When do state-dependent local projections work?

Sílvia Gonçalves[*], Ana María Herrera[†], Lutz Kilian[‡] and Elena Pesavento[§]
Preliminary - Please do not cite

March 4, 2022

Abstract

Many empirical studies estimate impulse response functions that depend on the state of the economy. Most of these studies rely on a variant of the local projection (LP) approach to estimate the state-dependent impulse response functions. Despite its widespread application, the asymptotic validity of the LP approach to estimating state-dependent impulse responses has not been established to date. We formally derive this result for a structural state-dependent vector autoregressive process. The model only requires the structural shock of interest to be identified. A crucial condition for the consistency of the state-dependent LP estimator of the response function is that current and future states are conditionally mean independent of the structural shocks, given the information available at the time the shock is realized. This rules out models in which the state of the economy is a function of current or future realizations of the outcome variable of interest, as is often the case in applied work. Even when the state is a function of past values of this variable only, consistency may hold only at short horizons.

---

[*]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: epesave@emory.edu.

1 Introduction

The recent empirical macroeconomics literature has emphasized the importance of allowing for nonlinearities when estimating the effects of exogenous shocks in macroeconomic variables of interest. A key question in empirical work is how impulse response functions depend on the state of the economy. For example, many studies estimating the government spending multiplier allow for the possibility that this multiplier may be different during recessions and expansions (e.g., Auerbach and Gorodnichenko (2012, 2013a,b), Bachmann and Sims (2012), Owyang, Ramey and Zubairy (2013), Caggiano, Castelnuovo, Colombo and Nodari (2015), Ramey and Zubairy (2018), Alloza (2019), and Ghassibe and Zanetti (2020)). There is also a related literature on the dependence of tax multipliers on the business cycle (e.g., Candelon and Lieb (2013), Alesina, Azzalini, Favero, Giavazzi and Miano (2018), Sims and Wolff (2018), Eskandari (2019), and Demirel (2021)). Similar questions arise in many other contexts including the analysis of monetary policy shocks. For example, Santoro, Petrella, Pfajfar and Gaffeo (2014), Tenreyro and Thwaites (2016), Angrist, Jordà and Kuersteiner (2018), Barnichon and Matthes (2018) and Klepacz (2020) allow the responses to monetary policy shocks to vary as a function of the state of the economy. Other studies allow these responses to vary depending on whether the zero lower bound is binding (e.g., Ramey and Zubairy 2018, Mavroeidis 2021). Yet another example of the estimation of state-dependent responses is the work of Caggiano, Castelnuovo and Groshenny (2014) who examine the dependence of the effects of uncertainty shocks on whether the economy is in recession or expansion.

Most of these studies rely on a variant of the local projection (LP) approach of Jordà (2005, 2009) to estimate the state-dependent impulse response functions. One argument for using state-dependent local projections rather than structural nonlinear vector autoregressive (VAR) models to estimate the impulse response functions is its computational simplicity. Estimating impulse responses in state-dependent VAR models by numerical methods tends to be computationally more challenging than the estimation of state-dependent local projections by the method of least squares.[1] Yet, despite its widespread application, the validity of the LP approach to estimating state-dependent impulse responses has not been established to date.[2]

In this paper, we clarify the conditions under which the state-dependent LP estimator can be

---

[1] For example, Ramey (2016, p. 87) stresses that "if one is interested in estimating state dependent models, the ... local projection method is a simple way to estimate such a model and calculate impulse response functions."

[2] LPs have become an increasingly popular alternative to VAR based estimators of impulse responses. The original LP estimator, as proposed by Jordà (2005, 2009) did not allow for the impulse response function to change depending on the state of the economy. For a review of the rationale underlying standard linear LPs the reader is referred to Plagborg-Møller and Wolf (2021). In this paper we are not concerned with linear approximations to nonlinear processes as in Plagborg-Møller and Wolf (2021), but with approximations that are explicitly state dependent and hence nonlinear.

expected to recover the population impulse responses in multivariate models. Our analysis only requires the structural shock of interest to be identified, allowing the user to remain agnostic about the identification of the remainder of the structural model. As it turns out, the crucial condition for the validity of the LP estimator in this context relates to the information set used to compute the state indicators. If this set only includes exogenous variables determined outside of the model, the state-dependent LP estimator is asymptotically valid and recovers the conditional IRF at any finite horizon. If instead the state indicator is a function of endogenous model variables, the asymptotic validity of the LP estimator depends on whether the state of the economy is a function of current, lagged or future realizations of the endogenous model variables. For example, if the state depends on current values of these variables, the LP estimator asymptotically recovers the impact response, but not the responses at horizons greater than zero. Basing the state only on lagged values instead allows the LP estimator to consistently estimate impulse responses at longer horizons. The longer the horizon of interest is, the more restrictive the lag structure needs to be when determining the current state of the economy as a function of endogenous variables in the model. In particular, to identify impulse responses up to order h_{max}, the minimum lag order should be h_{max}. Put differently, to be able to identify impulse responses at horizon $h = 0, 1, \ldots, h_{\max}$, the state indicator H_t has to be a function of $y_{t-h_{\max}}, y_{t-h_{\max}-1}, \ldots$.

While these results do not formally establish the inconsistency of the LP estimator when our sufficient conditions are violated, we show by simulation that the LP estimator of the response function tends to be asymptotically biased except for the impact response, when the state of the economy is endogenous. These asymptotic biases may become substantial when cumulating the level responses of the model variables, as required for computing fiscal or monetary multipliers, for example.

State-dependent local projections are extremely popular in macroeconomics because they are easily implemented and because they are believed to be more robust to dynamic model misspecification than numerical estimates of impulse response functions obtained from state-dependent structural VAR models. Our results suggest that researchers need to think carefully about the model specification underlying these local projections. In fact, assessing the validity of the state-dependent LP estimator requires the user to explicitly state the underlying structural data generating process.

Of particular concern is that in many macroeconomic applications one would expect exogenous shocks to affect not only the future realizations of the model variables, but also the future state of the economy, rendering the state of the economy endogenous with respect to the model variables. The implicit assumption in many empirical studies is that the state of the economy is exogenous with respect to the model variables. This assumption often is empirically implausible. For example, in models that

include log real GDP and express the state of the economy as a function of the unemployment rate, as in Ramey and Zubairy (2018), the unemployment rate changes systematically with the current log-level of real GDP. This renders the state of the economy endogenous with respect to the model variables.

The exogeneity assumption is also implausible when including log real GDP among the endogenous model variables, while measuring expansions and recessions of the economy based on the deviations of log real GDP from a two-sided HP filter trend, which makes the state of the economy dependent on past, current and future realizations of the endogenous model variables (e.g., Auerbach and Gorodnichenko 2013a). Similarly, exogenously imposing NBER business cycle dates, as suggested in Ramey and Zubairy (2018), is inconsistent with the state of the business cycle depending on the response of the model variables to an exogenous shock, since these model variables are correlated with the data underlying the NBER business cycle definition. Defining the state of the business cycle based on one-sided moving average filters, say, by defining a recession as two successive quarters of negative real GDP growth or by defining the business cycle based on the deviation from a one-sided HP filter trend, as in Alloza (2019), does not materially change this result. As in the case of endogenous states, only the impact response can be consistently estimated.

Although we make these points in the context of specific state-dependent VAR data generating processes, they apply more generally to other data generating processes as well. The remainder of the paper is organized as follows. In Section 2, we describe the state-dependent structural model of interest in this paper and define the conditional impulse response function. As is customary in applied work, this response function conditions on the state of the economy in the most recent period, but not on the state of the economy in the current period or in future periods. In Section 3, we define the state-dependent LP estimator of this response function and provide sufficient conditions for its consistency. Section 4 explains why this estimator is not expected to be asymptotically valid in general, when the state of the economy is endogenous with respect to the model variables. We show by simulation that in this case, the state-dependent LP estimator tends to be asymptotically valid in the impact period, but not at longer horizons. We also quantify the large-sample bias of the LP estimator of the response function for several DGPs. The concluding remarks are in section 5.

2 Framework

2.1 The model

Let $z_t \equiv (x_t, y_t')'$ denote an $n \times 1$ vector of strictly stationary time series, where y_t is $k \times 1$ with $k = n - 1$. We consider a structural state-dependent VAR process of the form

$$C_{t-1} z_t = B_{t-1}(L) z_{t-1} + \varepsilon_t, \tag{1}$$

where $\varepsilon_t = (\varepsilon_{1t}, \varepsilon_{2t}')'$ defines the vector of mutually independent structural shocks. Let

$$B_{t-1}(L) = B_{1,t-1} + B_{2,t-1} L + \ldots + B_{p,t-1} L^{p-1},$$

where p denotes the polynomial lag order. For later convenience, we partition $B_{t-1}(L)$ conformably with z_t as

$$B_{t-1}(L) = \begin{pmatrix} B_{11,t-1}(L) & B_{12,t-1}(L) \\ B_{21,t-1}(L) & B_{22,t-1}(L) \end{pmatrix}$$

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

All model coefficients evolve over time depending on the state of the economy. In the simplest case, there are only two states (such as a recession and an expansion). Let

$$\begin{aligned} C_{t-1} &= C_E H_{t-1} + C_R (1 - H_{t-1}), \\ B_{j,t-1} &= B_{jE} H_{t-1} + B_{jR} (1 - H_{t-1}) \text{ for } j = 1, \ldots, p, \end{aligned}$$

where H_{t-1} is a binary stationary time series that takes the value one if the economy is in expansion and 0 otherwise. Unlike in Markov switching models, H_{t-1} is observed.[3]

We are interested in the response of $\{y_{t+h} : h = 0, 1, \ldots, h_{\max}\}$ to a one-time shock in ε_{1t}, conditionally on observing $H_{t-1} = 1$ or $H_{t-1} = 0$. Here, $h_{\max}$ denotes the largest horizon of the impulse response function of interest. To identify this conditional impulse response, we need to impose further restrictions on the model coefficients. In particular, we postulate that

$$C_{t-1} = \begin{pmatrix} 1 & 0 \\ -C_{21,t-1} & C_{22,t-1} \end{pmatrix}, \tag{2}$$

where $C_{21,t-1}$ is $k \times 1$ and $C_{22,t-1}$ is a $k \times k$ non-singular matrix (we set its diagonal elements to 1, which is a standard normalization condition). Under these assumptions, x_t is predetermined with respect to y_t. Note that we do not restrict $C_{22,t-1}$ to be lower triangular, which allows C_{t-1} to be block recursive. Hence, the model is only partially identified in that only the responses to ε_{1t} are identified.

---

[3]Following the applied literature (see e.g. Auerbach and Gorodnichenko (2012, 2013a,b), Alloza (2009)), we index the parameters for the system at time t with the index $t - 1$, e.g. we write C_{t-1}. This reflects the fact that C_{t-1} depends on H_{t-1}, the value of the state indicator at time $t - 1$.

We further restrict x_t to be a directly observed exogenous shock such that

$$B_{j,t-1} = \begin{pmatrix} 0 & 0 \\ B_{j,t-1,21} & B_{j,t-1,22} \end{pmatrix}, \tag{3}$$

for all $j = 1, \ldots, p$. This implies that $x_t = \varepsilon_{1t}$ and corresponds to the empirically relevant case where ε_{1t} is identified using extraneous information outside the model. For instance, it is popular to use a narrative approach to identification when identifying monetary policy shocks (e.g., Romer and Romer (1989), Tenreyro and Thwaites (2016)) and fiscal policy shocks (e.g., Ramey and Shapiro (1998), Ramey (2011), Ramey (2016)).

With these restrictions, the structural model for z_t can be written as

$$\begin{cases} x_t = \varepsilon_{1t} \\ \\ C_{22,t-1} y_t = C_{21,t-1} x_t + B_{21,t-1}(L) x_{t-1} + B_{22,t-1}(L) y_{t-1} + \varepsilon_{2t}. \end{cases} \tag{4}$$

Without further restrictions (such as postulating that $C_{22,t-1}$ is lower triangular), the parameters in the equations for y_t are not identified. However, the fact that ε_{1t} is identified suffices to identify the conditional response function of y_t to a one-time shock in ε_{1t}.

We impose the following standard martingale difference sequence (m.d.s.) assumption on the structural errors ε_t.

Assumption 1 *Let $\mathcal{F}^{t-1} = \sigma(z_{t-1}, H_{t-1}, z_{t-2}, H_{t-2}, \ldots)$. Then, $\varepsilon_t | \mathcal{F}^{t-1} \sim (0, \Sigma)$, where Σ is a diagonal matrix with diagonal elements given by σ_i^2 for $i = 1, \ldots, n$.*

Assumption 1 stipulates that the structural errors ε_t are a martingale difference sequence (m.d.s.) with respect to $\mathcal{F}^{t-1}$, the information set generated by the past realizations of z_t and H_t. This standard assumption implies that ε_t is serially uncorrelated. Since $x_t = \varepsilon_{1t}$, this rules out serial correlation in x_t. We conjecture that these results could be generalized to allow x_t to depend on lags of z_t, as in Alloza et al. (2021) or Goncalves et al. (2021), at the cost of more complicated notation, but our focus is on the empirically most common setting where x_t is a directly observed structural shock. This setting facilitates the derivation of analytical results for the nonlinear response functions and makes our analysis more transparent. Assumption 1 also rules out conditional heteroskedasticity in ε_t by assuming that Σ is constant. This assumption is key to establishing the consistency of state-dependent local projections, as we will explain later. Finally, the assumption that Σ is diagonal implies that the structural errors are mutually uncorrelated, as is standard in the structural VAR literature.

2.2 Conditional impulse response function

Consistent with the empirical literature, our goal is to define the causal effect on y_{t+h} of a one-time shock in ε_{1t}, conditionally on H_{t-1}, the state of the economy at time $t-1$. The fact that our model

is state dependent is reflected in our definition of the conditional IRF. A common approach in the literature on nonlinear impulse response functions (e.g., Gallant, Rossi and Tauchen (1993), Koop, Pesaran and Potter (1996), Potter (2000), Gourieroux and Jasiak (2005), Kilian and Vigfusson (2011), Gonçalves et al. (2021)) is to compare, all else equal, two sample paths for the outcome variables of interest, one where ε_{1t} is subject to a one-time shock at time t and another one where no such shock is present. In a state-dependent model such as ours, this would require fixing ε_{2t} and H_t across the two sample paths. This thought experiment is not realistic when ε_t is correlated with current and future values of H_t because it ignores the possibility that a shock in ε_{1t} may change the states of the economy on impact and in the future.

Hence, we define the conditional IRF more generally as follows. We denote by $\{y_{t+h}\}$ the baseline path that corresponds to the observed data. This is implied by the sequence of structural disturbances and state indicators

$$\mathcal{E} \cup \mathcal{H} = \{\ldots, \varepsilon_{1t-1}, \varepsilon_{1t}, \varepsilon_{1t+1}, \ldots, \varepsilon_{2t-1}, \varepsilon_{2t}, \varepsilon_{2t+1}, \ldots\} \cup \{\ldots, H_{t-1}, H_t, H_{t+1}, \ldots\}.$$

The other sample path is $\{y_{t+h}^*\}$, which is the path implied by an alternative sequence of shocks and state indicators given by

$$\mathcal{E}^* \cup \mathcal{H}^* = \{\ldots, \varepsilon_{1t-1}, \varepsilon_{1t}^*, \varepsilon_{1t+1}, \ldots, \varepsilon_{2t-1}, \varepsilon_{2t}, \varepsilon_{2t+1}, \ldots\} \cup \{\ldots, H_{t-1}, H_t^*, H_{t+1}^*, \ldots\}.$$

With this choice of structural shocks and state indicators, the two sample paths are identical prior to shock in ε_{1t}. At time t, the shock hits ε_{1t}, yielding $\varepsilon_{1t}^* = \varepsilon_{1t} + 1$. All other shocks are kept the same. This choice of perturbation is consistent with the assumption that structural shocks are mutually uncorrelated. However, to accommodate the possibility that a shock to ε_{1t} may change current and future states, we allow for $H_s^* \neq H_s$ for $s \geq t$ when defining $\mathcal{H}^*$. If the states are exogenous (in a sense made precise in the next section), we can set $H_s^* = H_s$ for all s, in which case $\mathcal{H}^* = \mathcal{H}$.

Our definition of conditional IRF is given next.

Definition 1 *The conditional impulse response function of y_{t+h} to a one-time shock of size 1 in ε_{1t} is given by $CIRF_h(H_{t-1}) = E[y_{t+h}^* - y_{t+h}|H_{t-1}]$, for $h = 0, 1, 2, \ldots, h_{\max}$.*

Note that Definition 1 conditions only on H_{t-1}, the state of the economy in the period prior to the shock. Rather than conditioning on the current or future states of the economy, we average them out. This corresponds to the standard approach in applied macroeconomics, where two types of IRFs are often provided, depending on whether the economy was in an expansion or in a recession prior to the shock.

Although the counterfactual y_{t+h}^* is not observed, it may be recovered from the structural model given $\mathcal{E}^*$ and $\mathcal{H}^*$. The values of y_{t+h}^* obtained from solving the model given these sequences is related to the notion of potential outcomes, as defined by Angrist and Kuersteiner (2011) and Angrist, Jordá and Kuersteiner (2016). Further discussion of potential outcomes for time series processes can be found in White (2016) and Rambachan and Shephard (2021).

3   What happens when H_t is exogenous?

3.1   Expression for CIRF

In this section, we present an expression for $CIRF_h(H_{t-1})$ for the state-dependent structural model given in (4). We focus on a counterfactual that treats H_t as exogenous with respect to ε_t such that $\mathcal{H}^* = \mathcal{H}$ in Definition 1. To describe the population IRF, we evaluate the difference between y_{t+h}^* and y_{t+h}. Since C_{t-1} satisfies the identification condition (2), the inverse matrix of C_{t-1} exists and is given by

$$C_{t-1}^{-1} = \begin{pmatrix} 1 & 0 \\ C_{22,t-1}^{-1}C_{21,t-1} & C_{22,t-1}^{-1} \end{pmatrix} \equiv \begin{pmatrix} 1 & 0 \\ C_{t-1}^{21} & C_{t-1}^{22} \end{pmatrix},$$

where for any matrix $\mathcal{A}$, we let $\mathcal{A}^{ij}$ denote the block (i,j) of $\mathcal{A}^{-1}$.

Pre-multiplying (1) by C_{t-1}^{-1} yields

$$z_t = C_{t-1}^{-1}B_{t-1}(L)z_{t-1} + C_{t-1}^{-1}\varepsilon_t,$$

which we rewrite as

$$z_t = A_{t-1}(L)z_{t-1} + \eta_t, \tag{5}$$

where $\eta_t \equiv C_{t-1}^{-1}\varepsilon_t$ and

$$A_{t-1}(L) \equiv C_{t-1}^{-1}B_{t-1}(L) = A_{1,t-1} + A_{2,t-1}L + \ldots + A_{p,t-1}L^{p-1},$$

with $A_{j,t-1} \equiv C_{t-1}^{-1}B_{j,t-1}$. Note that the elements of the first row of $A_{j,t-1}$ are all zero when x_t is equal to ε_{1t} (i.e., when $B_{j,t-1}$ is as in (3)).

The value of y_{t+h} and y_{t+h}^* can be obtained from the companion-form representation of the reduced-form model (5). Let

$$\underset{np\times 1}{Z_t} = (z_t', z_{t-1}', \ldots, z_{t-p+1}')', \quad \underset{np\times 1}{\xi_t} = (\eta_t', 0')',$$

and

$$\underset{np\times np}{A_t} = \begin{pmatrix} A_{1,t-1} & A_{2,t-1} & \cdots & A_{p-1,t-1} & A_{p,t-1} \\ I_n & 0 & \cdots & 0 & 0 \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ 0 & 0 & \cdots & I_n & 0 \end{pmatrix}.$$

7

We can rewrite (5) as

$$Z_t = A_{t-1} Z_{t-1} + \xi_t. \tag{6}$$

To obtain y_t from Z_t, let

$$\underset{k \times np}{\mathbb{S}_k} = \left(\begin{array}{ccc} 0_{k \times 1} & I_k & 0_{k \times n(p-1)} \end{array} \right)$$

denote a $k \times np$ selection matrix (with $k = n - 1$ equal to the number of variables in y_t) which selects the subvector y_t from the vector Z_t. With this notation,

$$y_t = \mathbb{S}_k Z_t,$$

and, more generally, for any h,

$$y_{t+h} = \mathbb{S}_k Z_{t+h}.$$

Note that for $k = 1$ (i.e., for a bivariate system with $n = 2$), $\mathbb{S}_k = e'_{2,2p}$, where $e_{2,2p} = (0, 1, 0')$ is a $2p \times 1$ vector whose only non-zero element is equal to 1 and occurs in position 2. More generally, we let $e_{j,m}$ denote a $m \times 1$ vector with 1 in position j and 0 elsewhere.

Next, we use the companion form (6) to obtain the difference $y^*_{t+h} - y_{t+h}$ for different values of h. Starting with $h = 0$, and noting that the two sample paths coincide up to time $t - 1$, we have that

$$Z_t = A_{t-1} Z_{t-1} + \xi_t \quad \text{and} \quad Z_t^* = A_{t-1} Z_{t-1} + \xi_t^*.$$

Hence,

$$Z_t^* - Z_t = \xi_t^* - \xi_t = \left(\begin{array}{c} \eta_t^* - \eta_t \\ 0_{n(p-1) \times 1} \end{array} \right),$$

where $\eta_t^* - \eta_t = C_{t-1}^{-1} (\varepsilon_t^* - \varepsilon_t)$. Since we only perturb the first element of ε_t, the following decomposition of η_t is useful:

$$\eta_t \equiv C_{t-1}^{-1} \varepsilon_t = \left(\begin{array}{c} 1 \\ C_{t-1}^{21} \end{array} \right) \varepsilon_{1t} + \left(\begin{array}{c} 0 \\ C_{t-1}^{22} \end{array} \right) \varepsilon_{2t} \equiv C_{t-1}^{-1} e_{1n} \varepsilon_{1t} + C_{t-1}^{-1} I_{2:n} \varepsilon_{2t},$$

where $e_{1n} \equiv (1, 0')'$ is $n \times 1$ and $I_{2:n}$ is $k \times n$ and is equal to the $n \times n$ identity matrix with its first column removed:

$$I_{2:n} = \left(\begin{array}{ccc} e_{2n} & \cdots & e_{nn} \end{array} \right).$$

With this notation,

$$\eta_t^* - \eta_t = C_{t-1}^{-1} e_{1n} \underbrace{(\varepsilon_{1t}^* - \varepsilon_{1t})}_{=1} + C_{t-1}^{-1} I_{2:n} \underbrace{(\varepsilon_{2t}^* - \varepsilon_{2t})}_{=0} = C_{t-1}^{-1} e_{1n},$$

given our definition of $\mathcal{E}$ and $\mathcal{E}^*$. It follows that

$$Z_t^* - Z_t = \left(\begin{array}{c} \eta_t^* - \eta_t \\ 0_{n(p-1) \times 1} \end{array} \right) = \left(\begin{array}{c} C_{t-1}^{-1} e_{1n} \\ 0_{n(p-1) \times 1} \end{array} \right) = e_{1p} \otimes C_{t-1}^{-1} e_{1n},$$

8

and, consequently,

$$y_t^* - y_t = \mathbb{S}_k \left(Z_t^* - Z_t \right) = \mathbb{S}_k \left(e_{1p} \otimes C_{t-1}^{-1} e_{1n} \right).$$

The conditional response at $h = 0$ is given by

$$CIRF_0 \left(H_{t-1} \right) = E \left(y_t^* - y_t | H_{t-1} \right) = \mathbb{S}_k \left(e_{1p} \otimes C_{t-1}^{-1} e_{1n} \right), \tag{7}$$

since C_{t-1}^{-1} is known conditionally on H_{t-1}. In particular, the individual impact responses of each variable in y_t can be obtained as

$$CIRF_{0,j} \left(H_{t-1} \right) = E \left(y_{jt}^* - y_{jt} | H_{t-1} \right) = e_{j,np}' \left(e_{1p} \otimes C_{t-1}^{-1} e_{1n} \right),$$

for $j = 2, \ldots, n$. When $n = 2$ (bivariate system), $k = 1$, implying that $\mathbb{S}_k = e_{2,2p}'$.

The expression (7) shows that the conditional impact response can take on two different values, depending on whether $H_{t-1} = 1$ or $H_{t-1} = 0$,

$$CIRF_0 \left(H_{t-1} \right) = \begin{cases} \mathbb{S}_k \left(e_{1p} \otimes C_E^{-1} e_{1n} \right), & \text{if } H_{t-1} = 1 \\ \mathbb{S}_k \left(e_{1p} \otimes C_R^{-1} e_{1n} \right), & \text{if } H_{t-1} = 0, \end{cases}$$

since $C_{t-1} \equiv C_E H_{t-1} + C_R H_{t-1}$. It also shows that only the first column of C_{t-1}^{-1} (i.e., $C_{t-1}^{-1} e_{1n}$) matters for the identification of the conditional impact response.

For $h = 1$, we use the companion form to evaluate first $Z_{t+1}^* - Z_{t+1}$ and then $y_{t+1}^* - y_{t+1}$, as follows. In particular,

$$Z_{t+1} = A_t Z_t + \xi_{t+1},$$

where $\xi_{t+1} = \left(\eta_{t+1}', 0' \right)' = \left(\left(C_t \varepsilon_{t+1} \right)', 0' \right)'$, and where A_t and C_t depend on H_t. Similarly,

$$Z_{t+1}^* = A_t^* Z_t^* + \xi_{t+1}^*,$$

where $\xi_{t+1}^* = \left(\eta_{t+1}^{*\prime}, 0' \right)' = \left(\left(C_t^* \varepsilon_{t+1}^* \right)', 0' \right)'$ and A_t^* and C_t^* depend on H_t^*. Given our choice of structural shocks, $\varepsilon_{t+1}^* = \varepsilon_{t+1}$. Moreover, under the assumption that $H_t^* = H_t$, $A_t^* = A_t$ and $C_t^* = C_t$. This implies that

$$Z_{t+1}^* - Z_{t+1} = A_t \left(Z_t^* - Z_t \right) = A_t \left(e_{1p} \otimes C_{t-1}^{-1} e_{1n} \right),$$

given that $Z_t^* - Z_t = e_{1p} \otimes C_{t-1}^{-1} e_{1n}$. Thus, we have that

$$y_{t+1}^* - y_{t+1} = \mathbb{S}_k \left(Z_{t+1}^* - Z_{t+1} \right) = \mathbb{S}_k A_t \left(e_{1p} \otimes C_{t-1}^{-1} e_{1n} \right),$$

implying that

$$CIRF_1 \left(H_{t-1} \right) = \mathbb{S}_k E \left(A_t | H_{t-1} \right) \left(e_{1p} \otimes C_{t-1}^{-1} e_{1n} \right).$$

This expression generalizes to other values of h as follows.

9

Proposition 3.1 *Let $\mathcal{E}$, $\mathcal{H}$, $\mathcal{E}^*$ and $\mathcal{H}^*$ be as defined in Section 2.2. If $\mathcal{H}^* = \mathcal{H}$, the impulse response of y_t to a one-time shock in ε_{1t}, conditional on H_{t-1}, is*

$$CIRF_0(H_{t-1}) \equiv E\left(y_t^* - y_t | H_{t-1}\right) = \mathbb{S}_k\left(e_{1p} \otimes C_{t-1}^{-1} e_{1n}\right),$$

and for any $h \geq 1$,

$$CIRF_h(H_{t-1}) \equiv E\left(y_{t+h}^* - y_{t+h} | H_{t-1}\right) = \mathbb{S}_k E\left(A_{t+h-1} A_{t+h-2} \ldots A_t | H_{t-1}\right)\left(e_{1p} \otimes C_{t-1}^{-1} e_{1n}\right).$$

The impulse response function defined in Proposition 3.1 conditions only on H_{t-1}, the state of the economy in the period prior to the shock. It does not condition on the current or future states of the economy. Nor does it condition on the history of states prior to $t-1$ or on the histories of z_t.

To identify $CIRF_h(H_{t-1})$, we need to identify the first column of C_{t-1}^{-1}, $C_{t-1}^{-1} e_{1n}$, as well as the coefficients that enter the matrices A_{t+h-1} through A_t. Given that these matrices are linear in the state indicators, identification can be achieved from the reduced-form model (5). Even when the model is fully identified, evaluating $E\left(A_{t+h-1} A_{t+h-2} \ldots A_t | H_{t-1}\right)$ is challenging and requires knowledge of the conditional density of $H_{t+h-1}, \ldots, H_t$, given H_{t-1}. Local projections are a much simpler alternative. In the next section, we provide a set of sufficient conditions under which local projections are consistent.

3.2 Local projections

A state-dependent LP regression is a direct regression of y_{t+h} onto a constant, x_t and Z_{t-1}, each interacted with H_{t-1} and $1 - H_{t-1}$. The slope coefficients associated with $x_t H_{t-1}$ are usually interpreted as the CIRF of y_{t+h}, conditionally on $H_{t-1} = 1$, whereas the slope coefficients associated with $x_t(1 - H_{t-1})$ describe the CIRF of y_{t+h} when we condition on $H_{t-1} = 0$. The goal of this section is to provide a set of regularity conditions under which this interpretation is asymptotically valid.

Let $W_{t-1} \equiv (1, Z_{t-1}')'$ denote an $(np+1) \times 1$ vector of control variables which include a constant and p lags of z_t. A state-dependent LP for identifying the causal effect on y_{t+h} of a one-time shock in $x_t = \varepsilon_{1t}$ can be written as

$$y_{t+h} = b_{E,h} x_t H_{t-1} + \Pi_{E,h} W_{t-1} H_{t-1} + b_{R,h} x_t (1 - H_{t-1}) + \Pi_{R,h} W_{t-1}(1 - H_{t-1}) + v_{t+h}, \quad (8)$$

where the $k \times 1$ vectors $b_{E,h}$ and $b_{R,h}$ contain the main parameters of interest. In particular, $b_{E,h}$ is interpreted as the CIRF of y_{t+h} when $H_{t-1} = 1$, whereas $b_{R,h}$ contains the CIRF of y_{t+h}, conditionally on $H_{t-1} = 0$. The matrices $\Pi_{E,h}$ and $\Pi_{R,h}$ are of size $k \times (np+1)$; each row contains the constant and the slope coefficients associated with Z_{t-1} for the LP regression of each variable in y_{t+h}. The LP regression for variable $y_{j,t+h}$ is

$$y_{j,t+h} = b_{E,j,h} x_t H_{t-1} + \pi'_{E,j,h} W_{t-1} H_{t-1} + b_{R,j,h} x_t (1 - H_{t-1}) + \pi'_{R,j,h} W_{t-1}(1 - H_{t-1}) + v_{j,t+h}, \quad (9)$$

where $j = 2, \ldots, n$. The scalar coefficients $b_{E,j,h}$ and $b_{R,j,h}$ are the $(j-1)^{th}$ elements of $b_{E,h}$ and $b_{R,h}$, respectively. Similarly, $\pi'_{E,j,h}$ and $\pi'_{R,j,h}$ are the corresponding rows of $\Pi_{E,h}$ and $\Pi_{R,h}$.

Since H_t is observed, the coefficients in the multivariate state-dependent LP regression (8) can be obtained by running a multivariate LS regression of y_{t+h} onto $x_t H_{t-1}$, $W_{t-1} H_{t-1}$, $x_t (1 - H_{t-1})$ and $W_{t-1} (1 - H_{t-1})$. Note that this is equivalent to running a regression of $y_{j,t+h}$ onto $x_t H_{t-1}$, $W_{t-1} H_{t-1}$, $x_t (1 - H_{t-1})$ and $W_{t-1} (1 - H_{t-1})$, for each $j = 2, \ldots, n$. Put differently, the multivariate LS regression (8) is equivalent to the k univariate OLS regressions (9), equation-by-equation.

Let $\hat{b}_{E,h}$ and $\hat{b}_{R,h}$ denote the LS estimators of $b_{E,h}$ and $b_{R,h}$ in (8) based on a sample of size T given by $\{y_{t+h}, x_t, Z_{t-1}, H_{t-1} : t = 1, \ldots, T\}$. We can estimate each of these vectors separately, by restricting the sample to $H_{t-1} = 1$ and $H_{t-1} = 0$, respectively. For instance, $\hat{b}_{E,h}$ can be obtained from a regression of y_{t+h} on $x_t H_{t-1}$ and $W_{t-1} H_{t-1}$ (omitting $x_t (1 - H_{t-1})$ and $W_{t-1} (1 - H_{t-1})$ in the regression). This follows because the $H_{t-1} (1 - H_{t-1}) = 0$ for all t. Similarly, we can obtain $\hat{b}_{R,h}$ from a regression of y_{t+h} on $x_t (1 - H_{t-1})$ and $W_{t-1} (1 - H_{t-1})$ (omitting $x_t H_{t-1}$ and $W_{t-1} H_{t-1}$ in this regression).

As it turns out, Assumption 1 suffices to show the consistency of $\hat{b}_{E,h}$ and $\hat{b}_{R,h}$ when $h = 0$. To identify the CIRF at horizons $h = 1, \ldots, h_{\max}$ we add the following assumption.

Assumption 2 *Let $h_{\max} \geq 1$ denote the maximum horizon of the response function of interest. Then, for $h = 1, \ldots, h_{\max}$,*

(a) $E\left(\varepsilon_t | H_{t+h-1}, \ldots, H_t, \mathcal{F}^{t-1}\right) = E\left(\varepsilon_t | \mathcal{F}^{t-1}\right).$

(b) $E\left(\varepsilon_t \varepsilon'_t | H_{t+h-1}, \ldots, H_t, \mathcal{F}^{t-1}\right) = E\left(\varepsilon_t \varepsilon'_t | \mathcal{F}^{t-1}\right).$

Assumption 2 characterizes the relationship between the structural shocks $\{\varepsilon_t\}$ and the state indicators $\{H_t\}$. This condition is crucial for proving the validity of state-dependent local projections. A sufficient condition for Assumption 2 is to assume that $\{H_t\}$ is fully independent of $\{\varepsilon_s\}$. This assumption is satisfied if we construct H_t on the basis of variables not contained in z_t that are independent of the structural errors ε_t. Assumption 2 is a milder assumption than full independence between ε_t and H_t. It only requires the conditional first two moments of ε_t to be independent of $\{H_t, H_{t+1}, \ldots, H_{t+h-1}\}$, conditionally on $\mathcal{F}^{t-1}$, where $h \leq h_{\max}$. This allows for the possibility that H_t is obtained as a function of past values of z_t. How many lags of z_t can be included in H_t depends on the value of $h_{\max}$. For $h_{\max} = 1$, H_t can depend on z_{t-1} (and previous lags of z_{t-1}), but for $h_{\max} = 2$, H_t can depend only on z_{t-2} (or further lags of z_{t-2}). As $h_{\max}$ increases, the set of lags used to construct H_t shrinks. In the limit, if we are interested in the entire impulse response function, H_t

cannot be chosen as a function of $\{z_t\}$. We will further illustrate the content of Assumption 2 in the next section when we specialize H_t to be a deterministic function of z_t.

Under Assumptions 1 and 2, we can prove the following result.

Proposition 3.2 *Under Assumptions 1 and 2, as $T \to \infty$, for any $h = 0, 1, \ldots, h_{\max}$,*

$$\hat{b}_{E,h} \longrightarrow_p CIRF_h\left(H_{t-1} = 1\right) \ \text{and} \ \hat{b}_{R,h} \longrightarrow_p CIRF_h\left(H_{t-1} = 0\right),$$

where $CIRF_h\left(H_{t-1} = i\right) = \mathbb{S}_k E\left(A_{t+h-1}A_{t+h-2}\ldots A_t\left(e_{1p} \otimes C_{t-1}^{-1}e_{1n}\right) | H_{t-1} = i\right)$ for $i = 1, 0$.

The proof of Proposition 3.2 is in the Appendix. Proposition 3.2 shows that the LP regression (8) identifies the conditional IRF defined in Proposition 3.1. The latter corresponds to the CIRF of y_{t+h} derived under the counterfactual experiment that sets $\mathcal{H}^* = \mathcal{H}$, i.e., we assume that the shock of ε_{1t} does not change the state of the economy on impact nor in the future. This is consistent with Assumption 2, which imposes mean-independence conditions on ε_t and $H_{t+h-1}, \ldots, H_t$, conditionally on $\mathcal{F}^{t-1}$.

The model equation for y_{t+h} implied by the companion form representation of the structural model may be used to heuristically understand why the state-dependent LP works without Assumption 2 when $h = 0$, but not otherwise. More specifically, note that for any h, we can write

$$y_{t+h} = \mathbb{S}_k Z_{t+h},$$

where Z_{t+h} is obtained from (6). Consider first $h = 0$. Then,

$$Z_t = A_{t-1}Z_{t-1} + \xi_t,$$

where

$$\xi_t = \begin{pmatrix} \eta_t \\ 0 \end{pmatrix} = \begin{pmatrix} C_{t-1}^{-1}e_{1n}\varepsilon_{1t} + C_{t-1}^{-1}I_{2:n}\varepsilon_{2t} \\ 0 \end{pmatrix} = (e_{1p} \otimes C_{t-1}^{-1}e_{1n})\varepsilon_{1t} + e_{1p} \otimes C_{t-1}^{-1}I_{2:n}\varepsilon_{2t},$$

given that $\eta_t = C_{t-1}^{-1}\varepsilon_t$ and $\varepsilon_t = C_{t-1}^{-1}e_{1n}\varepsilon_{1t} + C_{t-1}^{-1}I_{2:n}\varepsilon_{2t}$, where e_{1n} and $I_{2:n}$ are as defined in Section 3.1. Hence,

$$y_t = \mathbb{S}_k Z_t = \mathbb{S}_k(e_{1p} \otimes C_{t-1}^{-1}e_{1n})\varepsilon_{1t} + \mathbb{S}_k A_{t-1}Z_{t-1} + \mathbb{S}_k(e_{1p} \otimes C_{t-1}^{-1}I_{2:n}\varepsilon_{2t}).$$

If we condition on $H_{t-1} = 1$, this equation implies that

$$y_t = \mathbb{S}_k Z_t = \underbrace{\mathbb{S}_k(e_{1p} \otimes C_E^{-1}e_{1n})\varepsilon_{1t}}_{=CIRF_0(H_{t-1}=1)} + \mathbb{S}_k A_E Z_{t-1} + \underbrace{\mathbb{S}_k(e_{1p} \otimes C_E^{-1}I_{2:n}\varepsilon_{2t})}_{\equiv u_t},$$

12

where the coefficient associated with ε_{1t} is the conditional IRF when $H_{t-1} = 1$. A linear regression of y_t on ε_{1t} and Z_{t-1} which conditions on $H_{t-1} = 1$ estimates consistently this coefficient if conditionally on $H_{t-1} = 1$ the regressors are uncorrelated with the error term given by u_t. This orthogonality condition holds under the m.d.s. assumption on ε_t (i.e., Assumption 1) without further restrictions on H_t.

For $h = 1$, the model equation for y_{t+1} is now

$$
\begin{aligned}
y_{t+1} &= \mathbb{S}_k Z_{t+1} = \mathbb{S}_k(A_t Z_t + \xi_{t+1}) \\
&= \mathbb{S}_k(A_t(A_{t-1} Z_{t-1} + \xi_t) + \xi_{t+1}) \\
&= \mathbb{S}_k A_t \xi_t + \mathbb{S}_k A_t A_{t-1} Z_{t-1} + \mathbb{S}_k \xi_{t+1},
\end{aligned}
$$

where $\xi_s = (e_{1p} \otimes C_{s-1}^{-1} e_{1n})\varepsilon_{1s} + e_{1p} \otimes C_{s-1}^{-1} I_{2:n}\varepsilon_{2s}$ for $s = t, t+1$. Hence, conditionally on $H_{t-1} = 1$,

$$
y_{t+1} = \underbrace{[\mathbb{S}_k A_t(e_{1p} \otimes C_E^{-1} e_{1n})]}_{\equiv \psi_t}\varepsilon_{1t} + \mathbb{S}_k A_t A_E Z_{t-1} + \underbrace{\mathbb{S}_k A_t(e_{1p} \otimes C_E^{-1} I_{2:n}\varepsilon_{2t}) + \mathbb{S}_k \xi_{t+1}}_{= u_{t+1}}. \tag{10}
$$

The coefficients in this equation depend on H_t (through A_t), conditionally on $H_{t-1} = 1$. This implies that a linear projection of y_{t+1} on ε_{1t} and Z_{t-1} which conditions on $H_{t-1} = 1$ estimates a conditional average of these state-dependent parameters. In particular, the estimand of the coefficient associated with ε_{1t} in this restricted OLS regression is $E(\psi_t | H_{t-1} = 1) \equiv CIRF_1(H_{t-1} = 1)$ provided the error term u_{t+1} is orthogonal to ε_{1t}, conditionally on $H_{t-1} = 1$. As u_{t+1} depends on H_t, this requires that H_t and ε_{1t} are uncorrelated, conditionally on $H_{t-1} = 1$.

For general values of h, we can write y_{t+h} as a function of ε_{1t} and Z_{t-1} and an error term that depends on $H_{t+h-1}, \ldots, H_{t-1}$. Conditionally on H_{t-1}, this is a state-dependent equation as it depends on $H_{t+h-1}, \ldots, H_t$. A linear local projection of y_{t+h} on ε_{1t} and Z_{t-1} which conditions only on H_{t-1} recovers the conditional IRF derived in Proposition 3.1 provided the error term is orthogonal to ε_{1t}, conditionally on H_{t-1}. Since this error depends on $H_{t+h-1}, \ldots, H_t$, we require that ε_{1t} (and more generally ε_t) be independent of $H_{t+h-1}, \ldots, H_t$, conditionally on H_{t-1}. Assumption 2 formalizes this independence condition. Because local projections are least squares estimates, it is natural that only first- and second-order conditional mean independence conditions on ε_t are required.

Note that the asymptotic validity of the state-dependent LP estimator does not depend on the full identification of the structural model parameters. The crucial condition is that the shock of interest ε_{1t} is identified. This shocks equals x_t in our setting, but one could instead allow it to be a function of lagged values of z_t. With this condition, and under Assumptions 1 and 2, the LP identifies the correct conditional IRF even though the contemporaneous state-dependent matrix C_{t-1} is only block recursive. This result is as expected. Proposition 3.1 shows that the conditional IRF depends on the

13

conditional expectation of a function of $A_{t+h-1}, \ldots, A_t$ and the first column of C_{t-1}^{-1}, which can all be identified from the reduced-form model (5).

4 What happens when H_t is endogenous?

In this section, we investigate the properties of state-dependent LPs when H_t does not satisfy Assumption 2. In particular, we consider the case when H_t depends on current values of the outcome variables y_t. To simplify the exposition and make the arguments clearer, we consider the special case of a bivariate structural model for $z_t \equiv (x_t, y_t)'$, where x_t is a directly observed shock and y_t has limited dynamics:

$$
\begin{cases}
x_t = \varepsilon_{1t}, \\
\\
y_t = \beta_{t-1} x_t + \gamma_{t-1} y_{t-1} + \varepsilon_{2t}.
\end{cases}
\tag{11}
$$

In terms of our previous notation, $n = 2$, $k = 1$, $C_{22,t-1} = 1$, $C_{21,t-1} = \beta_{t-1}$, $C_{21,t-1}(L) = 0$ and $B_{22,t-1}(L) = \gamma_{t-1}$. The state-dependent parameters β_{t-1} and γ_{t-1} depend on H_{t-1} as before. For instance, $\beta_{t-1} = \beta_E H_{t-1} + \beta_R (1 - H_{t-1})$. Crucially, we now endogenize H_t with respect to the structural shocks ε_t. In particular, we let $H_t = 1(y_t > 0)$. Given that the structural model sets the time t coefficients as a function of H_{t-1}, as is typically assumed in the empirical literature, setting $H_t = 1(y_t > 0)$ implies that β_{t-1} and γ_{t-1} are a function of y_{t-1}. A generalization of this scenario is to allow H_t to depend on current and lagged values of y_t, as in Alloza's (2019) study of the impact of a fiscal policy shock on output. Alloza sets $H_t = 1(y_t > 0 \text{ and } y_{t-1} > 0)$.

Next, we discuss the implications of this choice of H_t for the validity of the LP estimator. First, we show that the conditional IRF of interest is no longer given by the formula derived in Proposition 3.1. Next, we show that the LP estimand is not the same as the one derived in Proposition 3.2. Finally, because an analytical characterization of this estimand is infeasible, we numerically illustrate the magnitude of the asymptotic bias of the state-dependent LP estimator in this context.

4.1 Conditional IRF when H_t is endogenous

The goal is to obtain the response of y_{t+h} to a shock of size 1 in ε_{1t}. We follow the same approach as in Section 2.2 and compare the value of y_{t+h} with a counterfactual value y_{t+h}^* which corresponds to what we would have been observed if we had perturbed ε_{1t} by 1 without changing any of the other inputs to the system. Note that when H_t depends on y_t, the current and future values of H_t cannot be kept constant across these two sample paths. Thus, the counterfactual experiment that sets $\mathcal{H}^* = \mathcal{H}$ is not consistent with this choice of H_t. We need to account for the impact of the shock in ε_{1t} on the current and future values of the states of the economy such that $H_s^* \neq H_s$ for $s \geq t$.

14

Consider $h = 0$. Following the same steps as in Section 3.1, we can show that

$$y_t^* - y_t = \beta_{t-1}\left(x_t^* - x_t\right) = \beta_{t-1} \equiv \beta_E H_{t-1} + \beta_R\left(1 - H_{t-1}\right),$$

since $x_t^* = x_t + 1$, and importantly, $\beta_{t-1}^* = \beta_{t-1}$ and $\gamma_{t-1}^* = \gamma_{t-1}$. This follows because β_{t-1}^* and γ_{t-1}^* are defined as β_{t-1} and γ_{t-1}, but depend on $H_{t-1}^* = 1\left(y_{t-1}^* > 0\right) = H_{t-1} = 1\left(y_{t-1} > 0\right)$ since $y_{t-1}^* = y_{t-1}$. This implies that the conditional impact response defined in Proposition 3.1 is

$$CIRF_0\left(H_{t-1}\right) = \beta_{t-1} = \begin{cases} \beta_E & \text{if } H_{t-1} = 1 \\ \beta_R & \text{if } H_{t-1} = 0. \end{cases}$$

To see that this expression is a special case of Proposition 3.1, note that when $k = 1$, $\mathbb{S}_k = (0, 1)$, $e_{1p} = 1$ and $C_{t-1}^{-1}e_{1n} = \left(1, \beta_{t-1}\right)'$.

For $h = 1$, an important difference emerges. Now, β_t^* and γ_t^* depend on $H_t^* = 1\left(y_t^* > 0\right)$. Since y_t^* is not equal to y_t, we cannot set $\beta_t^* = \beta_t$ and $\gamma_t^* = \gamma_t$ when defining the counterfactual value of y_{t+1}. In particular, we now have

$$y_{t+1}^* = \beta_t^* x_{t+1}^* + \gamma_t^* y_t^* + \varepsilon_{2t+1}^* = \beta_t^* x_{t+1} + \gamma_t^* y_t^* + \varepsilon_{2t+1},$$

where the second equality follows because $\varepsilon_{2t+1}^* = \varepsilon_{2t+1}$, and $x_{t+1}^* = \varepsilon_{1t+1}^* = \varepsilon_{1t+1}$. The difference between y_{t+1}^* and y_{t+1} is

$$y_{t+1}^* - y_{t+1} = \left(\beta_t^* - \beta_t\right)x_{t+1} + \left(\gamma_t^* - \gamma_t\right)y_t^* + \gamma_t\left(y_t^* - y_t\right),$$

where $y_t^* - y_t = \beta_{t-1}$. The fact that H_t is a function of y_t implies that a shock at time t in ε_{1t} has an impact on y_t and hence an impact on the state-dependent coefficients β_t^* and γ_t^*. This explains the presence of the two extra terms in $y_{t+1}^* - y_{t+1}$.

The conditional impulse response at horizon $h = 1$ is the expectation of this difference, conditionally on H_{t-1}:

$$\begin{aligned} CIRF_1\left(H_{t-1}\right) &= E\left(y_{t+1}^* - y_{t+1}|H_{t-1}\right) \\ &= \underbrace{E[\left(\beta_t^* - \beta_t\right)x_{t+1}|H_{t-1}] + E[\left(\gamma_t^* - \gamma_t\right)y_t^*|H_{t-1}]}_{\text{Indirect effect}} + \underbrace{E\left(\gamma_t|H_{t-1}\right)\beta_{t-1}}_{\text{Direct effect}}. \end{aligned}$$

The second term corresponds to the CIRF derived in Proposition 3.1 under the assumption that the counterfactual value of H_t^* is equal to the observed value H_t, i.e., $\mathcal{H} = \mathcal{H}^*$. We interpret this as the direct effect as it captures the effect on y_{t+1} of the shock in ε_{1t} assuming that there is no change in the state H_t (and therefore no change in β_t and γ_t). The first term accounts for the effect on y_{t+1} that occurs because H_t has changed. When H_t is exogenous, this indirect effect is zero, but not otherwise.

15

Note that we can use the model equations to express the indirect effect as a function of observables. In particular, it can be shown that[4]

$$\text{Indirect effect} = (\gamma_E - \gamma_R) E[1\left(y_t + \beta_{t-1} > 0\right) - 1\left(y_t > 0\right)\left(y_t + \beta_{t-1}\right)|H_{t-1}].$$

The decomposition of the conditional response into a direct effect and an indirect effect generalizes to larger values of h. For instance, for $h = 2$,

$$
\begin{aligned}
CIRF_2\left(H_{t-1}\right) &= E\left(y_{t+2}^* - y_{t+2}|H_{t-1}\right) \\
&= \underbrace{E[\left(\beta_{t+1}^* - \beta_{t+1}\right) x_{t+2} + \left(\gamma_{t+1}^* - \gamma_{t+1}\right) y_{t+1}^*|H_{t-1}]}_{\text{Indirect effect due to time } t+1 \text{ change in parameters}} \\
&\quad + \underbrace{E[\gamma_{t+1}\left(\beta_t^* - \beta_t\right) x_{t+1} + \gamma_{t+1}\left(\gamma_t^* - \gamma_t\right) y_t^*|H_{t-1}]}_{\text{Indirect effect due to time } t \text{ change in parameters}} \\
&\quad + \underbrace{E[\gamma_{t+1}\gamma_t\beta_{t-1}|H_{t-1}]}_{\text{Direct effect if no change in parameters}} ,
\end{aligned}
$$

where the last term is the CIRF at $h = 2$ derived in Proposition 3.1 under the assumption that $\mathcal{H}^* = \mathcal{H}$. This term captures the direct effect for $h = 2$. The indirect effect is represented by the first two terms. Characterizing these expectations analytically becomes intractable, even under strong assumptions about the conditional distribution.

The overall message is that when H_t depends on y_t, the conditional IRF is no longer the same as the one defined in Proposition 3.1. It now contains additional terms that capture the indirect effect of the shock in ε_{1t} on y_{t+h} that operates through the effect of the shock on the transition path of H_t through H_{t+h-1}.

4.2   Asymptotic bias in the LP estimator when H_t is endogenous

We now investigate the effect of endogenizing H_t on the estimand of a state-dependent LP. For simplicity, we again focus on the simple bivariate model considered in (11) with $H_t = 1\left(y_t < 0\right)$. The state-dependent LP in this context is given by

$$y_{t+h} = b_{E,h}x_t H_{t-1} + \pi'_{E,h}W_{t-1}H_{t-1} + b_{R,h}x_t(1 - H_{t-1}) + \pi'_{R,h}W_{t-1}(1 - H_{t-1}) + v_{t+h}, \tag{12}$$

where $W'_{t-1} = (1, y_{t-1})$.

Proceeding as in Section 3.2, we can show that the LP estimate of $b_{E,h}$ is converges in probability to

$$b_{E,h} = \frac{E\left(x_t H_{t-1} y_{t+h}\right)}{E\left(x_t^2 H_{t-1}\right)} = \frac{E\left(x_t y_{t+h}|H_{t-1} = 1\right)}{E\left(x_t^2|H_{t-1} = 1\right)}.$$

---

[4]Further simplifying this expression involves computing truncated moments of $y_t + \beta_{t-1}$, conditionally on H_{t-1}. This can be done for $h = 1$ under parametric assumptions on the conditional distribution of y_t given H_{t-1}. However, this approach quickly becomes intractable as we increase the value of h. A simpler approach is to use numerical methods to approximate this expectation, which is the approach we use below to evaluate the asymptotic bias of the LP estimates.

The LS estimate of $b_{E,h}$ in (12) can be obtained from an OLS regression of y_{t+h} on x_t, conditionally on $H_{t-1} = 1$. This result does not depend on whether H_t is exogenous or endogenous. Rather it follows from the fact that $E(x_t H_{t-1} W_{t-1}) = 0$ by the m.d.s. assumption on ε_t (i.e., by Assumption 1).

When $h = 0$, one can easily show that $b_{E,0} = \beta_E = CIRF_0(H_{t-1})$. Thus, the state-dependent LP recovers the correct impact conditional IRF even when H_t depends on y_t. This follows because conditionally on $H_{t-1} = 1$, the structural model equation for y_t is

$$y_t = \beta_E x_t + \gamma_E y_{t-1} + \varepsilon_{2t},$$

and a linear local projection of y_t onto x_t and y_{t-1} that conditions on observations with $H_{t-1} = 1$ recovers β_E provided ε_{2t} is orthogonal to x_t and y_{t-1} (conditionally on $H_{t-1} = 1$). Assumption 1 alone suffices for this result.

When $h = 1$, evaluating $b_{E,1} = p\lim \hat{b}_{E,1}$ is more challenging when H_t is endogenous. To see this, note that conditionally on $H_{t-1} = 1$, the equation for y_{t+1} can be described as[5]

$$y_{t+1} = \gamma_t \beta_E x_t + \gamma_t \gamma_E y_{t-1} + u_{t+1},$$

where

$$u_{t+1} = \beta_t \varepsilon_{1t+1} + \gamma_t \varepsilon_{2t} + \varepsilon_{2t+1}.$$

Thus, conditionally on H_{t-1}, the model for y_{t+1} is state-dependent because it depends on H_t through the parameters γ_t and β_t.

As explained in Section 3.2, under exogeneity of H_t (as stated in Assumption 2), the LP estimand of the slope coefficient associated with $x_t H_{t-1}$ is equal to $E(\gamma_t | H_{t-1} = 1)\beta_E$. This is the correct CIRF when H_t satisfies Assumption 2 and it is the direct effect contained in the population CIRF when H_t is endogenous. The LP estimand of this coefficient is not necessarily equal to $E(\gamma_t | H_{t-1} = 1)\beta_E$, when $H_t = 1(y_t > 0)$. The main reason is that Assumption 2 is not satisfied. For this choice, H_t and ε_t are no longer mean independent, conditionally on H_{t-1}. For instance,

$$E\left(\varepsilon_{1t} | H_t, \mathcal{F}^{t-1}\right) = E\left(\varepsilon_{1t} | y_t > 0, \mathcal{F}^{t-1}\right) = E\left(\varepsilon_{1t} | \beta_{t-1}\varepsilon_{1t} + \varepsilon_{2t} > -\gamma_{t-1}y_{t-1}, \mathcal{F}^{t-1}\right),$$

which is a conditional truncated moment of ε_{1t}. Although under Assumption 1, ε_{1t} has mean zero conditionally on $\mathcal{F}^{t-1}$, adding information on H_t in the form of the restriction $\beta_{t-1}\varepsilon_{1t} + \varepsilon_{2t} > -\gamma_{t-1}y_{t-1}$ makes this mean not zero. The same is true for the second conditional moment of ε_t. For instance, $E\left(\varepsilon_{1t}^2 | H_t, \mathcal{F}^{t-1}\right)$ is no longer equal to $\sigma_1^2 = E\left(\varepsilon_{1t} | \mathcal{F}^{t-1}\right)$.

Next, we evaluate the limit of the LP estimator by simulations and compare it with the population CIRF and its decomposition into a direct and indirect effect.

---

[5] Note that this is the analogue of (10) for the bivariate model (11).

4.3 Does the LP Estimator Converge to the Population Response?

The literature has taken for granted that the LP estimator asymptotically recovers the population response, when the DGP is a state-dependent structural VAR model (see, e.g., Auerbach and Gorodnichenko 2013a; Alloza 2019). Our analysis shows that this conclusion is indeed correct, when H_t is exogenous with respect to z_t. In the empirically more relevant case when H_t directly or indirectly depends on y_t, however, our sufficient conditions ensure the asymptotic validity of the state-dependent LP estimator only for the impact response. Although our theoretical analysis does not formally prove that the LP estimator of the response function is invalid at other horizons in this case, there is no presumption that it does recover the population response function. In this section, we explore this question based on several stylized bivariate DGPs and show that the LP estimator of the response function indeed appears to be inconsistent when H_t is endogenous. We consider three DGPs.

Let

$$
\begin{aligned}
x_t &= \varepsilon_{1t} \\
y_t &= \beta_{t-1} x_t + \gamma_{t-1} y_{t-1} + \varepsilon_{2t}
\end{aligned}
\tag{13}
$$

where

$$
\begin{aligned}
\beta_{t-1} &= \beta_E H_{t-1} + \beta_R (1 - H_{t-1}), \\
\gamma_{t-1} &= \gamma_E H_{t-1} + \gamma_R (1 - H_{t-1}),
\end{aligned}
\tag{14}
$$

$\varepsilon_t = (\varepsilon_{1t}, \varepsilon_{2t})' \sim N(0, I_2)$ and H_t is an indicator function that determines the state of the economy. When $H_t = 1$ the economy is in an expansion, E, and when $H_t = 0$ the economy is in recession, R. In DGP1, $H_t = F(q_t) \equiv 1(q_t > 0)$, where q_t follows an exogenous process

$$
q_t = 0.6 q_{t-1} + u_t,
$$

and $u_t \sim N(0, 1)$ is independent of ε_t. DGP2 and DGP3 differ from DGP1 in that the indicator function is given by $H_t = F(y_t) \equiv 1(y_t > 0)$, so that the state of the economy is determined endogenously. DGP2 sets $\beta_E = 2.4$, $\beta_R = 1.6$, $\gamma_E = 0.7$ and $\gamma_R = 0.1$, whereas DGP1 sets $\beta_E = 2.5$, $\beta_R = 3.5$, $\gamma_E = 0.9$ and $\gamma_R = -0.1$. In all DGPs the intercepts have been normalized to 0.

We consider the effect on y_{t+h} of a shock of size 1 in ε_{1t}. The conditional impulse response function is evaluated as $E(y_{t+h}^* - y_{t+h} | H_{t-1})$, whereas the LP estimands are evaluated as (12). We also compute the direct effect (given by the formula in Proposition 3.1) and the indirect effect (which we obtain as the difference between $E(y_{t+h}^* - y_{t+h} | H_{t-1})$ and the direct effect). The number of draws used to compute all these conditional expectations is equal to 50 millions. In addition to reporting

18

results of the effect of the shock on the level of y_{t+h}, we also compute the cumulative effects. These are obtaining by summing the individual CIRFs and the corresponding LP coefficients. For instance, the cumulative CIRF at horizon $h = 1$ equals $\sum_{h=0}^{1} E\left(y_{t+h}^{*} - y_{t+h}|H_{t-1}\right)$ and the LP estimand is $\sum_{h=0}^{1} b_{i,h}$ with $b_{i,h} = b_{E,h}$ if $H_{t-1} = 1$ and $b_{i,h} = b_{R,h}$ if $H_{t-1} = 0$.

Figures 1 and 2 contain the results when H_t is exogenous (DGP1) whereas Figures 3 through 6 contain results for the endogenous case (DGP2 and DGP3). Starting with DGP1, Figure 1 shows that the CIRF is equal to the LP estimand at all horizons. In addition, the indirect effect is zero, making the CIRF equal to the direct effect. This is consistent with our theoretical results (cf. Proposition 3.2). Because the LP estimand coincides with the CIRF for y_{t+h}, LP also recovers the cumulative effect, as shown by Figure 2.

Figures 3 through 6 show that these results change when $H_t = 1\left(y_t > 0\right)$, making H_t endogenous with respect to ε_{1t}. These figures show that the LP estimands no longer coincide with the population response function of interest (both in levels and as a sum). In particular, although the impact effect is well recovered by the state-dependent LP, this is no longer true at intermediate values of h (as h increases, the CIRF and the corresponding LP estimand both tend to zero, making the bias disappear; this is no longer true for the cumulative LP bias, which remains non-zero for all values of h). The decomposition of the CIRF into the direct and indirect effect shows that the LP estimand follows closely the direct effect, missing the indirect effect. Thus, the bias in the LP estimator is close to the indirect effect in these simulations.

The size of the asymptotic bias depends on the parameter values we choose. In this simple model, the bias increases with $\gamma_E - \gamma_R$, implying that it is larger in absolute value in DGP3 than in DGP2 (compare Figures 3 and 5 for the CIRF and Figures 4 and 6 for the cumulative CIRF). Although Figures 3 and 5 seem to suggest that the bias of the LP estimator is modest relative to the value of the impulse response function, this bias is significant when measured as a function of the population CIRF of interest. For instance, for DGP1, the bias of LP relative to the CIRF is equal -10%, -13%, -14%, and -15% for $h = 1, 2, 3, 4$, when in expansions, and $-20\%, -20\%, -20\%, -21\%$ when in recessions. These numbers imply a relative bias for the cumulative response function that varies between -4% and -7% in expansions and -6% and -10% in recessions. For DGP2, the relative bias of the LP estimator for the CIRF varies between -6% and -14% in expansions and -29% and -40% in recessions. This translates to a relative bias in the cumulative CIRF that varies between -3% and -8% for expansions and -11% and -23% for recessions.

5   Conclusion

State-dependent LP impulse response estimators have become one of the most commonly used tools in empirically macroeconomics in recent years. The idea that the effects of economic shocks may differ depending on the state of the economy has a long tradition, but apparent nonlinearities in recent macroeconomic data such as the zero lower bound on interest rates have, if anything, further heightened interest among applied researchers in such state dependencies. Much of what we know about the state dependence of fiscal multipliers and the state-dependent effects of monetary policy shocks, for example, is based on this LP approach, yet its validity has never been formally established.

Although there has been much discussion of the perceived advantages of this approach in the literature compared to the estimation of state-dependent structural VAR models, including its apparent simplicity and its potential robustness to possible dynamic misspecification of nonlinear VAR models, the question remains under what conditions the state-dependent LP estimator recovers the population impulse response functions of interest. It also remains unclear what impulse response function the LPs are estimating, among many competing impulse response concepts. In this paper, we made precise the nature of the state-dependent impulse responses captured by the LP estimator, and we provided sufficient conditions, under which this estimator is consistent. These conditions tend to be violated in many empirical studies. Our analysis suggests that, when the state of the economy is endogenously determined, the LP estimator tends to be valid only for the impact response. This is a concern not only for impulse response analysis but also for the construction of fiscal and monetary multipliers that are often computed at higher horizons (or relative to the peak in the response function).

While our theoretical analysis does not formally establish the inconsistency of the LP estimator (and the multipliers derived from those responses) in cases not covered by our theoretical analysis, we showed by simulation that in practice the LP estimator of the response function tends to be asymptotically biased, unless the state of the economy is exogenous. The fact that many applications of the state-dependent LP estimator implicitly treat the state of the economy as exogenous with respect to the model variables, when it clearly is endogenous, calls into question their substantive conclusions. This result is important not only from an econometric point of view, but also for the ongoing debate about the magnitude of fiscal and monetary multipliers. Our analysis highlights the need to be specific about nature of the state dependence when applying the state-dependent LP estimator. LP estimators cannot be interpreted and their validity cannot be assessed without taking a stand on the data generating process.

References

[1] Alesina, A., Azzalini, G., Favero, C., Giavazzi, F., Miano, A., 2018. Is it the 'how' or the 'when' that matters in fiscal adjustments? IMF Economic Review 66, 144-88.

[2] Alloza, M., 2019. Is fiscal policy more effective in uncertain times or during recessions? Manuscript, Bank of Spain.

[3] Alloza, M., Gonzalo, J., Sanz, C., 2021. Dynamic effects of persistent shocks. Manuscript, Banco de España.

[4] Angrist, J.D., Kuersteiner, G.M., 2011. Causal effects of monetary policy shocks: Semiparametric conditional independence tests with a multinomial propensity score. Review of Economics and Statistics 93, 725-747.

[5] Angrist, J.D., Jordà, O., Kuersteiner, G.M., 2018. Semiparametric estimates of monetary policy effects: String theory revisited. Journal of Business and Economic Statistics 36, 371-387.

[6] Auerbach, A.J., Gorodnichenko, Y., 2012. Measuring the output responses to fiscal policy. American Economic Journal: Economic Policy 4, 1-27.

[7] Auerbach, A.J., Gorodnichenko, Y., 2013a. Fiscal multipliers in recession and expansion. In: Alesina, A., Giavazzi, F. (eds.): Fiscal Policy after the Financial Crisis. University of Chicago Press, pp. 19-55.

[8] Auerbach, A.J., Gorodnichenko, Y., 2013b. Output spillovers from fiscal policy. American Economic Review Papers & Proceedings 103, 141-46.

[9] Bachmann, R., Sims, E.R., 2012. Confidence and the transmission of government spending shocks. Journal of Monetary Economics 59, 235-249.

[10] Barnichon, R., Matthes, C., 2018. Functional approximation of impulse responses. Journal of Monetary Economics, 99, 41-55.

[11] Caggiano, G., Castelnuovo, E., Colombo, V., Nodari, G., 2015. Estimating fiscal multipliers: News from a non-linear world. Economic Journal 125, 746-776.

[12] Caggiano, G., Castelnuovo, E., Groshenny, N., 2014. Uncertainty shocks and unemployment dynamics in U.S. recessions. Journal of Monetary Economics 67, 78-92.

[13] Candelon, B., Lieb, L., 2013. Fiscal policy in good and bad times. Journal of Economic Dynamics and Control 37, 2679-2694.

[14] Demirel, U.D., 2021. The short-term effects of tax changes: The role of state dependence. Journal of Monetary Economics 117, 918-934.

[15] Eskandari, R., 2019. State-dependent macroeconomic effects of tax changes. Manuscript, University of York.

[16] Gallant, A.R., Rossi, P.E., Tauchen, G., 1993. Nonlinear dynamic structures. Econometrica 61, 871-907.

[17] Gourieroux, C., Jasiak, J., 2005. Nonlinear innovations and impulse responses with application to VaR sensitivity. Annales d'Economie et de Statistique 78, 1-31.

[18] Ghassibe, M., Zanetti, F., 2020. State dependence of fiscal multipliers: The source of fluctuations matters. Manuscript, University of Oxford.

[19] Gonçalves, S., Herrera, A.M., Kilian, L,. Pesavento, E., 2021. Impulse response analysis for structural dynamic models with nonlinear regressors. Journal of Econometrics, forthcoming.

[20] Jordà, O., 2005. Estimation and inference of impulse responses by local projections. American Economic Review 95, 161-182.

[21] Jordà, O., 2009. Simultaneous confidence regions for impulse responses. Review of Economics and Statistics 91, 629-647.

[22] Kilian, L., Vigfusson, R.J., 2011. Are the responses of the U.S. economy asymmetric in energy price increases and decreases? Quantitative Economics 2, 419-453.

[23] Klepacz, M. 2020. Price setting and volatility: Evidence from oil price volatility shocks. Manuscript, Federal Reserve Board.

[24] Koop, G., Pesaran, M.H., Potter, S.M., 1996. Impulse response analysis in nonlinear multivariate models. Journal of Econometrics 74, 119-147.

[25] Mavroeidis, S. 2021. Identification at the zero lower bound. Econometrica 89, 2855-2885.

[26] Owyang, M. T., Ramey, V.A., Zubairy, S., 2013. Are government spending multipliers greater during periods of slack? Evidence from twentieth-century historical data. American Economic Review Papers & Proceedings 103, 129-134.

[27] Plagborg-Møller, M., Wolf, C.K., 2021. Local projections and VARs estimate the same impulse responses. Econometrica 89, 955-980.

[28] Potter, S.M., 2000. Nonlinear impulse response functions. Journal of Economic Dynamics and Control 24, 1425-1446.

[29] Rambachan, A., Shephard, N., 2020. Econometric analysis of potential outcomes time series: Instruments, shocks, linearity and the causal response function. Manuscript, Harvard University.

[30] 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.

[31] 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.

[32] Santoro, E., Petrella, I., Pfajfar, D., Gaffeo, E. 2014. Loss aversion and the asymmetric transmission of monetary policy. Journal of Monetary Economics 68, 19-36.

[33] Sims, E., Wolff, J., 2018. The state-dependent effects of tax shocks. European Economic Review 107, 57-85.

[34] Tenreyro, S., and Thwaites, G., 2016. Pushing on a string: U.S. montary policy is less powerful in recessions. American Economic Journal: Macroeconomics 8, 43-74.

[35] White, H., 2006. Time series estimation of the effects of natural experiments. Journal of Econometrics 135, 527-566.

A Appendix

Proof of Proposition 3.1. The proof for $h = 0$ and $h = 1$ is in the text. We omit the proof for general h since it follows from similar arguments.

Proof of Proposition 3.2. To define $\hat{b}_{E,h}$, let

$$\underset{T \times k}{Y} = \begin{pmatrix} y'_{1+h} \\ \vdots \\ y'_{T+h} \end{pmatrix}, \quad \underset{T \times 1}{X_1} = \begin{pmatrix} x_1 H_0 \\ \vdots \\ x_T H_{T-1} \end{pmatrix}, \quad \text{and} \quad \underset{T \times (np+1)}{X_2} = \begin{pmatrix} W'_0 H_0 \\ \vdots \\ W'_{T-1} H_{T-1} \end{pmatrix},$$

and define $M_2 = I_T - X_2 \left(X'_2 X_2 \right)^{-1} X'_2$.

By the Frish-Waugh-Lovell (FWL) Theorem, $\hat{b}'_{E,h} = \left(X'_1 M_2 X_1 \right)^{-1} X'_1 M_2 Y$, or

$$\hat{b}_{E,h} = T^{-1}(Y' M_2 X_1) \left(T^{-1} X'_1 M_2 X_1 \right)^{-1} \equiv \hat{Q}_{1y.2,h} \hat{Q}_{11.2}^{-1}.$$

23

A similar expression holds for $\hat{b}_{R,h}$ with the difference that the regressors x_t and W_{t-1} are interacted with $1 - H_{t-1}$ rather than H_{t-1}.

Our goal is to derive the probability limit of $\hat{b}_{E,h}$ (and $\hat{b}_{R,h}$) as $T \to \infty$. We can write

$$
\begin{aligned}
\hat{Q}_{11.2} &= T^{-1}X_1'X_1 - T^{-1}X_1'X_2 \left(T^{-1}X_2'X_2\right)^{-1} T^{-1}X_2'X_1, \text{ and} \\
\hat{Q}_{1y.2,h} &= T^{-1}Y'X_1 - T^{-1}Y'X_2 \left(T^{-1}X_2'X_2\right)^{-1} T^{-1}X_2'X_1.
\end{aligned}
$$

If a law of large numbers applies to each term[6],

$$
\begin{aligned}
\hat{Q}_{11.2} \xrightarrow{p} Q_{11.2} &\equiv E\left(x_t^2 H_{t-1}\right) - E\left(x_t H_{t-1} W_{t-1}'\right) \left[E\left(W_{t-1} W_{t-1}' H_{t-1}\right)\right]^{-1} E\left(W_{t-1} H_{t-1} x_t\right), \text{ and} \\
\hat{Q}_{1y.2,h} \xrightarrow{p} Q_{1y.2,h} &\equiv E\left(y_{t+h} x_t H_{t-1}\right) - E\left(y_{t+h} H_{t-1} W_{t-1}'\right) \left[E\left(W_{t-1} W_{t-1}' H_{t-1}\right)\right]^{-1} E\left(W_{t-1} H_{t-1} x_t\right).
\end{aligned}
$$

It is easy to see that $E\left(x_t H_{t-1} W_{t-1}'\right) = 0$ under the assumption that $x_t = \varepsilon_{1t}$ is a m.d.s. Thus,

$$
Q_{11.2} = E\left(x_t^2 H_{t-1}\right) \text{ and } Q_{1y.2} = E\left(y_{t+h} x_t H_{t-1}\right),
$$

implying that[7]

$$
\hat{b}_{E,h} \xrightarrow{p} b_{E,h} \equiv E\left(y_{t+h} x_t H_{t-1}\right) \left[E\left(x_t^2 H_{t-1}\right)\right]^{-1}.
$$

A similar argument implies that

$$
\hat{b}_{R,h} \xrightarrow{p} b_{R,h} \equiv E\left(y_{t+h} x_t (1 - H_{t-1})\right) \left[E\left(x_t^2 (1 - H_{t-1})\right)\right]^{-1}.
$$

Next, we express $Q_{11.2}$ and $Q_{1y.2,h}$ in terms of the model's parameters. Under the conditional homoskedasticity assumption on $x_t = \varepsilon_{1t}$,

$$
Q_{11.2} = E\left(\varepsilon_{1t}^2 | H_{t-1}\right) \Pr\left(H_{t-1} = 1\right) = \sigma_1^2 \Pr\left(H_{t-1} = 1\right).
$$

To evaluate $Q_{1y.2,h} = E\left(y_{t+h} x_t H_{t-1}\right)$, we use the fact that for any h,

$$
y_{t+h} = \mathbb{S}_k Z_{t+h},
$$

where Z_{t+h} is obtained from the companion form representation of the model given by (6). Consider first $h = 0$. Then,

$$
Z_t = A_{t-1} Z_{t-1} + \xi_t,
$$

---

[6] We assume that the data are strictly sationary and ergodic and that the usual moment and rank conditions on the regressors are satisfied. We leave these as implicit high level assumptions since our focus here is on the conditions that H_t needs to satisfy in order for the LP estimator to be consistent.

[7] This result is consistent with the fact that when x_t is a directly observed shock we can simply regress y_{t+h} onto $x_t H_{t-1}$ to obtain a consistent estimator of $b_{E,h}$. When $x_t = \varepsilon_{1t}$, adding the controls $W_{t-1} H_{t-1}$ is not required for consistency, but can be important for efficiency. More generally, when x_t is only predetermined with respect to y_t, adding controls becomes important for consistency.

where

$$\xi_t = \begin{pmatrix} \eta_t \\ 0 \end{pmatrix} = \begin{pmatrix} C_{t-1}^{-1} e_{1n} \varepsilon_{1t} + C_{t-1}^{-1} I_{2:n} \varepsilon_{2t} \\ 0 \end{pmatrix} = (e_{1p} \otimes C_{t-1}^{-1} e_{1n}) \varepsilon_{1t} + e_{1p} \otimes C_{t-1}^{-1} I_{2:n} \varepsilon_{2t},$$

given that $\eta_t = C_{t-1}^{-1} \varepsilon_t$ and $\varepsilon_t = C_{t-1}^{-1} e_{1n} \varepsilon_{1t} + C_{t-1}^{-1} I_{2:n} \varepsilon_{2t}$, where e_{1n} and $I_{2:n}$ are as defined in Section 3.1. Hence,

$$y_t = \mathbb{S}_k Z_t = \mathbb{S}_k (e_{1p} \otimes C_{t-1}^{-1} e_{1n}) \varepsilon_{1t} + \mathbb{S}_k A_{t-1} Z_{t-1} + \mathbb{S}_k (e_{1p} \otimes C_{t-1}^{-1} I_{2:n} \varepsilon_{2t}). \tag{15}$$

We use this expression to evaluate $Q_{1y.2,0} = E(y_t x_t H_{t-1})$. More specifically, using the fact that $x_t = \varepsilon_{1t}$ and the above decomposition of y_t, we can write

$$Q_{1y.2,0} \equiv E(y_t x_t H_{t-1}) = Q_{1,0} + Q_{2,0} + Q_{3,0},$$

where

$$\begin{aligned} Q_{1,0} &= E[\mathbb{S}_k (e_{1p} \otimes C_{t-1}^{-1} e_{1n}) H_{t-1} \varepsilon_{1t}^2], \\ Q_{2,0} &= E[\mathbb{S}_k A_{t-1} Z_{t-1} H_{t-1} \varepsilon_{1t}], \text{ and} \\ Q_{3,0} &= E[\mathbb{S}_k (e_{1p} \otimes C_{t-1}^{-1} I_{2:n} \varepsilon_{2t}) H_{t-1} \varepsilon_{1t}]. \end{aligned}$$

Under Assumption 1, it can be shown that $Q_{2,0} = Q_{3,0} = 0$, implying that $Q_{1y.2,0} = Q_{1,0}$. More specifically, by the LIE,

$$Q_{2,0} = E[\mathbb{S}_k A_{t-1} Z_{t-1} H_{t-1} \underbrace{E(\varepsilon_{1t}|\mathcal{F}^{t-1})}_{=0}] = 0,$$

given the m.d.s. condition on ε_{1t}. Similarly, we have that

$$Q_{3,0} = E[\mathbb{S}_k (e_{1p} \otimes C_{t-1}^{-1} I_{2:n} \varepsilon_{2t} H_{t-1} \varepsilon_{1t})] = E[\mathbb{S}_k (e_{1p} \otimes C_{t-1}^{-1} I_{2:n} H_{t-1} \underbrace{E(\varepsilon_{2t} \varepsilon_{1t}|\mathcal{F}^{t-1})}_{=0})] = 0,$$

where the last equality holds by the assumption that $\Sigma = E(\varepsilon_t \varepsilon_t'|\mathcal{F}^{t-1})$ is diagonal. Finally, note that

$$\begin{aligned} Q_{1,0} &= E[\mathbb{S}_k (e_{1p} \otimes C_{t-1}^{-1} e_{1n}) H_{t-1} \underbrace{E(\varepsilon_{1t}^2|\mathcal{F}^{t-1})}_{=\sigma_1^2}] \\ &= \sigma_1^2 E[\mathbb{S}_k (e_{1p} \otimes C_{t-1}^{-1} e_{1n}) H_{t-1}] \\ &= E[\mathbb{S}_k (e_{1p} \otimes C_{t-1}^{-1} e_{1n})|H_{t-1} = 1] \underbrace{\sigma_1^2 \Pr(H_{t-1} = 1)}_{=Q_{11.2}}. \end{aligned}$$

Since $b_{E,0} = Q_{1y.2,0} Q_{11.2}^{-1}$, and $Q_{1y.2,0} = Q_{1,0} = E[\mathbb{S}_k (e_{1p} \otimes C_{t-1}^{-1} e_{1n})|H_{t-1} = 1] Q_{11.2}$, we conclude that

$$\hat{b}_{E,0} \xrightarrow{p} b_{E,0} = E[\mathbb{S}_k (e_{1p} \otimes C_{t-1}^{-1} e_{1n})|H_{t-1} = 1] = \mathbb{S}_k (e_{1p} \otimes C_E^{-1} e_{1n}).$$

25

A similar argument shows that

$$\hat{b}_{R,0} \xrightarrow{p} b_{R,0} = E[\mathbb{S}_k(e_{1p} \otimes C_{t-1}^{-1}e_{1n})|H_{t-1} = 0] = \mathbb{S}_k(e_{1p} \otimes C_R^{-1}e_{1n}).$$

These results show that the state-dependent LP regression (8) recovers the conditional IRF obtained in Proposition 3.1 with $h = 0$ under Assumption 1. No further assumptions are required (provided a law of large numbers can be applied to $\hat{Q}_{11.2}$ and $\hat{Q}_{1y.2,0}$).

As we will show next, this is no longer the case when $h > 0$. To illustrate this, consider $h = 1$. Now,

$$\hat{b}_{E,1} \xrightarrow{p} b_{E,1} \equiv E\left(y_{t+1}x_t H_{t-1}\right)\left[E\left(x_t^2 H_{t-1}\right)\right]^{-1} \equiv Q_{1y.2,1}Q_{11.2}^{-1},$$

where $Q_{11.2} = \sigma_1^2 \Pr\left(H_{t-1} = 1\right)$, as for $h = 0$. To obtain $Q_{1y.2,1}$, we can use the fact that

$$
\begin{aligned}
y_{t+1} &= \mathbb{S}_k Z_{t+1} = \mathbb{S}_k(A_t Z_t + \xi_{t+1}) \\
&= \mathbb{S}_k(A_t(A_{t-1}Z_{t-1} + \xi_t) + \xi_{t+1}) \\
&= \mathbb{S}_k A_t \xi_t + \mathbb{S}_k A_t A_{t-1} Z_{t-1} + \mathbb{S}_k \xi_{t+1},
\end{aligned}
\tag{16}
$$

where $\xi_s = (e_{1p} \otimes C_{s-1}^{-1}e_{1n})\varepsilon_{1s} + e_{1p} \otimes C_{s-1}^{-1}I_{2:n}\varepsilon_{2s}$ for $s = t, t+1$. This implies that

$$Q_{1y.2,1} \equiv E\left(y_{t+1}x_t H_{t-1}\right) = Q_{1,1} + Q_{2,1} + Q_{3,1},$$

where

$$
\begin{aligned}
Q_{1,1} &= E(\mathbb{S}_k A_t \xi_t H_{t-1}\varepsilon_{1t}), \\
Q_{2,1} &= E[\mathbb{S}_k A_t A_{t-1} Z_{t-1} H_{t-1}\varepsilon_{1t}], \text{ and} \\
Q_{3,1} &= E[\mathbb{S}_k \xi_{t+1} H_{t-1}\varepsilon_{1t}].
\end{aligned}
$$

Given the definition of ξ_{t+1}, we can easily see that $Q_{3,1} = 0$ by Assumption 1, since it implies that $E\left(\xi_{t+1}|\mathcal{F}^t\right) = 0$. However, to conclude that $Q_{2,1} = 0$, we need further assumptions. More specifically, this term now depends on $A_t \equiv A_E H_t + A_R\left(1 - H_t\right)$ and ε_{1t} (in addition to A_{t-1}, Z_{t-1} and H_{t-1}). Conditionally on $\mathcal{F}^{t-1}$, H_t and ε_{1t} may be correlated, implying that $Q_{2,1} \neq 0$. Indeed, by the law of iterated expectations,

$$Q_{2,1} = E[\mathbb{S}_k E\left(A_t\varepsilon_{1t}|\mathcal{F}^{t-1}\right) A_{t-1} Z_{t-1} H_{t-1}],$$

where

$$
\begin{aligned}
E\left(A_t\varepsilon_{1t}|\mathcal{F}^{t-1}\right) &= A_E E\left(H_t\varepsilon_{1t}|\mathcal{F}^{t-1}\right) + A_R E\left((1 - H_t)\varepsilon_{1t}|\mathcal{F}^{t-1}\right) \\
&= (A_E - A_R)E\left(H_t\varepsilon_{1t}|\mathcal{F}^{t-1}\right) + A_R E\left(\varepsilon_{1t}|\mathcal{F}^{t-1}\right) \\
&= (A_E - A_R)E\left(H_t\varepsilon_{1t}|\mathcal{F}^{t-1}\right),
\end{aligned}
$$

since $E\left(\varepsilon_{1t}|\mathcal{F}^{t-1}\right) = 0$. Thus, unless $E\left(H_t\varepsilon_{1t}|\mathcal{F}^{t-1}\right) = 0$ (or $A_E = A_R$), $Q_{2,1} \neq 0$. Assumption 2(a) (with $h = 1$) implies that $E\left(H_t\varepsilon_{1t}|\mathcal{F}^{t-1}\right) = 0$, ensuring that $Q_{2,1} = 0$.

Additional conditions are also required to simplify $Q_{1,1} = E(\mathbb{S}_k A_t\xi_t H_{t-1}\varepsilon_{1t})$ and show that the LP recovers the conditional IRF when $H_{t-1} = 1$. Using the definition of ξ_t, this term can be decomposed as follows:

$$Q_{1,1} = E[\mathbb{S}_k A_t(e_{1p} \otimes C_{t-1}^{-1}e_{1n})H_{t-1}\varepsilon_{1t}^2] + E[\mathbb{S}_k A_t(e_{1p} \otimes C_{t-1}^{-1}I_{2:n}H_{t-1}\varepsilon_{2t}\varepsilon_{1t})].$$

The presence of A_t (which depends on H_t) again complicates the evaluation of these expectations. For instance, the second term is not zero if $E\left(\varepsilon_{1t}\varepsilon_{2t}|H_t, \mathcal{F}^{t-1}\right) \neq 0$ even if Σ is diagonal. Similarly, we cannot conclude that the first term simplifies to $E[\mathbb{S}_k A_t(e_{1p} \otimes C_{t-1}^{-1}e_{1n})|H_{t-1} = 1]\sigma_1^2 \Pr\left(H_{t-1} = 1\right)$ unless we impose the condition that $E\left(\varepsilon_{1t}^2|H_t, \mathcal{F}^{t-1}\right) = \sigma_1^2$. This is Assumption 2(b) with $h = 1$, which together with Assumption 1 and 2(b) ensures the consistency of the LP estimator for $h = 1$. The proof for other values of h follows from similar arguments provided Assumption 1 is strengthening by Assumption 2.

Figure 1: DGP1: Exogenous H_t, Level Effects

Figure 2: DGP1: Exogenous H_t, Cumulative Effects

Figure 3: DGP2: Endogenous H_t, Level Effects

Figure 4: DGP2: Endogenous H_t, Cumulative Effects

Figure 5: DGP3: Endogenous H_t, Level Effects

Figure 6: DGP3: Endogenous H_t, Cumulative Effects