

Impulse response analysis for structural dynamic models with nonlinear regressors*

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

June 18, 2020

Abstract

We study the construction of nonlinear impulse responses in structural dynamic models that include nonlinearly transformed regressors. 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. The conventional approach to estimating nonlinear responses is by Monte Carlo integration. We show that the population impulse responses in this class of models may instead be derived analytically from the structural model. We use this insight to study under what conditions linear projection (LP) estimators may be used to recover the population impulse responses. We find that, unlike in vector autoregressive models, the asymptotic equivalence between estimators based on the structural model and LP estimators breaks down. Only in one important special case can the population impulse response be consistently estimated by LP methods. The construction of this LP estimator, however, differs from the LP approach currently used in the literature. Simulation evidence suggests that the modified LP estimator is less accurate in finite samples than estimators based on the structural model, when both are valid.

JEL codes : C22, C32, C51

Key Words: local projection, structural model, censored regressor, nonlinear transformation, nonlinear responses, Monte Carlo integration

*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 Atsushi Inoue and Mikkel Plagborg-Møller for helpful discussions.

[†]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

There has been much interest in the relationship between linear vector autoregressive (VAR) models and linear local projection (LP) models. Linear LP estimators of impulse responses were first proposed by Jordà (2005, 2009) and popularized by Ramey (2016). Kilian and Lütkepohl (2017) showed that the impulse responses implied by a linear VAR process can, in principle, be recovered by a linear local projection. Plagborg-Møller and Wolf (2019) proved that both linear VAR models and LP regression models may be viewed as a linear approximation to an unknown possibly nonlinear process. The latter result, however, allows this approximation to be arbitrarily poor. In this paper, we therefore examine the ability of LP regressions to recover nonlinear 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 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 wide range of censored regressors that have been employed in the empirical literature including increases, net increases, net changes and large increases in model variables as well as powers of regressors. We follow a large existing literature in postulating that $(y_t \ x_t)'$ is a bivariate structural dynamic process such that, in general, x_t may linearly depend on its own lags or 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 (e.g., Kilian and Lütkepohl 2017). Thus, x_t is assumed to be predetermined with respect to y_t . We are interested in estimating the response 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.

It is well known that this nonlinear impulse response function may be estimated by Monte Carlo integration based on estimates of the structural model and based on the data, but this approach tends to be computationally demanding. Our paper makes four main contributions. First, we derive exact analytical solutions for the population impulse responses in this class of structural models. This result not only allows the evaluation of the finite-sample accuracy of alternative impulse response estimators by simulation, but it enables us to derive a novel plug-in estimator of the population impulse responses that dispenses with the need for simulations. The availability of fast analytic solutions is an important advantage when bootstrapping the distribution of the impulse response estimator or when constructing critical values under the null of symmetric response functions.

¹Examples include Hamilton (2011), Kilian and Vigfusson (2011a,b, 2017), Venditti (2013), Herrera, Lagalo and Wada (2011, 2015), Herrera and Karaki (2015), Alsalman and Herrera (2015), Hussain and Malik (2016), Tenreyro and Thwaites (2016), Hwa, Kapinos and Ramirez (2018), and Barnichon, Matthes and Ziegenbein (2019).

Second, there has been increasing interest in applied work in recent years in estimating the response function in this class of models by local projections. There has been no theoretical work justifying the use of LP estimators for models involving nonlinearly transformed regressors, however. We show that LP estimators may be used to recover the population impulse responses in this class of models, provided that x_t is an observable exogenous i.i.d. shock. This case has become the empirically most important setting in macroeconomics in recent years. More generally, when x_t is exogenous, but serially correlated, or when x_t is endogenous, LP estimators are not valid. Thus, the equivalence between impulse response estimators based on the structural model and LP estimators breaks down in this class of models.

Third, when the LP approach may be used, the construction of the LP estimator differs from the LP estimators currently used in the literature, which effectively ignore the nonlinearity of the impulse responses (e.g., Tenreyro and Thwaites 2016, Barnichon, Matthes and Ziegenbein 2019). This result has important implications for applied work on the asymmetry of responses to positive and negative shocks of the same magnitude in a wide range of macroeconomic contexts.

Fourth, we prove that both the plug-in estimator and the modified LP estimator proposed in this paper are consistent when x_t is an observable i.i.d. shock. Monte Carlo simulations, however, suggest that in finite samples estimators based on the structural model tend to be systematically more accurate than the LP estimator when both are valid. Regardless of the model specification, there is little to choose between the plug-in estimator and the Monte Carlo integration estimator based on their finite-sample accuracy, but the analytical approach is 200-300 times faster.

The remainder of the paper is organized as follows. In section 2, we introduce matrix notation for our model and discuss the definition of 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. Section 3 discusses how to recover impulse response functions when x_t is restricted to be an observable exogenous i.i.d. shock. We show that these responses may be identified by local projections and in closed-form based on the structural model. Section 4 focuses on the general model, including the special case when x_t is an exogenous autoregressive process. We analytically derive the impulse response function and demonstrate the inability of local projections to recover the impulse responses in this general model. Section 5 introduces the plug-in estimator and a modified LP estimator and proves the consistency of these estimators for the important special case when x_t is an observable exogenous i.i.d. shock. Section 6 summarizes what we have learned. In section 7, we use simulations to assess the relative accuracy of the modified LP estimator and of the simulation-based estimator and the plug-in estimator based on the structural model. We also provide

evidence on the computational advantages of the plug-in estimator. The concluding remarks are in section 8. Details of the proofs and a review of the conventional impulse response estimator based on Monte Carlo integration can be found in the appendix.

2 Nonlinear impulse response functions in models with nonlinearly transformed regressors

2.1 Structural dynamic models with nonlinearly transformed regressors

For expository purposes, let $(x_t, y_t)'$ denote a bivariate stationary time series.² For expositional purposes, we assume that all variables have been suitably demeaned. A widely used structural data generating process that allows for the inclusion of nonlinear regressors is

$$\begin{cases} \phi(L)x_t = \alpha(L)y_t + \varepsilon_{1t} \\ \rho(L)y_t = \beta(L)x_t + \gamma(L)f(x_t) + \varepsilon_{2t}, \end{cases} \quad (1)$$

where

$$\phi(L) = 1 - \phi_1L - \phi_2L^2 - \dots - \phi_pL^p, \quad \rho(L) = 1 - \rho_1L - \dots - \rho_pL^p,$$

and

$$\begin{aligned} \alpha(L) &= \alpha_0 + \alpha_1L + \dots + \alpha_pL^p, \\ \beta(L) &= \beta_0 + \beta_1L + \dots + \beta_pL^p, \quad \text{and} \quad \gamma(L) = \gamma_0 + \gamma_1L + \dots + \gamma_pL^p. \end{aligned}$$

We assume that all lag polynomials have the same order p . This is without loss of generality because we can set $p = \max(p_\phi, p_\alpha, p_\beta, p_\rho, p_\gamma)$ and zero out the elements that exceed the true lag order. For instance, if p_ϕ is the number of lags in $\phi(L)$, set $\phi_i = 0$ for all $i > p_\phi$. For convenience, we collect all parameters in the vector $\boldsymbol{\theta} = (\phi', \alpha', \rho', \beta', \gamma)'$.³ It should be noted that it is important that $f(\cdot)$ is a function of x_t rather than ε_{1t} . This assumption is dictated by the economic reasoning underlying this class of models in the literature.

Model (1) can be written in matrix form as

$$B_0 z_t = B(L) z_{t-1} + C(L) f(x_t) + \varepsilon_t, \quad (2)$$

where $z_t \equiv (x_t, y_t)'$, $\varepsilon_t = (\varepsilon_{1t}, \varepsilon_{2t})'$, and

$$B_0 = \begin{pmatrix} 1 & -\alpha_0 \\ -\beta_0 & 1 \end{pmatrix}, \quad B(L) = \begin{pmatrix} \bar{\phi}(L) & \bar{\alpha}(L) \\ \bar{\beta}(L) & \bar{\rho}(L) \end{pmatrix} \quad \text{and} \quad C(L) = \begin{pmatrix} 0 \\ \gamma(L) \end{pmatrix},$$

²Although there is no reason why one could not consider higher-dimensional processes, the empirical literature has tended to focus on bivariate processes.

³This specification avoids the pitfalls that arise when specifying a VAR model for $(f(x_t), y_t)'$ and treating this model as linear. Kilian and Vigfusson (2011a) prove that responses constructed from such VAR models are invalid.

with

$$\begin{aligned}\bar{\phi}(L) &= \phi_1 + \phi_2 L + \dots + \phi_p L^{p-1}, & \bar{\alpha}(L) &= \alpha_1 + \alpha_2 L + \dots + \alpha_p L^{p-1}, \\ \bar{\beta}(L) &= \beta_1 + \beta_2 L + \dots + \beta_p L^{p-1}, & \bar{\rho}(L) &= \rho_1 + \rho_2 L + \dots + \rho_p L^{p-1}.\end{aligned}$$

When there are no nonlinearities, $\gamma(L) = 0$ and hence $C(L) = 0$. In this case, (2) is a recursive structural linear VAR model for z_t . When $C(L) \neq 0$, y_t depends on $f(x_t)$, a measurable nonlinear function of x_t , and lags of $f(x_t)$. We will consider several economically interesting examples of $f(\cdot)$ at the end of this section.

In order to derive our results, we make the following assumptions.

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

$$\varepsilon_t \equiv \begin{pmatrix} \varepsilon_{1t} \\ \varepsilon_{2t} \end{pmatrix} \sim \text{i.i.d.} \left(\begin{pmatrix} 0 \\ 0 \end{pmatrix}, \begin{pmatrix} \sigma_1^2 & 0 \\ 0 & \sigma_2^2 \end{pmatrix} \right).$$

Assumption P. The process $z_t = (x_t, y_t)'$ is generated by (1) (or (2)) with $\alpha_0 = 0$.

Assumption S. The roots of the equation $\det(D(L)) = 0$ are outside the unit circle, where $D(L) \equiv I_2 - B_0^{-1}B(L)L$, and the process z_t is stationary and ergodic.

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

Under Assumption E, the structural errors ε_{1t} and ε_{2t} are i.i.d. over time and follow mutually independent processes. These assumptions are stronger than the white noise assumption, which is not enough to justify the validity of our impulse response estimators in the presence of $f(x_t)$. Assumption P postulates that $\alpha_0 = 0$, implying that the matrix B_0 (which describes the instantaneous relations between x_t and y_t) is lower triangular. This assumption allows us to identify the structural parameter vector θ . It means that x_t is predetermined with respect to y_t . This assumption is reasonable in typical applications (e.g., Romer and Romer 2004, 2010; Kilian and Vega 2011).

Assumption S contains two parts. First, we assume that the roots of the determinantal equation associated with the matrix polynomial $D(L) = I_2 - B_0^{-1}B(L)L$ are all outside the unit circle. This standard 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 $\gamma(L) = 0$, it is not enough when $\gamma(L) \neq 0$. For this reason, we assume stationarity and ergodicity of z_t as a high level condition. When both $\gamma(L) \neq 0$ and $\alpha(L) \neq 0$, providing more primitive conditions under which this high-level condition is verified is challenging. This question is outside the scope of this paper. When $\alpha(L) = 0$, the following assumption (together with Assumption E) suffices for Assumption S.

Assumption S'

- (i) $\alpha(L) = 0$ and the roots of $\phi(L) = 0$ and $\rho(L) = 0$ are all outside the unit circle.
- (ii) The function $f(\cdot)$ is nonperiodic and bounded on compact sets such that $f(x) = O(|x|^{\gamma_1})$ as $x \rightarrow \infty$ for some $\gamma_1 \in \mathbb{R}$.
- (iii) $E\left(|\varepsilon_{1t}|^{\max(1, \gamma_1 + \gamma_2)}\right) < \infty$ for some $\gamma_2 > 0$.

When $\alpha(L) = 0$, z_t given by (2) is a special case of the nonlinear bivariate ARX model studied by Masry and Tjostheim (1997), where the only additive nonlinear term is itself an additive function of current and lagged values of $f(x_t)$. We can show that Assumptions E and S' suffice for their Assumption 3.3, implying that z_t is stationary and strongly mixing with exponentially decaying coefficients (cf. their Lemma 3.1). The condition on $f(\cdot)$ allows for functions that grow faster than a linear function, as $x \rightarrow \infty$, and is satisfied for the empirically motivated examples of $f(\cdot)$ we discuss below. Finally, Assumption M postulates that x_t and $f(x_t)$ have finite first-order moments. This assumption together with our remaining assumptions also implies that $\sup_t E|y_t| < \infty$. The moment condition on x_t follows from the moment condition on ε_t when $\gamma(L) = 0$, but not necessarily when $\gamma(L) \neq 0$ unless $\alpha(L) = 0$.

One special case of model (2) arises when $\bar{\phi}(L) = \bar{\alpha}(L) = 0$, or, equivalently, when in model (1) $\phi(L) = 1$ and $\alpha(L) = 0$, implying that $x_t = \varepsilon_{1t}$ is a serially uncorrelated directly observable exogenous shock, 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., Kilian 2008, Hamilton 2011, Bastianini and Manera 2018), news shocks (e.g., Ramey 2011; Kilian and Vega 2011; Kilian and Hicks 2013), financial market shocks (e.g., Barnichon, Matthes and Ziegenbein 2019) and shocks to market expectations (e.g., Kuttner 2001; Cochrane and Piazzesi 2002; Piazzesi and Swanson 2008; Baumeister and Kilian 2016). Another special case emerges when $\bar{\alpha}(L) = 0$ but $\bar{\phi}(L) \neq 0$, implying that x_t is a strictly exogenous autoregressive process. In the latter case, x_t is independent of ε_{2t} but is allowed to be serially correlated. Thus, (2) can be viewed as a bivariate ARX model for $(x_t, y_t)'$, where y_t is augmented with exogenous regressors given by $\beta(L)x_t = \beta_0 x_t + \bar{\beta}(L)x_{t-1}$ and $\gamma(L)f(x_t)$. We will discuss these two special cases in more detail below.

Note that under our assumptions, x_t (and hence $f(x_t)$) does not depend on the current value of ε_{2t} , given that it is only a function of ε_{1t} and its lags as well as lags of ε_{2t} when $\alpha(L) \neq 0$. Since we assume that ε_{1t} and ε_{2t} are mutually independent i.i.d. processes, $\beta(L)x_t + \gamma(L)f(x_t)$ is independent

of ε_{2t} in model (2). Hence, we can estimate the structural parameters in (2) by equation-by-equation OLS. The estimation of the model is discussed in more detail in Sections 5.1 and 5.2.

We are interested in estimating the responses of y_{t+h} at horizon $h = 0, 1, \dots, H$ to a shock of magnitude δ in ε_{1t} . Note that even though model (2) is linear in the parameters, it is nonlinear in the variables. Hence, the impulse responses are inherently nonlinear.

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 the model 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.

More specifically, we trace out the effect of a shock in ε_{1t} at time t on the future value of the outcome variable 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 (e.g., Gouriéroux and Jasiak 2005).

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

Definition 1 (Unconditional IRF) *The unconditional nonlinear impulse response function (IRF) 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 $h = 0, 1, 2, \dots$.*

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 whole 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 $\{y_{t+h} : h = 0, 1, \dots, H\}$. In doing so, we are effectively using the expected value of the difference between these two sample paths as a summary measure of their whole joint distribution (see Definition 4 in Gouriéroux and Jasiak (2005) who proposed this approach in the context of nonlinear univariate reduced-form models). In this paper, we are interested in the expected value of the difference between the two sample paths of y_{t+h} , but one could consider other functionals of their joint distribution.

Second, unlike the studies by Gallant, Rossi and Tauchen (1993), Koop, Pesaran and Potter (1996), Potter (2000), Gouriéroux and Jasiak (2005) and Kilian and Vigfusson (2017), we do not condition on the history of the process up to time $t - 1$, denoted Ω_{t-1} . We note, however, that it is also possible to define a version of our *IRF* that conditions on Ω_{t-1} .

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 $h = 0, 1, 2, \dots$*

This definition is routinely employed in estimating the unconditional response function by Monte Carlo integration. Details of this algorithm can be found in Appendix 1.

Third, a further difference with respect to some of the previous literature (e.g. Koop et al. (1996), Potter (2000) and Kilian and Vigfusson (2011a)) is the type of shocks that are presumed to occur between t and $t + h$. The aforementioned papers fix $\varepsilon_{1t}(\delta)$ equal to δ in the perturbed model and let ε_{1t} denote a random draw of the shock at t in the baseline model. Instead, 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 model (1) with $\phi(L) = 1$ and $\alpha(L) = 0$), it may potentially 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, $(\varepsilon_{1t}, \varepsilon_{2t})'$ 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

Next we discuss several economically interesting examples of nonlinear transformations $f(\cdot)$ of x_t (and possibly additional lags of x_t).

Example 1: Increases

A leading 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 2019), 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).

Example 2: Net increases

Concerns about the ability of models including oil price increases and decreases to match the data prompted Hamilton (2003) to argue for an alternative specification in which oil price increases matter only to the extent that they exceed their maximum value over the preceding three years, denoted by $x_t^* \equiv \max(x_{t-1}, \dots, x_{t-36})$. The resulting net oil price increase measure $x_t^{net,+} \equiv \max(0, x_t - x_t^*)$ is another example of a censored regressor. This specification, which has been motivated on behavioral grounds, is widely used in empirical work (e.g., Kilian and Vigfusson 2011a,b; Hamilton 2011; Herrera, Lagalo and Wada 2011, 2015).

Example 3: Net changes

In related work, Kilian and Vigfusson (2013) allow real GDP growth to depend on net changes in the price of oil, defined as the sum of net increases and net decreases in the price of oil, consistent with the behavioral arguments advanced in Hamilton (2011). Their net change measure is defined as $x_t^{net} \equiv x_t^{net,-} + x_t^{net,+}$, where $x_t^{net,-} \equiv \min(0, x_t - x_t^{**})$ with x_t^{**} defined as the minimum over the last 36 months.

Example 4: Large increases

Yet another hypothesis is that economic agents respond more to large changes in the model variables than to small changes. For example, Goldberg (1988) conjectures that the presence of costs to monitoring energy costs and of adjusting consumption patterns might make households reluctant to respond to small changes in energy prices. A gasoline price shock of typical magnitude might not evoke any reaction from consumers at all. In fact, it might go unnoticed by many households. Davis and Kilian (2011) make an analogous argument with respect to gasoline tax changes. Similarly, Alsalman and Herrera (2015) explore threshold effects in the response of stock returns to oil price shocks. Such threshold dynamics may be modelled, for example, by replacing $f(\cdot)$ in the process (1) by $x_t^{large} = x_t 1(|x_t| > \delta)$, where commonly $\delta = n\sigma$ and n denotes a multiple of the standard deviation σ of x_t , typically set to 1 or 2. Similar model structures may also arise in other economic contexts when modeling trigger points or thresholds.

Example 5: Large changes

Another empirically relevant specification involves higher 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 for several versions of our model. This will allow us to obtain an analytical estimator of the IRF that does not require Monte Carlo integration. We then investigate the ability of local projections to estimate these impulse responses.

3 The population IRF: x_t is an i.i.d. observed shock

The goal of this section is to derive a closed-form expression for $IRF_{h,\delta}$ for a special case of model (1), or equivalently, model (2). In particular, we assume not only that x_t is strictly exogenous (so that $\alpha(L) = 0$), but that it is serially uncorrelated (so, $\phi(L) = 1$). Hence, $x_t = \varepsilon_{1t}$.

There are three reasons for starting with this simpler model. First, this model corresponds to the case when x_t is an observed sequence of i.i.d. shocks, as is often the case in applied work. Second, the derivation of the population IRF is much simpler in this case, which helps build intuition for the general results that follow in Section 4. Third, we demonstrate that in this model one can, alternatively,

recover the population IRF with a modified local projection (LP) approach. This result only holds for the special case of $x_t = \varepsilon_{1t}$.

3.1 A closed-form expression for the IRF

Let

$$\begin{cases} x_t = \varepsilon_{1t}, \\ \rho(L) y_t = \beta(L) x_t + \gamma(L) f(x_t) + \varepsilon_{2t}, \end{cases} \quad (3)$$

Define the following inverse and product filters,

$$\begin{aligned} \psi(L) &= (1 - \rho_1 L - \rho_2 L^2 - \dots - \rho_p L^p)^{-1} = (1 + \psi_1 L + \psi_2 L^2 + \dots), \\ \psi_\beta(L) &= \psi(L) \beta(L) = \psi_{\beta,0} + \psi_{\beta,1} L + \psi_{\beta,2} L^2 + \dots, \text{ and} \\ \psi_\gamma(L) &= \psi(L) \gamma(L) = \psi_{\gamma,0} + \psi_{\gamma,1} L + \psi_{\gamma,2} L^2 + \dots \end{aligned}$$

Note that we can express the coefficients of these filters as a function of the structural parameters $\theta = (\rho', \beta', \gamma)'$. Under the assumption that the roots of $\rho(L)$ are outside the unit circle, these inverse filters have absolutely summable coefficients and, the process $(x_t, y_t)'$ is stationary under Assumptions S' and E.

We can write

$$y_t = \psi_\beta(L) x_t + \psi_\gamma(L) f(x_t) + \psi(L) \varepsilon_{2t}, \quad (4)$$

where $x_t = \varepsilon_{1t}$.

In order to obtain the population IRF for y_{t+h} , we rely on Definition 1 and on equation (4) to evaluate the difference $y_{t+h}(\delta) - y_{t+h}$. Specifically, using the definitions of $\psi_\beta(L)$ and $\psi_\gamma(L)$ and the two sequences of shocks ε^∞ and $\varepsilon^\infty(\delta)$, we have that

$$y_{t+h} = \psi_{\beta,0} \varepsilon_{1t+h} + \dots + \psi_{\beta,h} \varepsilon_{1t} + \dots + \psi_{\gamma,0} f(\varepsilon_{1t+h}) + \dots + \psi_{\gamma,h} f(\varepsilon_{1t}) + \dots + \psi(L) \varepsilon_{2t+h}$$

and

$$y_{t+h}(\delta) = \psi_{\beta,0} \varepsilon_{1t+h} + \dots + \psi_{\beta,h} [\varepsilon_{1t} + \delta] + \dots + \psi_{\gamma,0} f(\varepsilon_{1t+h}) + \dots + \psi_{\gamma,h} f(\varepsilon_{1t} + \delta) + \dots + \psi(L) \varepsilon_{2t+h}.$$

Hence,

$$y_{t+h}(\delta) - y_{t+h} = \psi_{\beta,h} \delta + \psi_{\gamma,h} [f(\varepsilon_{1t} + \delta) - f(\varepsilon_{1t})]. \quad (5)$$

As (5) shows, this difference depends on the realization of the shock $x_t = \varepsilon_{1t}$ at time t , except if there are no nonlinearities in the model (in which case $\psi_{\gamma,h} = 0$). In order to obtain the population IRF, we average over the realizations of $x_t = \varepsilon_{1t}$.

The following proposition formalizes this result.

Proposition 3.1 *Assume Assumptions E, P, S' and M hold with $\phi(L) = 1$ and $\alpha(L) = 0$. Then, for any $h = 0, 1, 2, \dots$,*

$$IRF_{h,\delta} = \psi_{\beta,h}\delta + \psi_{\gamma,h}A_{0,\delta}, \quad \text{where } A_{0,\delta} \equiv E[f(x_t + \delta) - f(x_t)].$$

This formula indicates that, unless $\gamma(L) = 0$ (implying that $\psi_{\gamma,h} = 0$), the nonlinear $IRF_{h,\delta}$ is different from the traditional linear IRF given by $\psi_{\beta,h}\delta$. The nonlinearity is captured by a second term that depends on $A_{0,\delta}$, the expected value of the difference between the nonlinear function $f(\cdot)$ evaluated at $x_t + \delta$ and x_t , respectively.

3.2 Can local projections identify the IRF?

Next, we investigate the ability of linear LP regressions to recover the IRF. In order to develop intuition, we first consider a special case of model (3), where $\rho(L) = 1 - \rho L$, $\beta(L) = \beta$, $\gamma(L) = \gamma$. For this special case,

$$IRF_{h,\delta} = \beta\rho^h + \gamma\rho^h E(f(x_t + \delta) - E(f(x_t))),$$

i.e. $\psi_{\beta,h} = \beta\rho^h$ and $\psi_{\gamma,h} = \gamma\rho^h$. Note also that $\psi_h = \rho^h$ in this case.

The value of y_{t+h} is

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

which we can write as

$$y_{t+h} = \beta\rho^h x_t + \gamma\rho^h f(x_t) + \rho^{h+1} y_{t-1} + u_{t+h}, \quad (6)$$

where u_{t+h} is a function of $\{\varepsilon_{1t+h}, \dots, \varepsilon_{1t+1}, f(\varepsilon_{1t+h}), \dots, f(\varepsilon_{1t+1}), \varepsilon_{2t+h}, \dots, \varepsilon_{2t+1}, \varepsilon_{2t}\}$. Specifically,

$$\begin{aligned} u_{t+h} &= \beta \left(1 + \rho L + \dots + \rho^{h-1} L^{h-1}\right) \varepsilon_{1t+h} + \gamma \left(1 + \rho L + \dots + \rho^{h-1} L^{h-1}\right) f(\varepsilon_{1t+h}) \\ &\quad + \left(1 + \rho L + \dots + \rho^{h-1} L^{h-1}\right) \varepsilon_{2t+h} + \rho^h \varepsilon_{2t}. \end{aligned}$$

If $\gamma = 0$, $IRF_{h,\delta} = \psi_{\beta,h} = \beta\rho^h$ and expression (6) can be written as

$$y_{t+h} = \pi_{x,h} x_t + \pi_{y,h} y_{t-1} + u_{t+h}, \quad (7)$$

where

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

using the fact that ε_{1t} and ε_{2t} are mutually uncorrelated sequences of i.i.d. zero mean random variables. Hence, (7) is a local projection where $\pi_{x,h} = \psi_{\beta,h}$ and $\pi_{y,h} = \psi_{h+1}$. 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} .

When $\gamma \neq 0$, this result no longer holds. In this case, as we will see, one may still use an LP regression to obtain $\psi_{\beta,h}$ and $\psi_{\gamma,h}$, but one needs to add an estimate of $A_{0,\delta} \equiv E[f(x_t + \delta) - f(x_t)]$.

The LP regression (6) in this case contains x_t , $f(x_t)$ and y_{t-1} as regressors, and the error term u_{t+h} depends on the nonlinear functions $f(x_t)$, evaluated at $t+h, \dots, t+1$. As these functions are not mean zero (even if $x_t = \varepsilon_{1t}$ has mean zero), we need to subtract $E(u_{t+h})$ from the error term. This introduces a constant in equation (6) given by

$$\pi_{c,h} \equiv E(u_{t+h}) = \gamma E \left[\left(1 + \rho L + \dots \rho^{h-1} L^{h-1} \right) f(\varepsilon_{1t+h}) \right] = \gamma \left(1 + \rho + \dots + \rho^{h-1} \right) E(f(\varepsilon_{1t})),$$

where we used the stationarity of ε_{1t} to simplify the expression. Hence, the LP regression becomes

$$y_{t+h} = \pi_{c,h} + \pi_{x,h} x_t + \pi_{f,h} f(x_t) + \pi_{y,h} y_{t-1} + v_{t+h}, \quad (8)$$

where $\pi_{x,h} = \beta \rho^h = \psi_{\beta,h}$, $\pi_{f,h} = \gamma \rho^h = \psi_{\gamma,h}$ and $\pi_{y,h} = \rho^{h+1} = \psi_{h+1}$. The error term is equal to

$$\begin{aligned} v_{t+h} &= \beta \left(1 + \rho L + \dots \rho^{h-1} L^{h-1} \right) \varepsilon_{1t+h} + \\ &\quad + \gamma \left(1 + \rho L + \dots \rho^{h-1} L^{h-1} \right) [f(\varepsilon_{1t+h}) - E(f(\varepsilon_{1t+h}))] \\ &\quad + \left(1 + \rho L + \dots \rho^{h-1} L^{h-1} \right) \varepsilon_{2t+h} + \rho^h \varepsilon_{2t}, \end{aligned}$$

and satisfies the orthogonality conditions

$$\begin{aligned} E(v_{t+h}) &= u_{t+h} - E(u_{t+h}) = 0, \text{ and} \\ E(x_t v_{t+h}) &= E(f(x_t) v_{t+h}) = E(y_{t-1} v_{t+h}) = 0 \end{aligned}$$

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} .

Equation (8) is the local projection equation we need to estimate to recover $IRF_{h,\delta}$ when $\gamma \neq 0$. The coefficients $\pi_{x,h}$ and $\pi_{f,h}$ associated with x_t and $f(x_t)$, respectively, are equal to $\psi_{\beta,h}$ and $\psi_{\gamma,h}$, the coefficients that enter the IRF.

As the next proposition shows, these results generalize to model (3).

Proposition 3.2 *Assume Assumptions E, P, S' and M hold with $\phi(L) = 1$ and $\alpha(L) = 0$. 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

$$\begin{aligned}
\pi_{x,h}(L) &= \pi_{x,1,h} + \pi_{x,2,h}L \dots + \pi_{x,p,h}L^{p-1}, \\
\pi_{f,h}(L) &= \pi_{f,1,h} + \pi_{f,2,h}L + \dots + \pi_{f,p,h}L^{p-1}, \\
\pi_{f,h}(L) &= \pi_{f,1,h} + \pi_{f,2,h}L + \dots + \pi_{f,p,h}L^{p-1}, \text{ and} \\
\pi_{y,h}(L) &= \pi_{y,1,h} + \pi_{y,2,h}L + \dots + \pi_{y,p,h}L^{p-1}, \text{ with} \\
\pi_{x,h} &= \psi_{\beta,h} \quad \text{and} \quad \pi_{f,h} = \psi_{\gamma,h}.
\end{aligned}$$

Moreover, v_{t+h} is mean zero and is orthogonal to all the regressors in the equation, i.e.

$$E(\mathbf{p}_t v_{t+h}) = 0,$$

where

$$\mathbf{p}'_t = (1 \quad x_t \quad \dots \quad x_{t-p} \quad f(x_t) \quad \dots \quad f(x_{t-p}) \quad y_{t-1} \quad \dots \quad y_{t-p}).$$

Proposition 3.2. implies that the IRF can be estimated by h single regressions of y_{t+h} on $x_t, f(x_t)$, and p lags of $x_t, f(x_t)$ and y_t , where the coefficients associated with x_t and $f(x_t)$ are equal to

$$\pi_{x,h} = \psi_{\beta,h} \text{ and } \pi_{f,h} = \psi_{\gamma,h},$$

the inputs necessary to compute $IRF_{h,\delta}$ given in Proposition 3.1. Note that the error term v_{t+h} is an $MA(h)$ process which satisfies the orthogonality conditions $E(\mathbf{p}_t v_{t+h}) = 0$, thus implying that the model $y_{t+h} = \mathbf{p}'_t \pi_h + v_{t+h}$ is a local projection. We prove the consistency of this modified LP-based IRF estimator in Section 5.1.

4 The population IRF for the general model

4.1 A closed-form expression for the IRF

To describe the population IRF for the general model, we proceed as in the previous section and evaluate the difference between $y_{t+h}(\delta)$ and y_{t+h} . We rely on the matrix equation

$$B_0 z_t = B(L) z_{t-1} + C(L) f(x_t) + \varepsilon_t, \tag{9}$$

Pre-multiplying (9) by B_0^{-1} yields

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

where $B_0^{-1} C(L) = C(L)$. Note that given the form of B_0 , B_0^{-1} exists for any value of β_0 and is given by

$$B_0^{-1} = \begin{pmatrix} 1 & 0 \\ \beta_0 & 1 \end{pmatrix}.$$

Hence, we obtain⁴

$$z_t = A(L) z_{t-1} + C(L) f(x_t) + u_t, \quad (10)$$

where

$$A(L) = B_0^{-1} B(L) \text{ and } u_t = B_0^{-1} \varepsilon_t.$$

We can easily show that $A(L)$ is given by

$$A(L) = \sum_{i=1}^p A_i L^{i-1},$$

where

$$A_i \equiv \begin{pmatrix} \phi_i & \alpha_i \\ \beta_0 \phi_i + \beta_i & \beta_0 \alpha_i + \rho_i \end{pmatrix}.$$

We can write

$$\underbrace{(I_2 - A(L)L)z_t}_{=D(L)} = C(L) f(x_t) + u_t, \quad (11)$$

where

$$D(L) \equiv I_2 - A(L)L = I_2 - \sum_{i=1}^p A_i L^i$$

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 S. Using the definition of the inverse filter, i.e. $\Psi(L)D(L) = I_2$, we can show that $\Psi_0 = I_2$, $\Psi_1 = \Psi_0 A_1, \dots$, and $\Psi_j = \Psi_{j-1} A_1 + \dots + \Psi_{j-p} A_p$, for any $j \geq p$.

By pre-multiplying (11) by $\Psi(L)$ and using the fact that $u_t = B_0^{-1} \varepsilon_t$, we can write

$$z_t = \Psi_\beta(L) \varepsilon_t + \Psi_\gamma(L) f(x_t), \quad (12)$$

where

$$\Psi_\beta(L) \equiv \Psi(L) B_0^{-1} \text{ and } \Psi_\gamma(L) \equiv \Psi(L) C(L).$$

The second equation of (12) is the analogue of equation (4) for our general model. In particular, we can write this equation as

$$y_t = \Psi_\beta^{(2,1)}(L) \varepsilon_{1t} + \Psi_\beta^{(2,2)}(L) \varepsilon_{2t} + \Psi_\gamma^{(2,1)}(L) f(x_t), \quad (13)$$

where for any matrix \mathcal{A} , we let $\mathcal{A}^{(i,j)}$ denote its (i,j) element.

⁴If $C(L) = 0$, (10) is the standard reduced-form version of the structural model (2). When $C(L) \neq 0$, (10) is not quite a reduced-form model as $f(x_t)$ still appears in the second equation of the system. In fact, given the form of B_0^{-1} , the ‘‘reduced-form’’ errors for this equation are equal to $u_{2t} = \beta_0 \varepsilon_{1t} + \varepsilon_{2t}$, implying that $f(x_t)$ is correlated with u_{2t} (since it depends on x_t which in turn depends on ε_{1t}). Thus, we cannot estimate the second equation of (10) by OLS (the first equation is the same as the that of (2) since $u_{1t} = \varepsilon_{1t}$ and can be estimated by OLS). We will only use (10) to define the population IRF.

We rely on equation (13) to evaluate the two sample paths of the outcome variable y_{t+h} corresponding to the sequences of shocks $\varepsilon^\infty(\delta)$ and ε^∞ . This yields

$$y_{t+h}(\delta) - y_{t+h} = \Psi_{\beta,h}^{(2,1)}\delta + \Psi_\gamma^{(2,1)}(L)[f(x_{t+h}(\delta)) - f(x_{t+h})],$$

where the second term reflects the contribution of the nonlinearities to the IRF.

We can further simplify this term by noting that $x_{t+h}(\delta)$ must be equal to x_{t+h} for all $h < 0$. This is true because we can use the first equation of (12) to see that x_t depends only on ε_{1t} and lags of $\varepsilon_t = (\varepsilon_{1t}, \varepsilon_{2t})'$ and their nonlinear transforms. Since the two sequences of shocks are the same prior to t , we must have $x_{t+h}(\delta) = x_{t+h}$ for $h < 0$. However, for $h > 0$, $x_{t+h}(\delta) \neq x_{t+h}$. As a result, $y_{t+h}(\delta) - y_{t+h}$ is equal to

$$\Psi_{\beta,h}^{(2,1)}\delta + \Psi_{\gamma,0}^{(2,1)}[f(x_{t+h}(\delta)) - f(x_{t+h})] + \Psi_{\gamma,1}^{(2,1)}[f(x_{t+h-1}(\delta)) - f(x_{t+h-1})] + \dots + \Psi_{\gamma,h}^{(2,1)}[f(x_t(\delta)) - f(x_t)].$$

In contrast to the model considered in Section 3, where $x_t = \varepsilon_{1t}$, a shock to ε_{1t} propagates to future values of x_t , either because x_t is serially correlated or because it depends on lags of y_t , inducing a feedback effect. For these reasons, $f(x_{t+h}(\delta)) - f(x_{t+h}) \neq 0$ for all $h \geq 0$ and not just for $h = 0$, as in Section 3.

The population $IRF_{h,\delta}$ is the unconditional expectation of $y_{t+h}(\delta) - y_{t+h}$. The next proposition formalizes this result and describes an algorithm for computing $A_{h,\delta} \equiv E[f(x_{t+h}(\delta)) - f(x_{t+h})]$.

Proposition 4.1 *Under Assumptions E, P, S, and M, for any $h = 0, 1, 2, \dots$,*

$$IRF_{h,\delta} = \Psi_{\beta,h}^{(2,1)}\delta + \Psi_{\gamma,0}^{(2,1)}A_{h,\delta} + \Psi_{\gamma,1}^{(2,1)}A_{h-1,\delta} + \dots + \Psi_{\gamma,h}^{(2,1)}A_{0,\delta},$$

where

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

The following steps can be used to calculate $A_{j,\delta}$ 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

$$x_{t+j}(\delta) = x_{t+j} + \Psi_{\beta,j}^{(1,1)}\delta + \Psi_{\gamma,1}^{(1,1)}[f(x_{t+j-1}(\delta)) - f(x_{t+j-1})] + \dots + \Psi_{\gamma,j}^{(1,1)}[f(x_t(\delta)) - f(x_t)].$$

For each j , denote

$$x_{t+j}(\delta) = g_{j,\delta}(x_{t+j}, x_{t+j-1}, \dots, x_t; \boldsymbol{\theta}),$$

where the function $g_{j,\delta}$ is defined implicitly by the iteration above.

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

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

As Proposition 4.1 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 $g_{j,\delta}(x_{t+j}, \dots, x_t; \boldsymbol{\theta})$ of the random variables (x_{t+j}, \dots, x_t) , the magnitude of the shock δ and the vector of structural parameters $\boldsymbol{\theta}$. Step (iii) defines $A_{j,\delta}$ as the expectation of the difference of $f(g_{j,\delta}(x_{t+j}, \dots, x_t; \boldsymbol{\theta}))$ and $f(x_{t+j})$.

To illustrate the algorithm described in Proposition 4.1, 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

$$x_{t+1}(\delta) = x_{t+1} + \Psi_{\beta,1}^{(1,1)}\delta + \Psi_{\gamma,1}^{(1,1)}[f(x_t + \delta) - f(x_t)] \equiv g_{1,\delta}(x_{t+1}, x_t; \boldsymbol{\theta}),$$

where $x_t(\delta) = x_t + \delta$ from step (i). Step (iii) then implies that

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

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 $j = 1$.

It is important to note that the iterative algorithm in Proposition 4.1 may be implemented in practice without relying on Monte Carlo integration, making this approach computationally attractive. In particular, estimating the $A_{j,\delta}$ terms, only requires evaluating the sample average of the difference between $f(g_{j,\delta}(x_{t+j}, x_{t+j-1}, \dots, x_t; \hat{\boldsymbol{\theta}})) - f(x_{t+j})$ at each horizon $j = 1, 2, \dots, h$, where $\hat{\boldsymbol{\theta}}$ is an estimator of $\boldsymbol{\theta}$. For instance, in the preceding example for $h = 1$, we compute $x_t(\delta) = x_t + \delta$ for $t = 1, \dots, T$ to obtain the estimate $\hat{A}_{0,\delta} = \frac{1}{T} \sum_{t=1}^T (f(x_t + \delta) - f(x_t))$. Similarly, $\hat{A}_{1,\delta} = \frac{1}{T-1} \sum_{t=2}^T (f(g_{1,\delta}(x_{t+1}, x_t; \hat{\boldsymbol{\theta}})) - f(x_{t+1}))$. Further details of the estimation procedure are discussed in section 5.

4.2 Revisiting the special case when $x_t = \varepsilon_{1t}$

Next, we show that the computation of $g_{j,\delta}$ and hence of $A_{j,\delta}$ can be simplified in two special cases of our model. One special case of our results is the model studied in Section 3. This corresponds to model (2) with

$$B(L) = \begin{pmatrix} 0 & 0 \\ \bar{\beta}(L) & \bar{\rho}(L) \end{pmatrix}.$$

For this model, $A(L) = B(L)$ and

$$D(L) = \begin{pmatrix} 1 & 0 \\ -\bar{\beta}(L)L & 1 - \bar{\rho}(L)L \end{pmatrix} = \begin{pmatrix} 1 & 0 \\ -\bar{\beta}(L)L & \rho(L) \end{pmatrix}.$$

It follows that

$$\Psi(L) = D(L)^{-1} = \begin{pmatrix} 1 & 0 \\ \rho^{-1}(L)\bar{\beta}(L)L & \rho^{-1}(L) \end{pmatrix}.$$

Hence,

$$\Psi_\beta(L) = \Psi(L)B_0^{-1} = \begin{pmatrix} 1 & 0 \\ \rho^{-1}(L)\beta(L) & \rho^{-1}(L) \end{pmatrix} \equiv \begin{pmatrix} 1 & 0 \\ \psi_\beta(L) & \psi(L) \end{pmatrix},$$

using the fact that $\beta(L) = \bar{\beta}(L)L + \beta_0$ and the definition of $\psi_\beta(L) \equiv \rho^{-1}(L)\beta(L)$. Similarly,

$$\Psi_\gamma(L) = \Psi(L)C(L) = \begin{pmatrix} 0 \\ \rho^{-1}(L)\gamma(L) \end{pmatrix} \equiv \begin{pmatrix} 0 \\ \psi_\gamma(L) \end{pmatrix}.$$

We conclude that for any $h = 0, 1, 2, \dots$

$$\Psi_{\beta,h}^{(2,1)} = \psi_{\beta,h} \text{ and } \Psi_{\gamma,h}^{(2,1)} = \psi_{\gamma,h}.$$

Moreover, since for $j = 1, \dots, h$,

$$\Psi_{\beta,j}^{(1,1)} = \Psi_{\gamma,j}^{(1,1)} = 0,$$

we conclude that $x_{t+j}(\delta) = x_{t+j}$ for any $j \geq 1$. Thus, $A_{h,\delta} = 0$ for all $h > 0$ and we obtain

$$IRF_{h,\delta} = \Psi_{\beta,h}^{(2,1)}\delta + \underbrace{\Psi_{\gamma,0}^{(2,1)}A_{h,\delta}}_{=0} + \underbrace{\Psi_{\gamma,1}^{(2,1)}A_{h-1,\delta}}_{=0} + \dots + \underbrace{\Psi_{\gamma,h}^{(2,1)}A_{0,\delta}}_{\neq 0} = \psi_{\beta,h}\delta + \psi_{\gamma,h}A_{0,\delta},$$

where

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

This shows that Proposition 3.1 is a special case of Proposition 4.1.

Because only $A_{0,\delta}$ enters in the computation of the IRF at horizon h , for any $h = 0, 1, 2, \dots$, and

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

the formula for the $IRF_{h,\delta}$ is considerably simpler in this model. In particular, $A_{0,\delta}$ does not depend on any unknown parameters, which simplifies the estimation of $IRF_{h,\delta}$.

4.3 Special case: x_t is an exogenous $AR(p)$ process

Another special case emerges when x_t is exogenous (i.e. $\alpha(L) = 0$) but serially correlated (i.e. $\bar{\phi}(L) \neq 0$). In particular, suppose that

$$x_t = \bar{\phi}(L)x_{t-1} + \varepsilon_{1t},$$

or equivalently,

$$\phi(L)x_t = \varepsilon_{1t}, \quad \text{where } \phi(L) = 1 - \bar{\phi}(L)L.$$

This corresponds to model (2) with

$$B(L) = \begin{pmatrix} \bar{\phi}(L) & 0 \\ \bar{\beta}(L) & \bar{\rho}(L) \end{pmatrix}.$$

For this model,

$$A(L) = \begin{pmatrix} \bar{\phi}(L) & 0 \\ \beta_0 \bar{\phi}(L) + \bar{\beta}(L) & \bar{\rho}(L) \end{pmatrix}$$

and

$$D(L) = \begin{pmatrix} \phi(L) & 0 \\ \beta_0 \phi(L) - \beta(L) & \rho(L) \end{pmatrix}.$$

Note that $A(L)$ and $D(L)$ coincide with the expressions obtained in the previous section when we set $\bar{\phi}(L) = 0$ and $\phi(L) = 1$. The inverse matrix filter of $D(L)$ is equal to

$$\Psi(L) = D(L)^{-1} = \begin{pmatrix} \phi^{-1}(L) & 0 \\ -\beta_0 \rho^{-1}(L) + \phi^{-1}(L) \rho^{-1}(L) \beta(L) & \rho^{-1}(L) \end{pmatrix}$$

implying that

$$\Psi_\beta(L) = \Psi(L)B_0^{-1} = \begin{pmatrix} \phi^{-1}(L) & 0 \\ \phi^{-1}(L) \rho^{-1}(L) \beta(L) & \rho^{-1}(L) \end{pmatrix},$$

and

$$\Psi_\gamma(L) = \Psi(L)C(L) = \begin{pmatrix} 0 \\ \rho^{-1}(L) \gamma(L) \end{pmatrix} \equiv \begin{pmatrix} 0 \\ \psi_\gamma(L) \end{pmatrix}.$$

Hence, $\Psi_\gamma(L)$ is the same as when $x_t = \varepsilon_{1t}$, but the first column of $\Psi_\beta(L)$ is multiplied by $\phi^{-1}(L)$ when $\phi(L)x_t = \varepsilon_{1t}$. The implication is that for any $j = 0, 1, 2, \dots, h$,

$$\begin{aligned} x_{t+j}(\delta) &= x_{t+j} + \Psi_{\beta,j}^{(1,1)} \delta + \underbrace{\Psi_{\gamma,1}^{(1,1)}}_{=0} [f(x_{t+j-1}(\delta)) - f(x_{t+j-1})] + \dots + \underbrace{\Psi_{\gamma,j}^{(1,1)}}_{=0} [f(x_t(\delta)) - f(x_t)] \\ &= x_{t+j} + \Psi_{\beta,j}^{(1,1)} \delta, \end{aligned}$$

where $\Psi_{\beta,j}^{(1,1)}$ is the j^{th} coefficient of the infinite order lag polynomial,

$$\Psi_\beta^{(1,1)}(L) = \phi^{-1}(L) = 1 + \psi_{\phi,1}L + \psi_{\phi,2}L^2 + \dots$$

Consequently,

$$x_{t+j}(\delta) = x_{t+j} + \psi_{\phi,j}\delta \equiv g_{j,\delta}(x_{t+j}; \boldsymbol{\theta}), \text{ for } j = 0, 1, \dots, h, \quad (14)$$

and

$$IRF_{h,\delta} = \Psi_{\beta,h}^{(2,1)}\delta + \Psi_{\gamma,0}^{(2,1)}A_{h,\delta} + \Psi_{\gamma,1}^{(2,1)}A_{h-1,\delta} + \dots + \Psi_{\gamma,h}^{(2,1)}A_{0,\delta},$$

where now

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

Two remarks are in order. First, compared to the previous special case, $x_{t+j}(\delta) \neq x_{t+j}$, for $j > 0$. Thus, the $IRF_{h,\delta}$ depends not only on $A_{0,\delta}$ but also on $A_{1,\delta}$ through $A_{h,\delta}$. This is due to the fact that a shock in ε_{1t} propagates into future values of x_{t+j} due to serial correlation of $\{x_t\}$. Moreover, (14) shows that the function $g_{j,\delta}$ that governs how $x_{t+j}(\delta)$ compares to x_{t+j} is very easy to characterize because it depends only on x_{t+j} and the coefficient $\psi_{\phi,j}$. Step (ii) thus is much simpler in this case, as compared to the general case, where we need to iterate on the equation for $x_{t+j}(\delta)$.

4.4 The inability of local projections to identify the IRF in our general model

A natural question is whether we can use an LP approach to obtain an estimate of $IRF_{h,\delta}$, as we did in Section 3. The answer is no. To explain this result, we consider a simplified version of our model, where $\phi(L) = 1 - \phi L$, $\rho(L) = 1 - \rho L$, $\beta(L) = \beta$ and $\gamma(L) = \gamma$. Given Proposition 4.1, we can show that

$$IRF_{h,\delta} = \beta\rho^h + \gamma A_{h,\delta} + \gamma\rho A_{h-1,\delta} + \dots + \gamma\rho^h A_{0,\delta},$$

where

$$A_{h,\delta} = E\left[f\left(x_t + \phi^h\delta\right) - f\left(x_t\right)\right].$$

We can write

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

where now $x_t = \phi x_{t-1} + \varepsilon_{1t}$. Suppose first $\gamma = 0$. Proceeding as in Section 3, we can decompose y_{t+h} as

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

Now, for any $h \geq 0$, we can write

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

This implies that

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

Letting $\{z_{t+1} \dots z_{t+h}\}$ denote a linear combination of the variables between z_{t+1} and z_{t+h} , we can then write

$$y_{t+h} = \psi_{\beta,h} x_t + \rho^{h+1} y_{t-1} + u_{t+h},$$

where

$$u_{t+h} = \left(\varepsilon_{2t+h} + \rho \varepsilon_{2t+h-1} + \dots + \rho^{h-1} \varepsilon_{2t+1} + \rho^h \varepsilon_{2t} \right) + \{\varepsilon_{1t+1} \dots \varepsilon_{1t+h}\} \equiv \{\varepsilon_{1t+1} \dots \varepsilon_{1t+h}; \varepsilon_{2t} \dots \varepsilon_{1t+h}\}.$$

Given the definition of u_{t+h} and our assumptions on ε_{1t} and ε_{2t} , it is clear that $E(x_t u_{t+h}) = E(y_{t-1} u_{t+h}) = 0$, implying that we can recover $\psi_{\beta,h}$ from a local projection given by the regression of y_{t+h} onto x_t and y_{t-1} .

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

$$y_{t+h} = \pi_{c,h} + \psi_{\beta,h} x_t + \rho^{h+1} y_{t-1} + \gamma \rho^h f(x_t) + v_{t+h}, \quad (15)$$

where the constant $\pi_{c,h} = \gamma (1 + \rho + \rho^2 + \dots + \rho^{h-1}) E(f(x_t))$ accounts for the fact that $\mu_f \equiv E(f(x_t)) \neq 0$ if $f(\cdot)$ is nonlinear. The error term is equal to

$$v_{t+h} = \{\varepsilon_{1t+1} \dots \varepsilon_{1t+h}; \varepsilon_{2t} \dots \varepsilon_{1t+h}\} + \gamma \left[(f(x_{t+h}) - \mu_f) + \rho (f(x_{t+h-1}) - \mu_f) + \dots + \rho^{h-1} (f(x_{t+1}) - \mu_f) \right].$$

It is the presence of the nonlinear term in v_{t+h} that complicates things. Even though we can still write

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

it is not the case in general that we can write

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

for some function $g(\cdot)$. Since all we can say is that

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

the presence of ε_{1t+h} through ε_{1t+1} inside the function $f(x_{t+h})$ implies that the error term v_{t+h} is not orthogonal to x_t nor y_{t-1} since $f(x_{t+h})$ depends on x_t for all $h \geq 0$. Thus, equation (15) is not a local projection. Estimates of this equation do not provide consistent estimates of $\psi_{\beta,h}$ or of the IRF. There are no easy solutions of this problem such as including $f(x_{t+1})$ in the LP.⁵

5 Estimation of the IRF

This section introduces a new estimator of the population *IRF* based on the closed-form expressions obtained in the previous sections. This estimator is an alternative to the Monte Carlo integration method proposed by Kilian and Vigfusson (2011a) for this class of models (see Appendix 1). The new estimator is computationally more attractive, as it does not involve any simulations. Since it is based on a plug-in approach, with the unknown parameters entering $IRF_{h,\delta}$ replaced by sample estimates, we will refer to it as the plug-in estimator, denoted by $IRF_{h,\delta}^{Plug-in}$.

Section 5.1 provides results for the special case when $x_t = \varepsilon_{1t}$, in which case the form of the population IRF simplifies. We also show that in this special case it is possible to estimate the population *IRF* using a modified local projection approach. We then discuss impulse response estimation in the general model in Section 5.2.

5.1 Special case: x_t is an i.i.d. shock

5.1.1 Plug-in estimator

A natural estimator of the population IRF given by Proposition 3.1 is a plug-in estimator that replaces the structural parameters in the lag polynomials $\beta(L)$, $\rho(L)$ and $\gamma(L)$ by consistent estimators $\hat{\beta}(L)$, $\hat{\rho}(L)$ and $\hat{\gamma}(L)$, and estimates the nonlinear term $E[f(x_t + \delta) - f(x_t)]$ by its sample analogue. Let $\mathbf{y}_{t-1} = (y_{t-1}, \dots, y_{t-p})'$, $\mathbf{x}_t = (x_t, \dots, x_{t-p}, f(x_t), \dots, f(x_{t-p}))'$ and $\mathbf{p}_t = (1, \mathbf{y}'_{t-1}, \mathbf{x}'_t)'$.

The steps used to estimate $IRF_{h,\delta}$ using the plug-in approach based on the structural model (3) are described below.

Algorithm for the IRF plug-in estimator when x_t is i.i.d.

⁵If 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 need to estimate the residual by regressing x_t on lags of x_t and y_t . Of course, the $f(\varepsilon_{1t})$ specification is not what applied researchers have been interested in, making this point moot.

i) Regress y_t onto $\mathbf{w}_t = (\mathbf{y}'_{t-1}, \mathbf{x}'_t)'$ and obtain

$$\hat{\boldsymbol{\theta}} = (\hat{\boldsymbol{\rho}}', \hat{\boldsymbol{\beta}}', \hat{\boldsymbol{\gamma}}')' = \left(\sum_{t=1+p}^T \mathbf{w}_t \mathbf{w}'_t \right)^{-1} \sum_{t=1+p}^T \mathbf{w}_t y_t.$$

ii) Compute $\hat{\psi}_{\beta,h}$ and $\hat{\psi}_{\gamma,h}$ using the definitions of $\hat{\psi}_{\beta}(L) = \hat{\beta}(L) \hat{\psi}(L)$ and $\hat{\psi}_{\gamma}(L) = \hat{\gamma}(L) \hat{\psi}(L)$, where $\hat{\psi}(L) = \hat{\rho}(L)^{-1}$.

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

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

iv) Then

$$IRF_{h,\delta}^{Plug-in} = \hat{\psi}_{\beta,h} \delta + \hat{\psi}_{\gamma,h} \hat{A}_{0,\delta}.$$

5.1.2 Modified LP estimator

As shown in Proposition 3.2 of Section 3, we can, alternatively, identify the parameters $\psi_{\beta,h}$ and $\psi_{\gamma,h}$ entering $IRF_{h,\delta}$ via local projections, but, as in the plug-in approach above, consistent estimation of the IRF by LP methods requires an estimate of $A_{0,\delta}$. The algorithm for constructing this modified LP estimator is:

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

i) Regress y_{t+h} onto $\mathbf{p}_t = (1, \mathbf{y}'_{t-1}, \mathbf{x}'_t)'$ and get

$$\hat{\boldsymbol{\pi}} = \left(\sum_{t=1+p}^{T-h} \mathbf{p}_t \mathbf{p}'_t \right)^{-1} \sum_{t=1+p}^{T-h} \mathbf{p}_t y_{t+h}.$$

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) Then

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

Remark 1 *The proof of Proposition 3.2 shows that if the lag polynomials $\rho(L)$, $\beta(L)$, and $\gamma(L)$ have different orders, then we must modify the local projection appropriately. In particular, the local projection should include p_{β} lags of x_t , p_{γ} lags of $f(x_t)$ and p_{ρ} lags of y_t .*

Remark 2 *It should be noted that existing studies in the literature working with (3) have not used the LP algorithm described here (e.g., Tenreyro and Thwaites 2016, Barnichon et al. 2019). Consider the example of $f(x_t) = x_t^+$. The conventional approach has been to fit local projections,*

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

or, equivalently,

$$y_{t+h} = \pi_{c,h} + \pi_{x^+,h}x_t^+ + \pi_{x^-,h}x_t^- + \dots + v_{t+h},$$

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 $\{\pi_{x^-,h}\}_{h=0}^H$, after normalizing the impact coefficients to unity, are 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 response functions even asymptotically.

5.1.3 Consistency

To conclude this section, we establish the consistency of $IRF_{h,\delta}^{Plug-in}$ and $IRF_{h,\delta}^{LP}$ for the population IRF . The plug-in estimator is consistent, provided (i) $\hat{\theta}$ is consistent for θ (which implies the consistency of $\hat{\psi}_{\beta,h}$ and $\hat{\psi}_{\gamma,h}$ given that these are continuous functions of $\hat{\theta}$) and (ii) $\hat{A}_{0,\delta}$ is consistent for $A_{0,\delta}$. Establishing the consistency of the LP estimator also follows standard arguments by showing (ii) and the consistency of the LP coefficient estimators $\hat{\pi}_{x,h}$ and $\hat{\pi}_{f,h}$.

In order to establish these results, we add the following assumptions.

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

Assumption R. $\Sigma_w \equiv E(\mathbf{w}_t\mathbf{w}_t')$ and $\Sigma_p = E(\mathbf{p}_t\mathbf{p}_t')$ are positive definite matrices.

Theorem 5.1 *Under Assumptions E, S', M' and R with $\alpha(L) = 0$ and $\phi(L) = 1$, 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 when $\alpha(L) = 0$ and $\phi(L) = 1$, Assumptions E and S' suffice for Assumption S, as discussed above. We strengthen Assumption M by requiring the existence of finite fourth moments rather than first-order moments of x_t and $f(x_t)$. This is required to ensure the existence of Σ_w and Σ_p . Specifically, we use Assumption M' and the Cauchy Schwarz inequality to show that $\sup_t E |\mathbf{w}_t'\mathbf{w}_t| < \infty$ and $\sup_t E |\mathbf{p}_t'\mathbf{p}_t| < \infty$. Assumption R is a standard rank condition which ensures that Σ_w^{-1} and Σ_p^{-1}

exist. Under these assumptions, one can show that $\hat{\boldsymbol{\theta}}$ is consistent for $\boldsymbol{\theta}$ (and $\hat{\boldsymbol{\pi}}$ is consistent for $\boldsymbol{\pi}$) using standard arguments. This implies the consistency of $\hat{\psi}_{\beta,h}$ and $\hat{\psi}_{\gamma,h}$, given that these are continuous functions of $\hat{\boldsymbol{\theta}}$. Moreover, we can show that $\hat{A}_{0,\delta}$ is consistent for $A_{0,\delta}$. Since $\hat{A}_{0,\delta}$ is the sample average of $f(x_t + \delta) - f(x_t)$, where x_t is i.i.d., this follows by a standard law of large numbers for i.i.d. data. No further restrictions on f beyond those imposed by Assumption S'(ii) are required in this case.

5.2 General model

Next, we extend the analytical estimator introduced in Section 5.1 to the general model. Note that we do not consider estimation results based on LP regressions because these cannot be used to identify the *IRF* in the general model, as discussed in Section 4.

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

- i) Regress y_t onto $\mathbf{w}_t = (\mathbf{y}'_{t-1}, \mathbf{x}'_t)'$. Collect the estimated parameters in $\hat{\boldsymbol{\rho}}, \hat{\boldsymbol{\beta}}$, and $\hat{\gamma}$.
- ii) Regress x_t onto p lags of x_t and p lags of y_t and obtain $\hat{\boldsymbol{\phi}}$ and $\hat{\boldsymbol{\alpha}}$. Let $\hat{\boldsymbol{\theta}} = (\hat{\boldsymbol{\rho}}', \hat{\boldsymbol{\beta}}', \hat{\gamma}', \hat{\boldsymbol{\phi}}', \hat{\boldsymbol{\alpha}})'$.
- iii) Use $\hat{\boldsymbol{\theta}}$ to form

$$\hat{B}_0^{-1} = \begin{pmatrix} 1 & 0 \\ \hat{\boldsymbol{\beta}}_0 & 1 \end{pmatrix}, \quad \hat{C}(L) = \begin{pmatrix} 0 \\ \hat{\gamma}(L) \end{pmatrix},$$

and

$$\hat{\Psi}(L) = I_2 + \hat{\Psi}_1 L + \hat{\Psi}_2 L^2 + \dots + \hat{\Psi}_h L^h + \dots,$$

where

$$\begin{aligned} \hat{\Psi}_0 &= I_2, \quad \hat{\Psi}_1 = \hat{\Psi}_0 \hat{A}_1, \quad \hat{\Psi}_2 = \hat{\Psi}_1 \hat{A}_1 + \hat{\Psi}_0 \hat{A}_2, \quad \text{and} \\ \hat{\Psi}_h &= \hat{\Psi}_{h-1} \hat{A}_1 + \dots + \hat{\Psi}_{h-p} \hat{A}_p, \quad \text{for any } h \geq p, \end{aligned}$$

with

$$\hat{A}_i \equiv \begin{pmatrix} \hat{\phi}_i & \hat{\alpha}_i \\ \hat{\beta}_0 \hat{\phi}_i + \hat{\beta}_i & \hat{\beta}_0 \hat{\alpha}_i + \hat{\rho}_i \end{pmatrix}, \quad i = 1, 2, \dots$$

- iii) For $j = 0, 1, \dots, h$, compute the j^{th} lag coefficients of the elements of

$$\hat{\Psi}_\beta(L) = \hat{\Psi}(L) \hat{B}_0^{-1} \quad \text{and} \quad \hat{\Psi}_\gamma(L) = \hat{\Psi}(L) \hat{C}(L).$$

- iv) Let

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

For $j = 1, \dots, h$, estimate $A_{j,\delta} \equiv E[f(x_{t+j}(\delta)) - f(x_{t+j})]$ with

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

where $g_{j,\delta}(\cdot; \hat{\boldsymbol{\theta}})$ is defined by the recursion in step ii) of Proposition 4.1. For instance, for $h = 1$,

$$g_{1,\delta}\left(x_{t+1}, x_t; \hat{\boldsymbol{\theta}}\right) = x_{t+1} + \hat{\Psi}_{\beta,1}^{(1,1)}\delta + \hat{\Psi}_{\gamma,1}^{(1,1)}[f(x_t + \delta) - f(x_t)]$$

and

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

v) Set

$$IRF_{h,\delta}^{Plug-in} = \hat{\Psi}_{\beta,h}^{(2,1)}\delta + \hat{\Psi}_{\gamma,0}^{(2,1)}\hat{A}_{h,\delta} + \hat{\Psi}_{\gamma,1}^{(2,1)}\hat{A}_{h-1,\delta} + \dots + \hat{\Psi}_{\gamma,h}^{(2,1)}\hat{A}_{0,\delta}.$$

As for the special case considered in Section 5.1, establishing the consistency of the plug-in estimator in the general model requires two steps: (i) consistency of $\hat{\boldsymbol{\theta}}$ for $\boldsymbol{\theta}$; and (ii) consistency of the nonlinear terms $\hat{A}_{j,\delta}$ for $j = 0, 1, 2, \dots, h$. The first step follows the same arguments as in Section 5.1, once we assume that $z_t = (x_t, y_t)'$ is stationary and ergodic. As we discussed in Section 2.1, providing primitive conditions for Assumption S in the general model is more involved because of the presence of $f(x_t)$ in the structural equation for y_t and the fact that $\alpha(L) \neq 0$. This task is beyond the scope of this paper and left for future research.

Another complication compared to Section 5.1 is that we now need to show consistency of the additional nonlinear terms $\hat{A}_{j,\delta}$ for $j = 1, 2, \dots, h$. As the algorithm above makes clear, these terms are no longer a simple sample mean of $f(x_t + \delta) - f(x_t)$, as is the case for $\hat{A}_{0,\delta}$. Instead, they depend on $g_{j,\delta}\left(x_{t+j}, x_{t+j-1}, \dots, x_t; \hat{\boldsymbol{\theta}}\right)$, a continuous function of $\hat{\boldsymbol{\theta}}$ that is recursively defined for $j = 1, 2, \dots, h$. Consequently, showing the consistency of $\hat{A}_{j,\delta}$ towards $A_{j,\delta}$ for $j > 0$ involves more than a simple application of a law of large numbers. If we define $f_{\theta}\left(x_{t+j}, x_{t+j-1}, \dots, x_t; \boldsymbol{\theta}\right) = f\left(g_{j,\delta}\left(x_{t+j}, x_{t+j-1}, \dots, x_t; \boldsymbol{\theta}\right)\right)$ and assume that $f_{\theta}(\cdot, \boldsymbol{\theta})$ is a differentiable function of $\boldsymbol{\theta}$, a standard application of the Taylor expansion implies the result. If f_{θ} is not differentiable as a function of $\boldsymbol{\theta}$, as in some of our examples, the result follows by an application of a uniform law of large numbers for stationary ergodic data (see e.g. Theorem A.2.2. of White, 1996). This in turn requires more structure on the function $f(\cdot)$ such as a Lipschitz continuity condition.

6 Discussion

Table 1 summarizes the main results. Regardless of the specification of model (1), it is always possible to evaluate the unconditional response function of interest numerically by Monte Carlo integration.

Typically, this involves constructing response functions conditional on randomly drawn histories Ω_{t-1} and averaging the response functions over these histories.

Alternatively, one can use the plug-in estimator proposed in this paper whose computational cost is negligible. This approach also works in the special cases when x_t is an observable exogenous shock or when x_t follows an exogenous autoregressive process. The availability of plug-in estimators is an important advantage when bootstrapping the distribution of this estimator or when constructing critical values under the null of symmetric response functions (see Kilian and Vigfusson 2011a).

LP estimators of the impulse response function can only be used when x_t is an observable exogenous i.i.d. shock. The latter situation has become the empirically most important setting in macroeconomics in recent years. More generally, when x_t is exogenous, but serially correlated, or when x_t is endogenous, LP estimators are not valid in this class of models, contrary to conjectures in the existing literature. Thus, the equivalence between impulse response estimators based on the structural model and LP estimators breaks down in this class of models.

It should be noted, however, that the construction of the LP estimator differs from the standard LP estimators currently used in the literature, which effectively ignore the nonlinearity of the impulse responses (e.g., Tenreyro and Thwaites 2016, Barnichon, Matthes and Ziegenbein 2019). Conventional LP estimators are invalid. This result has important implications for applied work on the asymmetry of responses to positive and negative shocks of the same magnitude in a wide range of macroeconomic contexts and for studies that evaluate the differential impact of large and small shocks.

7 Simulation evidence

In this section, we examine the finite-sample accuracy of the plug-in estimator, the Monte Carlo integration (MCI) estimator and (where applicable) the modified LP estimator. The evaluation criteria are the bias, variance and mean squared error (MSE) of the impulse response estimators.

7.1 Simulation design

We consider three data generating processes (DGPs). DGP 1 is the population model (1) with x_t restricted to an observed i.i.d. shock:

$$\begin{cases} x_t = \varepsilon_{1t}, \\ y_t = 0.5y_{t-1} + 0.5x_t + 0.3x_{t-1} - 0.4 \max(0, x_t) + 0.3 \max(0, x_{t-1}) + \varepsilon_{2t}, \end{cases} \quad (16)$$

In DGP 2, x_t is instead restricted to follow an exogenous AR(1) process:

$$\begin{cases} x_t = 0.5x_{t-1} + \varepsilon_{1t}, \\ y_t = 0.5y_{t-1} + 0.5x_t + 0.3x_{t-1} - 0.4 \max(0, x_t) + 0.3 \max(0, x_{t-1}) + \varepsilon_{2t}, \end{cases} \quad (17)$$

DGP 3 corresponds to the unrestricted model (1):

$$\begin{cases} x_t = 0.3x_{t-1} + 0.2y_{t-1} + \varepsilon_{1t}, \\ y_t = 0.5y_{t-1} + 0.5x_t + 0.3x_{t-1} - 0.4 \max(0, x_t) + 0.2 \max(0, x_{t-1}) + \varepsilon_{2t} \end{cases} . \quad (18)$$

In all DGPs, the intercept has been normalized to 0 in population and the population innovations are Gaussian. 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 $\delta = 1$. The MCI method is implemented, as discussed in the appendix, with $M_1 = M_2 = 1,000$. The plug-in method and LP method are implemented as discussed in section 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. Qualitatively similar results are obtained for $T = 120$ and $T = 480$, corresponding to ten years of monthly (or 30 years of quarterly) data and 40 years of monthly data.

7.2 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 slightly lower bias at longer horizons, it has much higher variance and hence a much higher MSE. The performance of the plug-in estimator and the MCI estimator is identical, as expected for this DGP. In the second row of Figure 1, the bias of the MCI and plug-in estimators tends to be very small and their variance is 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 last 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, 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, as shown in Table 2. Even for these simple DGPs, the plug-in estimator is 200-300 times faster than the MCI estimator.

8 Conclusions

Our analysis in this paper has focused on structural dynamic models with nonlinearly transformed regressors that are widely used in applied work. We compared existing approaches to estimating

impulse responses in this setting based on Monte Carlo integration methods and linear projections (LPs) to a new approach based on analytical derivations.

We first derived exact solutions for population impulse responses in this class of models. This result not only is required for evaluating the finite-sample accuracy of alternative impulse response estimators by simulation, but it allows us to derive a plug-in estimator based on estimates of the structural model that avoids the high computational cost of Monte Carlo integration methods. We showed that this estimator has a closed-form solution when the forcing variable is exogenous. It can be computed iteratively in more general model specifications without resorting to Monte Carlo integration.

Second, we examined the theoretical support for the use of linear projections in dynamic models that allow for various forms of nonlinear responses, while maintaining a linear structure that facilitates the estimation of the model. We showed that, among the estimators considered, in all but one important special case of this model there is no alternative to the use of impulse response estimators based on the structural model. When the LP estimator may be applied, its construction differs from current implementations of the LP estimator in applied work. This result has important implications for existing empirical work on the role of asymmetries and other nonlinearities in the transmission of shocks to macroeconomic aggregates.

Third, in addition to establishing the identifiability of the population *IRF*, we proved the consistency of the modified LP estimator and the plug-in estimator for the special case of the nonlinear response to an observable exogenous i.i.d. shock.

Fourth, we compared the finite-sample accuracy of the modified LP estimator to plug-in and Monte Carlo integration estimators based on the structural model. Simulations suggest that, when all estimators are valid, the modified LP estimator tends to be less reliable for realistic sample sizes. In contrast, there is little to choose between the Monte Carlo integration estimator and the plug-in estimator based on their finite-sample accuracy for any of the model specifications considered. The analytical approach, however, is 200-300 times faster in practice, which facilitates the use of bootstrap methods of inference in this class of models.

An interesting question for future research will be whether LP estimators of impulse responses are theoretically justified in 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 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. Thus, the implied impulse responses presumably can only be estimated by nonlinear local

projections, whereas the current paper focused on more conventional linear local projections.

9 References

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. 2019. Are the effects of financial market disruptions big or small? Manuscript, Federal Reserve Bank of San Francisco.

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

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.

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.

Jorda, O., 2005. Estimation and inference of impulse responses by local projections. *American Economic Review* 95, 161-182.

Jorda, 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., 2013. Do oil prices help forecast U.S. real GDP? The role of nonlinearities and asymmetries. *Journal of Business and Economic Statistics* 31, 78-93.

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.

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., 2019. Local projections and VARs estimate the same impulse responses. Manuscript, Princeton University.

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.

A Appendix 1: Algorithm for Computing Responses by Monte Carlo Integration

The algorithm involves first constructing the conditional response function, $IRF_{h,\delta,\Omega_{t-1}}$ at horizons $h = 0, \dots, H$ from model (1), or any of its variants, and then the unconditional response function, $IRF_{h,\delta}$. We focus on the most general model specification, but the algorithm below applies with the obvious changes to any of the models considered in this paper.

1. Consider a block of p consecutive values of x_τ , $f(x_\tau)$, and y_τ for $\tau = t-p, \dots, t-1$. This sequence defines a history, Ω_{t-1} .
2. Simulate two alternative time paths for x_{t+h} , $f(x_{t+h})$, and y_{t+h} , $h = 0, \dots, H$, from model (1) by iterating the model forward, given Ω_{t-1} and the estimated model coefficients. In simulating data from the model, we need to take a stand on the future realizations of the structural shocks. Each realization of $\varepsilon_{1,t+h}$, $h = 0, \dots, H$, is drawn with replacement from the set of structural residuals, $\{\widehat{\varepsilon}_{1,t+h}\}_{p+1}^T$. Realizations of $\varepsilon_{2,t+h}$, $h = 0, \dots, H$, are independently generated by drawing with replacement from the set of structural residuals $\{\widehat{\varepsilon}_{2,t+h}\}_{p+1}^T$. The only difference between the baseline and the alternative shock sequence is that for the alternative we replace $\varepsilon_{1,t}$ by $\varepsilon_{1,t} + \delta$, where the constant δ measures the magnitude of the shock of interest.
3. Subtract the simulated baseline path of y_{t+h} , $h = 0, \dots, H$, from the simulated alternative path. The difference across the two paths measures the impact of a shock of size δ on y_{t+h} , $h = 0, \dots, H$, conditional on Ω_{t-1} .
4. Repeat steps 2 and 3 M_1 times and average this difference to obtain an estimate of the impulse response function conditional on Ω_{t-1} , denoted by $IRF_{h,\delta,\Omega_{t-1}}$.

The corresponding unconditional response function can be computed as the average value of the conditional response functions over M_2 randomly drawn histories, Ω_{t-1}^r , each of which is generated by drawing with replacement a block of p consecutive observations from the observed data:

$$IRF_{h,\delta} = \int IRF_{h,\delta,\Omega_{t-1}^r} d\Omega_{t-1}^r.$$

B Appendix 2: Proofs

Proof of Proposition 3.1. The proof is in the text. ■

Proof of Proposition 3.2. We prove the result for $h \geq p$. The proof for $0 \leq h < p$ is similar, but requires some adjustments. Given model (3), we have that

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

Since $\psi_\beta(L) = \psi(L)\beta(L)$, the first term can be written as

$$\begin{aligned}\psi_\beta(L)x_{t+h} &= \psi(L)\beta(L)x_{t+h} = \psi(L)(\beta_0 + \beta_1L + \dots + \beta_pL^p)x_{t+h} \\ &= \psi(L)\beta_0x_{t+h} + \psi(L)\beta_1x_{t+h-1} + \dots + \psi(L)\beta_px_{t+h-p}.\end{aligned}$$

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

$$\begin{aligned}\psi_\beta(L)x_{t+h} &= \beta_0\psi_hx_t + \beta_0\left\{1 + \psi_1L + \dots + \psi_{h-1}L^{h-1}\right\}x_{t+h} \\ &\quad + \beta_0\left\{\psi_{h+1}L^{h+1} + \psi_{h+2}L^{h+2} + \dots\right\}x_{t+h} + \beta_1\psi_{h-1}x_t \\ &\quad + \beta_1\left\{1 + \psi_1L + \dots + \psi_{h-2}L^{h-2}\right\}x_{t+h-1} + \beta_1\left\{\psi_hL^h + \psi_{h+1}L^{h+1} + \dots\right\}x_{t+h-1} \\ &\quad + \dots \\ &\quad + \beta_px_{h-p}x_t + \beta_p\left\{1 + \psi_1L + \dots + \psi_{h-p-1}L^{h-p-1}\right\}x_{t+h-p} + \beta_p\left\{\psi_{h-p+1}L^{h-p+1} + \dots\right\}x_{t+h-p} \\ &= \underbrace{(\beta_0\psi_h + \beta_1\psi_{h-1} + \dots + \beta_p\psi_{h-p})}_{=\psi_{\beta,h}}x_t + u_{1h} + u_{2h},\end{aligned}$$

where

$$\begin{aligned}u_{1h} &\equiv \beta_0\underbrace{\left\{1 + \psi_1L + \dots + \psi_{h-1}L^{h-1}\right\}}_{=\{x_{t+h}, \dots, x_{t+1}\}}x_{t+h} + \beta_1\underbrace{\left\{1 + \psi_1L + \dots + \psi_{h-2}L^{h-2}\right\}}_{=\{x_{t+h-1}, \dots, x_{t+1}\}}x_{t+h-1} \\ &\quad + \dots + \beta_p\underbrace{\left\{1 + \psi_1L + \dots + \psi_{h-p-1}L^{h-p-1}\right\}}_{=\{x_{t+h-p}, \dots, x_{t+1}\}}x_{t+h-p}\end{aligned}$$

is a (linear) function only of future values of x , $\{x_{t+1}, \dots, x_{t+h}\}$.⁶ Note that because $x_t = \varepsilon_{1t}$, an i.i.d. shock over time, we have that $E(u_{1h}) = 0$. In addition, u_{1h} is orthogonal to x_t , $f(x_t)$, their lags as well as lags of y_t . Instead, the term u_{2h} is defined as

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

⁶Henceforth, 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.

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

$$\begin{aligned}
u_{2h} &= \beta_0 \{ \psi_{h+1}x_{t-1} + \psi_{h+2}x_{t-2} + \dots \} + \beta_1 \{ \psi_h x_{t-1} + \psi_{h+1}x_{t-2} + \dots \} \\
&+ \beta_2 \{ \psi_{h-1}x_{t-1} + \psi_h x_{t-2} + \psi_{h+1}x_{t-3} + \dots \} \\
&+ \dots \\
&+ \beta_p \{ \psi_{h-p+1}x_{t-1} + \psi_{h-p+2}x_{t-2} + \dots + \psi_h x_{t-p} + \psi_{h+1}x_{t-p-1} + \dots \} \\
&= \psi_{h+1} \underbrace{\{ \beta_0 + \beta_1 L + \dots + \beta_p L^p \}}_{=\beta(L)} x_{t-1} + \underbrace{(\beta_1 \psi_h + \beta_2 \psi_{h-1} + \dots + \beta_p \psi_{h-p+1})}_{\equiv \pi_{x,1,h}} x_{t-1} \\
&+ \psi_{h+2} \underbrace{\{ \beta_0 + \beta_1 L + \dots + \beta_p L^p \}}_{=\beta(L)} \underbrace{x_{t-2}}_{Lx_{t-1}} + \underbrace{(\beta_2 \psi_h + \dots + \beta_p \psi_{h-p+2})}_{\equiv \pi_{x,2,h}} x_{t-2} \\
&+ \dots \\
&+ \psi_{h+p} \underbrace{\{ \beta_0 + \beta_1 L + \dots + \beta_p L^p \}}_{=\beta(L)} \underbrace{x_{t-p}}_{=L^{p-1}x_{t-1}} + \underbrace{(\beta_p \psi_h)}_{\equiv \pi_{x,p,h}} x_{t-p} \\
&+ \psi_{h+p+1} \{ \beta_0 + \beta_1 L + \dots + \beta_p L^p \} x_{t-p-1} + 0 \\
&+ \dots \\
&= \underbrace{\pi_{x,1,h}x_{t-1} + \dots + \pi_{x,p,h}x_{t-p}}_{=\pi_x(L)x_{t-1}} + (\psi_{h+1} + \psi_{h+2}L + \dots) \beta(L) x_{t-1}
\end{aligned}$$

Assembling these results yields

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

where

$$\pi_{x,h} = \psi_{\beta,h}$$

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

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

where

$$\begin{aligned}
\pi_{f,h} &= \psi_{\gamma,h}, \\
\pi_f(L) &= \pi_{f,1,h} + \dots + \pi_{f,p,h}L^{p-1},
\end{aligned}$$

and

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

and it can be shown that

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

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

$$\begin{aligned}
y_{t+h} &= \pi_{c,h} + \psi_{\beta,h} x_t + \psi_{\gamma,h} f(x_t) + \pi_{x,h}(L) x_{t-1} + \pi_{f,h}(L) f(x_{t-1}) \\
&+ [(\psi_{h+1} + \psi_{h+2}L + \dots) \beta(L) x_{t-1}] \\
&+ (\psi_{h+1} + \psi_{h+2}L + \dots) \gamma(L) f(x_{t-1}) \\
&+ (\psi_{h+1} + \psi_{h+2}L + \dots) \varepsilon_{2t-1}] \\
&+ v_{t+h} \\
&= \pi_{c,h} + \psi_{\beta,h} x_t + \psi_{\gamma,h} f(x_t) + \pi_{x,h}(L) x_{t-1} + \pi_{f,h}(L) f(x_{t-1}) \\
&+ a_h + v_{t+h}
\end{aligned}$$

where a_h is the term in square brackets and v_{t+h} is the error term. Thus,

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

and

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

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 $a_h \equiv (\psi_{h+1} + \psi_{h+2}L + \dots) \rho(L) y_{t-1}$ 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,

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

This concludes the proof. ■

Proof of Proposition 4.1. Let $\mathcal{A}^{(\cdot,j)}$ denote the j^{th} column of any matrix \mathcal{A} . With this notation, we can write (12) as

$$\begin{pmatrix} x_t \\ y_t \end{pmatrix} = \Psi_\beta^{(\cdot,1)}(L) \varepsilon_{1t} + \Psi_\beta^{(\cdot,2)}(L) \varepsilon_{2t} + \Psi_\gamma(L) f(x_t). \quad (19)$$

The elements of the vectors $\Psi_\beta^{(\cdot,1)}(L)$, $\Psi_\beta^{(\cdot,2)}(L)$ and $\Psi_\gamma(L)$ are infinite order lag polynomials whose coefficients are a function of the structural model parameters. Importantly, we can show that

$$\Psi_{\beta,0}^{(\cdot,1)} = (1, \beta_0)', \quad \Psi_{\beta,0}^{(\cdot,2)} = (0, 1)' \quad \text{and} \quad \Psi_{\gamma,0} = (0, \gamma_0)'.$$

This implies that x_t does not depend on ε_{2t} nor it depends on $f(x_t)$ (although it can depend on lags of ε_{2t} and $f(x_t)$). For the benchmark economy, using (19),

$$z_{t+h} = \Psi_\beta^{(\cdot,1)}(L) \varepsilon_{1t+h} + \Psi_\beta^{(\cdot,2)}(L) \varepsilon_{2t+h} + \Psi_\gamma(L) f(x_{t+h}),$$

whereas for the δ -economy,

$$z_{t+h}(\delta) = \Psi_\beta^{(\cdot,1)}(L) \varepsilon_{1t+h}(\delta) + \Psi_\beta^{(\cdot,2)}(L) \varepsilon_{2t+h}(\delta) + \Psi_\gamma(L) f(x_{t+h}(\delta)).$$

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

$$z_{t+h}(\delta) - z_{t+h} = \Psi_{\beta,h}^{(\cdot,1)}\delta + \Psi_{\gamma}(L)[f(x_{t+h}(\delta)) - f(x_{t+h})].$$

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

$$\begin{aligned} z_{t+h}(\delta) - z_{t+h} &= \Psi_{\beta,h}^{(\cdot,1)}\delta + \Psi_{\gamma,0}[f(x_{t+h}(\delta)) - f(x_{t+h})] \\ &\quad + \Psi_{\gamma,1}[f(x_{t+h-1}(\delta)) - f(x_{t+h-1})] + \dots + \Psi_{\gamma,h}[f(x_t(\delta)) - f(x_t)]. \end{aligned} \quad (20)$$

This describes a system of two equations which we can use to obtain the population IRF in closed form. In particular, the second equation is

$$y_{t+h}(\delta) - y_{t+h} = \Psi_{\beta,h}^{(2,1)}\delta + \Psi_{\gamma,0}^{(2,1)}[f(x_{t+h}(\delta)) - f(x_{t+h})] + \dots + \Psi_{\gamma,h}^{(2,1)}[f(x_t(\delta)) - f(x_t)],$$

implying that

$$\begin{aligned} IRF_{h,\delta} &\equiv E(y_{t+h}(\delta) - y_{t+h}) \\ &= \Psi_{\beta,h}^{(2,1)}\delta + \Psi_{\gamma,0}^{(2,1)} \underbrace{E[f(x_{t+h}(\delta)) - f(x_{t+h})]}_{\equiv A_{h,\delta}} + \dots + \Psi_{\gamma,h}^{(2,1)} \underbrace{E[f(x_t(\delta)) - f(x_t)]}_{\equiv A_{0,\delta}}. \end{aligned}$$

To end the proof, we explain why the proposed algorithm can be used to evaluate $A_{j,\delta}$ for any $j = 0, 1, \dots, h$ (these are the terms needed to obtain $IRF_{h,\delta}$ for any $h = 0, 1, 2, \dots, H$).

Suppose we care about $IRF_{0,\delta}$, which requires $A_{0,\delta}$. The first equation of (20) with $h = 0$ implies

$$x_t(\delta) = x_t + \Psi_{\beta,0}^{(1,1)}\delta = x_t + \delta,$$

since $\Psi_{\beta,0}^{(1,1)} = 1$. Hence, $A_{0,\delta} = E(f(x_t + \delta) - f(x_t))$ as in step (i). Next, consider $h = 1$. We require now $A_{j,\delta}$ for $j = 0, 1$. Step (i) gives $A_{0,\delta}$. For $j = 1$, steps (i) and (ii) imply that

$$\begin{aligned} x_{t+1}(\delta) &= x_{t+1} + \Psi_{\beta,1}^{(1,1)}\delta + \Psi_{\gamma,1}^{(1,1)}[f(x_t(\delta)) - f(x_t)] \\ &= x_{t+1} + \Psi_{\beta,1}^{(1,1)}\delta + \Psi_{\gamma,1}^{(1,1)}[f(x_t + \delta) - f(x_t)], \end{aligned}$$

using the fact that $x_t(\delta) = x_t + \delta$. This shows that we can express $x_{t+1}(\delta)$ as a function of x_{t+1} and x_t , as well as δ and the model parameters $\boldsymbol{\theta}$. We call this function $x_{t+1}(\delta) = g_{1,\delta}(x_{t+1}, x_t; \boldsymbol{\theta})$. Step (iii) defines

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

By iterating on steps (ii) and (iii), we can obtain $A_{j,\delta}$ for any $j = 0, 1, \dots, h$ and any $h = 0, 1, 2, \dots$

■

Proof of Theorem 5.1. The proof follows standard arguments and therefore we only provide a sketch of the proof. For the plug-in estimator, to show that $\hat{\boldsymbol{\theta}}$ is consistent for $\boldsymbol{\theta}$, note that we can write

$$\hat{\boldsymbol{\theta}} - \boldsymbol{\theta} = \left(\frac{1}{T} \sum_{t=1+p}^T \mathbf{w}_t \mathbf{w}_t' \right)^{-1} \frac{1}{T} \sum_{t=1+p}^T \mathbf{w}_t \varepsilon_{2t}.$$

Under the stated assumptions, $\{\mathbf{w}_t \mathbf{w}_t'\}$ is stationary and ergodic, which implies that $\frac{1}{T} \sum_{t=1+p}^T \mathbf{w}_t \mathbf{w}_t' \xrightarrow{P} \Sigma_w > 0$. Similarly, we can show $\frac{1}{T} \sum_{t=1+p}^T \mathbf{w}_t \varepsilon_{2t} \xrightarrow{P} E(\mathbf{w}_t \varepsilon_{2t}) = 0$, where the orthogonality conditions can be verified under Assumption E. In particular, \mathbf{w}_t depends on \mathbf{y}_{t-1} and \mathbf{x}_t , which are only a function of ε_{1t} and lags of $\varepsilon_t = (\varepsilon_{1t}, \varepsilon_{2t})$ and their nonlinear transforms. Since \mathbf{w}_t is not a function of ε_{2t} , Assumption E implies that $E(\mathbf{w}_t \varepsilon_{2t}) = 0$, delivering the result. A similar argument can be used to show the consistency of $\hat{\boldsymbol{\pi}}$ towards $\boldsymbol{\pi}$. The consistency of $\hat{A}_{0,\delta}$ towards $A_{0,\delta}$ follows by a standard LLN for i.i.d. observations. ■

Table 1: Validity of Alternative Impulse Response Estimators

Data Generating Process	Based on Structural Model		Based on LP	
	Plug-in	Monte Carlo Integration	Modified	Conventional
$\begin{cases} x_t = \varepsilon_{1t} \\ \rho(L)y_t = \beta(L)x_t + \gamma(L)f(x_t) + \varepsilon_{2t} \end{cases}$	YES	YES	YES	NO
$\begin{cases} \phi(L)x_t = \varepsilon_{1t} \\ \rho(L)y_t = \beta(L)x_t + \gamma(L)f(x_t) + \varepsilon_{2t} \end{cases}$	YES	YES	NO	NO
$\begin{cases} \phi(L)x_t = \alpha(L)y_t + \varepsilon_{1t} \\ \rho(L)y_t = \beta(L)x_t + \gamma(L)f(x_t) + \varepsilon_{2t} \end{cases}$	YES	YES	NO	NO

Table 2: Ratio of Computation Time for Monte Carlo Integration Relative to Plug-in Estimator

	$T = 120$	$T = 240$	$T = 480$
DGP 1	278.57	272.30	272.97
DGP 2	302.87	295.83	285.13
DGP 3	306.07	321.65	227.03




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