively, one could consider a version of our response function that conditions on Ω_{t-1} as in Kilian and Vigfusson (2017).

Definition 2 (Conditional IRF) *The conditional nonlinear impulse response function of y_{t+h} to a shock of size δ in ε_{1t} is given by $IRF_{h,\delta,\Omega_{t-1}} = E[(y_{t+h}(\delta) - y_{t+h}) | \Omega_{t-1}]$, for $h = 0, 1, 2, \dots, H$.*

The special challenges of estimating the latter impulse response function in our context are briefly discussed in Section 3 and in Appendix B.

Third, a further difference with respect to some of the previous literature is the type of shocks that are presumed to occur between t and $t+h$. Koop et al. (1996), Potter (2000) and Kilian and Vigfusson (2011a), for example, set $\varepsilon_{1t}(\delta)$ equal to δ in the perturbed model and let ε_{1t} denote a random draw of

the shock at t in the baseline model. In contrast, we consider the impact of $\varepsilon_{1t}(\delta) = \varepsilon_{1t} + \delta$ versus ε_{1t} , averaged over the possible realizations of ε_{1t} . Although this difference does not matter for some model specifications (such as when $B_{11}(L) = 0$ and $B_{12}(L) = 0$), it may matter for other specifications.

Finally, our approach is closely related to that of Gallant, Rossi and Tauchen (1993) and Gouriéroux and Jasiak (2005), with the difference that Gallant, Rossi and Tauchen (1993) consider shocks to the outcome variable rather than shocks to innovations, and Gouriéroux and Jasiak (2005) consider shocks to a sequence of nonlinear innovations within the context of univariate reduced-form models. Our premise is that the shocks $\varepsilon_t = (\varepsilon_{1t}, \varepsilon'_{2t})'$ in the structural model (1), are i.i.d. over time and mutually independent. This allows us to perturb one of these structural shocks, namely ε_{1t} , without perturbing the other structural shocks ($\{\varepsilon_{2s}\}$ and $\{\varepsilon_{1s}, \text{ for } s \neq t\}$).

2.3 Examples of nonlinearly transformed variables

Our analysis focuses on two leading examples of economically interesting nonlinear transformations $f(\cdot)$ of x_t . One example is the censored variable $x_t^+ \equiv \max(0, x_t)$. Note that model (1) with $f(x_t) = x_t^+$ is equivalent to a model that includes both x_t^+ and $x_t^- \equiv \min(0, x_t)$ with potentially different coefficients. This specification was originally proposed by Mork (1989) and allows for asymmetries in the response of the economy to positive and negative oil price shocks. Mork argued that increases in oil prices matter more than decreases. This proposition has been explored by Kilian and Vigfusson (2011a, 2017), Herrera, Lagalo and Wada (2011, 2015), Alsalman and Herrera (2015), and Herrera and Karaki (2015), among others. Other applications of this framework include the potentially asymmetric passthrough of oil price shocks to gasoline prices (Venditti 2013), the differential effects of positive and negative tax changes on U.S. real GDP (Hussain and Malik 2016), the effects of positive and negative shocks to financial regulation on inflation and industrial output growth (Barnichon, Matthes and Ziegenbein 2020), the effect of positive and negative shocks to regulatory bank oversight (Hwa, Kapinos and Ramirez 2018), and the effects of contractionary and expansionary monetary policy shocks on the economy (Cover 1992, Tenreyro and Thwaites 2016).

The other example involves powers of x_t . For example, Tenreyro and Thwaites (2016) consider an exogenous shock series x_t and include $f(x_t) = x_t^3$ in the regression in addition to x_t . This specification allows larger values of x_t of either sign to have more powerful effects on the outcome variable. A similar approach has also been employed by Hwa, Kapinos and Ramirez (2018) to study the impact of large exogenous changes in ratings by bank supervisors on economic activity.

Next, we derive closed-form expressions for the nonlinear impulse response functions. This will allow us to obtain an estimator of the IRF that does not require Monte Carlo integration.

3 A closed-form expression for the population IRF

To describe the population IRF, we evaluate the difference between $y_{t+h}(\delta)$ and y_{t+h} . Recall that

$$B_0 z_t = b + B(L) z_{t-1} + C(L) f(x_t) + \varepsilon_t. \quad (4)$$

Since B_0 satisfies (2), the inverse matrix of B_0 exists and is given by

$$B_0^{-1} = \begin{pmatrix} 1 & 0 \\ B_{0,22}^{-1} B_{0,21} & B_{0,22}^{-1} \end{pmatrix} \equiv \begin{pmatrix} 1 & 0 \\ B_0^{21} & B_0^{22} \end{pmatrix}.^4$$

Pre-multiplying (4) by B_0^{-1} yields

$$z_t = B_0^{-1} b + B_0^{-1} B(L) z_{t-1} + B_0^{-1} C(L) f(x_t) + B_0^{-1} \varepsilon_t,$$

which we rewrite as

$$z_t = k + A(L) z_{t-1} + G(L) f(x_t) + u_t, \quad (5)$$

where $k = B_0^{-1} b$, and

$$\begin{aligned} A(L) &= B_0^{-1} B(L) \equiv \sum_{i=1}^p A_i L^{i-1}, \text{ where } A_i = B_0^{-1} B_i, \\ G(L) &= B_0^{-1} C(L) \equiv \sum_{i=0}^p G_i L^i, \text{ where } G_i = B_0^{-1} C_i, \text{ and } u_t = B_0^{-1} \varepsilon_t. \end{aligned}$$

If $C(L) = 0$, then (5) is the reduced-form version of the structural model (1). When $C(L) \neq 0$, (5) is not quite a reduced-form model because $f(x_t)$ still appears in the second block of equations of the system and this variable is correlated with u_t . We therefore refer to (5) as the ‘‘pseudo-reduced form’’ model. We will propose a method of estimating the parameters of this model in Section 4, but for now we use this model to define the population IRF.

It follows from (5) that

$$\underbrace{(I_n - A(L)L)}_{=D(L)} z_t = k + G(L) f(x_t) + u_t, \quad (6)$$

where $D(L)$ is a p^{th} degree lag matrix polynomial whose inverse matrix filter $\Psi(L) \equiv D(L)^{-1}$ exists and is absolutely summable under Assumption 2. Using the definition of the inverse filter, i.e. $\Psi(L) D(L) = I_n$, we can show that $\Psi_0 = I_n$, $\Psi_1 = \Psi_0 A_1, \dots$, and $\Psi_j = \Psi_{j-1} A_1 + \dots + \Psi_{j-p} A_p$, for any $j \geq p$.

By pre-multiplying (6) by $\Psi(L)$ and using the fact that $u_t = B_0^{-1} \varepsilon_t$, we can write

$$z_t = \mu + \Theta(L) \varepsilon_t + \Gamma(L) f(x_t), \quad (7)$$

⁴For any matrix A , we let A^{ij} denote the block (i, j) of A^{-1} .

where

$$\begin{aligned}\mu &= \Psi(1)b = \Psi(1)B_0^{-1}k, \\ \Theta(L) &\equiv \Psi(L)B_0^{-1}, \\ \Gamma(L) &\equiv \Psi(L)G(L) = \Psi(L)B_0^{-1}C(L).\end{aligned}$$

If $C(L) = 0$, then $\Gamma(L) = 0$, and (7) is the vector moving average representation of z_t . In this case, we can obtain the IRF from the first column of $\Theta(L)$. With $\Gamma(L) \neq 0$, additional terms must be added to the IRF. Let

$$\Theta_{\cdot 1}(L) = \begin{pmatrix} \Theta_{11}(L) \\ \Theta_{21}(L) \end{pmatrix}_{n \times 1} \text{ and } \Theta_{\cdot 2}(L) = \begin{pmatrix} \Theta_{12}(L) \\ \Theta_{22}(L) \end{pmatrix}_{n \times n_1},$$

where $\Theta_{\cdot 1}(L)$ denotes the first column of $\Theta(L)$ and $\Theta_{\cdot 2}(L)$ denotes its remaining columns. Note that for any i, j , $\Theta_{ij}(L) = \Theta_{0,ij} + \Theta_{1,ij}L + \dots + \Theta_{h,ij}L^h + \dots$, implying that, for example,

$$\Theta_{\cdot 1}(L) = \Theta_{0,\cdot 1} + \Theta_{1,\cdot 1}L + \dots + \Theta_{h,\cdot 1}L^h + \dots,$$

with $\Theta_{h,\cdot 1} = (\Theta_{h,11}, \Theta'_{h,21})'$.

From (7), for the baseline model, we can write

$$z_{t+h} = \mu + \Theta_{\cdot 1}(L)\varepsilon_{1t+h} + \Theta_{\cdot 2}(L)\varepsilon_{2t+h} + \Gamma(L)f(x_{t+h}),$$

whereas for the δ -perturbed model,

$$z_{t+h}(\delta) = \mu + \Theta_{\cdot 1}(L)\varepsilon_{1t+h}(\delta) + \Theta_{\cdot 2}(L)\varepsilon_{2t+h}(\delta) + \Gamma(L)f(x_{t+h}(\delta)).$$

Since $\varepsilon_{1t}(\delta) = \varepsilon_{1t} + \delta$, $\varepsilon_{1s}(\delta) = \varepsilon_{1s}$ for all $s \neq t$, and $\varepsilon_{2s}(\delta) = \varepsilon_{2s}$ for all s , we have that

$$z_{t+h}(\delta) - z_{t+h} = \Theta_{h,\cdot 1}\delta + \Gamma(L)[f(x_{t+h}(\delta)) - f(x_{t+h})].$$

Our definition of the two sequences of shocks implies that $x_{t+h}(\delta) = x_{t+h}$ for any $h < 0$. Hence,

$$\begin{aligned}z_{t+h}(\delta) - z_{t+h} &= \Theta_{h,\cdot 1}\delta + \Gamma_0[f(x_{t+h}(\delta)) - f(x_{t+h})] \\ &\quad + \Gamma_1[f(x_{t+h-1}(\delta)) - f(x_{t+h-1})] + \dots + \Gamma_h[f(x_t(\delta)) - f(x_t)].\end{aligned}\quad (8)$$

We rely on equation (8) to evaluate the two sample paths of the outcome variables y_{t+h} corresponding to the sequences of shocks $\varepsilon^\infty(\delta)$ and ε^∞ . This yields

$$y_{t+h}(\delta) - y_{t+h} = \Theta_{h,21}\delta + \Gamma_{0,21}[f(x_{t+h}(\delta)) - f(x_{t+h})] + \dots + \Gamma_{h,21}[f(x_t(\delta)) - f(x_t)],$$

where the last $h+1$ terms reflect the contribution of the nonlinearities to the IRF. These terms depend on the differences $f(x_{t+h}(\delta)) - f(x_{t+h})$, which are random functions of the path of x_t up to time $t+h$.

Our approach is to integrate out this randomness, by defining $IRF_{h,\delta}$ as the unconditional expectation of $y_{t+h}(\delta) - y_{t+h}$. The next proposition formalizes this result and describes an algorithm for evaluating $A_{j,\delta} = E(f(x_{t+j}(\delta))) - E(f(x_{t+j}))$, where $x_{t+j}(\delta)$ is written as a function of $\{x_{t+j}, \dots, x_t\}$ and the model parameters.

Proposition 3.1 *Under Assumptions 1, 2, and 3, for any $h = 0, 1, 2, \dots, H$,*

$$IRF_{h,\delta} = \Theta_{h,21}\delta + \Gamma_{0,21}A_{h,\delta} + \Gamma_{1,21}A_{h-1,\delta} + \dots + \Gamma_{h,21}A_{0,\delta},$$

where

$$A_{j,\delta} = E(f(x_{t+j}(\delta))) - E(f(x_{t+j})).$$

The following steps can be used to calculate $x_{t+j}(\delta)$ as a function of $\{x_{t+j}, \dots, x_t\}$ for $j = 0, 1, \dots, h$:

i) For $j = 0$, set $x_t(\delta) = x_t + \delta$ and $A_{0,\delta} = E(f(x_t + \delta) - f(x_t))$.

ii) For $j = 1, 2, \dots, h$, let

$$\begin{aligned} x_{t+j}(\delta) &= x_{t+j} + \Theta_{j,11}\delta + \Gamma_{1,11}[f(x_{t+j-1}(\delta)) - f(x_{t+j-1})] + \dots + \Gamma_{j,11}[f(x_t(\delta)) - f(x_t)] \\ &\equiv g_{j,\delta}(x_{t+j}, x_{t+j-1}, \dots, x_t; \beta_j), \end{aligned}$$

where $g_{j,\delta}(\cdot, \beta_j)$ is implicitly defined by this recursion and $\beta_j = (\Theta_{1,11}, \dots, \Theta_{j,11}, \Gamma_{1,11}, \dots, \Gamma_{j,11})'$.

iii) For $j = 1, 2, \dots, h$, let

$$A_{j,\delta} = E(f(g_{j,\delta}(x_{t+j}, x_{t+j-1}, \dots, x_t; \beta_j))) - E(f(x_{t+j})).$$

As Proposition 3.1 shows, computing the IRF at horizon h involves evaluating $h + 1$ expectation terms $A_{j,\delta}$ for $j = 0, 1, \dots, h$. Each of these evaluations requires computing $x_{t+j}(\delta)$, where $x_{t+j}(\delta)$ is the value of the variable x at time $t+j$ in the perturbed version of the model. For $j = 0$, $x_t(\delta) = x_t + \delta$ as given by step (i). For $j > 0$, we use step (ii) to obtain $x_{t+j}(\delta)$ recursively. This defines $x_{t+j}(\delta)$ as an implicit function $g_{j,\delta}(x_{t+j}, \dots, x_t; \beta_j)$ of the random variables (x_{t+j}, \dots, x_t) , the magnitude of the shock δ and β_j , a vector of parameters that depend continuously on the structural parameters θ . In particular, we can show that $\beta_j = (\Theta_{1,11}, \dots, \Theta_{j,11}, \Gamma_{1,11}, \dots, \Gamma_{j,11})'$.

To illustrate the algorithm described in Proposition 3.1, suppose we want to evaluate $IRF_{h,\delta}$ for $h = 1$. This requires evaluating the terms $A_{0,\delta}$ and $A_{1,\delta}$. By step (i),

$$x_t(\delta) = x_t + \delta \text{ and } A_{0,\delta} = E(f(x_t + \delta) - f(x_t)),$$

where the expectation is with respect to the marginal distribution of x_t . To obtain $A_{1,\delta}$, we use step (ii) with $j = 1$ to write

$$x_{t+1}(\delta) = x_{t+1} + \Theta_{1,11}\delta + \Gamma_{1,11} [f(x_t + \delta) - f(x_t)] \equiv g_{1,\delta}(x_{t+1}, x_t; \beta_1),$$

where $x_t(\delta) = x_t + \delta$ from step (i) and $\beta_1 = (\Theta_{1,11}, \Gamma_{1,11})'$. This implies that

$$A_{1,\delta} = E(f(g_{1,\delta}(x_{t+1}, x_t; \beta_1)) - f(x_{t+1})),$$

where the expectation is with respect to the joint distribution of (x_{t+1}, x_t) . We can proceed in this manner to compute $IRF_{h,\delta}$ for any value of h . As we increase h , more terms $A_{j,\delta}$ need to be computed since $IRF_{h,\delta}$ requires $A_{j,\delta}$ for $j = 0, 1, \dots, h$. The functions $g_{j,\delta}$ that implicitly define $x_{t+j}(\delta)$ as a function of the observables $\{x_{t+j}, x_{t+j-1}, \dots, x_t\}$ can be computed recursively as we did for $j = 1$.

Note that the computation of $g_{j,\delta}$ and hence of $A_{j,\delta}$ can be simplified in two special cases of our model. One is when $x_t = \varepsilon_{1t}$, in which case $x_{t+h}(\delta) = x_{t+h}$ for all $h \neq 0$ since $\varepsilon_{1s}(\delta) = \varepsilon_{1s}$ for all $s \neq t$. In this case, $A_{0,\delta} = E(f(x_t + \delta) - f(x_t))$, as in step (i) above, but $A_{j,\delta} = 0$ for all $j \neq 0$. We will consider a bivariate version of this special case in Section 5 to illustrate the properties of local projections in our framework. Another special case where the computation of $A_{j,\delta}$ is simplified is when x_t is an exogenous strictly stationary $AR(p)$ process. In this case, $B_{11}(L) \neq 0$ but $B_{12}(L) = 0$, which implies that $\Gamma_{11}(L) = 0$. Thus, step (ii) simplifies to

$$x_{t+j}(\delta) = x_{t+j} + \Theta_{j,11}\delta,$$

showing that $x_{t+j}(\delta)$ is a function of x_{t+j} and $\beta_j \equiv \Theta_{j,11}$, the j^{th} coefficient of the lag polynomial $\Theta_{11}(L)$.

The closed-form expressions for the IRF given by Proposition 3.1 can be evaluated in practice by replacing the unknown parameters with consistent estimates. This yields a novel “plug-in” estimator of the IRF which can be used as an alternative to the MCI approach proposed by Kilian and Vigfusson (2011a) for fully recursive models. Note in particular that the iterative algorithm in Proposition 3.1 does not require any simulations, contrary to the MCI method, making this approach computationally attractive. To estimate the $A_{j,\delta}$ terms, we only need to evaluate the sample average of the difference $f\left(g_{j,\delta}\left(x_{t+j}, x_{t+j-1}, \dots, x_t; \hat{\beta}_j\right)\right) - f(x_{t+j})$ at each horizon $j = 1, 2, \dots, h$, where $\hat{\beta}_j$ is an estimator of β_j . More importantly, whereas the MCI approach is feasible only in the special case of fully recursive models, the plug-in estimator can be used without taking a stand on the identification of ε_{2t} .

When x_t is i.i.d., there is no difference between the unconditional IRF and the IRF conditional on Ω_{t-1} as in Definition 2. This equivalence breaks down when x_t depends on lags of x_t or y_t . Although it is feasible to obtain the closed-form solutions for the conditional IRF in the latter case, these analytical

expressions are specific to the functional form of $f(x_t)$ and require imposing further restrictions on the model, in particular in the form of assumptions about the distribution of the error term ε_{1t} , as illustrated in Appendix B.

4 Estimation of the IRF

The closed-form expressions for $IRF_{h,\delta}$ obtained in Proposition 3.1 show that the IRF depends on two sets of coefficients: those from the first column of $\Theta(L)$ given by

$$\Theta_{\cdot 1}(L) = \Psi(L) \begin{pmatrix} 1 \\ B_0^{21} \end{pmatrix},$$

where $\Psi(L) = (I_n - A(L)L)^{-1}$, and those from

$$\Gamma(L) = \Psi(L) \begin{pmatrix} 0 \\ G_{21}(L) \end{pmatrix}.$$

Thus, in order to identify $IRF_{h,\delta}$ we need to identify the parameters in $\Psi(L)$, $G_{21}(L)$ and B_0^{21} . When the model is fully recursive (i.e. $B_{0,22}$ is lower triangular with 1's along the main diagonal), OLS estimation of each equation of the structural model yields consistent estimates of θ , the structural parameters entering $B_0, B(L)$ and $C(L)$. These in turn imply consistent estimates of the parameters in $\Psi(L)$, $G_{21}(L)$ and B_0^{21} . Full identification of the structural model, however, is not required to identify the dynamic responses of y_t to a shock in ε_{1t} . The only shock we need to identify is ε_{1t} . This only requires B_0 to be block recursive. Under this partial identification scheme, we can rely on the pseudo-reduced form equation for z_t to identify $IRF_{h,\delta}$, as shown next.

Recall our pseudo-reduced form model (5):

$$z_t = k + A(L)z_{t-1} + G(L)f(x_t) + u_t,$$

where $G(L) = (0, G_{21}(L))'$ and $A(L) = A_1 + A_2L + \dots + A_pL^{p-1}$. In the following, let $A_{i,1}$ denote the $1 \times n$ vector containing the first row of A_i , and let $A_{i,2}$ denote the $n_1 \times n$ matrix containing the remaining rows of A_i . With this notation, we can write $A(L)$ as

$$A(L) = \begin{pmatrix} A_{1\cdot}(L) \\ A_{2\cdot}(L) \end{pmatrix} = \begin{pmatrix} A_{1,1} + A_{2,1}L + \dots + A_{p,1}L^{p-1} \\ A_{1,2} + A_{2,2}L + \dots + A_{p,2}L^{p-1} \end{pmatrix}.$$

The pseudo-reduced form model then is

$$\begin{cases} x_t = k_1 + A_{1\cdot}(L)z_{t-1} + u_{1t} \\ y_t = k_2 + A_{2\cdot}(L)z_{t-1} + G_{21}(L)f(x_t) + u_{2t}, \end{cases} \quad (9)$$

where

$$u_t = B_0^{-1}\varepsilon_t = \begin{pmatrix} 1 & 0 \\ B_0^{21} & B_0^{22} \end{pmatrix} \begin{pmatrix} \varepsilon_{1t} \\ \varepsilon_{2t} \end{pmatrix} = \begin{pmatrix} \varepsilon_{1t} \\ B_0^{21}\varepsilon_{1t} + B_0^{22}\varepsilon_{2t} \end{pmatrix}.$$

Since $u_{1t} = \varepsilon_{1t}$, u_{1t} is independent of z_{t-1} and its lags, and $A_1(L)$ can be estimated by OLS. However, the presence of nonlinearities in the second block of equations implies that $f(x_t)$ is correlated with u_{2t} . Hence, when $G_{21}(L) \neq 0$, OLS estimation of (9) will not yield consistent estimates, except for the first equation.

The source of the endogeneity is the presence of ε_{1t} in u_{2t} . Suppose for a moment that we could observe ε_{1t} . Then we could effectively remove this source of endogeneity by including ε_{1t} as an additional control in the equations for y_t . This can be seen from rewriting (9) as

$$\begin{cases} x_t = k_1 + A_1(L) z_{t-1} + \varepsilon_{1t} \\ y_t = k_2 + A_2(L) z_{t-1} + G_{21}(L) f(x_t) + B_0^{21} \varepsilon_{1t} + u_{2t}^*, \end{cases} \quad (10)$$

where $u_{2t}^* \equiv B_0^{22} \varepsilon_{2t}$. Since by Assumption 1 u_{2t}^* is orthogonal to all the regressors in the equation for y_t (including ε_{1t}), the regression of y_t on a constant, z_{t-1} , $f(x_t)$, their lags, and ε_{1t} would provide consistent estimates of $A_2(L)$, $G_{21}(L)$ and B_0^{21} . This together with $A_1(L)$ would identify $A(L)$, and therefore $\Psi(L)$, implying that all the parameters required for identifying $IRF_{h,\delta}$ could be consistently estimated.

In practice, we do not observe ε_{1t} , but we can estimate ε_{1t} from the first equation in (10). Replacing ε_{1t} with $\hat{\varepsilon}_{1t}$ before applying OLS to the second equation amounts to using a control function approach (e.g. Wooldridge 2010, p. 268-269). For this purpose, we introduce the following notation. First, we write the first equation of (10) as

$$x_t = \pi_1' w_{1t} + \varepsilon_{1t},$$

where $w_{1t} = (1, Z_{t-1}')'$ is $d_1 \times 1$ with $Z_{t-1} = (z_{t-1}', \dots, z_{t-p}')'$, $d_1 = 1 + np$, and

$$\pi_1 = (k_1, A_{1,1}', \dots, A_{p,1}')'.$$

Furthermore, let $F_t = (f(x_t), f(x_{t-1}), \dots, f(x_{t-p}))'$ denote a $(p+1) \times 1$ vector containing the current and lagged values of the nonlinear function $f(x_t)$. Let $w_{2t} = (w_{1t}', F_t', \varepsilon_{1t})'$ be a $d_2 \times 1$ vector containing the regressors in the second equation of (10), including the unobserved error ε_{1t} (implying that $d_2 = d_1 + p + 2$). Similarly, let

$$\Pi_2' = \begin{pmatrix} k_2 & A_{1,2} & \cdots & A_{p,2} & G_{0,21} & \cdots & G_{p,21} & B_0^{21} \end{pmatrix}.$$

Then the second equation of (10) can be expressed as

$$y_t = \Pi_2' w_{2t} + u_{2t}^*,$$

and the system of equations (10) can be compactly written as

$$\begin{aligned} x_t &= \pi_1' w_{1t} + \varepsilon_{1t}, \\ y_t &= \Pi_2' w_{2t} + u_{2t}^*, \end{aligned}$$

where the orthogonality conditions

$$E(w_{1t}\varepsilon_{1t}) = 0 \text{ and } E(u_{2t}^*w_{2t}') = 0,$$

hold by Assumption 1. This implies that

$$\pi_1 = (E(w_{1t}w_{1t}'))^{-1} E(w_{1t}x_t) \text{ and } \Pi_2 = (E(w_{2t}w_{2t}'))^{-1} E(w_{2t}y_t')$$

exist, provided the inverse of $\Sigma_{w_1} = E(w_{1t}w_{1t}')$ and $\Sigma_{w_2} = E(w_{2t}w_{2t}')$ exists. This suggests estimating π_1 and Π_2 by OLS. However, since w_{2t} depends on ε_{1t} , we do not observe w_{2t} . Our strategy is to replace w_{2t} with $\hat{w}_{2t} = (w_{1t}', F_t', \hat{\varepsilon}_{1t})'$, where $\hat{\varepsilon}_{1t} = x_t - \hat{\pi}_1'w_{1t}$. This introduces a generated regressor problem in the estimation of Π_2 , which affects inference but not the consistency of the estimator.⁵

Algorithm for the IRF plug-in estimator in the general structural model

i) Regress x_t onto $w_{1t} = (1, Z_{t-1}')'$, where $Z_{t-1} = (z_{t-1}', \dots, z_{t-p}')'$. Collect the estimated parameters in $\hat{\pi}_1 = (\hat{k}_1, \hat{A}_{1,1}', \dots, \hat{A}_{p,1}')'$ and let $\hat{\varepsilon}_{1t} = x_t - \hat{\pi}_1'w_{1t}$.

ii) Regress y_t onto $\hat{w}_{2t} = (w_{1t}', F_t', \hat{\varepsilon}_{1t})'$ and collect the estimated parameters in $\hat{\pi}_2 = \text{vec}(\hat{\Pi}_2)$, where

$$\hat{\Pi}_2' = \begin{pmatrix} \hat{k}_2 & \hat{A}_{1,2} & \dots & \hat{A}_{p,2} & \hat{G}_{0,21} & \dots & \hat{G}_{p,21} & \hat{B}_0^{21} \end{pmatrix}.$$

iii) Use $\hat{\pi} = (\hat{\pi}_1', \hat{\pi}_2')'$ to compute

$$\hat{\Psi}(L) = (I_n - \hat{A}(L)L)^{-1} = I_n + \hat{\Psi}_1L + \hat{\Psi}_2L^2 + \dots + \hat{\Psi}_hL^h + \dots,$$

where

$$\hat{A}(L) = \hat{A}_1 + \hat{A}_2L + \dots + \hat{A}_pL^{p-1},$$

implying that

$$\begin{aligned} \hat{\Psi}_0 &= I_n, \quad \hat{\Psi}_1 = \hat{\Psi}_0\hat{A}_1, \quad \hat{\Psi}_2 = \hat{\Psi}_1\hat{A}_1 + \hat{\Psi}_0\hat{A}_2, \quad \dots, \text{ and} \\ \hat{\Psi}_h &= \hat{\Psi}_{h-1}\hat{A}_1 + \dots + \hat{\Psi}_{h-p}\hat{A}_p, \text{ for any } h \geq p. \end{aligned}$$

Then compute

$$\hat{\Theta}_{\cdot 1}(L) = \hat{\Psi}(L) \begin{pmatrix} 1 \\ \hat{B}_0^{21} \end{pmatrix} \equiv \begin{pmatrix} \hat{\Theta}_{11}(L) \\ \hat{\Theta}_{21}(L) \end{pmatrix} \text{ and } \hat{\Gamma}(L) = \hat{\Psi}(L) \begin{pmatrix} 0 \\ \hat{G}_{21}(L) \end{pmatrix} \equiv \begin{pmatrix} \hat{\Gamma}_{11}(L) \\ \hat{\Gamma}_{21}(L) \end{pmatrix}.$$

⁵Our analysis does not take into account the possibility that x_t may also be estimated, as recently discussed in Breitung and Brüggemann (2020). Of course, x_t need not be a generated regressor. For example, the Kilian (2008) measure of exogenous OPEC oil supply shocks does not rely on regression analysis. Nor do measures of expectation shifts computed as changes in futures prices around policy announcement dates (see Piazzesi and Swanson 2008).

iv) Set

$$\hat{A}_{0,\delta} = \frac{1}{T} \sum_{t=1}^T (f(x_t + \delta) - f(x_t)),$$

and for $j = 1, \dots, h$, let

$$\hat{A}_{j,\delta} = \frac{1}{T-j} \sum_{t=1}^{T-j} f\left(g_{j,\delta}\left(x_{t+j}, x_{t+j-1}, \dots, x_t; \hat{\beta}_j\right)\right) - \frac{1}{T} \sum_{t=1}^T f(x_t),$$

where $g_{j,\delta}(\cdot; \hat{\beta}_j)$ is defined by the recursion in step ii) of Proposition 3.1 with

$$\hat{\beta}_j = \left(\hat{\Theta}_{1,11}, \dots, \hat{\Theta}_{j,11}, \hat{\Gamma}_{1,11}, \dots, \hat{\Gamma}_{j,11}\right)'.$$

v) Set

$$IRF_{h,\delta}^{Plug-in} = \hat{\Theta}_{h,21}\delta + \hat{\Gamma}_{0,21}\hat{A}_{h,\delta} + \hat{\Gamma}_{1,21}\hat{A}_{h-1,\delta} + \dots + \hat{\Gamma}_{h,21}\hat{A}_{0,\delta}.$$

To summarize, in steps i) and ii) we construct the OLS estimator of $\pi = (\pi'_1, \pi'_2)$. Step iii) uses $\hat{\pi}$ to construct estimates of the coefficients $\Theta_{\cdot 1}(L)$ and $\Gamma(L)$, which enter the formula of the IRF. Step iv) computes an estimator of $A_{j,\delta} \equiv E[f(x_{t+j}(\delta)) - f(x_{t+j})]$ as a function of the data and the estimates obtained in step iii). For instance, for $j = 1$,

$$\hat{A}_{1,\delta} = \frac{1}{T-1} \sum_{t=1}^{T-1} f\left(g_{1,\delta}\left(x_{t+1}, x_t; \hat{\theta}\right)\right) - \frac{1}{T} \sum_{t=1}^T f(x_t),$$

where

$$g_{1,\delta}\left(x_{t+1}, x_t; \hat{\beta}_1\right) = x_{t+1} + \hat{\Theta}_{1,11}\delta + \hat{\Gamma}_{1,11}[f(x_t + \delta) - f(x_t)].$$

Step v) evaluates the population expression for the IRF at the estimated parameter values.

We conclude this section by proving the consistency of the plug-in IRF estimator. The following additional assumptions jointly with Assumptions 1-3 are sufficient to establish this result.

Assumption 4. $\sup_t E|x_t|^4 < \infty$ and $\sup_t E|f(x_t)|^4 < \infty$.

Assumption 5. $\Sigma_{w_1} = E(w_{1t}w'_{1t})$ and $\Sigma_{w_2} \equiv E(w_{2t}w'_{2t})$ are positive definite.

Assumption 6. The function f is continuous and is such that for each $j = 1, \dots, h$,

$$|f(g_{j,\delta}(x_{t+j}, \dots, x_t; \beta_j))| < b_j(x_{t+j}, \dots, x_t)$$

where β_j lies in a compact space and $E|b_j(x_{t+j}, \dots, x_t)| < \infty$.

Theorem 4.1 Under Assumptions 1 through 6, $IRF_{h,\delta}^{Plug-in} \xrightarrow{P} IRF_{h,\delta}$ as $T \rightarrow \infty$.

The proof of Theorem 4.1 follows from Lemma A.1 and A.2 in the Appendix. Lemma A.1 establishes the consistency of $\hat{\pi}$ for π . This result follows under Assumptions 1-5 by standard arguments that rely on laws of large numbers for stationary ergodic processes. The consistency of $\hat{\Theta}_{h,21}$, $\hat{\Theta}_{j,11}$, $\hat{\Gamma}_{j,11}$ and $\hat{\Gamma}_{j,21}$ follows from the consistency of $\hat{\pi}$ given that these estimators are continuous functions of $\hat{\pi}$. Note that we do not need to identify the structural parameters θ to prove the consistency of the IRF. Indeed, the structural equations for y_t are not identified unless we put further restrictions (such as a triangular structure) on B_0^{22} . Hence, ε_{2t} is not identified in this model. This is the sense in which the model is only partially identified (see Kilian and Lütkepohl 2017). However, the structural shock of interest, ε_{1t} , is identified and this is sufficient to identify π and hence the IRF through the control function approach described above.

Assumption 6 serves to show that $\hat{A}_{j,\delta}$ is a consistent estimator of $A_{j,\delta}$ for $j = 1, \dots, h$ (cf. Lemma A.2). This assumption ensures that we can apply a uniform law of large numbers for stationary ergodic data. The continuity assumption covers the empirically motivated examples in Section 2.3, but does not cover, for example, threshold dynamics of the form $x_t^{l\text{arg}e} = x_t 1(|x_t| > \kappa\sigma)$, where κ denotes a multiple of the standard deviation σ of x_t , typically set to 1 or 2, as discussed in Goldberg (1988), Davis and Kilian (2011) and Alsalman and Herrera (2015).

5 Can local projections identify the IRF?

The main goal of this section is to investigate the ability of local projections to identify the unconditional IRF in our framework. To do so, we focus on a bivariate version of our model where $x_t = \varepsilon_{1t}$. Our main result is that even in this simplified model, the conventional LP approach fails to identify the population IRF when there are nonlinearities. However, we show that it is possible to modify the LP approach for this simple model in order to recover the nonlinear IRF. We establish the consistency of this modified LP estimator of the IRF when $x_t = \varepsilon_{1t}$. We show that this result does not extend to situations when x_t is serially correlated.

In order to develop intuition, consider a simplified version of our model with $b_1 = 0$, $b_2 = b$, $B_{0,21} = \beta$, $B_{21}(L) = 0$, $B_{22}(L) = \rho$, and $C_{21}(L) = c$. Thus,

$$\begin{cases} x_t = \varepsilon_{1t}, \\ y_t = b + \beta x_t + \rho y_{t-1} + cf(x_t) + \varepsilon_{2t}. \end{cases} \quad (11)$$

In this special case,

$$IRF_{h,\delta} = \beta\rho^h\delta + c\rho^h \underbrace{E(f(x_t + \delta) - E(f(x_t)))}_{=A_{0,\delta}},$$

such that $\Theta_{h,21} \equiv \beta\rho^h$ and $\Gamma_{h,21} \equiv c\rho^h$. The value of y_{t+h} is

$$y_{t+h} = b + \beta x_{t+h} + \rho y_{t+h-1} + cf(x_{t+h}) + \varepsilon_{2t+h},$$

which we can write as

$$y_{t+h} = k_h + \beta\rho^h x_t + c\rho^h f(x_t) + \rho^{h+1} y_{t-1} + u_{t+h}, \quad (12)$$

where $k_h = \mu(1 - \rho^{h+1})$ and u_{t+h} is a function of $\{\varepsilon_{1t+h}, \dots, \varepsilon_{1t+1}, f(\varepsilon_{1t+h}), \dots, f(\varepsilon_{1t+1}), \varepsilon_{2t+h}, \dots, \varepsilon_{2t+1}, \varepsilon_{2t}\}$.

If $c = 0$, expression (12) can be written as

$$y_{t+h} = k_h + \pi_{x,h} x_t + \pi_{y,h} y_{t-1} + u_{t+h}, \quad (13)$$

where $\pi_{x,h} = \Theta_{h,21}$ and

$$E(u_{t+h}) = E(x_t u_{t+h}) = E(y_{t-1} u_{t+h}) = 0,$$

using the fact that ε_{1t} and ε_{2t} are mutually uncorrelated sequences of i.i.d. zero mean random variables. This implies the usual result that we can recover the IRF at lag h by setting $IRF_{h,\delta} = \pi_{x,h}$, where $\pi_{x,h}$ is the slope coefficient associated with x_t in the regression of y_{t+h} onto x_t and y_{t-1} (and a constant).

When $c \neq 0$, this result no longer holds. In this case, one may still use an LP regression to obtain $\Theta_{h,21}$ and $\Gamma_{h,21}$, but one needs to add an estimate of $A_{0,\delta} \equiv E[f(x_t + \delta) - f(x_t)]$. The LP regression is

$$y_{t+h} = \pi_{c,h} + \pi_{x,h} x_t + \pi_{f,h} f(x_t) + \pi_{y,h} y_{t-1} + v_{t+h}, \quad (14)$$

where $\pi_{x,h} = \beta\rho^h \equiv \Theta_{h,21}$, $\pi_{f,h} = c\rho^h \equiv \Gamma_{h,21}$ and $\pi_{y,h} = \rho^{h+1}$. The error term v_{t+h} has mean zero and satisfies the orthogonality conditions $E(x_t v_{t+h}) = E(f(x_t) v_{t+h}) = E(y_{t-1} v_{t+h}) = 0$ given that v_{t+h} only depends on $\{\varepsilon_{1s}, f(\varepsilon_{1s}) - E(f(\varepsilon_{1s})) : s = t+1, \dots, t+h\}$ and $\{\varepsilon_{2s} : s = t, \dots, t+h\}$ and given that these shocks are independent of $x_t = \varepsilon_{1t}$ and $f(x_t)$ as well as of y_{t-1} .⁶

Equation (14) is the local projection equation we need to estimate to recover $IRF_{h,\delta}$ when $c \neq 0$. The coefficients $\pi_{x,h}$ and $\pi_{f,h}$ associated with x_t and $f(x_t)$, respectively, are equal to $\Theta_{h,21}$ and $\Gamma_{h,21}$. Estimates of these coefficients together with an estimate of $A_{0,\delta}$ can be used to estimate $IRF_{h,\delta} = \Theta_{h,21} + \Gamma_{h,21} A_{0,\delta}$.

As the next proposition shows, these results generalize to the bivariate version of model (1) where $x_t = \varepsilon_{1t}$ and $B_{21}(L)L$, $B_{22}(L)L$ and $C_{21}(L)$ are general p^{th} order polynomials in L .

⁶Note that we redefine the error term of the LP regression as $v_{t+h} = u_{t+h} - E(u_{t+h})$. The reason is that $E(u_{t+h})$ is not necessarily zero because $E(f(x_t)) = E(f(\varepsilon_{1t}))$ may not be zero (even though $E(\varepsilon_{1t}) = 0$). This transformation is without loss of generality since we include a constant in the LP regression.

Proposition 5.1 *Assume Assumptions 1, 2 and 3 hold with $n_1 = 1$, and $b_1 = B_{11}(L) = B_{12}(L) = 0$. Then, for $h = 0, 1, 2, \dots$, we have that*

$$y_{t+h} = \pi_{c,h} + \pi_{x,h}x_t + \pi_{x,h}(L)x_{t-1} + \pi_{f,h}f(x_t) + \pi_{f,h}(L)f(x_{t-1}) + \pi_{y,h}(L)y_{t-1} + v_{t+h}$$

where

$$\pi_{x,h} = \Theta_{h,21} \quad \text{and} \quad \pi_{f,h} = \Gamma_{h,21},$$

and v_{t+h} is mean zero and is orthogonal to all the regressors.

Proposition 5.1 implies that the coefficients associated with x_t and $f(x_t)$ in the regression of y_{t+h} on $w'_{2t} = (1 \quad Z'_{t-1} \quad F'_t \quad x_t)'$ are equal to $\pi_{x,h} = \Theta_{h,21}$ and $\pi_{f,h} = \Gamma_{h,21}$. These coefficients are required to compute $IRF_{h,\delta}$ in Proposition 3.1.

The algorithm for constructing this modified LP estimator is:

Algorithm for estimating IRF based on LPs when x_t is i.i.d.

i) For each h , regress y_{t+h} onto $w'_{2t} = (1 \quad Z'_{t-1} \quad F'_t \quad x_t)'$ and let $\hat{\pi}_{x,h}$ and $\hat{\pi}_{f,h}$ denote the slope coefficients associated with x_t and $f(x_t)$, respectively.

ii) Obtain an estimate of $A_{0,\delta} \equiv E[f(x_t + \delta) - f(x_t)]$ as

$$\hat{A}_{0,\delta} = \frac{1}{T} \sum_{t=1}^T (f(x_t + \delta) - f(x_t)).$$

iii) Set

$$IRF_{h,\delta}^{LP} = \hat{\pi}_{x,h}\delta + \hat{\pi}_{f,h}\hat{A}_{0,\delta}.$$

The modified LP estimator $IRF_{h,\delta}^{LP}$ is an alternative to the plug-in estimator $IRF_{h,\delta}^{Plug-in}$ given in Section 4 when $x_t = \varepsilon_{1t}$. The modified LP estimator generates estimates of $\pi_{x,h} = \Theta_{h,21}$ and $\pi_{f,h} = \Gamma_{h,21}$ by running a regression of y_{t+h} onto w_{2t} for each horizon h , whereas the plug-in estimator obtains these coefficients by iterating on the OLS estimator from the regression of y_t onto w_{2t} . Both estimators rely on $\hat{A}_{0,\delta}$, the sample average of $A_{0,\delta}$, to obtain the final estimate of the IRF.

Remark 1 *The modified LP estimator described here differs from the LP estimator used in existing studies that allow for nonlinear transformations of x_t (e.g., Tenreyro and Thwaites 2016, Barnichon et al. 2020). Consider the example of $f(x_t) = x_t^\dagger$. The conventional approach has been to fit local projections,*

$$y_{t+h} = \pi_{c,h} + \pi_{x,h}x_t + \pi_{f,h}x_t^\dagger + v_{t+h}, \tag{15}$$

or, equivalently,

$$y_{t+h} = \pi_{c,h} + \pi_{x^+,h}x_t^+ + \pi_{x^-,h}x_t^- + v_{t+h}, \quad (16)$$

for $h = 0, 1, 2, \dots, H$, where, for simplicity, we have dropped all lagged regressors. Then the sequences $\{\pi_{x^+,h}\}_{h=0}^H$ and $\{\pi_{x^-,h}\}_{h=0}^H$ have been interpreted as the response functions with respect to a positive and a negative ε_{1t} shock of size $\delta = 1$ and $\delta = -1$, respectively. Our analysis shows that this approach does not recover the population unconditional response functions even asymptotically.

Remark 2 The coefficients $\pi_{x^+,h}$ and $\pi_{x^-,h}$ in the standard LP regression (16) can be re-interpreted as capturing a conditional version of the IRF at horizon h , where we condition on $x_t = \varepsilon_{1t} \geq 0$ and $\delta = 1$ and $x_t = \varepsilon_{1t} < 0$ and $\delta = -1$, respectively. To see this, let $x_t \geq 0$ and perturb it by $\delta = 1$. In this case, it follows that

$$y_{t+h}(1) - y_{t+h} = \Theta_{h,21}\delta + \Gamma_{h,21} \underbrace{(f(x_t+1) - f(x_t))}_{=1} = \Theta_{h,21} + \Gamma_{h,21} = \pi_{x^+,h},$$

where $\pi_{x^+,h} = E(y_{t+h}(1) - y_{t+h} | x_t \geq 0)$. This result does not necessarily apply to other specifications of $f(x_t)$, however, and does not apply when x_t is serially correlated.⁷

We can establish the consistency of $IRF_{h,\delta}^{LP}$ (as well as $IRF_{h,\delta}^{Plug-in}$) in the bivariate case with $B_{12}(L) = 0$ under a set of more primitive conditions than those used in Theorem 4.1, if we replace Assumption 2 by the following assumption.

Assumption 2'

- (i) $B_{12}(L) = 0$ and the roots of $1 - B_{11}(L)L = 0$ and $1 - B_{22}(L)L = 0$ are all outside the unit circle.
- (ii) The function f is nonperiodic and bounded on compact subsets of \mathbb{R} such that $|f(x)| \leq M|x|^{\gamma_1}$ for some $\gamma_1 \in \mathbb{R}$, for all $|x| > x_0$, for some $x_0 > 0$ and $M < \infty$.
- (iii) The i.i.d. random variables ε_{1t} and ε_{2t} have continuous probability density functions that are strictly positive on \mathbb{R} such that $E\left(|\varepsilon_{1t}|^{\max(1,\gamma_1+\gamma_2)}\right) < \infty$ for some $\gamma_2 > 0$.

When $B_{12}(L) = 0$, z_t given by (1) is a special case of the nonlinear bivariate ARX model studied by Masry and Tjøstheim (1997). Assumptions 2'(ii) and 2'(iii) correspond to their Assumptions 3(b) and 3(c) respectively. It can be shown that Assumption 2'(i) suffices for their Assumption 3.3(c). Thus, under Assumption 2', we can appeal to Masry and Tjøstheim (1997)'s Lemma 3.1 to conclude

⁷A similar conditional interpretation of LP coefficients when x_t is serially correlated would require conditioning on the sign of future values of x_t (e.g. $x_t \geq 0$ and $x_{t+1} \geq 0$ for $h = 1$ if $x_t = \phi x_{t-1} + \varepsilon_{1t}$) and would further require taking a stand on the sign of the dynamic coefficients driving x_t (e.g. assume that $\phi > 0$ in the AR(1) example).

that z_t is stationary and strongly mixing with exponentially decaying coefficients. Note that part (ii) of Assumption 2' allows f to grow at most at a polynomial rate, where γ_1 denotes the polynomial degree. This covers the empirically motivated examples discussed previously, but excludes other nonlinear functions whose growth rate is larger than the polynomial rate such as the exponential function. There is a trade-off between how fast $f(x_t)$ can grow as a function of x_t (as dictated by γ_1) and the number of finite moments of ε_{1t} , as required by Assumption 2' (iii).

Theorem 5.1 *Under Assumptions 1, 2', 3, 4 and 5 with $b_1 = B_{11}(L) = B_{12}(L) = 0$, as $T \rightarrow \infty$,*

$$IRF_{h,\delta}^{Plug-in} \xrightarrow{P} IRF_{h,\delta} \quad \text{and} \quad IRF_{h,\delta}^{LP} \xrightarrow{P} IRF_{h,\delta},$$

for any $h = 0, 1, \dots$ and fixed δ .

Note that we do not require Assumption 6 to prove Theorem 5.1 since we do not need to estimate $A_{j,\delta}$ for $j > 0$. The only nonlinear term that appears in the IRF is $A_{0,\delta} \equiv E(f(x_t + \delta) - f(x_t))$, which we estimate as the sample average of $f(x_t + \delta) - f(x_t)$. Since x_t is i.i.d., the consistency of $\hat{A}_{0,\delta}$ follows by a standard law of large numbers for i.i.d. data. An application of the uniform law of large numbers is not required, thus dispensing with Assumption 6.

A natural question is whether we can use the modified LP approach to obtain an estimate of $IRF_{h,\delta}$ when x_t is serially correlated. The answer is no. To explain this result, we again consider a simplified version of our model, where now

$$\begin{cases} x_t = b_1 + \phi x_{t-1} + \varepsilon_{1t}, \\ y_t = b_2 + \beta x_t + \rho y_{t-1} + c f(x_t) + \varepsilon_{2t}. \end{cases}$$

Given Proposition 3.1, we can show that

$$IRF_{h,\delta} = \Theta_{h,21} + \Gamma_{0,21} A_{h,\delta} + \Gamma_{1,21} A_{h-1,\delta} + \dots + \Gamma_{h,21} A_{0,\delta},$$

where $\Theta_{h,21} = \beta(\phi^h + \rho\phi^{h-1} + \dots + \rho^{h-1}\phi + \rho^h)$, $\Gamma_{h,21} = c\rho^h$ and $A_{h,\delta} = E[f(x_{t+h} + \phi^h\delta) - f(x_{t+h})]$, where the expression for $A_{h,\delta}$ follows from the fact that $x_{t+h}(\delta) = x_{t+h} + \phi^h\delta$. We can write

$$y_{t+h} = \mu + \beta(1 - \rho L)^{-1} x_{t+h} + c(1 - \rho L)^{-1} f(x_{t+h}) + (1 - \rho L)^{-1} \varepsilon_{2t+h},$$

where $\mu = b_2(1 - \rho)^{-1}$ and $x_t = \phi x_{t-1} + \varepsilon_{1t}$. Suppose first $c = 0$. We can decompose y_{t+h} as

$$\begin{aligned} y_{t+h} &= \mu \left(1 - \rho^{h+1}\right) + \beta \rho^h x_t + \rho^{h+1} y_{t-1} + \beta \left(x_{t+h} + \rho x_{t+h-1} + \dots + \rho^{h-1} x_{t+1}\right) \\ &\quad + \left(\varepsilon_{2t+h} + \rho \varepsilon_{2t+h-1} + \dots + \rho^{h-1} \varepsilon_{2t+1} + \rho^h \varepsilon_{2t}\right). \end{aligned}$$

For any $h \geq 0$, we can write

$$x_{t+h} = \pi_{x,c} + \phi^h x_t + \left(\phi^{h-1} \varepsilon_{1t+1} + \dots + \varepsilon_{1t+h} \right),$$

with $\pi_{x,c} = b_1 (1 + \phi + \dots + \phi^{h-1})$. This implies that

$$\begin{aligned} y_{t+h} &= \pi_{c,h} + \underbrace{\beta \left(\phi^h + \rho \phi^{h-1} + \dots + \rho^{h-1} \phi + \rho^h \right)}_{=\Theta_{h,21}} x_t + \rho^{h+1} y_{t-1} \\ &\quad + \beta \left(\phi^{h-1} \varepsilon_{1t+1} + \dots + \varepsilon_{1t+h} \right) + \beta \rho \left(\phi^{h-2} \varepsilon_{1t+1} + \dots + \varepsilon_{1t+h-1} \right) + \dots + \beta \rho^{h-1} \varepsilon_{1t+1} \\ &\quad + \left(\varepsilon_{2t+h} + \rho \varepsilon_{2t+h-1} + \dots + \rho^{h-1} \varepsilon_{2t+1} + \rho^h \varepsilon_{2t} \right), \end{aligned}$$

for some constant $\pi_{c,h}$. This equation can be written as

$$y_{t+h} = \pi_{c,h} + \Theta_{h,21} x_t + \rho^{h+1} y_{t-1} + u_{t+h},$$

where u_{t+h} is implicitly defined by the equation above and is a linear combination of $\varepsilon_{t+1} = (\varepsilon_{1t+1}, \varepsilon_{2t+1})'$ through $\varepsilon_{t+h} = (\varepsilon_{1t+h}, \varepsilon_{2t+h})'$ as well as ε_{2t} . Hence, $E(x_t u_{t+h}) = E(y_{t-1} u_{t+h}) = 0$. Thus, we can recover $\Theta_{h,21}$ from a local projection given by the regression of y_{t+h} onto x_t and y_{t-1} .

Now, suppose that $c \neq 0$. Then, a similar argument implies the following decomposition for y_{t+h} :

$$y_{t+h} = \pi_{c,h} + \Theta_{h,21} x_t + \rho^{h+1} y_{t-1} + \underbrace{c \rho^h}_{=\Gamma_{h,21}} f(x_t) + v_{t+h}, \quad (17)$$

where the error term now contains $f(x_{t+j})$ for $j > 0$ and is given by

$$v_{t+h} = u_{t+h} + c \left[f(x_{t+h}) + \rho f(x_{t+h-1}) + \dots + \rho^{h-1} f(x_{t+1}) \right].$$

While this error can be transformed to have mean zero (by subtracting $\mu_f = E f(x_{t+j})$ from each $f(x_{t+j})$), it is not orthogonal to the regressors in (17) when x_t is serially correlated. The reason is that each x_{t+j} can be written as a function of x_t and $\{\varepsilon_{1t+j}, \dots, \varepsilon_{1t+1}\}$, implying that $f(x_{t+j})$ is correlated with x_t and $f(x_t)$ as well as y_{t-1} . As a result, this equation is not a local projection and cannot be used to identify $\Theta_{h,21}$ nor $\Gamma_{h,21}$. Moreover, these are not the only coefficients needed to obtain $IRF_{h,\delta}$ when x_t is serially correlated. We also need the $\Gamma_{j,21}$ terms for $j = 0, \dots, h-1$, which are missing from (17)). One might consider including $f(x_{t+h}), \dots, f(x_{t+1})$ as additional regressors in (17) in order to capture all the required coefficients, as in

$$y_{t+h} = \pi_{c,h} + \Theta_{h,21} x_t + \rho^{h+1} y_{t-1} + \underbrace{c \rho^h}_{=\Gamma_{h,21}} f(x_t) + \underbrace{c}_{=\Gamma_{0,21}} f(x_{t+h}) + \underbrace{c \rho}_{=\Gamma_{1,21}} f(x_{t+h-1}) + \dots + \underbrace{c \rho^{h-1}}_{=\Gamma_{h-1,21}} f(x_{t+1}) + u_{t+h},$$

However, u_{t+h} would now be correlated with $f(x_{t+j})$ for $j > 0$.⁸

⁸If the nonlinear term were $f(\varepsilon_{1t})$ rather than $f(x_t)$ these complications would not arise because perturbing ε_{1t} by δ would only have an effect in the impact period. This means that we would effectively be back in the special case of $x_t = \varepsilon_{1t}$. The only difference would be that we would need to estimate the residual by regressing x_t on lags of x_t and y_t . We are not aware of any applications of this specification in empirical work. One reason is that the underlying economic arguments often dictate a nonlinearity in observables rather than shocks.

6 Simulation evidence

In this section, we compare the finite-sample accuracy of the plug-in estimator with that of the Monte Carlo integration (MCI) estimator and the modified LP estimator, where applicable. The evaluation criteria are the bias, variance and mean squared error (MSE) of the impulse response estimators. We also examine the robustness of the plug-in estimator to dynamic model misspecification and illustrate the convergence of the plug-in estimator in block recursive models.

6.1 Baseline simulation design

We consider three special cases of model (1). DGP 1 restricts x_t to an observed i.i.d. shock:

$$\begin{cases} x_t = \varepsilon_{1t}, \\ y_t = 0.5y_{t-1} + 0.5x_t + 0.3x_{t-1} - 0.4 \max(0, x_t) + 0.3 \max(0, x_{t-1}) + \varepsilon_{2t}, \end{cases} \quad (18)$$

In DGP 2, x_t is instead follows an exogenous AR(1) process:

$$\begin{cases} x_t = 0.5x_{t-1} + \varepsilon_{1t}, \\ y_t = 0.5y_{t-1} + 0.5x_t + 0.3x_{t-1} - 0.4 \max(0, x_t) + 0.3 \max(0, x_{t-1}) + \varepsilon_{2t}, \end{cases} \quad (19)$$

DGP 3 treats x_t as predetermined with respect to y_t :

$$\begin{cases} x_t = 0.3x_{t-1} + 0.2y_{t-1} + \varepsilon_{1t}, \\ y_t = 0.5y_{t-1} + 0.5x_t + 0.3x_{t-1} - 0.4 \max(0, x_t) + 0.2 \max(0, x_{t-1}) + \varepsilon_{2t} \end{cases} \quad (20)$$

In all DGPs, the intercept has been normalized to 0 in population and the population innovations are mutually independent and distributed $N(0, 1)$. Results for DGPs with other parameter values are qualitatively similar and hence are not reported.

The number of Monte Carlo trials is 10,000. For each draw from the DGP, we estimate the unconditional impulse response function of y_{t+h} , $h = 0, 1, \dots, H$ to a shock in ε_{1t} of magnitude δ . For expository purposes, we set $\delta = 1$. Very similar results are obtained for other choices of δ . The MCI method is implemented as the average value of the conditional response functions over 1,000 randomly drawn histories, where each conditional response function is based on 1,000 draws. The plug-in method and LP method are implemented as discussed in Sections 4 and 5. The lag order for all local projections is set to one, consistent with the assumption of a known lag order of $p = 1$ in the DGP. For expository purposes, we focus on samples of length $T = 240$, corresponding to 20 years of monthly data or 60 years of quarterly data. Very similar results are obtained for $T = 120$, corresponding to ten years of monthly data (or 30 years of quarterly data) and $T = 480$, corresponding to 40 years of monthly data. The complete results for all combinations of $T \in \{120, 240, 480\}$ and $\delta \in \{-2, -1, 1, 2\}$ are reported in the online appendix.

6.2 Baseline simulation results

The first row of Figure 1 shows that when the modified LP estimator is applicable, it tends to be less accurate in finite samples than the plug-in estimator or the MCI estimator. Although the modified LP estimator has about the same bias as the plug-in estimator, it has much higher variance and hence a higher MSE. The performance of the plug-in estimator and the MCI estimator is virtually identical, as expected for this DGP. In the second row of Figure 1, the bias and variance of the MCI and plug-in estimators again tends to be almost the same. Only for much smaller T , is there any evidence that these estimators are not effectively identical. Finally, for the unrestricted DGP in the third row, the plug-in estimator has slightly higher variance than the MCI estimator and slightly lower bias at short horizons. As a practical matter, these differences are negligible. Thus, in fully recursive models, there is little to choose between the plug-in estimator and the MCI estimator based on their finite-sample accuracy.

The plug-in estimator, however, is substantially less computationally demanding. It typically reduces the computational cost by 98% or more, which is a significant improvement when conducting bootstrap inference or when constructing critical values under the null of symmetric response functions. For example, when estimating the variance of the impulse response estimator by bootstrap, as required for tests of the symmetry of the response functions, the computation time for any of the models considered in Figure 1 drops from about one hour to under one minute. When allowing for six lags in the regression model, the computation time drops from close to four hours to about one minute.

6.3 Robustness to dynamic misspecification

We have assumed so far that the regression model is dynamically correctly specified. One of the perceived advantages of the conventional linear LP estimator compared to linear VAR models is its potential greater robustness to dynamic misspecification (Plagborg-Møller and Wolf 2021). This argument does not extend to our analysis, however, because the modified LP estimator assumes knowledge of the correct lag order much like the plug-in estimator.

In Figure 2, we examine how robust the ranking of these two estimators is to underfitting the DGP. We consider an empirically representative DGP, where $x_t = \varepsilon_{1t}$ and y_t depends on the current value and six lag of x_t and $\max(x_t, 0)$ as well as six lags of its own. Details of this DGP can be found in the online appendix. The regression models are based on $p \in \{4, 5, 6\}$. Since $p = 6$ in population, the estimated model is dynamically misspecified when its lag order is $p < 6$. The number of Monte Carlo trials is 10,000, and, for expository purposes, we set $T = 240$ and $\delta = 1$ as in Figure 1.

Figure 2 illustrates that there is no bias-variance trade-off between the modified LP estimator and

the plug-in estimator in our context. Even when underfitting by two lags, the plug-in estimator tends to have lower MSE than the modified LP estimator. Underfitting increases the bias, but lowers the variance of the plug-in estimator, as expected. For the modified LP estimator, there is no systematic change in the bias, but underfitting lowers the variance of the impulse responses. Figure 2 shows that, in finite samples, a modest degree of underfitting may lower the MSE of both estimators. The reason is that the increase in the variance from estimating population parameters at longer lags outweighs the reduction in the bias when these parameters are close to zero, as is typically the case when y_t is expressed as a growth rate in applied work.

6.4 Higher-dimensional block recursive models

In higher-dimensional block recursive models, there are no alternative estimators to compare the plug-in estimator to, but we can illustrate the convergence of the plug-in estimator, as $T \rightarrow \infty$. For expository purposes, we consider DGPs of the form

$$B_0 z_t = B_1 z_{t-1} + C_0 f(x_t) + C_1 f(x_{t-1}) + \varepsilon_t,$$

where ε_t is generated as a 3×1 vector of independent standard normal innovations and $f(x_t) = \max(0, x_t)$ or $f(x_t) = x_t^3$, respectively. DGP 4 treats x_t as an i.i.d. shock, DGP 5 allows x_t to be serially correlated, and DGP 6 only assumes that x_t is predetermined with respect to y_t . The parameter values are

$$B_0 = \begin{pmatrix} 1 & 0 & 0 \\ -0.45 & 1 & -0.3 \\ -0.05 & 0.1 & 1 \end{pmatrix}, C_0 = \begin{pmatrix} 0 \\ -0.2 \\ 0.08 \end{pmatrix}, C_1 = \begin{pmatrix} 0 \\ -0.1 \\ 0.2 \end{pmatrix},$$

and

$$B_1 = \begin{pmatrix} 0 & 0 & 0 \\ 0.15 & 0.17 & -0.18 \\ -0.08 & 0.03 & 0.6 \end{pmatrix} \text{ for DGP 4, } B_1 = \begin{pmatrix} -0.13 & 0 & 0 \\ 0.15 & 0.17 & -0.18 \\ -0.08 & 0.03 & 0.6 \end{pmatrix} \text{ for DGP 5,}$$

$$\text{and } B_1 = \begin{pmatrix} -0.13 & 0.05 & -0.01 \\ 0.15 & 0.17 & -0.18 \\ -0.08 & 0.03 & 0.6 \end{pmatrix} \text{ for DGP 6.}$$

Figures 3 and 4 show the MSE of the impulse response estimator based on 10,000 draws from each DGP for $T \in \{120, 240, 480, 960\}$. The plots illustrate that in all cases the pointwise MSE of the response estimates declines, as T increases.

7 Empirical application

There has been much interest in possible nonlinearities in the response of U.S. macroeconomic aggregates to monetary policy shocks (e.g., Cover 1992, Tenreyro and Thwaites 2016, Barnichon and

Matthes 2018). Our empirical application is motivated by the work of Tenreyro and Thwaites (2016) who employed a conventional LP estimator to examine the evidence for a nonlinear response of real GDP, inflation and the federal funds rate to an exogenous monetary policy shock. Their narrative measure of exogenous monetary policy shocks is similar to the Romer and Romer (2004) series.

The empirical model is given by the four-variable system of equations (1) where $z_t = (x_t \ i_t \ y_t \ \pi_t)'$, x_t is the narrative measure of U.S. monetary policy shocks, i_t is the federal funds rate, y_t is real GDP, and π_t is CPI inflation. The estimation period is 1969:Q1-2007:Q4. We follow Tenreyro and Thwaites in setting $p = 1$ and log-linearly detrending real GDP. Inflation and the federal funds rate are expressed in differences. Unlike Tenreyro and Thwaites' (2016) model, our model is block recursive. Since we are only interested in the responses to the monetary policy shock, the block recursive nature of B_0 and the assumption that $C_{11}(L) = 0$ suffice to identify these responses without having to take a stand on the identification of the remaining shocks. Two key questions raised in Tenreyro and Thwaites (2016) are whether large monetary policy shocks are disproportionately more powerful than small shocks and whether positive shocks have larger effects than negative shocks. They address the first question by defining $f(x_t) = x_t^3$ and the second question by allowing for separate coefficients for positive x_t and negative x_t , which is algebraically equivalent to defining $f(x_t) = \max(0, x_t)$ in our model. For expository purposes, we set $\delta = 1$.

Figure 5 reports the responses allowing for these nonlinearities. It shows that for both specifications of $f(x_t)$ an unexpected exogenous monetary tightening causes the federal funds rate to spike, followed by a gradual decline. Real GDP experiences a sustained, but temporary decline with a delay of about two quarters. Inflation shows a hump-shaped response, but that response is negligible relative to the variation in the federal funds rate. For comparison, we also include the responses based on the corresponding model excluding $f(x_t)$. As Figure 5 illustrates, the inclusion of $f(x_t) = \max(0, x_t)$ makes little difference for the responses of inflation and of the federal funds rate, but greatly reduces the economically implausible positive short-run response in real GDP associated with an unexpected monetary tightening in the model that does not allow for asymmetric responses. The inclusion of $f(x_t) = x_t^3$, in contrast, does not make much of a difference.

8 Conclusion

The paper has focused on the identification and consistent estimation of impulse responses functions in linear structural dynamic models with nonlinearly transformed regressors. We introduced a computationally efficient control function approach to estimating the population responses to a given shock without taking a stand on how the remainder of the model is identified. We provided sufficient condi-

tions for the asymptotic validity of this estimator that cover many of the model specifications currently in use in applied work. We contrasted our approach with local projection and Monte Carlo integration estimators.

An obvious next step would be to provide formal justification for the construction of bootstrap confidence intervals for these response functions and the use of bootstrap critical values in testing restrictions on the response functions. Another interesting question for future research would be the development of impulse response estimators for more general nonlinear models. One example is models with state dependence where the effect of a shock, for example, may depend on whether the economy is in recession or not. Another example is models with smooth threshold effects. Such nonlinear models are inherently different from the models studied in the current paper in that the regression models themselves are nonlinear in the parameters.

9 References

Alloza, M., Gonzalo, J., Sanz, C., 2020. Dynamic Effects of Persistent Shocks. Manuscript, Bank of Spain.

Alsalmán, Z., Herrera, A.M., 2015. Oil price shocks and the U.S. stock market: Do sign and size matter? *Energy Journal* 36, 171-188.

Barnichon, R., Matthes, C., Ziegenbein, A. 2020. Are the effects of financial market disruptions big or small? *Review of Economics and Statistics*, forthcoming.

Barnichon, R., Matthes, C. 2018. Functional Approximation of Impulse Responses. *Journal of Monetary Economics* 99, 41-55.

Bastianin, A., Manera, M., 2018. How does stock market volatility react to oil shocks? *Macroeconomic Dynamics* 22, 666-682.

Basu, S., Fernald, J.G., Kimball, M. 2006. Are Technology Improvements Contractionary? *American Economic Review* 96, 1418-1448.

Baumeister, C., Kilian, L., 2016. Forty years of oil price fluctuations: Why the price of oil may still surprise us. *Journal of Economic Perspectives* 30, 139-160.

Breitung, J., Brüggemann, R. 2020. Projection estimators for structural impulse responses. Manuscript, University of Cologne.

Cochrane, J., Piazzesi, M., 2002. The Fed and interest rates: A high-frequency identification. *American Economic Review* 92, 90-95.

Cover, J.P., 1992. Asymmetric effects of positive and negative money supply shocks. *Quarterly Journal of Economics* 107, 1261-1282.

- Davis, L. W., Kilian, L., 2011. Estimating the effect of a gasoline tax on carbon emissions. *Journal of Applied Econometrics* 16, 1187-1214.
- Gallant, A.R., Rossi, P.E., Tauchen, G., 1993. Nonlinear dynamic structures. *Econometrica* 61, 871-907.
- Gourieroux, C., Jasiak, J., 2005. Nonlinear innovations and impulse responses with application to VaR sensitivity. *Annales d'Economie et de Statistique* 78, 1-31.
- Goldberg, P.K., 1998. The effects of the corporate average fuel efficiency standards in the U.S. *Journal of Industrial Economics* 46, 1-33.
- Hamilton, J.D., 2003. What is an oil shock? *Journal of Econometrics* 113, 363-398.
- Hamilton, J.D., 2011. Nonlinearities and the macroeconomic effects of oil prices. *Macroeconomic Dynamics* 15, 364-378.
- Herrera, A.M., Karaki, M.B. 2015. The effect of oil price shocks on job reallocation. *Journal of Economic Dynamics and Control* 61, 95-113.
- Herrera, A.M., Lagalo, L.G., Wada, T., 2011. Oil price shocks and industrial production: Is the relationship linear? *Macroeconomic Dynamics* 15, 472-497.
- Herrera, A.M., Lagalo, L.G., Wada, T., 2015. Asymmetries in the response of economic activity to oil price increases and decreases? *Journal of International Money and Finance* 50, 108-133.
- Hussain, S.M., Malik, S., 2016. Asymmetric effects of exogenous tax changes. *Journal of Economic Dynamics and Control* 69, 268-300.
- Hwa, V., Kapinos, P., Ramirez, C.D., 2018. Does regulating bank oversight impact economic activity? A local projections approach. *Journal of Financial Stability* 39, 167-174.
- Jordà, O., 2005. Estimation and inference of impulse responses by local projections. *American Economic Review* 95, 161-182.
- Jordà, O., 2009. Simultaneous confidence regions for impulse responses. *Review of Economics and Statistics* 91, 629-647.
- Kilian, L., 2008. Exogenous oil supply shocks: How big are they and how much do they matter for the U.S. economy? *Review of Economics and Statistics* 90, 216-240.
- Kilian, L., Hicks, B. 2013. Did unexpectedly strong economic growth cause the oil price shock of 2003-2008? *Journal of Forecasting* 32, 385-394.
- Kilian, L., Lütkepohl, H., 2017. *Structural vector autoregressive analysis*. Cambridge University Press, New York.
- Kilian, L., Vega, C., 2011. Do energy prices respond to U.S. macroeconomic news? A test of the hypothesis of predetermined energy prices. *Review of Economics and Statistics* 93, 660-671.

- Kilian, L., Vigfusson, R.J., 2011a. Are the responses of the U.S. economy asymmetric in energy price increases and decreases? *Quantitative Economics* 2, 419-453.
- Kilian, L., Vigfusson, R.J., 2011b. Nonlinearities in the oil price-output relationship. *Macroeconomic Dynamics* 15, 337-363.
- Kilian, L., Vigfusson, R.J., 2017. The role of oil price shocks in causing U.S. recessions. *Journal of Money, Credit, and Banking* 49, 1747-1776.
- Koop, G., Pesaran, M.H., Potter, S.M., 1996. Impulse response analysis in nonlinear multivariate models. *Journal of Econometrics* 74, 119-147.
- Kuttner, K., 2001. Monetary policy surprises and interest rates: Evidence from the federal funds futures market. *Journal of Monetary Economics* 47, 523-544.
- Masry, E., Tjøstheim, D., 1997. Additive nonlinear ARX time series and projection estimates. *Econometric Theory* 13, 214-252.
- Mavroeidis, S. 2020. Identification at the zero lower bound. Manuscript, Oxford University.
- Mork, K. A., 1989. Oil and the macroeconomy. when prices go up and down: An extension of Hamilton's results. *Journal of Political Economy* 97, 740-744.
- Piazzesi, M., Swanson, E., 2008. Futures prices as risk adjusted forecasts of monetary policy. *Journal of Monetary Economics* 55, 677-691.
- Plagborg-Møller, M., Wolf, C.K., 2021. Local projections and VARs estimate the same impulse responses. *Econometrica*, forthcoming.
- Potter, S.M., 2000. Nonlinear impulse response functions. *Journal of Economic Dynamics and Control* 24, 1425-1446.
- Ramey, V. A., 2011. Identifying government spending shocks: It's all in the timing. *Quarterly Journal of Economics* 126, 1-50.
- Ramey, V. A., 2016. Macroeconomic shocks and their propagation. In J. B. Taylor and H. Uhlig (eds.): *Handbook of Macroeconomics*, volume 2, chapter 2, Elsevier, pp. 71-162.
- Ramey, V. A., Zubairy, S., 2018. Government spending multipliers in good times and in bad: Evidence from U.S. historical data. *Journal of Political Economy* 126, 850-901.
- Romer, C.D., Romer, D.H., 2004. A new measure of monetary shocks: Derivation and implications. *American Economic Review* 94, 1055-1084.
- Romer, C.D., Romer, D.H., 2010. The macroeconomic effects of tax changes: Estimates based on a new measure of fiscal shocks. *American Economic Review* 100, 763-801.
- Tenreyro, S., and Thwaites, G., 2016. Pushing on a string: U.S. monetary policy is less powerful in recessions. *American Economic Journal: Macroeconomics* 8, 43-74.

Venditti, F. 2013. From oil to consumer energy prices: How much asymmetry along the way? *Energy Economics* 40, 468-473.

Wooldridge, J. 1994. Estimation and inference for dependent processes. In: *Handbook of Econometrics*, Volume IV, Edited by R.F. Engle and D.L. McFadden, Elsevier Science.

Wooldridge, J. 2010. *Econometrics analysis of cross section and panel data*. MIT Press, Cambridge.

A Appendix: Proofs

Proof of Proposition 3.1. We can write (7) as

$$\begin{pmatrix} x_t \\ y_t \end{pmatrix} = \mu + \Theta_{\cdot,1}(L) \varepsilon_{1t} + \Theta_{\cdot,2}(L) \varepsilon_{2t} + \Gamma(L) f(x_t), \quad (21)$$

where

$$\Theta_{0,\cdot,1} = \begin{pmatrix} 1 \\ B_0^{21} \end{pmatrix}, \quad \Theta_{0,\cdot,2} = \begin{pmatrix} 0 \\ B_0^{22} \end{pmatrix} \quad \text{and} \quad \Gamma_0 = \begin{pmatrix} 0 \\ B_0^{21} C_{0,21} \end{pmatrix}.$$

To see this, note that $\Psi(L) = I_n + \Psi_1 L + \Psi_2 L^2 + \dots$, implying that

$$\Theta(L) = B_0^{-1} + \Psi_1 B_0^{-1} L + \Psi_2 B_0^{-1} L^2 + \dots \equiv \Theta_0 + \Theta_1 L + \Theta_2 L^2 + \dots$$

Hence,

$$\Theta_0 = B_0^{-1} = \begin{pmatrix} 1 & 0 \\ B_0^{21} & B_0^{22} \end{pmatrix} \equiv (\Theta_{0,\cdot,1} \quad \Theta_{0,\cdot,2}).$$

Similarly, we can write

$$\Gamma(L) = \Gamma_0 + \Gamma_1 L + \Gamma_2 L^2 + \dots,$$

where in particular

$$\Gamma_0 = B_0^{-1} C_0 = \begin{pmatrix} 1 & 0 \\ B_0^{21} & B_0^{22} \end{pmatrix} \begin{pmatrix} 0 \\ C_{0,21} \end{pmatrix} = \begin{pmatrix} 0 \\ B_0^{21} C_{0,21} \end{pmatrix} \equiv \begin{pmatrix} \Gamma_{0,11} \\ \Gamma_{0,21} \end{pmatrix}$$

The form of Θ_0 and Γ_0 implies that x_t does not depend on the current value of ε_{2t} (because $\Theta_{0,12} = 0$) nor does it depend on $f(x_t)$ (since $\Gamma_{0,1} = 0$). However, it may depend on past values of ε_{2t} and $f(x_t)$. Equation (21) and our definition of the shock sequences then imply

$$z_{t+h}(\delta) - z_{t+h} = \Theta_{h,\cdot,1} \delta + \Gamma_0 [f(x_{t+h}(\delta)) - f(x_{t+h})] + \dots + \Gamma_h [f(x_t(\delta)) - f(x_t)], \quad (22)$$

which describes a system of equations. The first equation is

$$\begin{aligned} x_{t+h}(\delta) - x_{t+h} &= \Theta_{h,11} \delta + \underbrace{\Gamma_{0,11}}_{=0} [f(x_{t+h}(\delta)) - f(x_{t+h})] + \dots + \Gamma_{h,11} [f(x_t(\delta)) - f(x_t)] \\ &= \Theta_{h,11} \delta + \Gamma_{1,11} [f(x_{t+h-1}(\delta)) - f(x_{t+h-1})] + \dots + \Gamma_{h,11} [f(x_t(\delta)) - f(x_t)], \end{aligned}$$

where we have used the fact that $\Gamma_{0,11} = 0$ due to our restrictions on B_0 and C_0 . The remaining equations of the system are given by

$$y_{t+h}(\delta) - y_{t+h} = \Theta_{h,21} \delta + \Gamma_{0,21} [f(x_{t+h}(\delta)) - f(x_{t+h})] + \dots + \Gamma_{h,21} [f(x_t(\delta)) - f(x_t)].$$

This implies

$$IRF_{h,\delta} = \Theta_{h,21} \delta + \Gamma_{0,21} \underbrace{E[f(x_{t+h}(\delta)) - f(x_{t+h})]}_{\equiv A_{h,\delta}} + \dots + \Gamma_{h,21} \underbrace{E[f(x_t(\delta)) - f(x_t)]}_{\equiv A_{0,\delta}}.$$

■

Next, we provide two lemmas that are instrumental in proving Theorem 4.1.

Lemma A.1 Under Assumptions 1-5, $\hat{\pi} - \pi \xrightarrow{P} 0$.

Proof. We first show the consistency of $\hat{\pi} = (\hat{\pi}'_1, \hat{\pi}'_2)'$ towards π . The proof that

$$\hat{\pi}_1 - \pi_1 = \left(\frac{1}{T} \sum_{t=1+p}^T w_{1t} w'_{1t} \right)^{-1} \frac{1}{T} \sum_{t=1+p}^T w_{1t} \varepsilon_{1t}$$

converges to 0 follows from standard arguments since $\{w_{1t} w'_{1t}\}$ is stationary and ergodic (which follows from Assumption 2). This implies that $\frac{1}{T} \sum_{t=1+p}^T w_{1t} w'_{1t} \xrightarrow{P} \Sigma_{w_1} > 0$ (given Assumptions 4 and 5). Similarly, we can show $\frac{1}{T} \sum_{t=1+p}^T w_{1t} \varepsilon_{1t} \xrightarrow{P} E(w_{1t} \varepsilon_{1t}) = 0$, where the orthogonality conditions can be verified under Assumption 1. In particular, w_{1t} depends on $Z_{t-1} = (z_{t-1}, \dots, z_{t-p})$, where $z_t = (x_t, y_t)'$, which is only a function of lags of $\varepsilon_t = (\varepsilon_{1t}, \varepsilon_{2t})$ and their nonlinear transforms. Assumption 1 then implies that $E(w_{1t} \varepsilon_{1t}) = 0$, delivering the result. To establish the consistency of $\hat{\pi}_2$ towards π_2 , let

$$\Pi_2^* = \left(\frac{1}{T} \sum_{t=1+p}^T w_{2t} w'_{2t} \right)^{-1} \frac{1}{T} \sum_{t=1+p}^T w_{2t} y'_t \equiv [\pi_{2,1}^*, \dots, \pi_{2,n_1}^*],$$

and note that the consistency of Π_2^* towards Π_2 follows from the orthogonality conditions $E(w_{2t} w'_{2t}) = 0$ provided $\Sigma_{w_2} \equiv E(w_{2t} w'_{2t})$ is nonsingular (which holds under Assumption 5). Next we show that replacing \hat{w}_{2t} with w_{2t} is asymptotically valid. We consider

$$\hat{\Pi}_2 = [\hat{\pi}_{2,1}, \dots, \hat{\pi}_{2,n_1}] = \left(\frac{1}{T} \sum_{t=1+p}^T \hat{w}_{2t} \hat{w}'_{2t} \right)^{-1} \frac{1}{T} \sum_{t=1+p}^T \hat{w}_{2t} y'_t,$$

where $\hat{w}_{2t} = w_{2t} + R_t$, with

$$R_t = \begin{pmatrix} 0 \\ \hat{\varepsilon}_{1t} - \varepsilon_{1t} \end{pmatrix} \text{ and } \hat{\varepsilon}_{1t} - \varepsilon_{1t} = -w'_{1t} (\hat{\pi}_1 - \pi_1).$$

We can write

$$\begin{aligned} \hat{\Pi}_2 - \Pi_2^* &= \zeta_1 + \zeta_2, \\ \zeta_1 &= \left\{ \left(\frac{1}{T} \sum_{t=1+p}^T \hat{w}_{2t} \hat{w}'_{2t} \right)^{-1} - \left(\frac{1}{T} \sum_{t=1+p}^T w_{2t} w'_{2t} \right)^{-1} \right\} \frac{1}{T} \sum_{t=1+p}^T \hat{w}_{2t} y'_t, \\ \zeta_2 &= \left(\frac{1}{T} \sum_{t=1+p}^T w_{2t} w'_{2t} \right)^{-1} \frac{1}{T} \sum_{t=1+p}^T (\hat{w}_{2t} - w_{2t}) y'_t. \end{aligned}$$

Next we show that $\zeta_1 = o_P(1)$ and $\zeta_2 = o_P(1)$. Starting with ζ_1 , note that

$$\frac{1}{T} \sum_{t=1+p}^T \hat{w}_{2t} \hat{w}'_{2t} = \frac{1}{T} \sum_{t=1+p}^T w_{2t} w'_{2t} + \frac{1}{T} \sum_{t=1+p}^T w_{2t} R'_t + \frac{1}{T} \sum_{t=1+p}^T R_t w'_{2t} + \frac{1}{T} \sum_{t=1+p}^T R_t R'_t. \quad (23)$$

The first term on the RHS of (23) converges to Σ_{w_2} , which is positive definite by assumption, so it suffices to show that the last three terms converge in probability to zero. Take for instance the second term:

$$\frac{1}{T} \sum_{t=1+p}^T w_{2t} R'_t = \frac{1}{T} \sum_{t=1+p}^T w_{2t} \begin{pmatrix} 0' & \hat{\varepsilon}_{1t} - \varepsilon_{1t} \end{pmatrix} = \begin{pmatrix} 0' & \frac{1}{T} \sum_{t=1+p}^T w_{2t} (\hat{\varepsilon}_{1t} - \varepsilon_{1t}) \end{pmatrix},$$

We need to show that

$$\frac{1}{T} \sum_{t=1+p}^T w_{2t} (\hat{\varepsilon}_{1t} - \varepsilon_{1t}) = -\frac{1}{T} \sum_{t=1+p}^T w_{2t} w'_{1t} (\hat{\pi}_1 - \pi_1) = o_P(1).$$

This follows provided $\hat{\pi}_1 - \pi_1 = o_P(1)$ and $\frac{1}{T} \sum_{t=1}^T w_{2t} w'_{1t} \xrightarrow{P} \Sigma_{w_1 w_2}$. Similarly, we can write

$$\frac{1}{T} \sum_{t=1+p}^T R_t R'_t = \frac{1}{T} \sum_{t=1+p}^T \begin{pmatrix} 0 & \\ \hat{\varepsilon}_{1t} - \varepsilon_{1t} & \end{pmatrix} \begin{pmatrix} 0' & \hat{\varepsilon}_{1t} - \varepsilon_{1t} \end{pmatrix} = \begin{pmatrix} 0 & 0 \\ 0 & \frac{1}{T} \sum_{t=1+p}^T (\hat{\varepsilon}_{1t} - \varepsilon_{1t})^2 \end{pmatrix},$$

where

$$\frac{1}{T} \sum_{t=1+p}^T (\hat{\varepsilon}_{1t} - \varepsilon_{1t})^2 = (\hat{\pi}_1 - \pi_1)' \frac{1}{T} \sum_{t=1+p}^T w_{1t} w'_{1t} (\hat{\pi}_1 - \pi_1) = o_P(1).$$

Thus, $\left(\frac{1}{T} \sum_{t=1+p}^T \hat{w}_{2t} \hat{w}'_{2t}\right)^{-1} - \left(\frac{1}{T} \sum_{t=1+p}^T w_{2t} w'_{2t}\right)^{-1} = o_P(1)$. To conclude that $\zeta_1 = o_P(1)$, it suffices that $\frac{1}{T} \sum_{t=1+p}^T \hat{w}_{2t} y'_t = o_P(1)$. To see this result, write

$$\frac{1}{T} \sum_{t=1}^T \hat{w}_{2t} y'_t = \frac{1}{T} \sum_{t=1}^T w_{2t} y'_t + \frac{1}{T} \sum_{t=1}^T (\hat{w}_{2t} - w_{2t}) y'_t = \frac{1}{T} \sum_{t=1}^T w_{2t} y'_t + \frac{1}{T} \sum_{t=1}^T R_t y'_t,$$

and note that the first term on the RHS is $O_P(1)$ whereas the second term is $o_P(1)$. The proof that $\zeta_2 = o_P(1)$ follows from similar arguments. This completes the proof of the consistency of $\hat{\Pi}_2$ towards Π_2 . ■

Lemma A.2 Under Assumptions 1-6, $\hat{A}_{j,\delta} - A_{j,\delta} \xrightarrow{P} 0$ for $j = 0, 1, \dots, h$.

Proof. For each $j = 1, \dots, h$, we can write $\hat{A}_{j,\delta} - A_{j,\delta} = I_1 + I_2$, where

$$I_1 = \frac{1}{T-h} \sum_{t=1}^{T-h} \left(f \left(g_{j,\delta} \left(x_{t+j}, x_{t+j-1}, \dots, x_t; \hat{\beta}_j \right) \right) - Ef \left(g_{j,\delta} \left(x_{t+j}, x_{t+j-1}, \dots, x_t; \beta_j \right) \right) \right) \text{ and}$$

$$I_2 = \frac{1}{T} \sum_{t=1}^T (f(x_{t+j}) - Ef(x_{t+j})).$$

Under Assumptions 2 and 3, by a standard law of large numbers for stationary ergodic processes, $I_2 = o_P(1)$. To show that $I_1 = o_P(1)$, we first apply a uniform law of large numbers for stationary ergodic processes. Specifically, we verify conditions (i) through (iii) of Theorem 4.1 of Wooldridge

(1994, p. 2651). By Assumption 6, β_j is in a compact set, thus verifying his condition (i). Since $g_{j,\delta}(x_{t+j}, x_{t+j-1}, \dots, x_t; \beta_j)$ is a continuous functions of β_j , it suffices to assume that f is a continuous function to conclude that $f(g_{j,\delta}(x_{t+j}, x_{t+j-1}, \dots, x_t; \beta_j))$ is continuous in β_j , which verifies his condition (ii). Finally, the dominance condition (iii) holds under our Assumption 6. The desired result then follows from Lemma A.1 of Wooldridge (1994), given the consistency of $\hat{\pi}$ for π and the fact that the estimators $\hat{\Theta}_{\cdot 1}(L)$ and $\hat{\Gamma}(L)$ are continuous functions of $\hat{\pi}$. ■

Proof of Theorem 5.1. The proof is omitted as it follows standard arguments. ■

Proof of Proposition 5.1. We prove the result for $h \geq p$. The proof for $0 \leq h < p$ is similar, but requires some adjustments. In order to simplify the notation, we rewrite the second equation of the bivariate structural model with $x_t = \varepsilon_{1t}$ as

$$y_t = k_2 + \beta_0 x_t + \beta(L) x_{t-1} + \rho(L) y_{t-1} + c(L) f(x_t) + \varepsilon_{2t}.$$

This implies that

$$y_t = \mu + \theta(L) x_t + \gamma(L) f(x_t) + \psi(L) \varepsilon_{2t},$$

where

$$\begin{aligned} \psi(L) &= (1 - \rho(L)L)^{-1} = (1 + \psi_1 L + \psi_2 L^2 + \dots), \\ \theta(L) &= \psi(L)(\beta_0 + \beta(L)L) = \theta_0 + \theta_1 L + \theta_2 L^2 + \dots, \text{ and} \\ \gamma(L) &= \psi(L)c(L) = \gamma_0 + \gamma_1 L + \gamma_2 L^2 + \dots, \end{aligned}$$

which yields

$$y_{t+h} = \mu + \theta(L)x_{t+h} + \gamma(L)f(x_{t+h}) + \psi(L)\varepsilon_{2t+h}.$$

Note that with this notation, $\theta(L) = \Theta_{21}(L)$ and $\gamma(L) = \Gamma_{21}(L)$. Given the definition of $\theta(L)$, we have that

$$\theta(L)x_{t+h} = \psi(L)(\beta_0 + \beta_1 L + \dots + \beta_p L^p) x_{t+h} = \psi(L)\beta_0 x_{t+h} + \psi(L)\beta_1 x_{t+h-1} + \dots + \psi(L)\beta_p x_{t+h-p}.$$

We can further decompose each of the terms above using the definition of $\psi(L)$. This yields

$$\theta(L)x_{t+h} = \underbrace{(\beta_0 \psi_h + \beta_1 \psi_{h-1} + \dots + \beta_p \psi_{h-p})}_{=\theta_h} x_t + u_{1,t+h} + u_{2,t+h},$$

where $\theta_h \equiv \Theta_{h,21}$, and

$$\begin{aligned} u_{1,t+h} &\equiv \underbrace{\beta_0 \{1 + \psi_1 L + \dots + \psi_{h-1} L^{h-1}\}}_{=\{x_{t+h}, \dots, x_{t+1}\}} x_{t+h} + \underbrace{\beta_1 \{1 + \psi_1 L + \dots + \psi_{h-2} L^{h-2}\}}_{=\{x_{t+h-1}, \dots, x_{t+1}\}} x_{t+h-1} + \dots \\ &\quad + \underbrace{\beta_p \{1 + \psi_1 L + \dots + \psi_{h-p-1} L^{h-p-1}\}}_{=\{x_{t+h-p}, \dots, x_{t+1}\}} x_{t+h-p} \end{aligned}$$

is a (linear) function only of future values of x , $\{x_{t+1}, \dots, x_{t+h}\}$.⁹ Note that because $x_t = \varepsilon_{1t}$, which is assumed to be i.i.d., we have that $E(u_{1,t+h}) = 0$. In addition, $u_{1,t+h}$ is orthogonal to x_t , $f(x_t)$, their lags as well as lags of y_t . Instead, the term $u_{2,t+h}$ is defined as

$$\begin{aligned} u_{2,t+h} &= \beta_0 \left\{ \psi_{h+1} L^{h+1} + \psi_{h+2} L^{h+2} + \dots \right\} x_{t+h} + \beta_1 \left\{ \psi_h L^h + \psi_{h+1} L^{h+1} + \dots \right\} x_{t+h-1} + \dots \\ &\quad + \beta_p \left\{ \psi_{h-p+1} L^{h-p+1} + \dots \right\} x_{t+h-p} \end{aligned}$$

and is a function only of past values of x . We can rewrite $u_{2,t+h}$ as

$$\begin{aligned} u_{2,t+h} &= \psi_{h+1} \underbrace{\{\beta_0 + \beta_1 L + \dots + \beta_p L^p\}}_{=\beta_0 + \beta(L)} x_{t-1} + \underbrace{(\beta_1 \psi_h + \beta_2 \psi_{h-1} + \dots + \beta_p \psi_{h-p+1})}_{\equiv \pi_{x,1,h}} x_{t-1} \\ &\quad + \dots \\ &\quad + \psi_{h+p} \underbrace{\{\beta_0 + \beta_1 L + \dots + \beta_p L^p\}}_{=\beta_0 + \beta(L)} \underbrace{x_{t-p}}_{=L^{p-1} x_{t-1}} + \underbrace{(\beta_p \psi_h)}_{\equiv \pi_{x,p,h}} x_{t-p} \\ &\quad + \psi_{h+p+1} \{\beta_0 + \beta_1 L + \dots + \beta_p L^p\} x_{t-p-1} + \dots \\ &= \underbrace{\pi_{x,1,h} x_{t-1} + \dots + \pi_{x,p,h} x_{t-p}}_{=\pi_x(L) x_{t-1}} + (\psi_{h+1} + \psi_{h+2} L + \dots) (\beta_0 + \beta(L)) x_{t-1}. \end{aligned}$$

Assembling these results yields

$$\theta(L)x_{t+h} = \pi_{x,h}x_t + \pi_{x,h}(L)x_{t-1} + (\psi_{h+1} + \psi_{h+2}L + \dots) (\beta_0 + \beta(L)) x_{t-1} + u_{1,t+h},$$

where $\pi_{x,h} = \theta_h \equiv \Theta_{h,21}$ and

$$\pi_{x,h}(L) = \pi_{x,1,h} + \pi_{x,2,h}L + \dots + \pi_{x,p,h}L^{p-1}.$$

Proceeding the same way, we have that

$$\gamma(L)f(x_{t+h}) = \pi_{f,h}f(x_t) + \pi_{f,h}(L)f(x_{t-1}) + (\psi_{h+1} + \psi_{h+2}L + \dots) \gamma(L)f(x_{t-1}) + w_{1,t+h},$$

where $\pi_{f,h} = \gamma_h \equiv \Gamma_{h,21}$, $\pi_f(L) = \pi_{f,1,h} + \dots + \pi_{f,p,h}L^{p-1}$, and

$$\begin{aligned} w_{1,t+h} &\equiv \underbrace{\gamma_0 \left\{ 1 + \psi_1 L + \dots + \psi_{h-1} L^{h-1} \right\} f(x_{t+h})}_{=\{f(x_{t+h}), \dots, f(x_{t+1})\}} + \underbrace{\gamma_1 \left\{ 1 + \psi_1 L + \dots + \psi_{h-2} L^{h-2} \right\} f(x_{t+h-1})}_{=\{f(x_{t+h-1}), \dots, f(x_{t+1})\}} \\ &\quad + \dots + \underbrace{\gamma_p \left\{ 1 + \psi_1 L + \dots + \psi_{h-p-1} L^{h-p-1} \right\} f(x_{t+h-p})}_{=\{f(x_{t+h-p}), \dots, f(x_{t+1})\}}, \end{aligned}$$

and it can be shown that

$$w_{2,t+h} = \pi_{f,1,h}f(x_{t-1}) + \dots + \pi_{f,p,h}f(x_{t-p}) + (\psi_{h+1} + \psi_{h+2}L + \dots) \gamma(L)f(x_{t-1}).$$

⁹Henceforth, and in order to simplify the notation, we will use $\{z_{t+1}, \dots, z_{t+h}\}$ to denote a linear combination of the variables inside the curly brackets.

Because $E(f(x_t)) \neq 0$, for any t (due to the nonlinearity of f), $E(w_{1h}) \neq 0$, so we need to add and subtract $\pi_{c,h} \equiv E(w_{1h})$ from the equation that defines y_{t+h} . Hence, we can write

$$y_{t+h} = \pi_{c,h} + \theta_h x_t + \gamma_h f(x_t) + \pi_{x,h}(L)x_{t-1} + \pi_{f,h}(L)f(x_{t-1}) + a_{t+h} + v_{t+h},$$

where

$$a_{t+h} = (\psi_{h+1} + \psi_{h+2}L + \dots) \underbrace{((\beta_0 + \beta(L))x_{t-1} + \gamma(L)f(x_{t-1}) + \varepsilon_{2t-1})}_{=\rho(L)y_{t-1}} = (\psi_{h+1} + \psi_{h+2}L + \dots) \rho(L) y_{t-1}$$

and

$$v_{t+h} = \psi_h \varepsilon_{2t} + u_{1,t+h} + (w_{1,t+h} - E(w_{1,t+h})) + \{\varepsilon_{1t+h}, \dots, \varepsilon_{2t+1}\}.$$

It can be shown that v_{t+h} has mean zero and is orthogonal to x_t , $f(x_t)$ and their lags using the properties of the structural errors ε_{1t} and ε_{2t} . To complete the proof, we show that $a_{t+h} \equiv (\psi_{h+1} + \psi_{h+2}L + \dots) \rho(L) y_{t-1}$ can be written as a linear function of y_{t-1}, \dots, y_{t-p} . This follows from proving that $(\psi_{h+1} + \psi_{h+2}L + \dots) \rho(L)$ is a lag polynomial of order $p-1$ using the definition of ψ_h . To show this, recall that for any $h \geq p$, $\psi_h = \psi_{h-1}\rho_1 + \psi_{h-2}\rho_2 + \dots + \psi_{h-p}\rho_p$. Thus,

$$\begin{aligned} & (\psi_{h+1} + \psi_{h+2}L + \dots) (1 - \rho_1 L - \dots - \rho_p L^p) \\ = & \psi_{h+1} - \rho_1 \psi_{h+1} L - \dots - \rho_p \psi_{h+1} L^p + \psi_{h+2} L - \rho_1 \psi_{h+2} L^2 - \dots - \rho_p \psi_{h+2} L^{p+1} \\ & + \dots + \psi_{h+p} L^{p-1} - \rho_1 \psi_{h+p} L^p - \dots - \rho_p \psi_{h+p} L^{2p-1} + \dots \\ = & \underbrace{\psi_{h+1}}_{\equiv \pi_{y,1,h}} + \underbrace{(\psi_{h+2} - \rho_1 \psi_{h+1}) L}_{\equiv \pi_{y,2,h} \neq 0} + \underbrace{(\psi_{h+3} - \rho_1 \psi_{h+2} - \rho_2 \psi_{h+1}) L^2}_{\equiv \pi_{y,3,h} \neq 0} \\ & + \dots + \underbrace{(\psi_{h+p} - \rho_1 \psi_{h+p-1} - \dots - \rho_{p-1} \psi_{h+1}) L^{p-1}}_{=\rho_p \psi_h \equiv \pi_{y,p,h} \neq 0} + \underbrace{(\psi_{h+p+1} - \rho_1 \psi_{h+p} - \dots - \rho_p \psi_{h+1}) L^p}_{=0} + \dots \\ = & \pi_{y,1,h} + \pi_{y,2,h} L + \dots + \pi_{y,p,h} L^{p-1}. \end{aligned}$$

This concludes the proof. ■

B Appendix: Conditional impulse response functions

For illustrative purposes, consider the simplified model

$$\begin{cases} x_t = b_1 + \phi x_{t-1} + \varepsilon_{1t}, \\ y_t = b_2 + \beta x_t + \rho y_{t-1} + cf(x_t) + \varepsilon_{2t}, \end{cases}$$

and the conditional IRF

$$IRF_{h,\delta}(\Omega_{t-1}) = E(y_{t+h}(\delta) - y_{t+h} | \Omega_{t-1}),$$

where $h = 0, 1, 2, \dots, H$ and $\Omega_{t-1} = \{x_{t-1}, y_{t-1}, \dots\}$. For this model,

$$y_{t+h}(\delta) - y_{t+h} = \Theta_{h,21}\delta + \Gamma_{0,21} \left[f(x_{t+h} + \phi^h \delta) - f(x_{t+h}) \right] + \dots + \Gamma_{h,21} [f(x_t + \delta) - f(x_t)],$$

where $\Theta_{h,21} = \beta(\phi^h + \rho\phi^{h-1} + \dots + \rho^h)$, $\Gamma_{h,21} = c\rho^h$. To obtain the conditional IRF, we need to evaluate the expected value of $f(x_{t+j} + \phi^j \delta) - f(x_{t+j})$, conditionally on Ω_{t-1} . We obtain

$$IRF_{h,\delta}(\Omega_{t-1}) = \Theta_{h,21}\delta + \Gamma_{0,21}C_{h,\delta} + \dots + \Gamma_{h,21}C_{h,\delta},$$

where now

$$C_{j,\delta} \equiv E[f(x_{t+j} + \phi^j \delta) - f(x_{t+j}) | \Omega_{t-1}] \quad \text{for } j \geq 0.$$

We can characterize analytically the nonlinear terms $C_{j,\delta}$, if we are willing to impose further assumptions on $f(x_t)$ and on the conditional distribution of x_{t+j} given Ω_{t-1} . Consider for instance the case where $f(x_t) = \max(x_t, 0)$ and suppose we care about $h = 0$. Then

$$C_{0,\delta} = E[f(x_t + \delta) - f(x_t) | \Omega_{t-1}] = E[f(x_t + \delta) - f(x_t) | x_{t-1}],$$

where $\Omega_{t-1} = \{x_{t-1}\}$ since $x_t = b_1 + \phi x_{t-1} + \varepsilon_{1t}$. If we assume in addition that $\varepsilon_{1t} | x_{t-1} \sim N(0, \sigma^2)$, then standard derivations imply that

$$\begin{aligned} C_{0,\delta} &= \left[\Phi\left(\frac{b_1}{\sigma} + \frac{\phi}{\sigma}x_{t-1} + \frac{\delta}{\sigma}\right) - \Phi\left(\frac{b_1}{\sigma} + \frac{\phi}{\sigma}x_{t-1}\right) \right] (b_1 + \phi x_{t-1}) + \delta \Phi\left(\frac{b_1}{\sigma} + \frac{\phi}{\sigma}x_{t-1} + \frac{\delta}{\sigma}\right) \\ &\quad + \sigma \left[\phi_f\left(\frac{b_1}{\sigma} + \frac{\phi}{\sigma}x_{t-1} + \frac{\delta}{\sigma}\right) - \phi_f\left(\frac{b_1}{\sigma} + \frac{\phi}{\sigma}x_{t-1}\right) \right], \end{aligned}$$

where ϕ_f and Φ denote the pdf and cdf of a standard normal distribution.¹⁰ This example illustrates that the conditional IRF at horizon $h = 0$ is a complicated nonlinear function of x_{t-1} which depends on the distribution of ε_{1t} .¹¹

¹⁰These derivations are similar to those used to analyze the Tobit regression model.

¹¹A similar approach in a different context has been taken by Mavroeidis (2020) who considers a model in which y_t depends on x_t and $f(x_t) = \max(x_t, 0)$, where x_t is latent. Mavroeidis derives the conditional response of y_{t+h} to x_t under the maintained assumption of Gaussian errors.

Although the derivations above could be extended to horizons $h > 0$, it is clear that the validity of a plug-in estimate of the conditional IRFs based on the resulting analytical expressions would depend on strong parametric assumptions on ε_{1t} and on the functional form of $f(x_t)$. If the model is fully recursive, our recommendation would be to rely on the Monte Carlo integration approach for constructing the conditional responses, which does not require restricting the distribution of ε_{1t} . This approach is not possible under the weaker assumption of a block recursive model, because the MCI approach requires an estimate of the full structural model.

Figure 1: The Accuracy of Alternative Impulse Response Estimators, $T = 240$

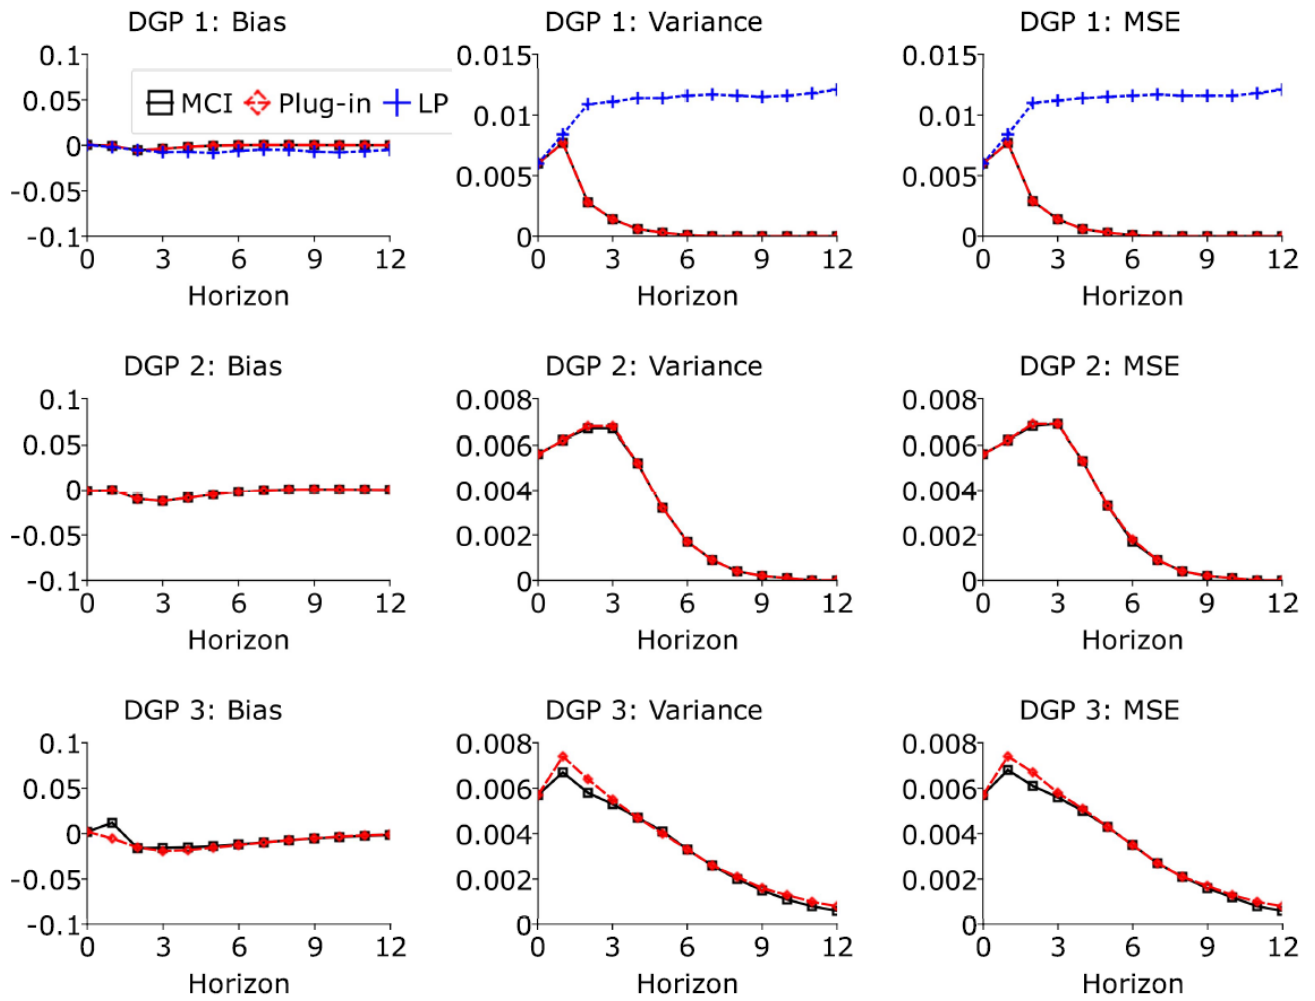


Figure 2: Accuracy of Alternative Impulse Response Estimators Under Dynamic Misspecification, $T = 240$

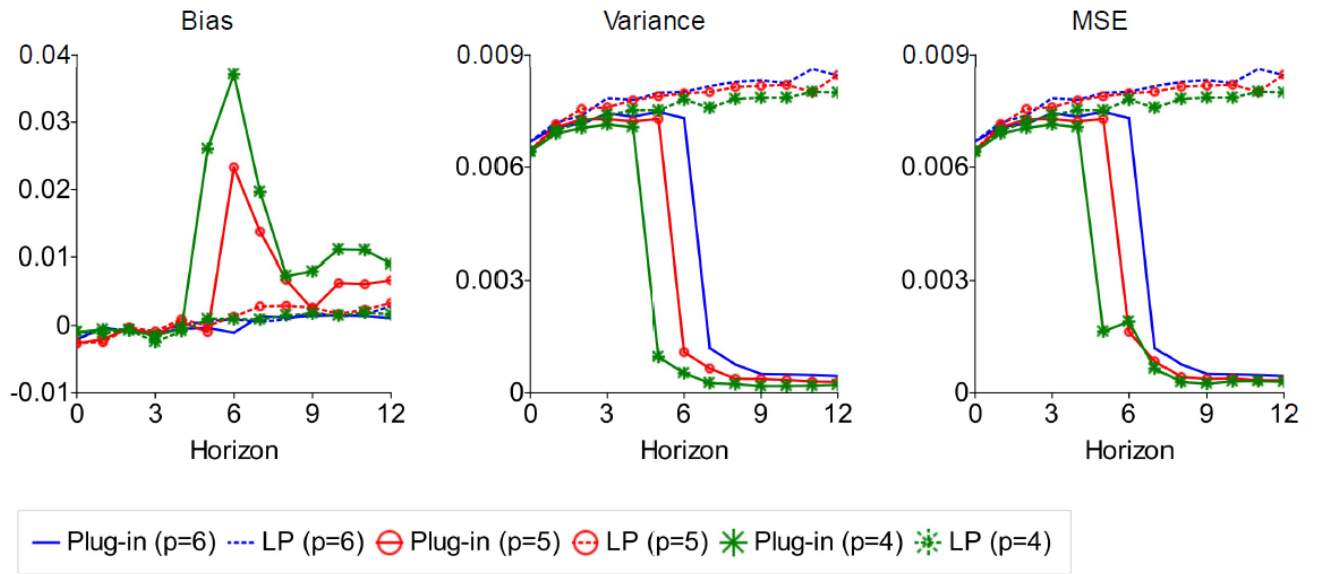


Figure 3: MSE Convergence of the Plug-in Estimator, $f(x_t) = \max(0, x_t)$

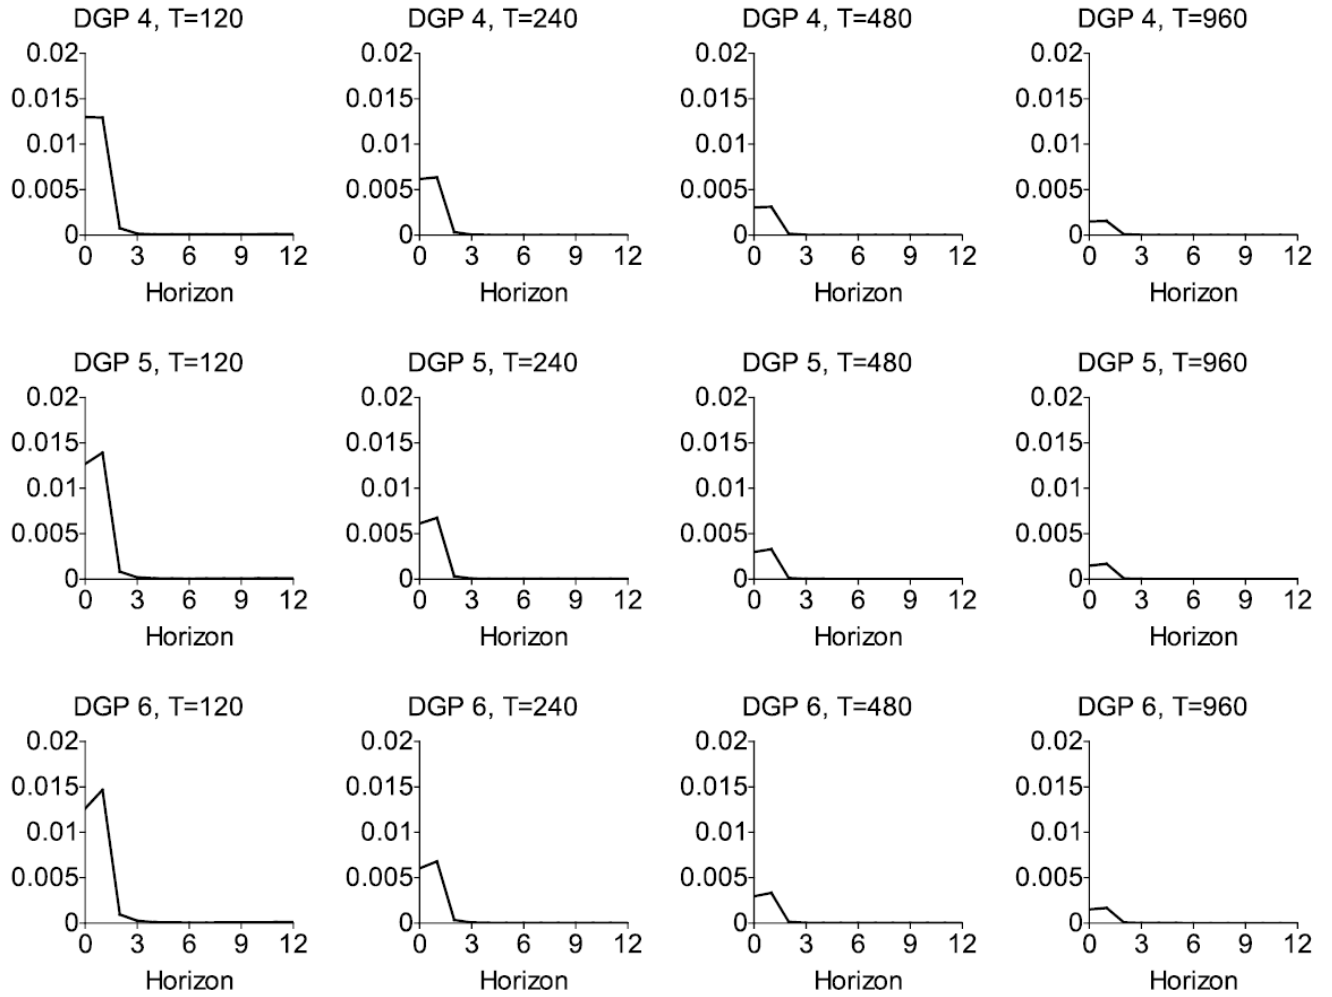


Figure 4: MSE Convergence of the Plug-in Estimator, $f(x_t) = x_t^3$

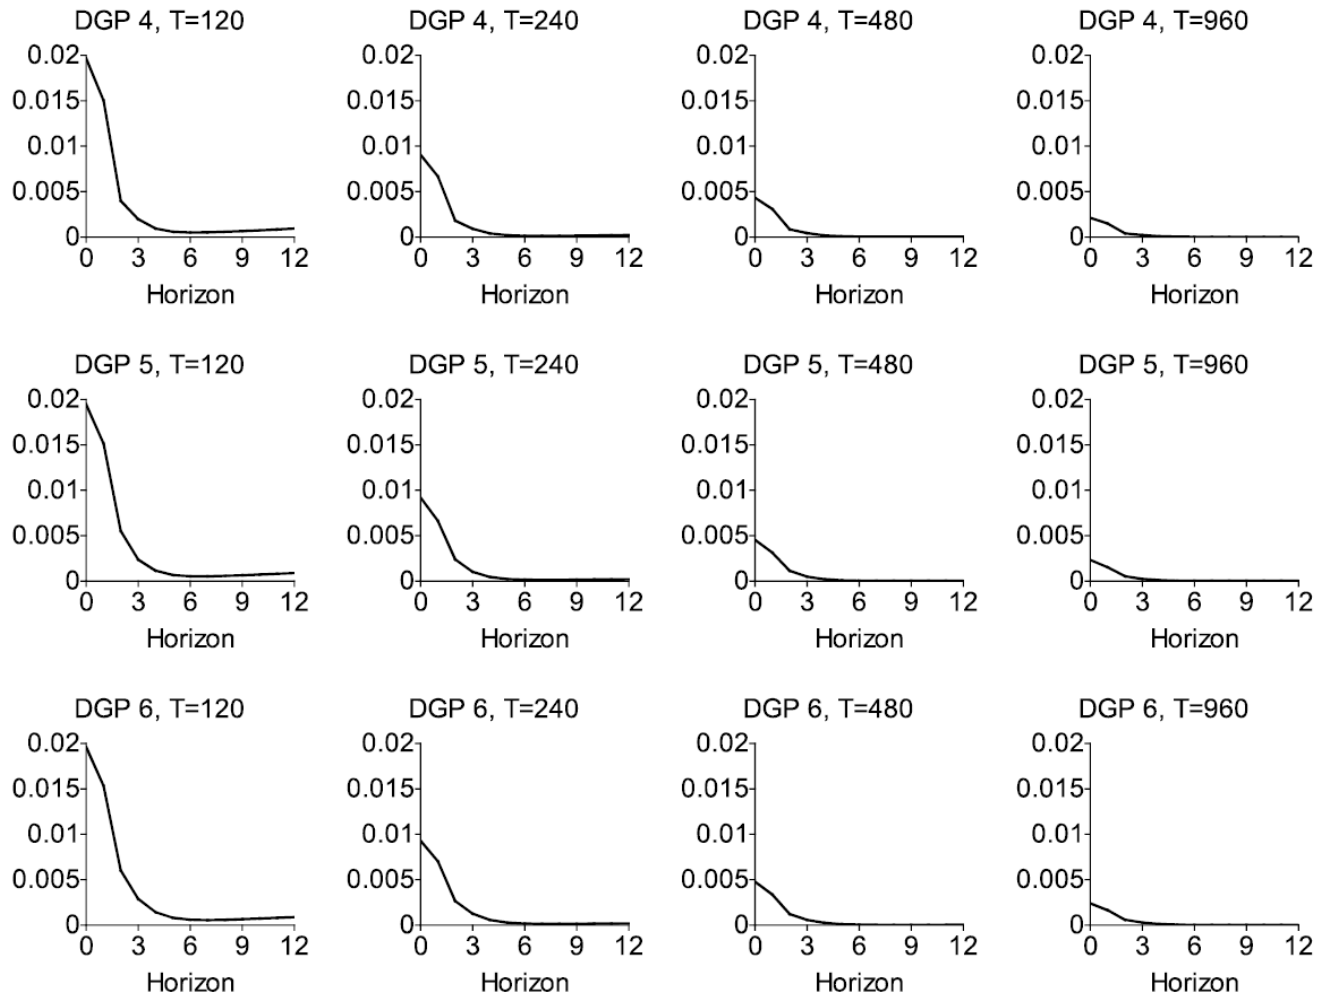


Figure 5: The Effect of an Unexpected U.S. Monetary Tightening When $f(x_t) = \max(0, x_t)$ (Upper Panel) and $f(x_t) = x_t^3$ (Lower Panel)

