

Local Projections and VARs Estimate the Same Impulse Responses*

Mikkel Plagborg-Møller
Princeton University

Christian K. Wolf
Princeton University

This version: November 3, 2019

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 belong to a spectrum of dimension reduction techniques with common estimand but different bias-variance properties. (ii) VAR-based structural estimation 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.

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

*Email: mikkelpm@princeton.edu and ckwolf@princeton.edu. We received helpful comments from Domenico Giannone, Marek Jarociński, Òscar Jordà, Peter Karadi, Lutz Kilian, Dake Li, Pepe Montiel Olea, Valerie Ramey, Giovanni Ricco, Neil Shephard, Chris Sims, Jim Stock, Mark Watson, 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.

– 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 such statements are found in Jordà (2005, p. 162), in the literature reviews of Ramey (2016, p. 83) and Nakamura & Steinsson (2018, pp. 80–81), and in the textbook treatment of Kilian & Lütkepohl (2017, ch. 12.8).¹ 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,” so that SVAR estimation is the only way to implement more exotic 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 on their relative merits reach conflicting conclusions and disagree on implementation details (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 is *nonparametric*, in that it essentially only requires the data to be weakly stationary and the lag structures in the two specifications to be unrestricted.³ 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, 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 agree out to horizon p (but

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

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

³Although linear LP and VAR estimators may be viewed as “parametric” procedures, we do not assume that the data generating process can be summarized by any finite-dimensional parametric model.

not further), again without imposing any parametric assumptions on the data generating process. In summary, if VAR and LP results differ in population or in sample, it is due to extraneous restrictions on the lag structure.

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

First, LP and VAR estimators are not conceptually different methods; instead, they belong to a spectrum of linear projection techniques that share the same estimand but differ in their finite-sample bias-variance properties. Standard LPs effectively provide no dimension reduction, while conventional low-order VARs extrapolate shock propagation from the first few autocorrelations of the data. The relative mean-square error of the two methods – and of intermediate dimension reduction techniques, such as shrinkage – necessarily depends on assumptions about the data generating process (DGP). VAR estimators are optimal if the true DGP is exactly a finite-order VAR, but this is rarely the case in theory or practice. The formal equivalence of LP and VAR impulse response estimation to direct and iterated forecasting, respectively, means that applied researchers can look to the existing forecasting literature for guidance on how to choose between the menu of available estimators (Schorfheide, 2005; Marcellino et al., 2006; Pesaran et al., 2011).

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 dimension reduction technique. 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 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).⁴ In particular, this

⁴In contemporaneous work, Noh (2018) also recommends including the IV as an internal instrument in a VAR; our result offers additional insights by drawing connections to LP-IV and to the general equivalence between LPs and VARs.

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 VARs. We show that their common estimand may be formally interpreted as a best linear approximation to the underlying, perhaps non-linear, data generating process.

In summary, in addition to clarifying misconceptions in the literature about the LP and VAR estimands, our results allow applied researchers to separate the choice of identification scheme from the choice of estimation technique. Researchers who prefer the intuitive regression interpretation and generally low bias 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 and generally low variance 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 contribution is to formally establish a nonparametric equivalence result and derive implications for estimation efficiency and structural identification. [Jordà \(2005\)](#) 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 this context, [Kilian & Lütkepohl](#) also discuss the relative efficiency of the two estimation methods and mention the literature on direct versus iterated forecasts. In contrast, our equivalence result is *non-parametric*, and we further demonstrate how structural VAR orderings map into particular choices of LP control variables, and *vice versa*.⁵ 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.⁶

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 when an instrument/proxy is available. We do not consider questions related to *inference*, and instead refer to the discussions in [Jordà \(2005\)](#), [Kilian & Lütkepohl \(2017\)](#), and [Stock & Watson \(2018\)](#).

⁵[Jordà et al. \(2019\)](#) informally discuss the connection between control variables and recursive SVARs.

⁶[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 nonparametric equivalence of the LP and VAR estimands, however.

OUTLINE. [Section 2](#) presents our core result on the population equivalence of local projections and VARs. Finite-sample estimation is discussed in [Section 3](#), while [Section 4](#) traces out implications for *structural* estimation. We illustrate our equivalence results with a practical application to IV-based identification of monetary policy shocks in [Section 5](#). [Section 6](#) concludes with several recommendations for empirical practice. Some proofs are relegated to [Appendix A](#).

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. Finally, we illustrate the results graphically. Our analysis in this section is “reduced form” in that it does not assume any specific underlying structural model; we merely work with linear projections of stationary time series. We will discuss implications for *structural* identification in [Section 4](#).

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. The distinction between them relates to whether they appear as *contemporaneous* controls or not, as will become clear in equations (1) and (2) below.

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

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. To simplify notation, we proceed as if $\{w_t\}$ were a (strictly stationary) jointly Gaussian vector time series.*

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

⁷The 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 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.⁸

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 .

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

$$y_{t+h} = \mu_h + \beta_h x_t + \gamma'_h r_t + \sum_{\ell=1}^{\infty} \delta'_{h,\ell} w_{t-\ell} + \xi_{h,t}, \quad (1)$$

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. The *LP impulse response function* of y_t with respect to x_t is given by $\{\beta_h\}_{h \geq 0}$. Notice that the projection (1) controls for the contemporaneous value of r_t but not of q_t .

2. VAR. Consider the multivariate linear “VAR(∞)” projection

$$w_t = c + \sum_{\ell=1}^{\infty} A_\ell w_{t-\ell} + u_t, \quad (2)$$

where $u_t \equiv w_t - E(w_t | \{w_\tau\}_{-\infty < \tau < t})$ 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. Define the lag polynomial

$$\sum_{\ell=0}^{\infty} C_\ell L^\ell = C(L) \equiv A(L)^{-1}.$$

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

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

⁸Throughout we write any linear projection 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.

since x_t and y_t are the $(n_r + 1)$ -th and $(n_r + 2)$ -th elements in w_t . The notation $C_{i,\bullet,h}$, say, means the i -th row of matrix C_h , while $B_{\bullet,j}$ is the j -th column of matrix B .

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](#), while all finite-sample considerations are relegated to [Section 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 controls in [Section 4](#).

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 [Assumption 1](#), 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 $p \rightarrow \infty$ is sufficiently flexible that it perfectly captures all covariance properties of the data. Thus, iterated forecasts based on the VAR coincide perfectly with direct forecasts $E[w_{t+h} \mid w_t, w_{t-1}, \dots]$. Although the intuition for the equivalence is simple, its implications do not appear to have been generally appreciated in the literature, as discussed 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

$$\beta_h = \frac{\text{Cov}(y_{t+h}, \tilde{x}_t)}{E(\tilde{x}_t^2)}. \quad (3)$$

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

$$\theta_h = C_{n_r+2, \bullet, h} B_{\bullet, n_r+1} = \text{Cov}(y_{t+h}, \eta_{x,t}), \quad (4)$$

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 have⁹

$$\eta_{x,t} = \frac{1}{\sqrt{E(\tilde{u}_{x,t}^2)}} \times \tilde{u}_{x,t}, \quad (5)$$

where we partition $u_t = (u'_{r,t}, u_{x,t}, u_{y,t}, u'_{q,t})'$ and define¹⁰

$$\tilde{u}_{x,t} \equiv u_{x,t} - E(u_{x,t} | u_{r,t}) = \tilde{x}_t. \quad (6)$$

From (4), (5), and (6) 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 (3). □

In conclusion, LPs and VARs offer two equivalent ways of arriving at the same population parameter (3), up to a scale factor that does not depend on the horizon h . Our argument was nonparametric and did not assume the validity of a specific structural model.

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

⁹ 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 $B_{n_r+1, n_r+1} \eta_{x,t} = u_{x,t} - E(u_{x,t} | \eta_{r,t}) = u_{x,t} - E(u_{x,t} | u_{r,t}) = \tilde{u}_{x,t}$. Expression (5) then follows from $E(\eta_{x,t}^2) = 1$.

¹⁰Observe that $u_{x,t} - \tilde{x}_t = E(x_t | r_t, \{w_\tau\}_{-\infty < \tau < t}) - E(x_t | \{w_\tau\}_{-\infty < \tau < t}) = E(u_{x,t} | r_t, \{w_\tau\}_{-\infty < \tau < t}) = E(u_{x,t} | u_{r,t}, \{w_\tau\}_{-\infty < \tau < t}) = E(u_{x,t} | u_{r,t})$.

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 (2), but now we shall study the propagation of *some* rotation of the reduced-form forecasting residuals,

$$\bar{\eta}_t \equiv b'u_t, \tag{7}$$

where b is a vector of the same dimension as w_t . Under [Assumption 1](#), 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

$$y_{t+h} = \bar{\mu}_h + \bar{\beta}_h(b'w_t) + \sum_{\ell=1}^{\infty} \bar{\delta}'_{h,\ell} w_{t-\ell} + \bar{\xi}_{h,t}, \tag{8}$$

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 (8) on a particular linear combination $b'w_t$ of the variables in the VAR (and their lags). For recursive orderings, this reduces to [Proposition 1](#).

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 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 finite-order VAR.

First, we define the finite-order LP and VAR estimands. We continue to impose the nonparametric [Assumption 1](#). 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 (2) 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 $A(L; p) \equiv I - \sum_{\ell=1}^p A_\ell(p)L^\ell$ 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 $p = \infty$.

Proposition 2. *Impose [Assumption 1](#). Define $\tilde{x}_t(\ell) \equiv x_t - E(x_t \mid r_t, \{w_\tau\}_{t-\ell \leq \tau < t})$ for all $\ell = 0, 1, 2, \dots$. Let the nonnegative integers h, p satisfy $h \leq p$. If $\tilde{x}_t(p) = \tilde{x}_t(p-h)$, then $\theta_h(p) = \sqrt{E(\tilde{x}_t(p)^2)} \times \beta_h(p)$.*

Proof. Please see [Appendix A.1](#). □

Thus, under the conditions of the proposition, the population LP and VAR impulse response estimands agree at all horizons $h \leq p$, 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](#), 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.¹¹

[Proposition 2](#) assumes $\tilde{x}_t(p) = \tilde{x}_t(p-h)$ to obtain an exact result, but the conclusion is likely to hold qualitatively under more general conditions. If x_t is a direct measure of a “shock” and thus uncorrelated with all past data, then $\tilde{x}_t(\ell) = x_t$ for all $\ell \geq 0$, so the conclusion of the proposition holds exactly. More generally, the LP estimand projects y_{t+h} onto $\tilde{x}_t(p)$ (and controls); thus, the projection depends on the first $p+h$ autocovariances of

¹¹[Baek & Lee \(2019\)](#) prove a similar result for the related but distinct setting of single-equation Autoregressive Distributed Lag models with a white noise exogenous regressor.

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 potentially project on different objects. However, at short horizons $h \ll p$, it will usually be the case in empirically relevant DGPs that $\tilde{x}_t(p) \approx \tilde{x}_t(p - h)$, since it is typically only the first few lags of the data that is useful for forecasting x_t . In this case, the conclusion of [Proposition 2](#) will hold approximately. We provide an illustration in [Section 2.4](#).

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

2.4 Graphical illustration

We finish the section by illustrating graphically the previous theoretical results. We do so 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.¹² 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 4](#).

The left panel of [Figure 1](#) shows LP and VAR impulse response estimands of the response of output to a government spending innovation. We assume the model’s government spending innovation is directly observed by the econometrician, who additionally controls for lags of output and government spending. This experiment is therefore 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 $\tilde{x}_t(\ell) = x_t$ for all $\ell \geq 0$. 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

¹²We 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(∞).

ILLUSTRATION: POPULATION EQUIVALENCE OF VAR AND LP ESTIMANDS

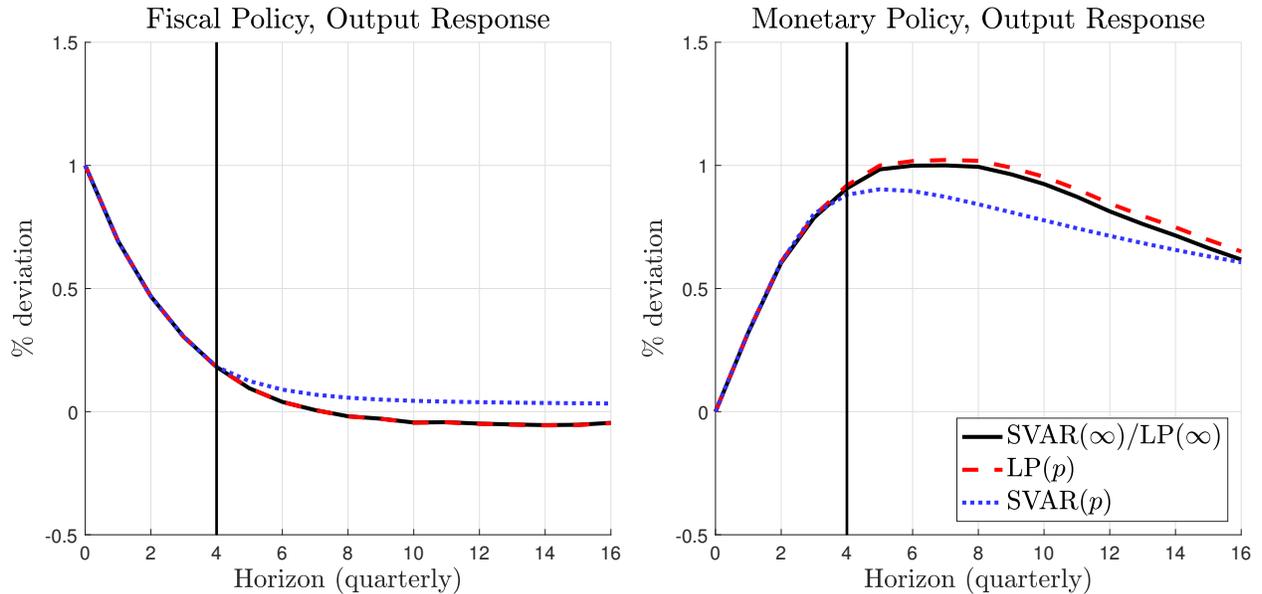


Figure 1: LP and VAR impulse response estimands in the structural model of [Smets & Wouters \(2007\)](#). Left panel: response of output to a government spending innovation. Right panel: response of output to an interest rate innovation. The horizontal line marks the horizon p after which the finite-lag-length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

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 $LP(\infty)$ and $VAR(\infty)$ estimands agree at all horizons. Now, however, the “impulse” $\tilde{x}_t(p)$ upon which the different methods project is different. Hence, $LP(p)$ and $VAR(p)$ estimands differ from each other, as well as from the population limit $LP(\infty)/VAR(\infty)$ estimands. Formally, [Proposition 2](#) only assures that the estimated *impact* impulse responses of $LP(p)$ and $VAR(p)$ agree exactly. Nevertheless, and consistent with the intuition offered in [Section 2.3](#), all impulse response estimands are *nearly* identical until the truncation horizon $p = 4$.

3 Efficient estimation of impulse responses

This section discusses our equivalence result in the context of finite-sample estimation of impulse responses. We first provide a sample analogue of our population equivalence result

when the lag length is large. Then we discuss the bias-variance trade-off associated with estimation of impulse response functions. While we maintain a reduced-form perspective in this section, in [Section 4](#) we will apply the insights to structural estimators.

3.1 Sample equivalence

In addition to being identical *conceptually* and in *population*, we show in the Online Appendix that local projection and VAR impulse response estimators are nearly identical *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)–(2) 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, 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.¹³

3.2 Bias-variance trade-off

Empirically relevant short sample sizes force researchers to economize on the number of lags, and the relative accuracy of LP and VAR estimators with a small/moderate number of lags invariably depends on the underlying data generating process (DGP). This is perfectly analogous to the choice between “direct” and “iterated” predictions in multi-step forecasting ([Marcellino et al., 2006](#); [Pesaran et al., 2011](#)). [Schorfheide \(2005\)](#) proves that the mean-square error ranking of LP (i.e., direct) and VAR (i.e., iterated) forecasts depends on how large in magnitude the partial autocorrelations of the DGP are at lags longer than the lag length used for estimation.¹⁴ Hence, although [Meier \(2005\)](#), [Kilian & Kim \(2011\)](#), and [Choi & Chudik \(2019\)](#) exhibit simulation evidence that VAR estimators (or other iterated estimators) outperform the LP estimator, this conclusion must necessarily depend on the choice of DGP. Indeed, [Brugnolini \(2018\)](#) and [Nakamura & Steinsson \(2018\)](#) exhibit DGPs where the LP estimator instead outperforms VARs.

The forecasting literature has generally found that LP (direct) methods tend to have relatively low bias, whereas VAR (iterated) methods have relatively low variance. The

¹³The appendix is available at: http://scholar.princeton.edu/mikkelpm/lp_var

¹⁴See also [Chevillon \(2007\)](#), [McElroy \(2015\)](#), and references therein.

trade-off is most relevant at longer response horizons, as shown by our finite- p equivalence result in [Proposition 2](#). The VAR(p) model extrapolates long-horizon impulse responses from the autocovariances at lags $0, 1, \dots, p$, and thus may potentially be substantially biased if p is not very large. For the same reason, though, VAR(p) estimators tend to deliver much smaller estimation variance than LPs at long horizons. [Hansen \(2010, 2016\)](#), [Pesaran et al. \(2011\)](#), and [Kilian & Lütkepohl \(2017, ch. 2.6\)](#) discuss methods for choosing the lag length p for VAR and LP estimators in a way that is informed by the bias-variance trade-off.

More generally, effective finite-sample estimation of impulse responses involves an unavoidable bias-variance trade-off, and many dimension reduction or penalization approaches may be sensible depending on the application. Bayesian VARs reduce effective dimensionality by imposing priors on longer-lag coefficients, e.g., through a Minnesota prior ([Giannone et al., 2015](#)); model averaging across restricted and unrestricted VARs has similar effects ([Hansen, 2016](#)). Dimension reduction can also be achieved through penalized local projection ([Plagborg-Møller, 2016, Ch. 3](#); [Barnichon & Brownlees, 2019](#)) or by shrinking unrestricted local projections towards low-order VAR estimates ([Miranda-Agrippino & Ricco, 2018b](#)). Alternatively, impulse response estimation could be based on plugging a shrinkage/regularized autocovariance function estimate into the explicit formula (3) for the LP/VAR estimand.

We believe that the different estimation methods in the literature are best viewed as sharing the *same* large-sample estimand but lying along a *spectrum* of small-sample bias-variance choices. Low-order VAR(p) models only have a conceptually special status insofar as we think the finite- p assumption is literally true, which is typically not the case. In general, the relative accuracy of the methods depends on smoothness/sparsity properties of the autocovariance function of the data. From the point of view of point estimation, no single method dominates for *all* empirically relevant data DGPs. In principle, standard VAR model diagnostic checks or pseudo-out-of-sample forecast performance can be used as a means to select between impulse response estimators. However, we recommend that researchers compare results from different methods, since any disparities may indicate that further thought about the DGP and/or the shrinkage procedure is warranted.

To summarize: Guided by the previously cited forecasting literature, the choice of estimation method should depend on (i) the researcher’s preferences over bias and variance and on (ii) features of the DGP. In contrast, in the next section we argue that the choice of *structural identification scheme* should *not* determine the choice between LPs and VARs (or other dimension reduction techniques).

4 Structural identification of impulse responses

We now show that our result on the equivalence of LP and VAR impulse response functions has important implications for structural identification. We have seen that LP and VAR methods only differ to the extent that they represent different approaches to finite-sample dimensionality reduction. The problem of *structural identification* is a population concept and is thus logically distinct from that of dimensionality reduction. In this section we apply our equivalence result to popular SVAR and local projection identification schemes – including short-run restrictions, long-run restrictions, sign restrictions, and external instruments – and we discuss how to think about non-linear models.

4.1 Structural model

To discuss structural identification, we now 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 (we address non-linear models in [Section 4.4](#)). SVMA models have been analyzed by [Stock & Watson \(2018\)](#), [Plagborg-Møller & Wolf \(2019\)](#), and many others. The class of SVMA models encompasses all discrete-time, linearized DSGE models as well as all stationary SVAR models.

Assumption 2. *The data $\{w_t\}$ are driven by an n_ε -dimensional vector $\varepsilon_t = (\varepsilon_{1,t}, \dots, \varepsilon_{n_\varepsilon,t})'$ of exogenous structural shocks,*

$$w_t = \mu + \Theta(L)\varepsilon_t, \quad \Theta(L) \equiv \sum_{\ell=0}^{\infty} \Theta_\ell L^\ell, \quad (9)$$

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:

$$\varepsilon_t \stackrel{i.i.d.}{\sim} N(0, I_{n_\varepsilon}). \quad (10)$$

Under these assumptions w_t is a nonsingular, strictly stationary jointly Gaussian time series, consistent with [Assumption 1](#) 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 $\Theta_{n_r+2,1,h}$, $h = 0, 1, 2, \dots$. 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.

4.2 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:

$$\varepsilon_{1,t} \in \text{span}(\{w_\tau\}_{-\infty < \tau \leq t}). \quad (11)$$

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

$$\tilde{\varepsilon}_{1,t} \equiv b'u_t, \quad (12)$$

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 (11) holds for all shocks in the system (9). 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 , all aggregates ordered before the Federal Funds Rate as the controls r_t , and all other variables collected in the vector q_t . [Christiano et al.](#) implement their structural analysis through the recursive VAR (2). By our main result, they could have equivalently estimated the regression (1) and collected the regression coefficients $\{\beta_h\}_{h \geq 0}$. The population estimand would have been the same, but in finite samples the mean-square error ranking of the two estimators is ambiguous, as discussed in [Section 3](#). \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 2](#) 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. (11). 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 a 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

$$gdp_{t+H} - gdp_{t-1} = \tilde{\mu}_H + \tilde{\beta}'_H w_t + \sum_{\ell=1}^{\infty} \tilde{\delta}'_{H,\ell} w_{t-\ell} + \tilde{\xi}_{H,t}. \quad (13)$$

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 (8) 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 $H \rightarrow \infty$, 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.¹⁵ To estimate relative impulse responses $\Theta_{i,2,h}/\Theta_{1,2,0}$ to the demand

¹⁵Absolute impulse responses can be identified by rescaling the identified shock so it has variance 1.

shock, the researcher can choose any vector \tilde{b} such that $\tilde{b}'b = 0$, and then implement the local projection (8) with \tilde{b} in lieu of b .

In finite samples, the mean-square 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.¹⁶ 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 VAR does. \square

Example 3 (Sign identification). [Uhlig \(2005\)](#) set-identifies the effects of monetary policy shocks by sign-restricting impulse responses. For concreteness, suppose we are interested in the impulse response of y_t (say, real GDP growth) to a monetary shock at horizon h . As before, assume that the full set of observed data $\{w_t\}$ follows an SVMA system (9) 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 $s = 0, 1, \dots, \bar{H}$.

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 $\varepsilon_{1,t} = \nu' u_t$, where $\nu \in \mathbb{R}^{n_w}$ is an unknown vector. If we knew ν , the structural impulse responses of any variable $w_{i,t}$ to $\varepsilon_{1,t}$ could be obtained as the linear combination ν 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 ν such that (i) the r_t impulse responses are positive at all horizons $s = 0, 1, \dots, \bar{H}$ and (ii) the impact r_t impulse response is normalized to 1 (other normalizations are also possible). Once we have determined the set of possible ν 's, we can then use the VAR to compute the corresponding set of possible impulse responses of y_t with respect to $\nu' u_t$.

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

¹⁶[Christiano et al. \(2006\)](#) and [Mertens \(2012\)](#) make the related point that SVAR-based long-run identification need not rely on the VAR-implied long-run variance matrix. Alternative nonparametric estimators of the latter may have attractive bias-variance properties, depending on the true DGP.

u_t . Exactly as in the VAR approach, we now seek the linear combination $\nu' \check{\beta}_h$ that equals the structural impulse response to the monetary shock $\varepsilon_{1,t} = \nu' u_t$. To find the set of ν '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 $\check{\beta}_s$ from the projection

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

The coefficients $\check{\beta}_s$ 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 program¹⁷

$$\begin{aligned} \sup_{\nu \in \mathbb{R}^{n_w}} \nu' \check{\beta}_h \quad \text{subject to} \quad & \check{\beta}'_0 \nu = 1, \\ & \check{\beta}'_s \nu \geq 0, \quad s = 1, \dots, \bar{H}. \end{aligned}$$

To compute the *smallest possible* impulse response, replace the supremum with an infimum.¹⁸ In population, this LP-based procedure recovers exactly the same identified set as analogous sign restrictions in an SVAR. It is straight-forward to implement more complicated identification schemes by adding additional equality or inequality constraints of the above type. \square

These three 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, as discussed in [Section 3](#).

BEYOND INVERTIBILITY. If the invertibility assumption (11) is violated, then identification strategies that erroneously assume invertibility – independent of whether they are

¹⁷To consider impulse responses to a *one-standard-deviation* monetary shock, replace the equality constraint in the linear program by the constraint $\nu' \text{Var}(u_t)^{-1} \nu = 1$. The resulting linear-quadratic program with inequality constraints is similar to those in [Gafarov et al. \(2018\)](#) and [Giacomini & Kitagawa \(2018\)](#).

¹⁸We 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](#)).

implemented using VARs, LPs, or any other dimensionality reduction technique – will not measure the true impulse responses.¹⁹ 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:

$$\tilde{\varepsilon}_{1,t} = \vartheta(L)\varepsilon_t. \quad (14)$$

The properties of the lag polynomial $\vartheta(L)$ are characterized in detail in [Fernández-Villaverde et al. \(2007\)](#) and [Wolf \(2019\)](#). Combining (9) and (14), we see that, in general, both LP and VAR impulse response estimands are linear combinations of contemporaneous and lagged 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 .²⁰ Instrumental variable identification, discussed next, is the leading example of this second case.

4.3 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](#)).²¹ More precisely, given [Assumption 2](#), the IV exclusion restrictions are that

$$\text{Cov}(z_t, \varepsilon_{j,s} \mid \{z_\tau, w_\tau\}_{-\infty < \tau < t}) \neq 0 \quad \text{if and only if} \quad \text{both } j = 1 \text{ and } t = s. \quad (15)$$

¹⁹Several 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.

²⁰In 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.

²¹We 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.

Stock & Watson (2018, p. 926) refer to this assumption as “LP-IV[⊥],” 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 (15) as follows.

Assumption 3.

$$z_t = c_z + \sum_{\ell=1}^{\infty} (\Psi_{\ell} z_{t-\ell} + \Lambda_{\ell} w_{t-\ell}) + \alpha \varepsilon_{1,t} + v_t, \quad (16)$$

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

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

$$y_{t+h} = \mu_{RF,h} + \beta_{RF,h} z_t + \sum_{\ell=1}^{\infty} \delta'_{RF,h,\ell} W_{t-\ell} + \xi_{RF,h,t} \quad (17)$$

for any $h \geq 0$. Consider also the “first-stage” IV projection²²

$$x_t = \mu_{FS} + \beta_{FS} z_t + \sum_{\ell=1}^{\infty} \delta'_{FS,\ell} W_{t-\ell} + \xi_{FS,t}. \quad (18)$$

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 $\beta_{LP-IV,h} \equiv$

²²As always, the coefficients and residuals in (17)–(18) should be interpreted as linear projections.

$\beta_{RF,h}/\beta_{FS}$ of reduced-form to first-stage coefficients (e.g. Angrist & Pischke, 2009, p. 122).²³

Stock & Watson (2018) show that, under Assumptions 2 and 3, the LP-IV estimand $\beta_{LP-IV,h}$ 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 (11).

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 is nonparametric and assumes nothing about the underlying structural model or about the IV z_t .

Corollary 1. *Let Assumption 1 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. (18). 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 $\tilde{\theta}_{y,h}/\tilde{\theta}_{x,0} = \beta_{LP-IV,h}$.

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. \square

This nonparametric result implies that, given the *structural* Assumptions 2 and 3, 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.²⁴ Importantly, under Assumptions 2 and 3, 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 (16).²⁵ But this independent measurement error merely leads to attenuation bias in the estimated impulse

²³In 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.

²⁴Plagborg-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.

²⁵Note that, even though Assumption 3 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 3).

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, 2018a](#)). 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 is valid to include an externally identified shock in a SVAR even if the shock is measured with (independent) error, as long as the noisily measured shock is ordered first.²⁶

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

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

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 2](#) and [3](#) and the invertibility condition [\(11\)](#) 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.²⁷ However, in the general non-invertible case, SVAR-IV mis-identifies the shock as $\tilde{\varepsilon}_{1,t} \neq \varepsilon_{1,t}$.²⁸ [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.

To summarize, the relative impulse responses obtained from the LP-IV procedure of [Stock](#)

²⁶[Romer & Romer \(2004\)](#) and [Barakchian & Crowe \(2013\)](#) include an externally identified monetary shock in a SVAR, but they order it last, which assumes additional exclusion restrictions. [Kilian \(2006\)](#), [Ramey \(2011\)](#), [Miranda-Agrippino \(2017\)](#), and [Jarociński & Karadi \(2019\)](#), among others, mention the strategy of ordering an IV first in a SVAR, but these papers do not consider the non-invertible case.

²⁷Consistent 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.

²⁸The VARX approach of [Paul \(2018\)](#) is equivalent with SVAR-IV under [Assumption 1](#).

& 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 – with their associated bias-variance properties discussed in Section 3 – while relying on the same 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.²⁹

4.4 Estimands in non-linear models

Our main result in Section 2.1 implies that linear local projections are exactly as “robust to non-linearities” as VAR methods, in population. We now show that the common LP/VAR estimand can be given a mathematically well-defined “best linear approximation” interpretation when the true underlying structural DGP is in fact non-linear.

Assume that the underlying structural DGP has the nonparametric causal structure

$$w_t = g(\varepsilon_t, \varepsilon_{t-1}, \varepsilon_{t-2}, \dots), \quad (19)$$

where $g(\cdot)$ is any non-linear function that yields a well-defined covariance stationary process $\{w_t\}$, and $\{\varepsilon_t\}$ is an n_ε -dimensional i.i.d. process with $\text{Cov}(\varepsilon_t) = I_{n_\varepsilon}$. The number of structural shocks ε_t may exceed the number of variables in w_t .

We show formally in Appendix A.3 that we can represent the process (19) as the *linear* Structural Vector Moving Average model

$$w_t = \mu^* + \sum_{\ell=0}^{\infty} \Theta_\ell^* \varepsilon_{t-\ell} + \sum_{\ell=0}^{\infty} \Psi_\ell^* \zeta_{t-\ell},$$

where ζ_t is an n_w -dimensional white noise process that is uncorrelated at all leads and lags with the structural shocks ε_t . The argument exploits the Wold decomposition of the residual of w_t after projecting on the structural shocks. Hence, the linear SVMA model (9) in Assumption 2 should not be thought of as restrictive, provided we do not restrict the

²⁹SVAR-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 3, cf. Stock & Watson (2018, sec. 2.1).

number of “shocks” relative to the number of variables.

The linear SVMA impulse responses Θ_ℓ^* corresponding to the structural shocks ε_t have a “best linear approximation” interpretation. Specifically,

$$(\Theta_0^*, \Theta_1^*, \dots) \in \underset{(\tilde{\Theta}_0, \tilde{\Theta}_1, \dots)}{\operatorname{argmin}} E \left[\left(g(\varepsilon_t, \varepsilon_{t-1}, \dots) - \sum_{\ell=0}^{\infty} \tilde{\Theta}_\ell \varepsilon_{t-\ell} \right)^2 \right]. \quad (20)$$

Thus, if a second-moment LP/VAR identification scheme is known to correctly identify the impulse responses in a linear SVMA model (9), and there is doubt about whether the true underlying DGP is in fact linear, the population estimand of the identification procedure can be given a formal “best linear approximation” interpretation. This is analogous to the “best linear predictor” property of Ordinary Least Squares in cross-sectional regression. In contrast, identification approaches that depart from standard linear projections – such as identification through higher moments or through heteroskedasticity – may not have a clear interpretation under functional form misspecification.

Of course, in some applications, the non-linearities of the true underlying DGP may be of interest *per se*. In such cases, non-linear VAR or LP estimators can be applied, for example by adding interaction or polynomial terms, regime switching, stochastic volatility, etc. Such issues are outside the scope of this paper, which deals exclusively with linear estimators.

5 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.³⁰ 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),

³⁰The 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.

RESPONSE OF BOND SPREAD TO MONETARY SHOCK: VAR AND LP ESTIMATES

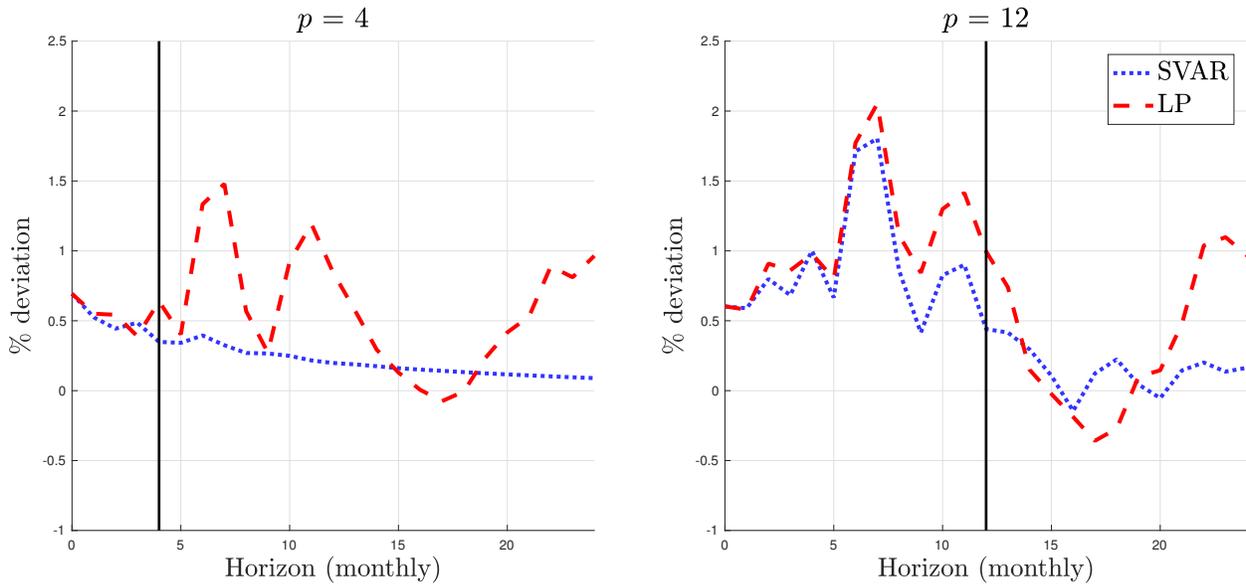


Figure 2: Estimated impulse response function of the Excess Bond Premium to a monetary policy shock, normalized to increase the 1-year bond rate by 100 basis points on impact. Left panel: lag length $p = 4$. Right panel: $p = 12$. The horizontal line marks the horizon p after which the VAR(p) and LP(p)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

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 2012.³¹

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 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 a concrete empirical illustration of our earlier claim that LP and

³¹The data were retrieved from: <https://www.aeaweb.org/articles?id=10.1257/mac.20130329>

VAR estimates are closely tied together at short horizons, not just in population but also in sample. The larger the lag length used for estimation, the more impulse response horizons will exhibit agreement between LP and VAR estimates. As this exercise is merely meant to illustrate our theoretical results, we refrain from conducting formal statistical tests of the relative finite-sample efficiency of the different estimation methods.

6 Conclusion

We demonstrated a general *nonparametric* equivalence of local projection and VAR impulse response function estimands. This result 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. Efficient finite-sample estimation requires navigating a bias-variance trade-off. Low-order VAR and local projection estimators resolve this trade-off differently, and several other recently proposed methods also lie on the continuum of possible dimension reduction or regularization approaches. Neither low-order VARs nor low-order local projections should be treated as having special status *generally*.
3. The bias-variance trade-off is equivalent to the well-known trade-off between direct and iterated forecasts. Thus, the finite-sample mean-square error ranking of different impulse response estimation methods depends on smoothness/sparsity properties of the autocovariance function of the data. The forecasting literature offers extensive guidance on the bias-variance trade-off (see references in [Section 3](#)). No single estimation method dominates for all empirically relevant data generating processes.
4. At short impulse response horizons, the various estimation methods are likely to approximately agree, but at longer horizons the bias-variance trade-off is unavoidable. A VAR estimator with large lag length will give similar results as a local projection, except at very long horizons.
5. It is a useful diagnostic to check if different estimation methods reach similar conclusions. If estimated impulse responses from VARs and local projections differ substantially at longer horizons, it must mean that the sample partial autocorrelations

at long lags are not small. This possibly calls into question the validity of the VAR approximation to the distribution of the data, depending on the standard errors.

6. 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 a finite-order SVAR model whenever the discussion turns to structural identification, as this conflates the population identification analysis and the dimension reduction technique of using a low-order VAR estimator.
7. 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.
8. 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.
9. Linear local projections are exactly as “robust to non-linearities” in the underlying data generating process as linear VARs.

We stress that this paper has focused entirely on identification and estimation of *impulse responses* using linear methods. Identification of other objects, such as variance/historical decompositions, is more involved, as shown in [Plagborg-Møller & Wolf \(2019\)](#).

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 near-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 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

$$\beta_h(p) = \frac{\text{Cov}(y_{t+h}, \tilde{x}_t(p))}{E(\tilde{x}_t(p)^2)}. \quad (21)$$

We now introduce 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

$$A(L; p)w_t = B(p)\bar{\eta}_t, \quad \bar{\eta}_t \sim WN(0, I), \quad (22)$$

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 (22) 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 (22) 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 (22) agree with the autocovariances of the actual data out to lag p , although generally not after lag p .

Under the counterfactual model (22), we have the moving average representation $w_t = C(L; p)B(p)\bar{\eta}_t$, and thus

$$\theta_h(p) = C_{n_r+2, \bullet, h}(p)B_{\bullet, n_r+1}(p) = \text{Cov}^p(y_{t+h}, \bar{\eta}_{x,t}), \quad (23)$$

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 (22) that

$$B_{n_r+1, n_r+1}(p)\bar{\eta}_{x,t} = x_t - E^p(x_t \mid r_t, \{w_\tau\}_{t-p \leq \tau < t}) = x_t - E(x_t \mid r_t, \{w_\tau\}_{t-p \leq \tau < t}) = \tilde{x}_t(p), \quad (24)$$

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 (24) 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 (23), (24), and the above equation, we have shown that

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

Under the stated assumption that $\tilde{x}_t(p) = \tilde{x}_t(p-h)$, the covariance on the right-hand side above depends only on autocovariances of the data w_t at lags $\ell = 0, 1, 2, \dots, p$. Hence, we can again appeal to the equivalence of $\text{Cov}^p(\cdot, \cdot)$ with the covariance function of the actual data, and the expression (21) yields the desired conclusion. \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 $u_t \equiv w_t - E(w_t | \{w_\tau\}_{-\infty < \tau < t})$ and Wold decomposition

$$w_t = \chi + C(L)u_t, \quad C(L) \equiv I_2 + \sum_{\ell=1}^{\infty} C_\ell L^\ell. \quad (25)$$

Since both structural shocks are assumed to be invertible, there exists a 2×2 matrix B such that $\varepsilon_t = Bu_t$. Comparing (9) and (25), 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

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

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

$$\lim_{H \rightarrow \infty} \tilde{\beta}'_H = e_1' C(1). \quad (26)$$

Define $\Sigma_u \equiv \text{Var}(u_t)$. Applying the Frisch-Waugh theorem to the projection (13), and using $w_{1,t} = \Delta gdp_t$, we find

$$\tilde{\beta}'_H = \text{Cov}(gdp_{t+H} - gdp_{t-1}, u_t) \Sigma_u^{-1} = \text{Cov} \left(\sum_{\ell=0}^H w_{1,t+\ell}, u_t \right) \Sigma_u^{-1} = \sum_{\ell=0}^H \text{Cov}(w_{1,t+\ell}, u_t) \Sigma_u^{-1}. \quad (27)$$

On the other hand, the Wold decomposition (25) implies (recall that u_t is white noise)

$$\sum_{\ell=0}^{\infty} \text{Cov}(w_{t+\ell}, u_t) \Sigma_u^{-1} = \sum_{\ell=0}^{\infty} C_\ell = C(1). \quad (28)$$

Comparing (27) and (28), we get the desired result (26). \square

A.3 Best linear approximation under non-linearity

Here we give the technical details behind the “best linear approximation” interpretation of a non-linear model, cf. Section 4.4. Assume the nonparametric model (19), and that $\{w_t\}$ is covariance stationary and purely nondeterministic. Let the linear projection of w_t on the orthonormal basis $(\varepsilon_t, \varepsilon_{t-1}, \varepsilon_{t-2}, \dots)$ be denoted $\sum_{\ell=0}^{\infty} \Theta_\ell^* \varepsilon_{t-\ell}$, with projection residual v_t . Assume v_t is either identically zero or purely non-deterministic. Then it has a Wold decomposition

$$v_t = \mu^* + \sum_{\ell=0}^{\infty} \Psi_\ell^* \zeta_{t-\ell},$$

where $\{\zeta_t\}$ is n_w -dimensional white noise with $\text{Cov}(\zeta_t) = I_{n_w}$. Since v_t is a function of $\{\varepsilon_\tau\}_{\tau \leq t}$, and $\{\varepsilon_t\}$ is i.i.d., we have $\text{Cov}(\varepsilon_{t+\ell}, v_t) = 0_{n_\varepsilon \times n_w}$ for all $\ell \geq 1$. Moreover, since v_t is a residual from a projection onto $\{\varepsilon_\tau\}_{\tau \leq t}$, we also have $\text{Cov}(\varepsilon_{t+\ell}, v_t) = 0_{n_\varepsilon \times n_w}$ for all $\ell \leq 0$. By the Wold decomposition theorem, ζ_t lies in the closed linear span of $\{v_\tau\}_{\tau \leq t}$, so we must have $\text{Cov}(\varepsilon_{t+\ell}, \zeta_t) = 0_{n_\varepsilon \times n_w}$ for all $\ell \in \mathbb{Z}$. Finally, the best linear approximation property (20) is a standard consequence of linear projection. We have thus verified all claims made in Section 4.4. \square

References

- 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.
- Baek, C. & Lee, B. (2019). ABCs of Understanding Single Equation Regressions. Manuscript, University of California, Berkeley.
- Barakchian, S. M. & Crowe, C. (2013). Monetary policy matters: Evidence from new shocks data. *Journal of Monetary Economics*, 60(8), 950–966.
- Barnichon, R. & Brownlees, C. (2019). Impulse Response Estimation by Smooth Local Projections. *Review of Economics and Statistics*, 101(3), 522–530.
- 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.
- 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.
- 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. (2018). Robust Bayesian Inference for Set-Identified Models. Manuscript, University College London.
- Giannone, D., Lenza, M., & Primiceri, G. E. (2015). Prior Selection for Vector Autoregressions. *Review of Economics and Statistics*, 97(2), 436–451.
- Gilchrist, S. & Zakrajšek, E. (2012). Credit Spreads and Business Cycle Fluctuations. *American Economic Review*, 102(4), 1692–1720.
- Hansen, B. E. (2010). Multi-Step Forecast Model Selection. Manuscript, University of Wisconsin-Madison.
- Hansen, B. E. (2016). Stein Combination Shrinkage for Vector Autoregressions. Manuscript, University of Wisconsin-Madison.

- Jarociński, M. & Karadi, P. (2019). Deconstructing Monetary Policy Surprises—The Role of Information Shocks. *American Economic Journal: Macroeconomics*. Forthcoming.
- Jordà, Ò. (2005). Estimation and Inference of Impulse Responses by Local Projections. *American Economic Review*, 95(1), 161–182.
- 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.
- 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.
- 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. (2018a). Identification with External Instruments in Structural VARs under Partial Invertibility. Manuscript, Bank of England.
- Miranda-Agrippino, S. & Ricco, G. (2018b). The Transmission of Monetary Policy Shocks. Manuscript, Bank of England.
- Nakamura, E. & Steinsson, J. (2018). Identification in Macroeconomics. *Journal of Economic Perspectives*, 32(3), 59–86.
- Noh, E. (2018). Impulse-response analysis with proxy variables. Manuscript, University of California San Diego.
- Paul, P. (2018). The Time-Varying Effect of Monetary Policy on Asset Prices. Federal Reserve Bank of San Francisco Working Paper 2017-09.
- 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. (2016). *Essays in Macroeconometrics*. PhD thesis, Harvard University.
- 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. Manuscript, Princeton University.