Local Projections and VARs Estimate the Same Impulse Responses*

Mikkel Plagborg-Møller
Princeton University

Christian K. Wolf
Princeton University

This version: October 6, 2020
First version: November 16, 2018

Abstract: We prove that local projections (LPs) and Vector Autoregressions (VARs) estimate the same impulse responses. This nonparametric result only requires unrestricted lag structures. We discuss several implications: (i) LP and VAR estimators are not conceptually separate procedures; instead, they are simply two dimension reduction techniques with common estimand but different finite-sample properties. (ii) VAR-based structural identification – including short-run, long-run, or sign restrictions – can equivalently be performed using LPs, and *vice versa*. (iii) Structural estimation with an instrument (proxy) can be carried out by ordering the instrument first in a recursive VAR, even under non-invertibility. (iv) Linear VARs are as robust to non-linearities as linear LPs.

Keywords: external instrument, impulse response function, local projection, proxy variable, structural vector autoregression. *JEL codes:* C32, C36.

*Email: `mikkelpm@princeton.edu` and `ckwolf@princeton.edu`. We received helpful comments from two anonymous referees, the editor Guido Imbens, Majid Al-Sadoon, Domenico Giannone, Marek Jarociński, Òscar Jordà, Peter Karadi, Lutz Kilian, Dake Li, Pepe Montiel Olea, Valerie Ramey, Giovanni Ricco, Ken Rogoff, Neil Shephard, Chris Sims, Jim Stock, Mark Watson, Jing Zeng, and seminar participants at various venues. The first draft of this paper was written while Wolf was visiting the European Central Bank, whose hospitality is gratefully acknowledged. Plagborg-Møller acknowledges that this material is based upon work supported by the NSF under Grant #1851665. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the NSF.

1 Introduction

Modern dynamic macroeconomics studies the propagation of structural shocks (Frisch, 1933; Ramey, 2016). Central to this impulse-propagation paradigm are impulse response functions – the dynamic response of a macro aggregate to a structural shock. Following Sims (1980), Bernanke (1986), and Blanchard & Watson (1986), Structural Vector Autoregression (SVAR) analysis remains the most popular empirical approach to impulse response estimation. Over the past decade, however, starting with Jordà (2005), local projections (LPs) have become an increasingly widespread alternative econometric approach.

How should we choose between SVAR and LP estimators of impulse responses? Unfortunately, so far there exists little theoretical guidance as to which method is preferable in practice. Conventional wisdom holds that SVARs are more efficient, while LPs are more robust to model misspecification. Examples of the former statement can be found in the textbook treatment of Kilian & Lütkepohl (2017, ch. 12.8) and the survey of Ramey (2016, p. 84), while the latter statement is expressed by Jordà (2005, p. 162), Ramey (2016, p. 83) and Nakamura & Steinsson (2018, pp. 80–81), among others.[1] Kilian & Lütkepohl (2017) and Stock & Watson (2018, p. 944), however, caution that these remarks are not based on formal analysis and call for further research. It is also widely believed that LPs invariably require a measure of a "shock" (perhaps obtained from an auxiliary SVAR model), so that SVAR estimation is required to implement non-recursive structural identification schemes such as long-run or sign restrictions.[2] Finally, when applied to the same empirical question, LP- and VAR-based approaches sometimes give substantively different results (Ramey, 2016). Existing simulation studies provide useful guidance on particular approaches to local projections or VARs, but differences in implementation details cause these studies to reach disparate conclusions (Meier, 2005; Kilian & Kim, 2011; Brugnolini, 2018; Nakamura & Steinsson, 2018; Choi & Chudik, 2019).

The central result of this paper is that linear local projections and VARs in fact estimate the exact same impulse responses in population. Specifically, any LP impulse response function can be obtained through an appropriately ordered recursive VAR, and any (possibly non-recursive) VAR impulse response function can be obtained through a LP with appropriate control variables. This result applies to all common implementations of local projections

[1]In the online postscript to her handbook chapter, Ramey corrects the claims and restates the relationship between LP and VAR estimands following the findings of this paper.

[2]See the reviews by Ramey (2016) and Kilian & Lütkepohl (2017, ch. 12.8).

used in the literature (Jordà, 2005, 2009; Ramey, 2016). While the result concerns linear estimators, we essentially only require the *nonparametric* assumption that the data are weakly stationary and that the lag structures in the two specifications are unrestricted. In particular, we do not impose restrictions on the linearity or dimensionality of the underlying data generating process (DGP). Intuitively, a VAR model with sufficiently large lag length captures all covariance properties of the data. Hence, iterated VAR(∞) forecasts coincide with direct LP forecasts. Since impulse responses are just forecasts conditional on specific innovations, LP and VAR impulse response estimands coincide in population. Furthermore, we prove that if only a fixed number p of lags are included in the LP and VAR, then the two impulse response estimands still approximately agree out to horizon p (but not further), again without imposing any parametric assumptions on the data generating process.

The equivalence of VAR and LP estimands has several implications for structural estimation in applied macroeconometrics.

First, LPs and VARs are not conceptually different methods; instead, they are simply two particular linear projection techniques that share the same estimand but differ in their finite-sample properties. At short impulse response horizons the two methods are likely to approximately agree if the same lag length is used for both methods. However, with finite lag lengths, the two methods may give substantially different results at long horizons.

Second, structural estimation with VARs can equally well be carried out using LPs, and *vice versa*. Structural identification – which is a population concept – is logically distinct from the choice of finite-sample estimation approach. In particular, we show concretely how various popular "SVAR" identification schemes – including recursive, long-run, and sign identification – can just as easily be implemented using local projection techniques. Ultimately, our results show that LP-based structural estimation can succeed if and only if SVAR estimation can succeed.

Third, valid structural estimation with an instrument (IV, also known as a proxy variable) can be carried out by ordering the IV first in a recursive VAR à la Kilian (2006) and Ramey (2011). This is because the LP-IV estimand of Stock & Watson (2018) can equivalently be obtained from a recursive (i.e., Cholesky) VAR that contains the IV. Importantly, the "internal instrument" strategy of ordering the IV first in a VAR yields valid impulse response estimates even if the shock of interest is non-invertible, unlike the well-known "external instrument" SVAR-IV approach (Stock, 2008; Stock & Watson, 2012; Mertens & Ravn, 2013).[3]

[3]In contemporaneous work, Noh (2018) also includes the IV as an internal instrument in a VAR; our result offers additional insights by drawing connections to LP-IV and to the general LP/VAR equivalence.

In particular, this result goes through even if the IV is contaminated with measurement error that is unrelated to the shock of interest.

Fourth, in population, linear local projections are exactly as "robust to non-linearities" in the DGP as linear VARs. We show that their common estimand may be formally interpreted as a best linear approximation to the underlying, perhaps non-linear, impulse responses.

In summary, in addition to clarifying misconceptions in the literature about the relationship between the LP and VAR estimands, our results allow applied researchers to separate the choice of identification scheme from the choice of estimation approach. Researchers who prefer the intuitive regression interpretation of the LP impulse response estimator can apply our methods for imposing "SVAR" identifying restrictions such as short-run, long-run, and sign restrictions. Researchers who instead prefer the explicit multivariate model of the VAR estimator can apply our results on how to use instruments/proxies without requiring invertible shocks, as in LP-IV.

LITERATURE. While the existing literature has pointed out connections between LPs and VARs, our contributions are to establish a formal equivalence result that does not require extraneous functional form assumptions and to derive implications for structural identification of impulse responses. Jordà (2005, Sec. I.B) and Kilian & Lütkepohl (2017, Ch. 12.8) show that, under the assumption of a finite-order VAR model, VAR impulse responses can be estimated consistently through LPs. In contrast, our equivalence result between these two linear estimation methods does not restrict the data generating process itself to be linear or finite-dimensional. While Dufour & Renault (1998, Eqn. 3.17) discuss a similar result in the context of testing for Granger causality, we go further by demonstrating how causal structural VAR orderings map into particular choices of LP control variables, and *vice versa*.[4] Moreover, to our knowledge, our results on long-run/sign identification, LP-IV, and best linear approximations have no obvious parallels in the preceding literature.[5]

In this paper we focus exclusively on identification and point estimation of impulse responses. Plagborg-Møller & Wolf (2019) provide identification results for variance/historical decompositions using IVs/proxies. We do not consider questions related to *inference*, and instead refer to Jordà (2005), Kilian & Lütkepohl (2017), and Stock & Watson (2018).

[4] Jordà et al. (2019) informally discuss the connection between control variables and recursive SVARs.

[5] Kilian & Lütkepohl (2017, Ch. 12.8) present alternative arguments for why it is a mistake to assert that finite-order LPs are generally more "robust to model misspecification" than finite-order VAR estimators. They do not appeal to the general equivalence of the LP and VAR estimands, however.

Outline. Section 2 presents our core result on the population equivalence of local projections and VARs in a reduced-form setting. Section 3 traces out the implications for *structural* estimation. We illustrate our equivalence results with a practical application to IV-based identification of monetary policy shocks in Section 4. Section 5 concludes by summarizing the takeaways for empirical practice. Some proofs are relegated to Appendix A, and supplementary details are presented in the Online Appendix.[6]

2 Equivalence between local projections and VARs

This section presents our core result: Local projections and VARs estimate the same impulse response functions in population. First we establish that local projections are equivalent with recursively identified VARs when the lag structure is unrestricted. Then we extend the argument to (i) non-recursive identification and (ii) finite lag orders, and we illustrate the results graphically. Finally, we discuss an in-sample asymptotic equivalence result that complements the population analysis.

Our analysis in this section is "reduced form" in that it does not assume any specific underlying structural/causal model; we merely manipulate linear projections of stationary time series. We will discuss implications for *causal* identification in Section 3.

2.1 Main result

Suppose the researcher observes data $w_t = (r_t', x_t, y_t, q_t')'$, where r_t and q_t are, respectively, $n_r \times 1$ and $n_q \times 1$ vectors of time series, while x_t and y_t are scalar time series. We are interested in the dynamic response of y_t after an impulse in x_t. The vector time series r_t and q_t (which may each be empty) will serve as control variables. Readers who wish to have a structural interpretation in mind may think of x_t as predetermined with respect to y_t and r_t as predetermined with respect to $\{x_t, y_t\}$. However, our reduced-form equivalence result below does not require any such predeterminedness assumptions. The precise roles of the controls r_t and q_t will become clear in equations (1) and (3) below.

For now, we only make the following standard nonparametric regularity assumption.[7]

[6] Online Appendix and replication files: `http://scholar.princeton.edu/mikkelpm/lp_var`

[7] The restriction to non-singular spectral density matrices rules out over-differenced data. We conjecture that this restriction could be relaxed using the techniques in Almuzara & Marcet (2017).

Assumption 1. *The data $\{w_t\}$ are covariance stationary and purely non-deterministic, with an everywhere nonsingular spectral density matrix and absolutely summable Wold decomposition coefficients.*

In particular, we assume nothing about the underlying causal structure of the economy, as this section is concerned solely with properties of linear projections.[8]

As an expositional device, we impose an additional assumption of joint Gaussianity.

Assumption 2. *$\{w_t\}$ is a jointly Gaussian vector time series.*

The Gaussianity assumption is made purely for notational simplicity, as this allows us to write conditional expectations instead of linear projections. If we drop the Gaussianity assumption, all calculations below hold with projections in place of conditional expectations.

We will show that, in population, the following two approaches estimate the *same* impulse response function of y_t with respect to an innovation in x_t.[9]

1. Local projection. Consider for each $h = 0, 1, 2, \dots$ the linear projection

$$y_{t+h} = \mu_h + \beta_h x_t + \gamma_h' r_t + \sum_{\ell=1}^{\infty} \delta_{h,\ell}' w_{t-\ell} + \xi_{h,t}, \tag{1}$$

where $\xi_{h,t}$ is the projection residual, and $\mu_h, \beta_h, \gamma_h, \delta_{h,1}, \delta_{h,2}, \dots$ the projection coefficients.

Definition 1. *The* LP impulse response function *of y_t with respect to x_t is given by $\{\beta_h\}_{h\geq 0}$ in equation (1).*

Effectively, this defines the LP impulse response estimand at horizon h as

$$\beta_h = E(y_{t+h} \mid x_t = 1, r_t, \{w_\tau\}_{\tau<t}) - E(y_{t+h} \mid x_t = 0, r_t, \{w_\tau\}_{\tau<t}). \tag{2}$$

Notice that the projection (1) controls for the contemporaneous value of r_t but not of q_t. Notice also that we do not require x_t to be a predetermined "shock" variable in this section, although such additional assumptions may be important for interpreting β_h

[8] Assumption 1 allows the time series to be discrete or censored, though *structural* interpretation of the linear impulse response estimand in such cases requires care, cf. the discussion below Proposition 1.

[9] We write linear projections on the span of infinitely many variables as an infinite sum. This is justified under Assumption 1, since we can invert the Wold representation to obtain a VAR(∞) representation.

structurally, as discussed in Section 3 below. Importantly, the formulation (1) is general enough to cover all common empirical implementations of local projections.[10]

2. VAR. Consider the multivariate linear "VAR(∞)" projection

$$w_t = c + \sum_{\ell=1}^{\infty} A_\ell w_{t-\ell} + u_t, \tag{3}$$

where $u_t \equiv w_t - E(w_t \mid \{w_\tau\}_{-\infty<\tau<t})$ is the projection residual, and $c, A_1, A_2, \dots$ the projection coefficients. Let $\Sigma_u \equiv E(u_t u_t')$, and define the Cholesky decomposition $\Sigma_u = BB'$, where B is lower triangular with positive diagonal entries. Consider the corresponding recursive SVAR representation

$$A(L)w_t = c + B\eta_t,$$

where $A(L) \equiv I - \sum_{\ell=1}^{\infty} A_\ell L^\ell$ and $\eta_t \equiv B^{-1}u_t$. Notice that r_t is ordered first in the VAR, while q_t is ordered last.[11] Define the lag polynomial $\sum_{\ell=0}^{\infty} C_\ell L^\ell = C(L) \equiv A(L)^{-1}$. Noting that x_t and y_t are the (n_r+1)-th and (n_r+2)-th elements in w_t, we now introduce the following familiar definition.

Definition 2. *The* VAR impulse response function *of y_t with respect to an innovation in x_t is given by $\{\theta_h\}_{h\geq 0}$, where*

$$\theta_h \equiv C_{n_r+2,\bullet,h} B_{\bullet,n_r+1},$$

and $\{C_\ell\}$ and B are defined above.

Here $C_{i,\bullet,h}$, say, refers to the i-th row of C_h, while $B_{\bullet,j}$ is the j-th column of B.

[10] This includes: i) estimating reduced-form impulse responses via LP and then rotating them using estimates of the impact impulse response matrix from an auxiliary SVAR as in Jordà (2005, 2009), (ii) projections on an exogenous shock $x_t = \varepsilon_{j,t}$ (e.g., Ramey, 2016; Nakamura & Steinsson, 2018) (in this case control variables are often omitted, though they may increase efficiency), and (iii) projections on an endogenous covariate x_t while controlling for confounding variables r_t (e.g., Jordà et al., 2013). The fact that options (ii) and (iii) are covered by (1) is immediate, while Chang & Sakata (2007) show that option (i) is equivalent in population to directly projecting on the shock x_t identified by the auxiliary SVAR as in option (ii).

[11] The relative ordering of y_t and q_t in the SVAR representation does not matter for our results, since it can be verified that the (n_r+1)-th column of B is equivariant with respect to this ordering. Similarly, if x_t is ordered after y_t in the SVAR representation, then the equivalence result below still obtains as long as we additionally control for y_t on the right-hand side of (1) (so in particular $\beta_0 = 0$).

Note that our definitions of the LP and VAR estimands include infinitely many lags of w_t in the relevant projections; we consider the case of finitely many lags in Section 2.3. Note also that we take the use of the control variables r_t and q_t as given in this section, as controls are common in applied work. We will discuss structural justifications for the use of such controls in Section 3.

Although LP and VAR approaches are often viewed as conceptually distinct in the literature, they in fact estimate the same population impulse response function.

Proposition 1. *Under Assumptions 1 and 2, the LP and VAR impulse response functions are equal, up to a constant of proportionality:* $\theta_h = \sqrt{E(\tilde{x}_t^2)} \times \beta_h$ *for all* $h = 0, 1, 2, \dots,$ *where* $\tilde{x}_t \equiv x_t - E(x_t \mid r_t, \{w_\tau\}_{-\infty<\tau<t})$.

That is, any LP impulse response function can equivalently be obtained as an appropriately ordered recursive VAR impulse response function. Conversely, any recursive VAR impulse response function can be obtained through a LP with appropriate control variables. We comment on non-recursive identification schemes below. The constant of proportionality in the proposition depends on neither the response horizon h nor on the response variable y_t. The reason for the presence of this constant of proportionality is that the implicit LP innovation $\tilde{x}_t$, after controlling for the other right-hand side variables, does not have variance 1. If we scale the innovation $\tilde{x}_t$ to have variance 1, or if we consider relative impulse responses θ_h/θ_0 (as further discussed below), the LP and VAR impulse response functions coincide.

The intuition behind the result is that a VAR(p) model with $p \to \infty$ is sufficiently flexible to perfectly capture all covariance properties of the data (Lewis & Reinsel, 1985; Inoue & Kilian, 2002). Thus, iterated forecasts based on the VAR coincide perfectly with direct forecasts $E(w_{t+h} \mid w_t, w_{t-1}, \dots)$. Since both recursive VAR and LP impulse responses are just linear functions of these direct reduced-form forecasts, they coincide. Although the intuition for this equivalence result is simple, its implications do not appear to have been generally appreciated in the literature on impulse response estimation, as discussed earlier in Section 1.

Proof. The proof of the proposition relies only on least-squares projection algebra. First consider the LP estimand. By the Frisch-Waugh theorem, we have that

$$\beta_h = \frac{\mathrm{Cov}(y_{t+h}, \tilde{x}_t)}{E(\tilde{x}_t^2)}. \tag{4}$$

For the VAR estimand, note that $C(L) = A(L)^{-1}$ collects the coefficient matrices in the

Wold decomposition

$$w_t = \chi + C(L)u_t = \chi + \sum_{\ell=0}^{\infty} C_\ell B \eta_t, \quad \chi \equiv C(1)c.$$

As a result, the VAR impulse responses equal

$$\theta_h = C_{n_r+2,\bullet,h} B_{\bullet,n_r+1} = \text{Cov}(y_{t+h}, \eta_{x,t}), \tag{5}$$

where we partition $\eta_t = (\eta_{r,t}', \eta_{x,t}, \eta_{y,t}, \eta_{q,t}')'$ the same way as $w_t = (r_t', x_t, y_t, q_t')'$. By $u_t = B\eta_t$ and the properties of the Cholesky decomposition, we have[12]

$$\eta_{x,t} = \frac{1}{\sqrt{E(\tilde{u}_{x,t}^2)}} \times \tilde{u}_{x,t}, \tag{6}$$

where we partition $u_t = (u_{r,t}', u_{x,t}, u_{y,t}, u_{q,t}')'$ and define[13]

$$\tilde{u}_{x,t} \equiv u_{x,t} - E(u_{x,t} \mid u_{r,t}) = \tilde{x}_t. \tag{7}$$

From (5), (6), and (7) we conclude that

$$\theta_h = \frac{\text{Cov}(y_{t+h}, \tilde{x}_t)}{\sqrt{E(\tilde{x}_t^2)}},$$

and the proposition now follows by comparing with (4). □

In the special case where x_t represents a "shock", in the sense that $E(x_t \mid r_t, \{w_\tau\}_{\tau<t}) = 0$, the LP estimand β_h coincides also with the impulse response estimand φ_h from a distributed lag regression $y_t = a + \sum_{\ell=0}^{\infty} \varphi_\ell x_{t-\ell} + \omega_t$ (understood as a linear projection), see Baek & Lee (2020). Note that in this special case, the LP estimand is unchanged if we drop all control variables in equation (1). However, the projection coefficient φ_h differs from the LP (and VAR) estimand if x_t correlates with r_t or with lags of the data (Alloza et al., 2019).

Proposition 1 implies that linear LPs are exactly as "robust to non-linearities" as linear

[12] B is lower triangular, so the $(n_r + 1)$-th equation in the system $B\eta_t = u_t$ is $B_{n_r+1,1:n_r}\eta_{r,t} + B_{n_r+1,n_r+1}\eta_{x,t} = u_{x,t}$, with obvious notation. Since $\eta_{x,t}$ and $\eta_{r,t}$ are uncorrelated, we find $B_{n_r+1,n_r+1}\eta_{x,t} = u_{x,t} - E(u_{x,t} \mid \eta_{r,t}) = u_{x,t} - E(u_{x,t} \mid u_{r,t}) = \tilde{u}_{x,t}$. Expression (6) then follows from $E(\eta_{x,t}^2) = 1$.

[13] Observe that $u_{x,t} - \tilde{x}_t = E(x_t \mid r_t, \{w_\tau\}_{-\infty<\tau<t}) - E(x_t \mid \{w_\tau\}_{-\infty<\tau<t}) = E(u_{x,t} \mid r_t, \{w_\tau\}_{-\infty<\tau<t}) = E(u_{x,t} \mid u_{r,t}, \{w_\tau\}_{-\infty<\tau<t}) = E(u_{x,t} \mid u_{r,t})$.

VAR methods, in population. This is because, while the equivalence result concerns linear estimation methods, our argument was nonparametric in that it did not rely on functional form assumptions on the true data generating process, such as linearity or finite dimensionality. In the Online Appendix we prove that the common LP/VAR estimand can be interpreted as a "best linear approximation" to the true, possibly non-linear, structural/causal impulse responses. Of course, this best linear approximation may bear little resemblance to the impulse responses in the underlying non-linear model, which will generally depend on the history and magnitudes of current and past shocks, unlike the linear impulse responses.[14]

In conclusion, LPs and VARs should not be thought of as conceptually different methods – they are simply two particular linear projection techniques with a shared estimand. LPs and VARs offer two equivalent ways of arriving at the same population parameter (4), or equivalently (2), up to a scale factor that does not depend on the horizon h.

2.2 Extension: Non-recursive specifications

Our equivalence result extends straightforwardly to the case of non-recursively identified VARs. Above we restricted attention to recursive identification schemes, as the VAR directly contains a measure of the impulse x_t. In a generic structural VAR identification scheme, the impulse is some – not necessarily recursive – rotation of reduced-form forecasting residuals. Thus, let us continue to consider the VAR (3), but now we shall study the propagation of *some* rotation of the reduced-form forecasting residuals,

$$\bar{\eta}_t \equiv b' u_t, \tag{8}$$

where b is a vector of the same dimension as w_t. Under Assumptions 1 and 2, we can follow the same steps as in Section 2.1 to establish that the VAR-implied impulse response at horizon h of y_t with respect to the innovation $\bar{\eta}_t$ equals – up to scale – the coefficient $\bar{\beta}_h$ of the linear projection

$$y_{t+h} = \bar{\mu}_h + \bar{\beta}_h (b' w_t) + \sum_{\ell=1}^{\infty} \bar{\delta}'_{h,\ell} w_{t-\ell} + \bar{\xi}_{h,t}, \tag{9}$$

where the coefficients are least-squares projection coefficients and the last term is the projection residual. Thus, any recursive or non-recursive SVAR(∞) identification procedure is equivalent with a local projection (9) on a particular linear combination $b' w_t$ of the variables

[14] See Kilian & Vigfusson (2011) as an example of a model in which the common *linear* estimand of local projections and VARs is not the structural object of interest.

in the VAR (and their lags). For recursive orderings, this reduces to Proposition 1. We give concrete examples of the mapping from non-recursive VAR to the rotation vector b in Section 3.2 as well as the Online Appendix.

2.3 Extension: Finite lag length

Whereas our main equivalence result in Section 2.1 relied on infinite lag polynomials, we now prove an equivalence result that holds when only finitely many lags are used. Specifically, when p lags of the data are included in the VAR and as controls in the LP, the impulse response estimands for the two methods agree approximately out to horizon p, but generally not at higher horizons. Importantly, this result is still entirely nonparametric, in the sense that we do not impose that the true DGP is a linear or finite-order VAR.

First, we define the finite-order LP and VAR estimands. We continue to impose Assumptions 1 and 2. Consider any lag length p and impulse response horizon h.

1. Local projection. The local projection impulse response estimand $\beta_h(p)$ is defined as the coefficient on x_t in a projection as in (1), except that the infinite sum is truncated at lag p. Again, we interpret all coefficients and residuals as resulting from a least-squares linear projection.

2. VAR. Consider a linear projection of the data vector w_t onto p of its lags (and a constant), i.e., the projection (3) except with the infinite sum truncated at lag p. Let $A_\ell(p)$, $\ell = 1, 2, \dots, p$, and $\Sigma_u(p)$ denote the corresponding projection coefficients and residual variance. Define $A(L; p) \equiv I - \sum_{\ell=1}^{p} A_\ell(p)$ and the Cholesky decomposition $\Sigma_u(p) = B(p)B(p)'$. Define also the inverse lag polynomial $\sum_{\ell=0}^{\infty} C_\ell(p)L^\ell = C(L; p) \equiv A(L; p)^{-1}$ consisting of the reduced-form impulse responses implied by $A(L; p)$. Then the VAR impulse response estimand at horizon h is defined as

$$\theta_h(p) \equiv C_{n_r+2,\bullet,h}(p)B_{\bullet,n_r+1}(p),$$

cf. the definition in Section 2.1 with $p = \infty$.

Note that the VAR(p) model used to define the VAR estimand above is "misspecified," in the sense that the reduced-form residuals from the projection of w_t on its first p lags are not white noise in general.

We now state the equivalence result for finite p. The statement of the result is a simple generalization of Proposition 1, which can be thought of as the case $p = \infty$. Define the projection residual $\tilde{x}_t(p) \equiv x_t - E(x_t \mid r_t, \{w_\tau\}_{t-p\le\tau<t}) = x_t - \sum_{\ell=0}^{p} \varrho_\ell(p)' w_{t-\ell}$ (where the last $n_q + 2$ elements of $\varrho_0(p)$ are zero). Let also the operator $\text{Cov}^p(\cdot,\cdot)$ denote the covariance between any variables in the VAR that *would hypothetically* obtain if the data in fact followed a VAR(p) model with the parameters $(A(L;p), \Sigma_u(p))$ defined above.

Proposition 2. *Impose Assumptions 1 and 2. Let the nonnegative integers h, p satisfy $h \le p$. Then $\theta_h(p) = \sqrt{E(\tilde{x}_t(p)^2)} \times \beta_h(p) + \phi_h(p)$, where the remainder is given by $\phi_h(p) = \{E(\tilde{x}_t(p)^2)\}^{-1/2} \sum_{\ell=p-h+1}^{p} \{\text{Cov}(y_{t+h}, w_{t-\ell}) - \text{Cov}^p(y_{t+h}, w_{t-\ell})\}\varrho_\ell(p)$.*

Proof. Please see Appendix A.1. □

Thus, if long lags of the data do not help to predict the impulse variable x_t – i.e., when $\varrho_\ell(p) = 0$ for all $\ell \ge p-h+1$ – then the population LP and VAR impulse response estimands agree at all horizons $h \le p$, although generally not at horizons $h > p$. This finding would not be surprising if the true DGP were assumed to be a finite-order VAR (as in Jordà, 2005, Sec. I.B, and Kilian & Lütkepohl, 2017, Ch. 12.8), but we allow for general covariance stationary DGPs. The reason why the result still goes through is that a VAR(p) obtained through least-squares projections perfectly captures the autocovariances of the data out to lag p (but not further), and these are precisely what determine the LP estimand.[15] For example, if $p = 2$, then $\text{Cov}^p(y_t, x_{t-2}) = \text{Cov}(y_t, x_{t-2})$, but generally $\text{Cov}^p(y_t, x_{t-3}) \neq \text{Cov}(y_t, x_{t-3})$.

Proposition 2 implies that LP and VAR impulse response estimands will agree approximately at short horizons for a wide range of empirically relevant DGPs. If, as in many applications, x_t is a direct measure of a "shock" and thus uncorrelated with r_t and all past data, then necessarily $\phi_h(p) = 0$ and so the LP/VAR equivalence holds exactly out to horizon h. More generally, the LP estimand projects y_{t+h} onto $\tilde{x}_t(p)$; thus, the projection depends on the first $p + h$ autocovariances of the data. The estimated VAR(p) generally does not precisely capture the autocovariances of the data at lags $p+1, \dots, p+h$, and so the LP and VAR may not agree exactly. However, as we illustrate in Section 2.4, empirically relevant DGPs often have $\varrho_\ell(p) \approx 0$ for long lags ℓ, since it is typically only the first few lags of the data w_t that contribute substantially to forecasting x_t. But then $\phi_h(p) \approx 0$ for $h \ll p$, so the LP/VAR equivalence holds approximately at short horizons (in particular, $\phi_0(p) = 0$).

In conclusion, even if we use "too short" a lag length p, the LP and VAR impulse response estimands only disagree materially at horizons longer than p. This is a comforting fact in applications where the questions of interest revolve around short-horizon impulse responses.

[15] Baek & Lee (2020) prove a similar result for the related but distinct setting of single-equation Autoregressive Distributed Lag models with a white noise exogenous regressor.

Illustration: Population equivalence of VAR and LP estimands

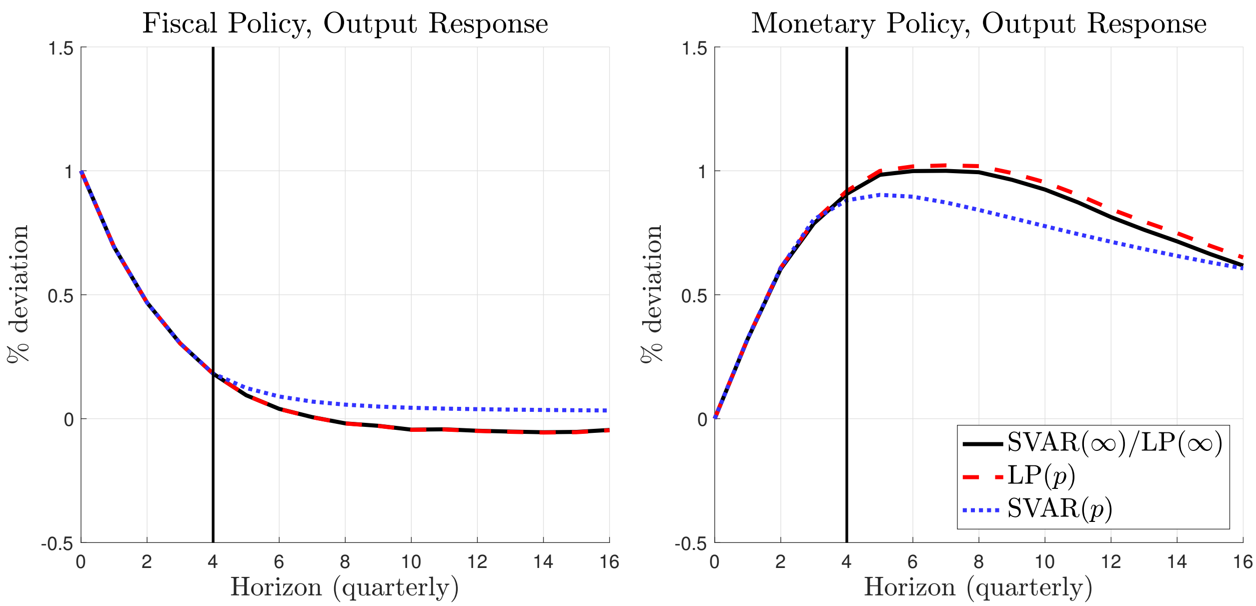

Figure 1: LP and VAR impulse response estimands in the structural model of Smets & Wouters (2007). Left panel: response of output to a government spending innovation. Right panel: response of output to a negative interest rate innovation. The horizontal line marks the horizon p after which the finite-lag-length LP(p) and VAR(p) estimands diverge.

2.4 Graphical illustration

In this section we graphically illustrate our previous theoretical results in the context of a particular data generating process: the structural macro model of Smets & Wouters (2007). We abstract from sampling uncertainty and throughout assume that the econometrician actually observes an infinite amount of data.[16] Since this section is merely intended to illustrate the properties of different projections, we do not comment on the relation of the projection estimands to true *structural* model-implied impulse responses. We formally discuss structural identification in Section 3.

The left panel of Figure 1 shows LP and VAR impulse response estimands of the response of output to an innovation in government spending. We assume the model's government spending innovation is directly observed by the econometrician, who additionally controls for lags of output and the fiscal spending innovation itself. This experiment is therefore

[16] We use the Dynare replication of Smets & Wouters (2007) kindly provided by Johannes Pfeifer. The code is available at: `https://sites.google.com/site/pfeiferecon/dynare`. We truncate the model-implied vector moving average representation at a large horizon ($H = 350$), and then invert to obtain a VAR(∞).

similar in spirit to that of Ramey (2011). As ensured by Proposition 1, the LP(∞) and VAR(∞) estimands – i.e., with infinitely many lags as controls – agree at all horizons. Since by assumption the "impulse" variable x_t is a direct measure of the government spending innovation, we have $\tilde{x}_t(p) = x_t$. Thus, any LP(p) estimand for finite p also agrees with the LP(∞) limit at all horizons. Finally, we observe that the impulse responses implied by a VAR(4) *exactly* agree with the true population projections up until horizon $h = 4$, as predicted by Proposition 2.

The right panel of Figure 1 shows LP and VAR impulse response estimands for the response of output to an innovation in the nominal interest rate. Here the model's innovation is not directly observed by the econometrician, only the interest rate. The LP specifications control for the contemporaneous value of output and inflation as well as lags of output, inflation, and the nominal interest rate; as discussed, this set of control variables is equivalent to ordering the interest rate last in the VAR. Thus, the experiment emulates the familiar monetary policy shock identification analysis of Christiano et al. (2005), although we, at least for the purposes of this section, interpret the projections purely in a reduced-form way. Again, the LP(∞) and VAR(∞) estimands agree at all horizons. Now, however, the "impulse" $\tilde{x}_t(p)$ upon which the LP(p) and VAR(p) methods project depends on the lag length p, so these estimands differ from the LP(∞)/VAR(∞) estimands. Furthermore, the LP(p) and VAR(p) estimands also differ from each other, as the remainder term $\phi_h(p)$ in Proposition 2 is not exactly zero. Nevertheless, because distant lags of w_t do not contribute substantially to forecasting x_t in this DGP, all impulse response estimands are *nearly* identical until the truncation horizon $p = 4$, consistent with the discussion in Section 2.3.

2.5 Sample equivalence

In addition to being identical *conceptually* and in *population*, we show in the Online Appendix that the difference between the local projection and VAR impulse response estimators converges to zero asymptotically *in sample* when large lag lengths are used in the regression specifications. Formally, let $\hat{\beta}_h(p)$ and $\hat{\theta}_h(p)$ denote the least-squares estimators of the LP and VAR specifications (1)–(3) if we include p lags of the data in the VAR and on the right-hand side of the local projection. Under standard nonparametric regularity conditions (Lewis & Reinsel, 1985), the sample analogue of the population equivalence result in Section 2.1 holds: There exists a constant of proportionality $\hat{\kappa}$ such that, at any fixed horizon h, the distance $|\hat{\theta}_h(p) - \hat{\kappa}\hat{\beta}_h(p)|$ tends to zero in probability asymptotically, provided that the lag length p tends to infinity with the sample size at an appropriate rate. We relegate the details of this result to the Online Appendix.

Combining Proposition 1 and the logic in Newey (1994), we conjecture that the least-squares local projection and VAR impulse response estimators $\hat{\beta}_h(p)$ and $\hat{\theta}_h(p)$ are equally asymptotically *efficient* (at a fixed horizon h) under weak regularity conditions, provided that the lag length p tends to infinity at an appropriate rate with the sample size. In finite samples, and with finite lag lengths, the two estimation methods are likely to agree approximately at short horizons due to Proposition 2, but the choice between the procedures at long horizons requires navigating a bias-variance trade-off.[17] Much more research is warranted on the optimal way to resolve the bias-variance trade-off in practice.[18]

3 Structural identification of impulse responses

We now show that our result on the population equivalence of LP and VAR impulse response functions has important implications for structural identification. The problem of structural identification is a population concept and thus logically distinct from the choice of finite-sample dimension reduction technique. We apply our equivalence result to show that popular "SVAR" identification schemes – including short-run restrictions, long-run restrictions, and sign restrictions – can equivalently be carried out using LPs. Conversely, invertibility-robust structural estimation with an external instrument (proxy) is also possible using VARs, not just LPs.

3.1 Structural model

To discuss structural identification, we impose a linear but otherwise general semiparametric Structural Vector Moving Average (SVMA) model. This model does not restrict the *linear* transmission mechanism of shocks to observed variables (see the Online Appendix for a discussion of non-linear models). SVMA models have been analyzed by Stock & Watson (2018), Plagborg-Møller & Wolf (2019), and many others. The class of SVMA models encompasses

[17] A VAR with p lags extrapolates its long-run responses from the first p sample autocovariances, which are estimated more precisely than the long-lag sample autocovariances used by the local projection estimator. Hence, the VAR impulse response estimates at long horizons typically have lower variance than local projections, but potentially higher asymptotic bias if the lag length is misspecified (though this depends on the magnitude of the misspecification, see Schorfheide, 2005, and Kilian & Kim, 2011).

[18] The forecasting literature offers extensive guidance on the choice between direct forecasts (analogous to LP) and iterated forecasts (analogous to VAR) (Schorfheide, 2005; Chevillon, 2007; Marcellino et al., 2006; Pesaran et al., 2011; McElroy, 2015).

all discrete-time, linearized DSGE models as well as all stationary SVAR models.

Assumption 3. *The data $\{w_t\}$ are driven by an n_ε-dimensional vector $\varepsilon_t = (\varepsilon_{1,t}, \ldots, \varepsilon_{n_\varepsilon,t})'$ of exogenous structural shocks,*

$$w_t = \mu + \Theta(L)\varepsilon_t, \quad \Theta(L) \equiv \sum_{\ell=0}^{\infty} \Theta_\ell L^\ell, \tag{10}$$

where $\mu \in \mathbb{R}^{n_w \times 1}$, $\Theta_\ell \in \mathbb{R}^{n_w \times n_\varepsilon}$, and L is the lag operator. $\{\Theta_\ell\}_\ell$ is assumed to be absolutely summable, and $\Theta(x)$ has full row rank for all complex scalars x on the unit circle. For notational simplicity, we further assume normality of the shocks:

$$\varepsilon_t \overset{i.i.d.}{\sim} N(0, I_{n_\varepsilon}). \tag{11}$$

Under these assumptions w_t is a nonsingular, strictly stationary jointly Gaussian time series, consistent with Assumptions 1 and 2 in Section 2. The (i, j) element $\Theta_{i,j,\ell}$ of the $n_w \times n_\varepsilon$ moving average coefficient matrix Θ_ℓ is the impulse response of variable i to shock j at horizon ℓ.

The researcher is interested in the propagation of the structural shock $\varepsilon_{1,t}$ to the observed macro aggregate y_t. Since y_t is the $(n_r + 2)$-th element in w_t, the parameters of interest are $\Theta_{n_r+2,1,h}$, $h = 0, 1, 2, \ldots$. In line with applied work, we also consider *relative* impulse responses $\Theta_{n_r+2,1,h}/\Theta_{n_r+1,1,0}$. This may be interpreted as the response in y_{t+h} caused by a shock $\varepsilon_{1,t}$ of a magnitude that raises x_t by one unit on impact.

3.2 Implementing "SVAR" identification using LPs

In this subsection we show that LP methods are as applicable as VAR methods when implementing common identification schemes. Our main result in Section 2.1 implies that LP-based causal estimation can succeed if and only if SVAR-based estimation can succeed. We will exhibit several concrete and easily implementable examples of this equivalence.

Identification under invertibility. Standard SVAR analysis assumes (partial) invertibility – that is, the ability to recover the structural shock of interest, $\varepsilon_{1,t}$, as a function of only current and past macro aggregates:

$$\varepsilon_{1,t} \in \operatorname{span}\left(\{w_\tau\}_{-\infty<\tau\leq t}\right). \tag{12}$$

A given SVAR identification scheme then identifies as the candidate structural shock a particular linear combination of the Wold forecast errors:

$$\tilde{\varepsilon}_{1,t} \equiv b' u_t, \tag{13}$$

where the chosen identification scheme gives the vector b as a function of the reduced-form VAR parameters $(A(L), \Sigma_u)$, or equivalently the Wold decomposition parameters $(C(L), \Sigma_u)$. Under invertibility, there must exist a vector b such that $\tilde{\varepsilon}_{1,t} = \varepsilon_{1,t}$, so SVAR identification can in principle succeed (Fernández-Villaverde et al., 2007; Wolf, 2019).

We now illustrate through three examples that common SVAR identification schemes are equally as simple to implement using LP methods. We first consider a standard recursive scheme covered by our benchmark analysis in Section 2.1. The second and third examples involve long-run and sign restrictions and require the general equivalence result of Section 2.2.

Example 1 (Recursive identification)**.** Christiano et al. (2005) identify monetary policy shocks through a recursive ordering. They assume that their observed data $\{w_t\}$ follow an invertible SVMA model, i.e. the condition (12) holds for all shocks in the system (10). They then additionally impose a temporal ordering on the set of variables w_t: Output, consumption, investment, wages, productivity, and the price deflator do not respond within the period to changes in the policy rate (Federal Funds Rate), which itself in turn does not react within the period to changes in profits and money growth. In the notation of Section 2.1, the assumed ordering corresponds to the Federal Funds Rate as the impulse variable x_t, real and price variables as the controls r_t, and financial variables collected in the vector q_t. Thus, for the purposes of structural interpretation, it is now explicitly assumed that x_t and r_t are predetermined with respect to y_t. Christiano et al. implement their structural analysis through the recursive VAR (3). By our main result, they could have equivalently estimated the LP regression (1) and collected the regression coefficients $\{\beta_h\}_{h \geq 0}$. The population estimand would have been the same. $\square$

Example 2 (Long-run identification)**.** Blanchard & Quah (1989) identify the effects of demand and supply shocks using long-run restrictions in a bivariate system. Let gdp_t and unr_t denote log real GDP (in levels) and the unemployment rate, respectively. Then $\Delta gdp_t \equiv gdp_t - gdp_{t-1}$ is log GDP growth. Blanchard & Quah impose that $w_t \equiv (\Delta gdp_t, unr_t)'$ follows the SVMA model in Assumption 3 with $n_\varepsilon = 2$ shocks, where the first shock is a supply shock, the second shock a demand shock, and both shocks are invertible, cf. (12). They then additionally impose the identifying restriction that the long-run effect of the demand shock on the *level* of output is zero, i.e., $\sum_{\ell=0}^{\infty} \Theta_{1,2,\ell} = 0$.

While Blanchard & Quah impose their long-run restriction on an SVAR model to estimate impulse responses, the extended equivalence result in Section 2.2 implies that the same restriction can be equivalently implemented using an LP approach. To see how, consider, for a large horizon H, the "long difference" projection

$$gdp_{t+H} - gdp_{t-1} = \tilde{\mu}_H + \tilde{\beta}_H' w_t + \sum_{\ell=1}^{\infty} \tilde{\delta}_{H,\ell}' w_{t-\ell} + \tilde{\xi}_{H,t}. \tag{14}$$

Intuitively, this projection uncovers the linear combination of the data that best explains long-run movements in GDP. By assumption, such explanatory power can only come from the supply shock. Thus, to estimate impulse responses with respect to the supply shock, we can run the local projection (9) with $b = \tilde{\beta}_H$ and with y_t given by the response variable of interest (either Δgdp_t or unr_t). Indeed, we show formally in Appendix A.2 that, as $H \to \infty$, this procedure correctly identifies the impulse responses $\Theta_{i,1,h}$ with respect to the supply shock, up to a constant scale factor. In this way, relative impulse responses $\Theta_{i,1,h}/\Theta_{1,1,0}$ are correctly identified.[19] To estimate relative impulse responses $\Theta_{i,2,h}/\Theta_{1,2,0}$ to the demand shock, the researcher can choose any vector $\tilde{b}$ such that $\tilde{b}'b = 0$, and then implement the local projection (9) with $\tilde{b}$ in lieu of b.

In finite samples, the mean squared error performance of the proposed procedure relative to the conventional SVAR(p) approach of Blanchard & Quah (1989) will depend on the tuning parameters H and p, and on whether the low-frequency properties of the data are well approximated by a low-order VAR model.[20] For researchers who prioritize bias over variance, the LP approach to long-run restrictions has the advantage that it does not extrapolate long-run impulse responses from short-run autocorrelations, as a finite-order VAR does. $\square$

Example 3 (Sign identification)**.** Uhlig (2005) and Rubio-Ramírez et al. (2010) consider set-identification via sign restrictions on impulse responses. For concreteness, suppose we are interested in the impulse response of real GDP growth y_t to a monetary shock $\varepsilon_{1,t}$ at horizon h. As before, assume that the full set of observed data $\{w_t\}$ follows an SVMA system (10) where all shocks are invertible. As a very simple example of sign restrictions, we may impose the identifying restriction that the scalar variable r_t (say, the nominal interest rate) responds positively to a monetary shock at all horizons $s = 0, 1, \ldots, \bar{H}$.

[19] Absolute impulse responses can be identified by rescaling the identified shock so it has variance 1.

[20] Christiano et al. (2006) and Mertens (2012) make the related point that SVAR-based long-run identification could employ nonparametric estimators of the long-run variance matrix instead of the VAR estimator.

The traditional SVAR approach to sign identification proceeds as follows. By invertibility, the monetary shock $\varepsilon_{1,t}$ is related to the Wold forecast errors u_t through $\varepsilon_{1,t} = b'u_t$, where $b \in \mathbb{R}^{n_w}$ is an unknown vector. If we knew b, the structural impulse responses of any variable $w_{i,t}$ to $\varepsilon_{1,t}$ could be obtained as the linear combination b of the reduced-form impulse responses of $w_{i,t}$ from a VAR in w_t. To impose the sign restrictions, we search over all possible vectors b such that (i) the r_t impulse responses are positive at all horizons $s = 0, 1, \dots, \bar{H}$ and (ii) the impact impulse response of r_t is normalized to 1 (other normalizations are also possible). Once we have determined the set of possible b's, we can then use the VAR to compute the corresponding set of possible impulse responses of y_t with respect to $b'u_t$.

By the logic in Section 2.2, we can alternatively impose sign restrictions using an LP approach. We simply estimate the reduced-form impulse responses using LPs instead of a VAR. Consider the coefficient vector $\check{\beta}_h$ obtained from the projection

$$y_{t+h} = \check{\mu}_h + \check{\beta}_h' w_t + \sum_{\ell=1}^{\infty} \check{\delta}_{h,\ell}' w_{t-\ell} + \check{\xi}_{h,t}.$$

The above LP yields the reduced-form impulse responses $\check{\beta}_h$ of y_t to the Wold forecast errors u_t. Exactly as in the VAR approach, we now seek the linear combination $b'\check{\beta}_h$ that equals the structural impulse response to the monetary shock $\varepsilon_{1,t} = b'u_t$. To find the set of b's consistent with the sign restrictions, the natural analogue of the VAR approach is as follows. For each horizon $s = 0, 1, \dots, \bar{H}$, store the coefficient vector $\ddot{\beta}_s$ from the projection

$$r_{t+s} = \ddot{\mu}_s + \ddot{\beta}_s' w_t + \sum_{\ell=1}^{\infty} \ddot{\delta}_{s,\ell}' w_{t-\ell} + \ddot{\xi}_{s,t}.$$

The coefficients $\ddot{\beta}_s$ measure the reduced-form impulse responses of r_t to u_t, so sign restrictions on the structural impulse responses of r_t amount to linear inequality restrictions on these coefficients. Consequently, the *largest possible* response of y_{t+h} to a monetary shock that raises r_t by one unit on impact can be obtained as the solution to the linear program[21]

$$\sup_{b \in \mathbb{R}^{n_w}} b'\check{\beta}_h \quad \text{subject to} \quad b'\ddot{\beta}_0 = 1,$$
$$b'\ddot{\beta}_s \geq 0, \; s = 1, \dots, \bar{H}.$$

[21]To consider impulse responses to a *one-standard-deviation* monetary shock, replace the equality constraint in the linear program by the constraint $b' \operatorname{Var}(u_t)^{-1} b = 1$. The resulting linear-quadratic program with inequality constraints is similar to those in Gafarov et al. (2018) and Giacomini & Kitagawa (2020).

To compute the *smallest possible* impulse response, replace the supremum with an infimum.[22] In population, this LP-based procedure recovers exactly the same identified set as analogous sign restrictions in an SVAR. □

In the Online Appendix we show how to perform local projection identification with a completely general system of sign, zero, and magnitude restrictions on multiple variables, shocks, and horizons, as in the SVAR frameworks of Rubio-Ramírez et al. (2010) and Kilian & Murphy (2012). There we also argue that narrative restrictions as in Antolín-Díaz & Rubio-Ramírez (2018) and Ludvigson et al. (2020b) can be exploited in a local projection framework, by recasting those restrictions in the instrumental variable framework analyzed in Section 3.3 below.

Our examples demonstrate that invertibility-based identification need not be thought of as "SVAR identification," contrary to standard practice in textbooks and parts of the literature. As a matter of identification (i.e., in population), the two methods succeed or fail together. Ideally, researchers ought to decide on the identification scheme separately from how they decide on the finite-sample dimension reduction technique. The former choice should be based on economic theory. The latter choice should be based on the researcher's preferences over bias and variance as well as on features of the DGP.

Beyond invertibility. If the invertibility assumption (12) is violated, then identification strategies that erroneously assume invertibility – independent of whether they are implemented using VARs or LPs – will not measure the true impulse responses.[23] Instead, these methods will measure the impulse responses to a white noise disturbance that is a linear combination of current and *lagged* true structural shocks:

$$\tilde{\varepsilon}_{1,t} = \vartheta(L)\varepsilon_t. \tag{15}$$

The properties of the lag polynomial $\vartheta(L)$ are characterized in detail in Fernández-Villaverde et al. (2007) and Wolf (2019). Combining (10) and (15), we see that, in general, both LP and VAR impulse response estimands are linear combinations of contemporaneous and lagged

[22] We focus on computing the *bounds* of the identified set. An alternative approach is to sample from the identified set, as is commonly done in the Bayesian SVAR literature (Rubio-Ramírez et al., 2010).

[23] Several recent papers have demonstrated how to perform valid semi-structural identification without assuming invertibility, cf. the references in Plagborg-Møller & Wolf (2019). Often such methods rely on LP or VAR techniques to compute relevant linear projections, without interpreting the VAR disturbances (i.e., Wold innovations) as linear combinations of the contemporaneous true shocks.

true impulse responses. Thus, projection on a given identified impulse $\tilde{\varepsilon}_{1,t}$ correctly identifies impulse response functions (up to scale) if and only if $\tilde{\varepsilon}_{1,t}$ affects the response variable y_t *only* through the contemporaneous true structural shock $\varepsilon_{1,t}$. Trivially, this is the case if $\tilde{\varepsilon}_{1,t}$ is a function only of $\varepsilon_{1,t}$ (the invertible case); less obviously, the same is also true if $\tilde{\varepsilon}_{1,t}$ is only contaminated by shocks that do not directly affect the response variable y_t.[24] Instrumental variable identification, discussed next, is the leading example of this second case.

3.3 Identification and estimation with instruments

Instruments (also known as proxy variables) are popular in semi-structural analysis. We here use our main result in Section 2 to show that the influential Local Projection Instrumental Variable estimation procedure is equivalent to estimating a VAR with the instrument ordered first. This is true irrespective of the underlying structural model.

An instrumental variable (IV) is defined as an observed variable z_t that is contemporaneously correlated *only* with the shock of interest $\varepsilon_{1,t}$, but not with other shocks that affect the macro aggregate y_t of interest (Stock, 2008; Stock & Watson, 2012; Mertens & Ravn, 2013).[25] More precisely, given Assumption 3, the IV exclusion restrictions are that

$$\text{Cov}(z_t, \varepsilon_{j,s} \mid \{z_\tau, w_\tau\}_{-\infty<\tau<t}) \neq 0 \quad \text{if and only if} \quad \text{both } j = 1 \text{ and } t = s. \tag{16}$$

Stock & Watson (2018, p. 926) refer to this assumption as "LP-IV$^\perp$," and it is routinely made in theoretical and applied work, as reviewed by Ramey (2016) and Stock & Watson (2018). The assumption requires that, once we control for all lagged data, the instrument is not contaminated by other structural shocks or by lags of the shock of interest.

Without loss of generality, we can use projection notation to phrase the IV exclusion restrictions (16) as follows.

Assumption 4.

$$z_t = c_z + \sum_{\ell=1}^{\infty}(\Psi_\ell z_{t-\ell} + \Lambda_\ell w_{t-\ell}) + \alpha\varepsilon_{1,t} + v_t, \tag{17}$$

where $\alpha \neq 0$, $c_z, \Psi_\ell \in \mathbb{R}$, $\Lambda_\ell \in \mathbb{R}^{1\times n_w}$, $v_t \overset{i.i.d.}{\sim} N(0, \sigma_v^2)$, *and* v_t *is independent of* ε_t *at all*

[24]In particular, this means that neither invertibility nor recoverability (as defined in Plagborg-Møller & Wolf, 2019) are *necessary* for successful semi-structural inference on impulse response functions.

[25]We focus on the case of a single IV. If multiple IVs for the same shock are available, Plagborg-Møller & Wolf (2019) show that (i) the model is testable, and (ii) all the identifying power of the IVs is preserved by collapsing them to a certain (single) linear combination.

leads and lags. The lag polynomial $1 - \sum_{\ell=1}^{\infty} \Psi_\ell L^\ell$ *is assumed to have all roots outside the unit circle, and* $\{\Lambda_\ell\}_\ell$ *is absolutely summable.*

Crucially, the assumption allows the IV to be contaminated by the independent measurement error v_t. In some applications, we may know by construction of the IV that the lag coefficients Ψ_ℓ and Λ_ℓ are all zero (so z_t satisfies assumption "LP-IV" of Stock & Watson, 2018, p. 924, without controls); obviously, such additional information will not present any difficulties for any of the arguments that follow.

LP-IV. The Local Projection Instrumental Variable (LP-IV) approach estimates the impulse responses to the first shock using a two-stage least squares version of LP. Loosely, Mertens (2015), Jordà et al. (2015, 2019), Leduc & Wilson (2017), Ramey & Zubairy (2018), and Stock & Watson (2018) propose to estimate the LP equation (1) using z_t as an IV for x_t. To describe the two-stage least-squares estimand in detail, define $W_t \equiv (z_t, w_t')'$ and consider the "reduced-form" IV projection

$$y_{t+h} = \mu_{RF,h} + \beta_{RF,h} z_t + \sum_{\ell=1}^{\infty} \delta'_{RF,h,\ell} W_{t-\ell} + \xi_{RF,h,t} \tag{18}$$

for any $h \geq 0$. Consider also the "first-stage" IV projection[26]

$$x_t = \mu_{FS} + \beta_{FS} z_t + \sum_{\ell=1}^{\infty} \delta'_{FS,\ell} W_{t-\ell} + \xi_{FS,t}. \tag{19}$$

Notice that the first stage does not depend on the horizon h. As in standard cross-sectional two-stage least-squares estimation, the LP-IV estimand is then given by the ratio $\beta_{LPIV,h} \equiv \beta_{RF,h}/\beta_{FS}$ of reduced-form to first-stage coefficients (e.g., Angrist & Pischke, 2009, p. 122).[27]

Stock & Watson (2018) show that, under Assumptions 3 and 4, the LP-IV estimand $\beta_{LPIV,h}$ correctly identifies the *relative* impulse response $\Theta_{n_r+2,1,h}/\Theta_{n_r+1,1,0}$. Importantly, this holds whether or not the shock of interest $\varepsilon_{1,t}$ is invertible in the sense of equation (12).

We now use our main result from Section 2.1 to show that the LP-IV impulse responses can equivalently be estimated from a recursive VAR that orders the IV first. As in Section 2, this result imposes no functional form assumptions on the underlying structural model and does not yet impose any exclusion restrictions on the IV z_t (although such assumptions are

[26] As always, the coefficients and residuals in (18)–(19) should be interpreted as linear projections.

[27] In the over-identified case with multiple IVs, the IV estimand can no longer be written as this simple ratio; we focus on a single IV as in most of the applied literature.

of course required to interpret the estimand causally, as discussed below).

Corollary 1. *Let Assumptions 1 and 2 hold for the expanded data vector $W_t \equiv (z_t, w_t')'$ in place of w_t. Assume also that $\beta_{FS} \neq 0$, cf. (19). Consider a recursively ordered SVAR(∞) in the variables $(z_t, w_t')'$, where the instrument is ordered first (the ordering of the other variables does not matter). Let $\tilde{\theta}_{y,h}$ be the SVAR-implied impulse response at horizon h of y_t with respect to the first shock. Let $\tilde{\theta}_{x,0}$ be the SVAR-implied impact impulse response of x_t with respect to the first shock.*

Then $\tilde{\theta}_{y,h}/\tilde{\theta}_{x,0} = \beta_{LPIV,h}$.

Proof. Let $\tilde{z}_t \equiv \alpha\varepsilon_{1,t} + v_t$ and $a \equiv \sqrt{E(\tilde{z}_t^2)} = \sqrt{\alpha^2 + \sigma_v^2}$. Proposition 1 states that $\tilde{\theta}_{y,h} = a \times \beta_{RF,h}$ for all h, and $\tilde{\theta}_{x,0} = a \times \beta_{FS}$. The claim follows. □

This essentially reduced-form result implies that, if we additionally impose the *structural* Assumptions 3 and 4, valid identification of relative structural impulse responses can equivalently be achieved through LP-IV or through an "internal instrument" recursive SVAR with the IV ordered first.[28] Importantly, under Assumptions 3 and 4, these equivalent estimation strategies are valid even when the shock of interest $\varepsilon_{1,t}$ is not invertible (Stock & Watson, 2018). Intuitively, although adding the IV z_t to the VAR does not render the shock $\varepsilon_{1,t}$ invertible, the only reason that the shock may be non-invertible with respect to the *expanded* information set $\{z_\tau, w_\tau\}_{-\infty<\tau\leq t}$ is the presence of the measurement error v_t in the IV equation (17).[29] But this independent measurement error merely leads to attenuation bias in the estimated impulse responses, and the bias (in percentage terms) is the same at all response horizons and for all response variables. Thus, it does not contaminate estimation of *relative* impulse responses.

IV identification is therefore an example of a setting where SVAR analysis works even though invertibility fails (including the partial invertibility notion of Forni et al., 2019, and Miranda-Agrippino & Ricco, 2019). The "internal instrument" recursive SVAR(∞) procedure estimates the right *relative* impulse responses despite the fact that no invertible structural VAR model generally exists under our assumptions. Our result implies that it

[28] Plagborg-Møller & Wolf (2019) show that point identification of *absolute* impulse responses – and thus variance decompositions – can be achieved under a further *recoverability* assumption that is mathematically and substantively weaker than assuming invertibility.

[29] Note that, even though Assumption 4 allows z_t to be correlated with lags of w_t, non-invertibility of $\varepsilon_{1,t}$ is entirely consistent with Theorem 1 of Stock & Watson (2018). That theorem states that if the shock is non-invertible, then it is *possible to construct an example* of an IV $\check{z}_t$ satisfying $E(\check{z}_t\varepsilon_{j,t}) = 0$ for all $j \neq 1$ and $E(\check{z}_t\varepsilon_{j,t-\ell} \mid \{w_\tau\}_{\tau<t}) \neq 0$ for some j and $\ell \geq 1$ (so $\check{z}_t$ does not satisfy Assumption 4).

is valid to include an externally identified shock in an SVAR even if the shock is measured with (independent) error, as long as the noisily measured shock is ordered first.[30]

SVAR-IV. Unlike the non-invertibility-robust procedure of ordering the IV first in a VAR, the popular SVAR-IV (also known as proxy-SVAR) procedure (Stock, 2008; Stock & Watson, 2012; Mertens & Ravn, 2013) is only valid under invertibility. This procedure uses an SVAR to identify the shock of interest as

$$\tilde{\varepsilon}_{1,t} \equiv \frac{1}{\sqrt{\mathrm{Var}(\tilde{z}_t^\dagger)}} \times \tilde{z}_t^\dagger,$$

where $\tilde{z}_t^\dagger$ is computed as a linear combination of the reduced-form residuals u_t from a VAR in w_t alone (i.e., excluding the IV from the VAR):

$$\tilde{z}_t^\dagger \equiv E(\tilde{z}_t \mid u_t) = E(\tilde{z}_t \mid \{w_\tau\}_{-\infty<\tau\leq t}).$$

If Assumptions 3 and 4 and the invertibility condition (12) hold, then SVAR-IV is valid. In fact, in this case SVAR-IV removes any attenuation bias, thus correctly identifying *absolute* (not just relative) impulse responses.[31] However, in the general non-invertible case, SVAR-IV mis-identifies the shock as $\tilde{\varepsilon}_{1,t} \neq \varepsilon_{1,t}$.[32] Plagborg-Møller & Wolf (2019, Appendix B.4) characterize the bias of SVAR-IV under non-invertibility and show that the invertibility assumption can be tested using the IV.

Summary. The relative impulse responses obtained from the LP-IV procedure of Stock & Watson (2018) are nonparametrically identical to the relative impulse responses from a recursive SVAR with the IV ordered first (an "internal instrument" approach). Assuming an SVMA model and the IV exclusion restrictions, these procedures correctly identify relative structural impulse responses, irrespective of the invertibility of the shock of interest. This allows researchers to exploit VAR estimation techniques while relying on the same

[30] Romer & Romer (2004) and Barakchian & Crowe (2013) include an externally identified monetary shock in an SVAR, but they order it last, which assumes additional exclusion restrictions. Kilian (2006), Ramey (2011), Miranda-Agrippino (2017), and Jarociński & Karadi (2020), among others, mention the strategy of ordering an IV first in an SVAR, but these papers do not consider the non-invertible case.

[31] Consistent with our analytical results, Carriero et al. (2015) observe in a calibrated simulation study that, under invertibility, SVAR-IV correctly identifies *absolute* impulse response functions, while direct projections on the IV suffer from attenuation bias.

[32] The VARX approach of Paul (2019) is equivalent with SVAR-IV under Assumptions 1 and 2.

invertibility-robust identifying restrictions as the popular two-stage least squares implementation of LP-IV. In contrast, the SVAR-IV procedure of Stock & Watson (2012) and Mertens & Ravn (2013) (an "external instrument" approach) requires invertibility.[33]

4 Empirical application

We finally illustrate our theoretical equivalence results by empirically estimating the dynamic response of corporate bond spreads to a monetary policy shock. We adopt the specification of Gertler & Karadi (2015), who, using high-frequency financial data, obtain an external instrument for monetary policy shocks.[34] Because of possible non-invertibility (Ramey, 2016; Plagborg-Møller & Wolf, 2019), we do not consider the external SVAR-IV estimator, but instead implement direct projections on the IV through (i) local projections and (ii) an "internal instrument" recursive VAR, following the logic of Corollary 1. In both cases, our vector of macro control variables exactly follows Gertler & Karadi (2015); it includes output growth (log growth rate of industrial production), inflation (log growth rate of CPI inflation), the 1-year government bond rate, and the Excess Bond Premium of Gilchrist & Zakrajšek (2012) as a measure of the non-default-related corporate bond spread. The data is monthly and spans January 1990 to June 2012.[35]

Figure 2 shows that LP-IV and "internal instrument" VAR impulse response estimates agree at short horizons, but diverge at longer horizons, consistent with Proposition 2. The figure shows point estimates of the response of the Excess Bond Premium to the monetary policy shock, for different projection techniques and different lag lengths. For all specifications, the Excess Bond Premium initially increases after a contractionary monetary policy shock, consistent with the results in Gertler & Karadi (2015). The left panel shows results for LP(4) and VAR(4) estimates. Up until horizon $h = 4$, the estimated impulse responses are closely aligned. At longer horizons, the iterated VAR structure enforces a smooth return to 0, while direct local projections give more erratic impulse responses. The right panel shows

[33] SVAR-IV does have one advantage over LP-IV (and thus also over the "internal instruments" VAR approach): Provided the shock is invertible, SVAR-IV does not require z_t to only be correlated with lagged shocks through *observed* lagged variables as in Assumption 4, cf. Stock & Watson (2018, sec. 2.1).

[34] The external IV z_t is constructed from changes in 3-month-ahead futures prices written on the Federal Funds Rate, where the changes are measured over short time windows around Federal Open Market Committee monetary policy announcement times. See Gertler & Karadi (2015) for details on the construction of the IV and a discussion of the exclusion restriction.

[35] The data were retrieved from: `https://www.aeaweb.org/articles?id=10.1257/mac.20130329`

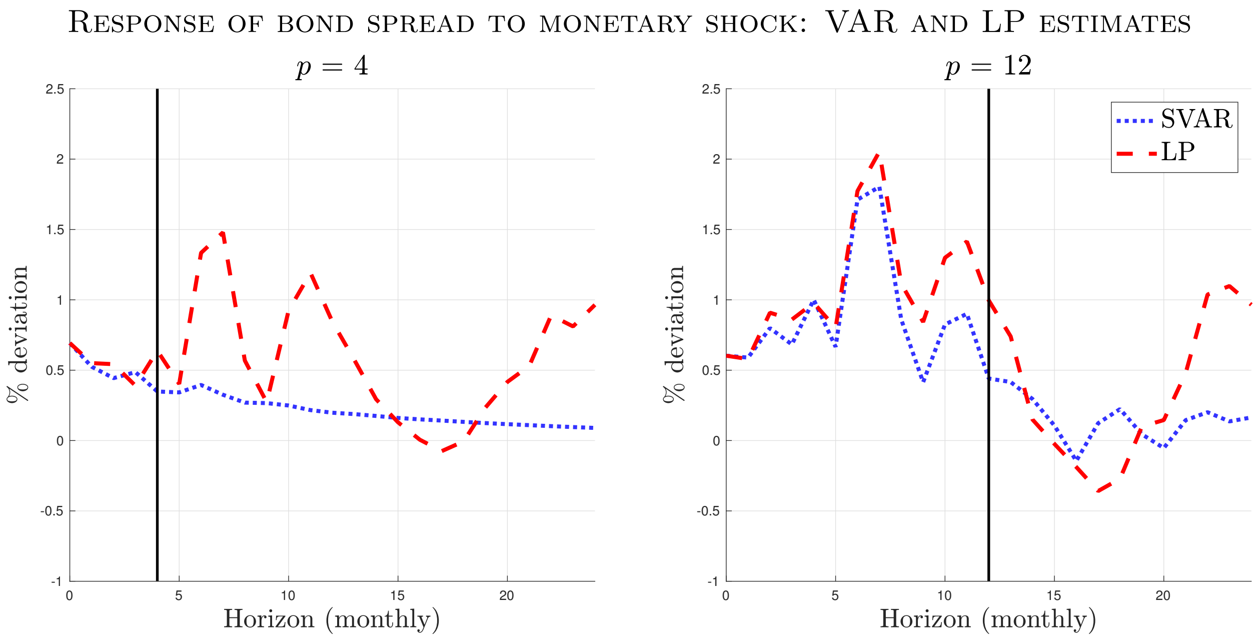

Figure 2: Estimated impulse response function of the Excess Bond Premium to a monetary policy shock, normalized to increase the 1-year bond rate by 100 basis points on impact. Left panel: lag length $p = 4$. Right panel: $p = 12$. The horizontal line marks the horizon p after which the VAR(p) and LP(p) estimates may diverge substantially.

an analogous picture for LP(12) and VAR(12) estimates: The estimated impulse responses agree closely until horizon $h = 12$, but they diverge at longer horizons.

These results provide an empirical illustration of our earlier claim that LP and VAR estimates are closely tied together at short horizons.[36] The results further illustrate that the larger the lag length used for estimation (in both the LP and VAR specifications), the more impulse response horizons will exhibit agreement between LP and VAR estimates, consistent with Section 2.3. As this exercise is merely meant to illustrate our theoretical results, we refrain from conducting statistical tests of the validity of the different regression models.

5 Conclusion

We demonstrated a general equivalence of the local projection and VAR impulse response function estimands. Although these estimation methods are linear, the equivalence between

[36]In unreported results, we confirm this finding for the impulse response functions of output growth, inflation, and the 1-year government bond rate. Ludvigson et al. (2020a), referencing our paper, reach a similar conclusion in a different empirical study.

them obtains nonparametrically, in the sense that we do not require linearity or other functional form assumptions on the data generating process. This result, which applies to all common implementations of local projections in the applied literature, has several implications for empirical practice:

1. VAR and local projection estimators of impulse responses should not be regarded as conceptually distinct methods – in population, they estimate the same thing, as long as we control flexibly for lagged data.

2. Local projections with p lags as controls and VAR(p) estimators agree approximately at impulse response horizons $h \leq p$. Hence, VARs that control for a large number of lags, as recommended by Kilian & Lütkepohl (2017), will tend to agree at short and medium-long horizons with local projections that also control for a rich set of lags. However, at long horizons the methods may disagree substantially.

3. Structural *identification* is logically distinct from the dimension reduction choices that must be made for *estimation* purposes. It may be counterproductive to follow standard practice in assuming an SVAR model whenever the discussion turns to structural identification, as this conflates the population identification analysis and the dimension reduction technique of using a VAR estimator.

4. Any structural estimation method that works for SVARs can be implemented with local projections, and *vice versa*. For example, if a paper already relies on local projections for parts of the analysis, then an additional sign restriction identification exercise, say, can also be implemented in a local projection fashion.

5. If an instrument/proxy for the shock of interest is available, structural impulse responses can be consistently estimated by ordering the instrument first in a recursive VAR (an "internal instrument" approach), even if the shock of interest is non-invertible. In contrast, the popular SVAR-IV estimator (an "external instrument" approach) is only consistent under invertibility.

6. Linear local projections and linear VARs will continue to estimate the same impulse responses even when the true data generating process is non-linear.

7. Although this paper has focused solely on population properties of estimators, the population equivalence result implies that no single estimation method should be expected to dominate in terms of mean squared error across every possible data generating

process. In finite samples, and with finite lag lengths, researchers must navigate a bias-variance trade-off at long horizons. Much more research is warranted on this key issue for applied research.

We conclude that each of the following common assertions is, in fact, *mistaken*: (i) VAR impulse response estimators are generally more efficient than local projection estimators; (ii) local projections are generally more robust to misspecification than VARs; (iii) SVAR analysis is required when implementing non-recursive, non-IV identification schemes; (iv) simple SVAR methods cannot be used when the structural shock of interest is noninvertible.

Our work points to several promising areas for future research. First, it would be useful to adapt the results in the present paper to *non-linear* estimators, such as regressions with interactions or polynomial terms. Second, future research could consider data with unit roots or cointegration. Third, we only discussed the population properties of IV estimators, and thus ignored weak IV issues. Fourth, it would be interesting to generalize our LP-IV equivalence result to settings with multiple instruments/proxies. Finally, we have deliberately avoided questions related to estimation and inference.

A Appendix

A.1 Equivalence result with finite lag length

We here prove Proposition 2 from Section 2.3. We proceed mostly as in the proof of Proposition 1. As a first step, the Frisch-Waugh theorem implies that

$$\beta_h(p) = \frac{\text{Cov}(y_{t+h}, \tilde{x}_t(p))}{E(\tilde{x}_t(p)^2)}. \tag{20}$$

We remind the reader of the notation $\text{Cov}^p(\cdot, \cdot)$, which denotes covariances of the data $\{w_t\}$ as implied by the (counterfactual) stationary "fitted" SVAR(p) model

$$A(L;p)w_t = B(p)\bar{\eta}_t, \quad \bar{\eta}_t \sim WN(0, I), \tag{21}$$

i.e., where $\bar{\eta}_t$ is truly white noise (unlike the residuals from the VAR(p) projection on the actual data). For example $\text{Cov}^p(y_t, x_{t-1})$ denotes the covariance of y_t and x_{t-1} that *would* obtain if $w_t = (r_t', x_t, y_t, q_t')'$ were generated by the model (21) with parameters $A(L;p)$ and $B(p)$ obtained from the projection on the *actual* data, as defined in Section 2.3. We similarly define any covariances that involve $\bar{\eta}_t$. Note that stationarity of the VAR model (21) follows from Brockwell & Davis (1991, Remark 2, pp. 424–425).

It follows from the argument in Brockwell & Davis (1991, p. 240) that $\text{Cov}^p(w_t, w_{t-h}) = \text{Cov}(w_t, w_{t-h})$ for all $h \leq p$ (see also Brockwell & Davis, 1991, Remark 2, pp. 424–425 for the multivariate generalization of the key step in the argument). In words, the autocovariances implied by the "fitted" SVAR(p) model (21) agree with the autocovariances of the actual data out to lag p, although generally not after lag p.

Under the counterfactual model (21), we have the moving average representation $w_t = C(L;p)B(p)\bar{\eta}_t$, and thus

$$\theta_h(p) = C_{n_r+2,\bullet,h}(p)B_{\bullet,n_r+1}(p) = \text{Cov}^p(y_{t+h}, \bar{\eta}_{x,t}), \tag{22}$$

where $\bar{\eta}_{x,t}$ is the (n_r+1)-th element of $\bar{\eta}_t$. Since $B(p)$ is lower triangular by definition, it is straight-forward to show from (21) that

$$B_{n_r+1,n_r+1}(p)\bar{\eta}_{x,t} = x_t - E^p(x_t \mid r_t, \{w_\tau\}_{t-p\leq\tau<t}) = x_t - E(x_t \mid r_t, \{w_\tau\}_{t-p\leq\tau<t}) = \tilde{x}_t(p), \tag{23}$$

where $E^p(\cdot \mid \cdot)$ denotes linear projection under the inner product $\text{Cov}^p(\cdot,\cdot)$, the second equality follows from the above-mentioned equivalence of $\text{Cov}^p(\cdot,\cdot)$ and $\text{Cov}(\cdot,\cdot)$ out to lag p, and the last equality follows by definition. Since $\text{Cov}^p(\bar{\eta}_{x,t}, \bar{\eta}_{x,t}) = 1$, equation (23) implies

$$B_{n_r+1,n_r+1}(p)^2 = \text{Cov}^p(\tilde{x}_t(p), \tilde{x}_t(p)) = E(\tilde{x}_t(p)^2),$$

where the last equality again uses the equivalence of $\text{Cov}^p(\cdot,\cdot)$ and $\text{Cov}(\cdot,\cdot)$ out to lag p. Putting together (22), (23), and the above equation, we have shown that

$$\theta_h(p) = \frac{1}{\sqrt{E(\tilde{x}_t(p)^2)}} \times \text{Cov}^p(y_{t+h}, \tilde{x}_t(p)).$$

If we compare with the expression (20), the desired conclusion follows from

$$\text{Cov}^p(y_{t+h}, \tilde{x}_t(p)) - \text{Cov}(y_{t+h}, \tilde{x}_t(p)) = \sum_{\ell=0}^{p} \underbrace{\{\text{Cov}(y_{t+h}, w_{t-\ell}) - \text{Cov}^p(y_{t+h}, w_{t-\ell})\}}_{=0 \text{ when } (t+h)-(t-\ell)\leq p} \varrho_\ell(p),$$

where we have used the notation for the projection coefficients $\varrho_\ell(p)$ defined immediately above Proposition 2, and we have again appealed to the equivalence of $\text{Cov}^p(\cdot,\cdot)$ with the covariance function of the actual data. $\square$

A.2 Long-run identification using local projections

Here we show that the LP-based long-run identification approach in Example 2 is valid. Define the Wold innovations $u_t \equiv w_t - E(w_t \mid \{w_\tau\}_{-\infty<\tau<t})$ and Wold decomposition

$$w_t = \chi + C(L)u_t, \quad C(L) \equiv I_2 + \sum_{\ell=1}^{\infty} C_\ell L^\ell. \tag{24}$$

Since both structural shocks are assumed to be invertible, there exists a 2×2 matrix B such that $\varepsilon_t = Bu_t$. Comparing (10) and (24), we then have $\Theta(1)B = C(1)$. Let $e_1 \equiv (1, 0)'$. Note that the Blanchard & Quah assumption $e_1'\Theta(1) = (\Theta_{1,1}(1), 0)$ implies

$$e_1'C(1) = e_1'\Theta(1)B = \Theta_{1,1}(1)e_1'B,$$

and therefore

$$e_1'C(1)u_t = \Theta_{1,1}(1) \times e_1'Bu_t = \Theta_{1,1}(1) \times \varepsilon_{1,t}.$$

By the result in Section 2.2, the claim in Example 2 follows if we show that

$$\lim_{H\to\infty} \tilde{\beta}_H' = e_1'C(1). \tag{25}$$

Define $\Sigma_u \equiv \text{Var}(u_t)$. Applying the Frisch-Waugh theorem to the projection (14), and using $w_{1,t} = \Delta gdp_t$, we find

$$\tilde{\beta}_H' = \text{Cov}(gdp_{t+H} - gdp_{t-1}, u_t)\Sigma_u^{-1} = \text{Cov}\left(\sum_{\ell=0}^{H} w_{1,t+\ell}, u_t\right)\Sigma_u^{-1} = \sum_{\ell=0}^{H} \text{Cov}(w_{1,t+\ell}, u_t)\Sigma_u^{-1}. \tag{26}$$

On the other hand, the Wold decomposition (24) implies (recall that u_t is white noise)

$$\sum_{\ell=0}^{\infty} \text{Cov}(w_{t+\ell}, u_t)\Sigma_u^{-1} = \sum_{\ell=0}^{\infty} C_\ell = C(1). \tag{27}$$

Comparing (26) and (27), we get the desired result (25). □

References

Alloza, M., Gonzalo, J., & Sanz, C. (2019). Dynamic Effects of Persistent Shocks. Banco de Espana Working Paper No. 1944.

Almuzara, T. & Marcet, A. (2017). To diff, or not to diff? Preprint.

Angrist, J. & Pischke, J. (2009). *Mostly Harmless Econometrics: An Empiricist's Companion*. Princeton University Press.

Antolín-Díaz, J. & Rubio-Ramírez, J. F. (2018). Narrative Sign Restrictions for SVARs. *American Economic Review*, *108*(10), 2802–29.

Baek, C. & Lee, B. (2020). A Guide to Single Equation Regressions for Impulse Response Estimations. Manuscript, Hong Kong University of Science and Technology.

Barakchian, S. M. & Crowe, C. (2013). Monetary policy matters: Evidence from new shocks data. *Journal of Monetary Economics*, *60*(8), 950–966.

Bernanke, B. S. (1986). Alternative explanations of the money-income correlation. In *Carnegie-Rochester Conference Series on Public Policy*, volume 25, (pp. 49–99). Elsevier.

Blanchard, O. J. & Quah, D. (1989). The Dynamic Effects of Aggregate Demand and Supply Disturbances. *American Economic Review*, *79*(4), 655–673.

Blanchard, O. J. & Watson, M. W. (1986). Are Business Cycles All Alike? In R. J. Gordon (Ed.), *The American Business Cycle: Continuity and Change* chapter 2, (pp. 123–180). University of Chicago Press.

Brockwell, P. & Davis, R. (1991). *Time Series: Theory and Methods* (2nd ed.). Springer Series in Statistics. Springer.

Brugnolini, L. (2018). About Local Projection Impulse Response Function Reliability. Manuscript, University of Rome "Tor Vergata".

Carriero, A., Mumtaz, H., Theodoridis, K., & Theophilopoulou, A. (2015). The Impact of Uncertainty Shocks under Measurement Error: A Proxy SVAR Approach. *Journal of Money, Credit and Banking*, *47*(6), 1223–1238.

Chang, P.-L. & Sakata, S. (2007). Estimation of impulse response functions using long autoregression. *The Econometrics Journal*, *10*(2), 453–469.

Chevillon, G. (2007). Direct Multi-Step Estimation and Forecasting. *Journal of Economic Surveys*, *21*(4), 746–785.

Choi, C.-Y. & Chudik, A. (2019). Estimating impulse response functions when the shock series is observed. *Economics Letters*, *180*, 71–75.

Christiano, L. J., Eichenbaum, M., & Evans, C. L. (2005). Nominal Rigidities and the Dynamic Effects of a Shock to Monetary Policy. *Journal of Political Economy*, *113*(1), 1–45.

Christiano, L. J., Eichenbaum, M., & Vigfusson, R. (2006). Alternative Procedures for Estimating Vector Autoregressions Identified with Long-Run Restrictions. *Journal of the European Economic Association*, *4*(2–3), 475–483.

Dufour, J.-M. & Renault, E. (1998). Short Run and Long Run Causality in Time Series: Theory. *Econometrica*, *66*(5), 1099–1125.

Fernández-Villaverde, J., Rubio-Ramírez, J. F., Sargent, T. J., & Watson, M. W. (2007). ABCs (and Ds) of Understanding VARs. *American Economic Review*, *97*(3), 1021–1026.

Forni, M., Gambetti, L., & Sala, L. (2019). Structural VARs and noninvertible macroeconomic models. *Journal of Applied Econometrics*, *34*(2), 221–246.

Frisch, R. (1933). Propagation Problems and Impulse Problems in Dynamic Economics. In *Economic Essays in Honor of Gustav Cassel*. George Allen & Unwin Ltd.

Gafarov, B., Meier, M., & Montiel Olea, J. L. (2018). Delta-Method Inference for a Class of Set-Identified SVARs. *Journal of Econometrics*, *203*(2), 316–327.

Gertler, M. & Karadi, P. (2015). Monetary Policy Surprises, Credit Costs, and Economic Activity. *American Economic Journal: Macroeconomics*, *7*(1), 44–76.

Giacomini, R. & Kitagawa, T. (2020). Robust Bayesian Inference for Set-Identified Models. cemmap Working Paper CWP13/20.

Gilchrist, S. & Zakrajšek, E. (2012). Credit Spreads and Business Cycle Fluctuations. *American Economic Review*, *102*(4), 1692–1720.

Inoue, A. & Kilian, L. (2002). Bootstrapping Smooth Functions of Slope Parameters and Innovation Variances in VAR(∞) Models. *International Economic Review*, *43*(2), 309–331.

Jarociński, M. & Karadi, P. (2020). Deconstructing Monetary Policy Surprises—The Role of Information Shocks. *American Economic Journal: Macroeconomics*, *12*(2).

Jordà, Ò. (2005). Estimation and Inference of Impulse Responses by Local Projections. *American Economic Review*, *95*(1), 161–182.

Jordà, O. (2009). Simultaneous Confidence Regions for Impulse Responses. *Review of Economics and Statistics*, *91*(3), 629–647.

Jordà, Ò., Schularick, M., & Taylor, A. M. (2013). When Credit Bites Back. *Journal of Money, Credit and Banking*, *45*(Supplement 2), 3–28.

Jordà, Ò., Schularick, M., & Taylor, A. M. (2015). Betting the house. *Journal of International Economics*, *96*(Supplement 1), S2–S18. 37th Annual NBER International Seminar on Macroeconomics.

Jordà, Ò., Schularick, M., & Taylor, A. M. (2019). The effects of quasi-random monetary experiments. *Journal of Monetary Economics*. Forthcoming.

Kilian, L. (2006). Not All Oil Price Shocks Are Alike: Disentangling Demand and Supply Shocks in the Crude Oil Market. CEPR Discussion Paper 5994.

Kilian, L. & Kim, Y. J. (2011). How Reliable Are Local Projection Estimators of Impulse Responses? *Review of Economics and Statistics*, *93*(4), 1460–1466.

Kilian, L. & Lütkepohl, H. (2017). *Structural Vector Autoregressive Analysis*. Cambridge University Press.

Kilian, L. & Murphy, D. P. (2012). Why agnostic sign restrictions are not enough: understanding the dynamics of oil market VAR models. *Journal of the European Economic Association*, *10*(5), 1166–1188.

Kilian, L. & Vigfusson, R. J. (2011). Are the responses of the US economy asymmetric in energy price increases and decreases? *Quantitative Economics*, *2*(3), 419–453.

Leduc, S. & Wilson, D. (2017). Are State Governments Roadblocks to Federal Stimulus? Evidence on the Flypaper Effect of Highway Grants in the 2009 Recovery Act. *American Economic Journal: Economic Policy*, *9*(2), 253–292.

Lewis, R. & Reinsel, G. C. (1985). Prediction of Multivariate Time Series by Autoregressive Model Fitting. *Journal of Multivariate Analysis*, *16*(3), 393–411.

Ludvigson, S. C., Ma, S., & Ng, S. (2020a). Covid19 and the Macroeconomic Effects of Costly Disasters. NBER Working Paper No. 26987.

Ludvigson, S. C., Ma, S., & Ng, S. (2020b). Shock Restricted Structural Vector-Autoregressions. Manuscript, New York University.

Marcellino, M., Stock, J. H., & Watson, M. W. (2006). A comparison of direct and iterated multistep AR methods for forecasting macroeconomic time series. *Journal of Econometrics*, *135*(1–2), 499–526.

McElroy, T. (2015). When are Direct Multi-step and Iterative Forecasts Identical? *Journal of Forecasting*, *34*(4), 315–336.

Meier, A. (2005). How Big is the Bias in Estimated Impulse Responses? A Horse Race between VAR and Local Projection Methods. Manuscript, European University Institute.

Mertens, E. (2012). Are spectral estimators useful for long-run restrictions in SVARs? *Journal of Economic Dynamics and Control*, *36*(12), 1831–1844.

Mertens, K. (2015). Advances in Empirical Macroeconomics, Lecture 2. Lecture slides, Bonn Summer School.

Mertens, K. & Ravn, M. O. (2013). The Dynamic Effects of Personal and Corporate Income Tax Changes in the United States. *American Economic Review*, *103*(4), 1212–1247.

Miranda-Agrippino, S. (2017). Unsurprising Shocks: Information, Premia, and the Monetary Transmission. Manuscript, Bank of England.

Miranda-Agrippino, S. & Ricco, G. (2019). Identification with External Instruments in Structural VARs under Partial Invertibility. Manuscript, Bank of England.

Nakamura, E. & Steinsson, J. (2018). Identification in Macroeconomics. *Journal of Economic Perspectives*, *32*(3), 59–86.

Newey, W. (1994). The Asymptotic Variance of Semiparametric Estimators. *Econometrica*, *62*(6), 1349–1382.

Noh, E. (2018). Impulse-response analysis with proxy variables. Manuscript, University of California San Diego.

Paul, P. (2019). The Time-Varying Effect of Monetary Policy on Asset Prices. *Review of Economics and Statistics*. Forthcoming.

Pesaran, M. H., Pick, A., & Timmermann, A. (2011). Variable selection, estimation and inference for multi-period forecasting problems. *Journal of Econometrics*, *164*(1), 173–187. Annals Issue on Forecasting.

Plagborg-Møller, M. & Wolf, C. K. (2019). Instrumental Variable Identification of Dynamic Variance Decompositions. Manuscript, Princeton University.

Ramey, V. A. (2011). Identifying Government Spending Shocks: It's All in the Timing. *Quarterly Journal of Economics*, *126*(1), 1–50.

Ramey, V. A. (2016). Macroeconomic Shocks and Their Propagation. In J. B. Taylor & H. Uhlig (Eds.), *Handbook of Macroeconomics*, volume 2 chapter 2, (pp. 71–162). Elsevier.

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*(2), 850–901.

Romer, C. D. & Romer, D. H. (2004). A New Measure of Monetary Shocks: Derivation and Implications. *American Economic Review*, *94*(4), 1055–1084.

Rubio-Ramírez, J. F., Waggoner, D. F., & Zha, T. (2010). Structural vector autoregressions: Theory of identification and algorithms for inference. *Review of Economic Studies*, *77*(2), 665–696.

Schorfheide, F. (2005). VAR forecasting under misspecification. *Journal of Econometrics*, *128*(1), 99–136.

Sims, C. A. (1980). Macroeconomics and Reality. *Econometrica*, *48*(1), 1–48.

Smets, F. & Wouters, R. (2007). Shocks and Frictions in US Business Cycles: A Bayesian DSGE Approach. *American Economic Review*, *97*(3), 586–606.

Stock, J. H. (2008). What's New in Econometrics: Time Series, Lecture 7. Lecture slides, NBER Summer Institute.

Stock, J. H. & Watson, M. W. (2012). Disentangling the Channels of the 2007–09 Recession. *Brookings Papers on Economic Activity*, *2012*(1), 81–135.

Stock, J. H. & Watson, M. W. (2018). Identification and Estimation of Dynamic Causal Effects in Macroeconomics Using External Instruments. *Economic Journal*, *128*(610), 917–948.

Uhlig, H. (2005). What are the effects of monetary policy on output? Results from an agnostic identification procedure. *Journal of Monetary Economics*, *52*(2), 381–419.

Wolf, C. K. (2019). SVAR (Mis-)Identification and the Real Effects of Monetary Policy Shocks. *American Economic Journal: Macroeconomics*. Forthcoming.