Instrumental Variable Identification of Dynamic Variance Decompositions*

Mikkel Plagborg-Møller
Princeton University

Christian K. Wolf
Princeton University

This version: June 9, 2019
First version: August 10, 2017

Abstract: Macroeconomists often estimate impulse response functions using external instruments (proxy variables) for the shocks of interest. However, existing methods are silent on the importance of the structural shocks for macroeconomic fluctuations. We provide tools for doing inference on variance decompositions in a general semiparametric moving average model, disciplined only by the availability of external instruments. The share of the variance that can be attributed to a shock is partially identified, albeit with informative bounds. Point identification of most parameters, including historical decompositions, can be achieved under additional assumptions that are weaker than invertibility, a condition imposed in conventional Structural Vector Autoregressive analysis. In fact, we prove that invertibility is testable in the presence of external instruments. We illustrate the practical usefulness of our methods by obtaining a tight upper bound on the importance of monetary policy shocks for U.S. inflation dynamics.

Keywords: external instrument, impulse response function, invertibility, proxy variable, variance decomposition. *JEL codes:* C32, C36.

*Email: mikkelpm@princeton.edu and ckwolf@princeton.edu. We received helpful comments from Isaiah Andrews, Tim Armstrong, Dario Caldara, Thorsten Drautzburg, Domenico Giannone, Yuriy Gorodnichenko, Ed Herbst, Marek Jarociński, Peter Karadi, Lutz Kilian, Michal Kolesár, Byoungchan Lee, Sophocles Mavroeidis, Pepe Montiel Olea, Ulrich Müller, Emi Nakamura, Giorgio Primiceri, Eric Renault, Giovanni Ricco, Luca Sala, Jón Steinsson, Jim Stock, Harald Uhlig, Mark Watson, and seminar participants at several venues. The first draft of this paper was written while Wolf was visiting the Bundesbank, whose hospitality is gratefully acknowledged. Wolf also acknowledges support from the Alfred P. Sloan Foundation and the Macro Financial Modeling Project.

1 Introduction

Empirical macroeconomists increasingly seek to estimate impulse response functions using easily interpretable and credible identification approaches. For example, local projections (LPs) have become a popular direct regression-based alternative to Structural Vector Autoregressions (SVARs) (Jordà, 2005; Angrist et al., 2018). Additionally, instrumental variable (IV, also known as proxy variable) methods are now routinely used to conduct structural analysis (Stock, 2008; Stock & Watson, 2012; Mertens & Ravn, 2013). Several recent papers have combined these two ideas, yielding an appealing semiparametric method, LP-IV, with a transparent framework for identification and estimation of impulse response functions (Mertens, 2015; Ramey, 2016; Stock & Watson, 2018).

However, researchers often care not just about impulse responses – they want to know how *important* different shocks are in driving economic fluctuations. In theoretical and applied macroeconomics, shock importance is usually quantified through variance decompositions and historical decompositions. Variance decompositions measure the fraction of the overall (unconditional or forecast) variance of a variable that can be attributed to each of the shocks, while historical decompositions measure the contributions of each shock to observed fluctuations at specific points in time. These decompositions have served as key tools for distinguishing between competing business cycle theories since Kydland & Prescott (1982).[1]

Although identifying shock importance is straight-forward in SVARs due to the implicit invertibility assumption, there exist no such methods in the less restrictive LP-IV framework. In the LP-IV model, existing results on how to identify (normalized) impulse responses do not yield identification of variance/historical decompositions. This is because the IV is usually assumed to be a noisy measure of the true shock, and we do not know the signal-to-noise ratio *a priori*. Traditional SVAR methods are able to identify the shock of interest because these methods assume *invertibility*: The shock must be obtainable as a linear function of current and past (but not future) values of the observed macro variables. This stark requirement implies the often unrealistic property that the econometrician's forecast errors must also be surprises to economic agents who observe the underlying shocks (Hansen & Sargent, 1991; Lippi & Reichlin, 1994; Nakamura & Steinsson, 2018). The invertibility restriction is

[1] For example, variance decompositions have been used to quantify the importance of productivity shocks (King et al., 1991), monetary shocks (Romer & Romer, 1989; Christiano et al., 1999), investment shocks (Justiniano et al., 2010), news shocks (Schmitt-Grohé & Uribe, 2012), financial shocks (Jermann & Quadrini, 2012; Christiano et al., 2014), and sentiment shocks (Angeletos et al., 2018). Cochrane (1994) and Smets & Wouters (2007) perform comprehensive shock accounting exercises.

especially implausible in applications where agents have (perhaps imperfect) foresight about economic fundamentals (Blanchard et al., 2013; Leeper et al., 2013). Stock & Watson (2018) highlight as a key attraction of the LP-IV framework that impulse responses are identified without assuming invertibility. It is therefore unfortunate that it has hitherto been unknown whether variance/historical decompositions are identifiable without assuming invertibility.

In this paper, we show precisely to what extent the data are informative about the importance of shocks in a general linear dynamic model disciplined by IVs. Hence, we expand the applied macroeconomist's toolkit to allow identification of not just impulse responses but also variance/historical decompositions, without requiring invertibility of shocks. We allow for an unrestricted moving average structure of shock transmission, consistent with essentially all linearized structural macroeconomic models. Assuming only validity of the instruments, we derive sharp – and informative – *bounds* on variance decompositions. *Point identification* of most parameters of interest, including historical decompositions, can be achieved under the additional assumption that the shock of interest is recoverable from the infinite past, present, and *future* values of the endogenous macro variables. This requirement is substantially weaker than the invertibility assumption. Finally, to perform inference, we develop easily computable, partial identification robust confidence intervals for variance decompositions and other objects of interest.

Throughout, we adopt the LP-IV model of Stock & Watson (2018), which, although linear, is semiparametric in the sense that we allow for a completely general infinite moving average structure for the transmission of shocks to observed variables.[2] Our sole assumption on the IVs is the usual exclusion restriction – the IVs correlate with the shock of interest, but not with the other macro shocks. Importantly, we allow the number of underlying exogenous shocks to be unknown and potentially exceed the number of observed endogenous variables. In particular, and unlike standard SVAR models, we do not restrict the shocks to be invertible. Stock & Watson (2018) show in this setting that *relative* impulse responses can be point-identified through simple two-stage least squares regressions; however, these do not pin down the *scale* of the shock of interest, which is crucial for identifying variance/historical decompositions.

In this baseline LP-IV model, we show that variance decompositions are only partially identified, albeit with informative bounds. Hence, even with an infinite sample, it would be impossible to pinpoint the exact importance of the shock of interest. Intuitively, as in the

[2]This Structural Vector Moving Average Model has been analyzed recently from a Bayesian viewpoint by Barnichon & Matthes (2018) and Plagborg-Møller (2019), although with little emphasis on IVs.

classical measurement error model, the challenge is that we do not know the signal-to-noise ratio in the IV equation *a priori*. We show, however, that the data are informative about this ratio. Specifically, the identified set of the variance decomposition is an interval, with nontrivial lower and upper bounds computable from the joint spectral density of the macro variables and the IV. The bounds depend on the strength of the external IV and the informativeness of the observed macro variables about the shock of interest. A *sufficient* condition for our upper bound to be tight is that the shock of interest is close to being invertible; yet, we show through theory and examples that our bounds can be highly informative about shock importance even when the shock is far from being invertible.

As the LP-IV model does not assume *a priori* that shocks are invertible, we are able to characterize the extent to which the data are informative about the "degree" of invertibility. Inference about invertibility is useful for gauging the ability of VAR models to perform valid structural analysis and for distinguishing between different classes of structural models, such as models with anticipated versus surprise shocks. The *degree* of invertibility of a shock is given by the R^2 in an (infeasible) regression of the shock on past and current values of the endogenous variables. We show that this R^2 measure is partially identified: The distribution of the data is inconsistent with invertibility if and only if the IV Granger-causes the observed endogenous variables. Without invertibility, the popular SVAR-IV estimator, which uses IVs to partially identify conventional SVARs, is inconsistent and overstates the (impact) forecast variance decomposition of the shock.[3] With invertibility, both SVAR-IV and our upper bound consistently estimate true shock importance.

Although the baseline model is partially identified, we additionally provide assumptions that guarantee point identification of certain variance decompositions and the degree of invertibility. We prove that point identification obtains if the shock of interest is *recoverable*, i.e., spanned by the infinite past, present, and future of the endogenous macro variables. This assumption also yields point identification of historical decompositions. The recoverability condition – although restrictive – is satisfied in certain classes of macro models, such as news and noise shock models, and it is substantially weaker than the invertibility condition that is automatically, if unintentionally, assumed in SVAR analysis. In particular, recoverability obtains if there are as many observed variables as shocks – a necessary, but not sufficient condition for invertibility. Still, we stress that researchers do not need to adopt any auxiliary assumptions to *partially* identify variance decompositions.

[3] The SVAR-IV estimator was developed by Stock (2008), Stock & Watson (2012), and Mertens & Ravn (2013). We characterize the population bias of SVAR-IV under non-invertibility in Online Appendix B.4.

To make our identification analysis practically useful, we develop partial identification robust confidence intervals for all objects of interest. In a first step, the researcher estimates a reduced-form VAR jointly in the macro variables and IVs. To be clear, this step merely uses the reduced-form VAR as a convenient tool for approximating the second moments of the data; it does not assume an underlying *structural* VAR model with invertible shocks.[4] The second step then constructs sample analogues of our population partial identification bounds and inserts these into the confidence procedure of Imbens & Manski (2004) and Stoye (2009). We prove that our confidence intervals have asymptotically valid frequentist coverage under weak *nonparametric* conditions on the data generating process. We also discuss a test of invertibility that has power against all falsifiable noninvertible alternatives.

We illustrate the usefulness of our identification bounds through the lens of the well-known structural business cycle model of Smets & Wouters (2007). We assume that the econometrician observes aggregate output, inflation, and a short-term policy interest rate, but she does not exploit the underlying structure of the model for inference. We separately consider external instruments for three different shocks: a standard monetary policy shock, a forward guidance (anticipated monetary) shock, and a technology shock. These three shocks vary greatly in terms of their degree of invertibility and recoverability, and we show that SVAR-based identification of the latter two shocks is severely biased. Nevertheless, our partial identification bounds are informative in all cases. This result is particularly striking for the technology shock, since the macro aggregates provide little information about its short- or medium-run cycles.

Finally, we apply our method to study the importance of monetary shocks in driving U.S. macroeconomic aggregates. Following Gertler & Karadi (2015), the external instrument for monetary policy shocks is constructed from changes in interest rate futures during short time windows around Federal Open Market Committee announcements. As discussed in Ramey (2016), the rising importance of forward guidance since the early 1990s is likely to invalidate the invertibility assumption and so threatens consistency of the standard SVAR-IV estimator; consistent with her concerns, we reject invertibility. Applying our novel methodology, we find that monetary shocks are almost irrelevant for aggregate inflation: Our partial identification robust 90% confidence intervals for the forecast variance contribution of monetary shocks rules out values above 8% at all horizons. This conclusion resonates with the recently

[4] We view the reduced-form VAR step as a dimension reduction technique, to address finite-sample concerns about unrestricted local projections raised by Kilian & Kim (2011). It is straight-forward to base inference on other first-step estimators of the joint spectrum of the data.

documented *unconditional* divorce of inflation from other macroeconomic aggregates (Hall, 2011). Thus, to the extent that inflation is a monetary phenomenon, it is so because of the systematic part of U.S. monetary policy, not because of its erratic conduct. Importantly, our results are obtained under much weaker identifying assumptions than existing methods.

LITERATURE. A growing literature has provided inference tools for the LP-IV model, although variance/historical decompositions have been neglected. External IVs relax the assumption that the shock in a local projection is directly observed (or consistently estimable). Theoretical results on LP-IV estimation of impulse response functions were established by Mertens (2015), Ramey (2016), Barnichon & Brownlees (2018), Jordà et al. (2018), Ramey & Zubairy (2018), and Stock & Watson (2018). These papers identify *relative* impulse responses (e.g., the responses of the macro variables to a shock which raises the first variable by 1 unit). We go further and derive the identified set of all LP-IV model parameters.

Our identification results for variance decompositions generalize several results in the literature. Variance decompositions are frequently reported in SVAR analysis, where identification is straight-forward due to the invertibility assumption (Kilian & Lütkepohl, 2017, Ch. 4). Stock & Watson (2018) assume invertibility of all shocks to identify forecast variance decompositions and historical decompositions in an LP-IV model; we substantially strengthen this result by showing that much weaker assumptions than invertibility (e.g., recoverability) are in fact sufficient for point identification of some of these objects. Our results are complementary to Gorodnichenko & Lee (2017), who consider finite-sample inference on what we call the "forecast variance ratio" in local projection models where the shock is assumed to be directly observed.

A key attraction of the LP-IV framework is that it allows for noninvertible shocks, unlike the standard SVAR model. Noninvertibility is now known to occur in many classes of structural models where economic agents observe better information than the econometrician, such as models with news shocks, private signals, or measurement error. Hence, the issue has received a lot of attention in the SVAR literature (see references in Plagborg-Møller, 2019, Sec. 2.3). Stock & Watson (2018) develop an LP-IV-based test of noninvertibility. Our contribution in this area is to sharply characterize the identified set for the *degree* of invertibility of the shocks, which in turn shows under what conditions the distribution of the data is consistent with invertibility. These conditions are related to Granger causality, as in the SVAR settings studied by Giannone & Reichlin (2006) and Forni & Gambetti (2014).

The weaker concept of "recoverability" has precursors in the literature. Our definition of recoverability has independently been proposed by Chahrour & Jurado (2018). Their generic framework does not specifically consider the LP-IV model, where the recoverability assumption is testable, as we show. As discussed below, recoverability is formally equivalent to recovering shocks from *dynamic* rotations of reduced-form VAR errors, as used by Lippi & Reichlin (1994), Mertens & Ravn (2010), and Forni et al. (2017a,b).

Our confidence interval procedure applies the generic methods for interval-identified parameters developed by Imbens & Manski (2004) and Stoye (2009). Our analysis is distinct from and complementary to the literature on inference in sign-identified SVARs (Gafarov et al., 2018; Giacomini & Kitagawa, 2018; Granziera et al., 2018).

Outline. Section 2 defines the LP-IV model and the parameters of interest. Section 3 contains our main identification results. Section 4 interprets the results through the lens of a structural macro model. Section 5 develops confidence intervals. Section 6 contains the empirical application. Section 7 concludes. Appendix A provides inference formulas and proofs of our main results. A supplementary appendix and Matlab code are available online.[5]

2 Model and parameters of interest

We begin by defining the Local Projection Instrumental Variable (LP-IV) model and its parameters of interest. The LP-IV model allows for an unrestricted linear shock transmission mechanism and, unlike standard SVAR analysis, does not assume shocks to be invertible. We assume the availability of valid external IVs (proxy variables) – variables that correlate with the shock of interest, but not with the other shocks. Although the model we study is identical to that of Stock & Watson (2018), our parameters of interest are entirely different: They study relative impulse responses, whereas we study variance decompositions, historical decompositions, and the degree of invertibility.

Model. We start out by describing the LP-IV model's semiparametric assumptions on shock transmission and the instrument exclusion restrictions. For notational clarity, we assume throughout that all time series below have mean zero and are strictly non-deterministic.

First, we specify the weak assumptions on shock transmission to endogenous variables.

[5] https://scholar.princeton.edu/mikkelpm/decomp_iv

Assumption 1. *The n_y-dimensional vector $y_t = (y_{1,t}, \dots, y_{n_y,t})'$ of observed macro variables is driven by an unobserved n_ε-dimensional vector $\varepsilon_t = (\varepsilon_{1,t}, \dots, \varepsilon_{n_\varepsilon,t})'$ of exogenous economic shocks,*

$$y_t = \Theta(L)\varepsilon_t, \quad \Theta(L) \equiv \sum_{\ell=0}^{\infty} \Theta_\ell L^\ell, \tag{1}$$

where L is the lag operator. The matrices Θ_ℓ are each $n_y \times n_\varepsilon$ and absolutely summable across ℓ. $\Theta(x)$ is assumed to have full row rank for all complex scalars x on the unit circle. The shocks are mutually orthogonal white noise processes:

$$\varepsilon_t \sim WN(0, I_{n_\varepsilon}),$$

where I_n denotes the n-dimensional identity matrix.

The (i, j) element $\Theta_{i,j,\ell}$ of the moving average coefficient matrix Θ_ℓ is the *impulse response* of variable i to shock j at horizon ℓ. The j-th column of Θ_ℓ is denoted by $\Theta_{\bullet,j,\ell}$ and the i-th row by $\Theta_{i,\bullet,\ell}$. The full-rank assumption guarantees a nonsingular stochastic process. This condition requires $n_\varepsilon \geq n_y$, but – crucially – we do not assume that the number of shocks n_ε is known. The mutual orthogonality of the shocks is the standard assumption in empirical macroeconomics. The model is semiparametric in that we place no *a priori* restrictions on the coefficients of the infinite moving average, except to ensure a valid stochastic process. In particular, we do not impose the usual invertibility conditions that point-identify $\Theta(L)$ in reduced-form time series analysis. It is well known that the infinite-order Structural Vector Moving Average model (1) is consistent with discrete-time Dynamic Stochastic General Equilibrium (DSGE) models as well as stable SVAR models for y_t. However, the appeal of LP-IV analysis is that it does not require any specific underlying structure.

Second, we assume the availability of one or more external IVs for the shock of interest. We specify the shock of interest to be the first one, $\varepsilon_{1,t}$. Each of the n_z IVs $z_t = (z_{1,t}, \dots, z_{n_z,t})'$ are assumed to correlate with the first shock but not the other shocks, after controlling for lagged variables: For all $i = 1, \dots, n_z$,

$$E(\tilde{z}_{i,t}\varepsilon_{1,t}) \neq 0, \quad E(\tilde{z}_{i,t}\varepsilon_{j,\tau}) = 0 \text{ for all } (j, \tau) \neq (i, t), \tag{2}$$

where $\tilde{z}_{i,t}$ is the population residual from projecting $z_{i,t}$ on all lags of $\{z_t, y_t\}$. The key exclusion restrictions are that the shock of interest $\varepsilon_{1,t}$ is the only contemporaneous shock to correlate with the IVs. Thus, $\tilde{z}_t$ is a proxy for $\varepsilon_{1,t}$ (up to scale) that is contaminated by classical measurement error. This is a strong assumption that must be carefully defended in applications. Ramey (2016) and Stock & Watson (2018) survey the extensive applied literature that has constructed plausibly valid external IVs for various shocks. In Online Appendix B.3 we discuss how our analysis changes if the exclusion restriction is relaxed.

Using linear projection notation, we can equivalently express the IV exclusion restrictions (2) as follows. $\|\cdot\|$ refers to the Euclidean norm.

Assumption 2. *The IVs* $z_t = (z_{1,t}, \dots, z_{n_z,t})'$ *satisfy*

$$z_t = \sum_{\ell=1}^{\infty}(\Psi_\ell z_{t-\ell} + \Lambda_\ell y_{t-\ell}) + \alpha\lambda\varepsilon_{1,t} + \Sigma_v^{1/2} v_t, \tag{3}$$

where Ψ_ℓ *is* $n_z \times n_z$*,* Λ_ℓ *is* $n_z \times n_y$*,* λ *is an* n_z*-dimensional vector normalized to unit length (*$\|\lambda\| = 1$*) and with its first nonzero element being positive,* $\alpha \geq 0$ *is a scalar, and* Σ_v *is a symmetric positive semidefinite* $n_z \times n_z$ *matrix. The elements of* Ψ_ℓ *and* Λ_ℓ *are absolutely summable across* ℓ*, and the polynomial* $x \mapsto \det(I_{n_z} - \sum_{\ell=1}^{\infty} \Psi_\ell x^\ell)$ *has all its roots outside the unit circle. The disturbance vector* v_t *is a white noise process that is dynamically uncorrelated with the structural shocks* ε_t*:*

$$v_t \sim WN(0, I_{n_z}), \quad \mathrm{Cov}(\varepsilon_t, v_\tau) = 0_{n_\varepsilon \times n_z} \text{ for all } t, \tau.$$

The scale parameter α (along with the residual variance-covariance matrix Σ_v) measures the overall strength of the IVs, while the unit-length vector λ determines which IVs are stronger than others. We emphasize that the linearity of equation (3) is not a structural assumption; it arises from a linear projection (as in the "first stage" of cross-sectional IV).[6]

Since we restrict attention to identification from second moments, we simplify notation by further assuming that the structural shocks and IV disturbances are Gaussian.

Assumption 3. $(\varepsilon_t', v_t')'$ *is i.i.d. jointly Gaussian.*

The Gaussianity assumption is strictly for notational convenience. We could instead have maintained the above white noise assumptions and phrased all our results using linear projection notation.[7] The sole meaningful restriction is that we only consider identification

[6] We allow for lagged values of z_t and y_t on the right-hand side of (3) because this is precisely enough to ensure that the LP-IV model is untestable (using second moments) in the case of a single IV, cf. Proposition 1 below. The model with multiple instruments is testable, as further discussed in Section 3.3. If lagged terms can be excluded *a priori*, it presents no difficulties for identification or inference.

[7] Simply replace conditional expectations by linear projections and replace conditional variances by variances of projection residuals.

from the second-moment properties of the data, as is standard in the applied macro literature (and without loss of generality for Gaussian data).[8] We will drop the Gaussianity assumption when considering inference in Section 5.

Note that we have normalized the variances of all shocks to 1, without loss of generality, as this simplifies our notation. Stock & Watson (2018) study the same LP-IV model as us, but they instead normalize certain impact impulse responses to equal 1, while letting the the shock variances be unrestricted. Hence, when Stock & Watson discuss identification of "impulse responses", this translates into our notation as identification of *relative* impulse responses of the type $\Theta_{i,1,\ell}/\Theta_{1,1,0}$. The choice of normalization of course does not matter for the identification analysis. Finally note also that Assumptions 1 to 3 together imply that the $(n_y + n_z)$-dimensional data vector $(y_t', z_t')'$ is strictly stationary.

Invertibility and recoverability. We now define invertibility, the degree of invertibility, and recoverability.

The shock $\varepsilon_{1,t}$ is said to be *invertible* if it is spanned by past and current (but not future) values of the endogenous variables y_t: $\varepsilon_{1,t} = E(\varepsilon_{1,t} \mid \{y_\tau\}_{-\infty<\tau\leq t})$. Invertibility of all structural shocks is assumed automatically by SVAR models, but the condition may or may not hold in a given moving average model (1), depending on the impulse response parameters Θ_ℓ. A sufficient condition for invertibility of all shocks is that $n_\varepsilon = n_y$ and the polynomial $x \mapsto \det(\Theta(x))$ has all its roots outside the unit circle. In many structural macro models, at least some of the shocks cannot be recovered from only past and current observed macro variables, i.e., the moving average representation is noninvertible. For example, this is often the case in models with news (anticipated) shocks or noise (signal extraction) shocks (Blanchard et al., 2013; Leeper et al., 2013). Furthermore, if $n_\varepsilon > n_y$, it is impossible for all shocks to be invertible.

A continuous measure of the *degree* of invertibility is the R^2 value in a population regression of the shock on past and current observed variables. More generally, define

$$R_\ell^2 \equiv \frac{\mathrm{Var}(\varepsilon_{1,t}) - \mathrm{Var}(\varepsilon_{1,t} \mid \{y_\tau\}_{-\infty<\tau\leq t+\ell})}{\mathrm{Var}(\varepsilon_{1,t})} = 1 - \mathrm{Var}(\varepsilon_{1,t} \mid \{y_\tau\}_{-\infty<\tau\leq t+\ell})$$

as an R^2 measure of invertibility of the shock of interest using data up to time $t + \ell$. If

[8] If we were to take the assumption of i.i.d. shocks seriously, and the shocks were *not* Gaussian, higher-order moments of the data would be informative about the parameters. However, we agree with most of the literature that the assumption of i.i.d. shocks is too strong due to the likely presence of stochastic volatility (SV). Our identification results allow for SV since we only require shocks to be white noise.

the shock is invertible in the sense of the previous paragraph, then $R_\ell^2 = 1$ for all $\ell \geq 0$. If $R_\ell^2 < 1$ for some $\ell \geq 0$, then the model is noninvertible and thus no SVAR model could generate the impulse responses $\Theta(L)$, although the model may be *nearly* consistent with an SVAR structure if the R^2 values are close to 1, as we discuss further in Online Appendix B.4 (Sims & Zha, 2006, pp. 243–245; Forni et al., 2018; Wolf, 2018). For noninvertible models, a plot of R_ℓ^2 for $\ell = 0, 1, 2, \ldots$ reveals how quickly the econometrician learns about the structural shock over time. To illustrate, we derive the R_ℓ^2 value for an MA(1) model in Online Appendix B.5.

A weaker condition than invertibility is that the shock of interest is *recoverable* from all leads and lags of the endogenous variables – that is, if $E(\varepsilon_{1,t} \mid \{y_\tau\}_{-\infty<\tau<\infty}) = \varepsilon_{1,t}$, or equivalently if $R_\infty^2 = 1$.[9] This property will become important when we consider assumptions that guarantee point identification.

Variance decompositions. Variance decompositions are the key parameters of interest in this paper. We now define several variance decomposition objects, including forecast variance decompositions (either conditioning on past observables or on past shocks) and a frequency-specific unconditional variance decomposition.

We consider two forecast variance decomposition concepts. First, define the *forecast variance ratio* (FVR) for the shock of interest for variable i at horizon ℓ as

$$FVR_{i,\ell} \equiv 1 - \frac{\mathrm{Var}(y_{i,t+\ell} \mid \{y_\tau\}_{-\infty<\tau\leq t}, \{\varepsilon_{1,\tau}\}_{t<\tau<\infty})}{\mathrm{Var}(y_{i,t+\ell} \mid \{y_\tau\}_{-\infty<\tau\leq t})} = \frac{\sum_{m=0}^{\ell-1} \Theta_{i,1,m}^2}{\mathrm{Var}(y_{i,t+\ell} \mid \{y_\tau\}_{-\infty<\tau\leq t})}$$

The FVR measures the reduction in the forecast variance that would come from knowing the entire path of future realizations of the first shock. The larger this measure is, the more important is the first shock for forecasting variable i at horizon ℓ. The FVR is always between 0 and 1. An unappealing feature, however, is that the FVR conflates two different sources of uncertainty: fundamental forecasting uncertainty (uncertainty related to future shock realizations) and noninvertibility-induced uncertainty (uncertainty related to imperfect knowledge about *past* shocks).[10] This means that, when $\varepsilon_{1,t}$ is noninvertible, the FVR does

[9] Chahrour & Jurado (2018) independently propose this definition. Recoverability is formally equivalent to the assumption that the structural shock is spanned by current and *future* reduced-form forecast errors $u_t \equiv y_t - E(y_t \mid \{y_\tau\}_{-\infty<\tau\leq t-1})$. Such *dynamic rotations* of u_t have been exploited for identification by Lippi & Reichlin (1994), Mertens & Ravn (2010), and Forni et al. (2017a,b).

[10] $\mathrm{Var}(y_{i,t+\ell} \mid \{y_\tau\}_{-\infty<\tau\leq t}) = \sum_{m=0}^{\ell-1} \Theta_{i,\bullet,m}\Theta_{i,\bullet,m}' + \mathrm{Var}\left(\sum_{m=\ell}^{\infty} \Theta_{i,\bullet,m}\varepsilon_{t+l-m} \,\middle|\, \{y_\tau\}_{-\infty<\tau\leq t}\right)$.

not equal 1 even if $\varepsilon_{1,t}$ is solely responsible for driving the i-th variable in equation (1).

The second variance decomposition concept is the *forecast variance decomposition* (FVD) for the shock of interest for variable i at horizon ℓ,

$$FVD_{i,\ell} \equiv 1 - \frac{\operatorname{Var}(y_{i,t+\ell} \mid \{\varepsilon_\tau\}_{-\infty<\tau\leq t}, \{\varepsilon_{1,\tau}\}_{t<\tau<\infty})}{\operatorname{Var}(y_{i,t+\ell} \mid \{\varepsilon_\tau\}_{-\infty<\tau\leq t})} = \frac{\sum_{m=0}^{\ell-1} \Theta_{i,1,m}^2}{\sum_{j=1}^{n_\varepsilon} \sum_{m=0}^{\ell-1} \Theta_{i,j,m}^2}. \tag{4}$$

The FVD measures the reduction in forecast variance that arises from learning the path of future realizations of the shock of interest, supposing that we already had the history of structural shocks ε_t available when forming our forecast. Because the econometrician generally does not observe the structural shocks directly, the FVD is best thought of as reflecting forecasts of economic agents who observe the underlying shocks. The FVD always lies between 0 and 1, purely reflects fundamental forecasting uncertainty, and equals 1 if the first shock is the only shock driving variable i in equation (1). The software package Dynare reports FVDs after having estimated a DSGE model.

While the FVR and FVD concepts generally differ, they coincide in the case where all shocks are invertible, since in that case the information set $\{y_\tau\}_{-\infty<\tau\leq t}$ equals the information set $\{\varepsilon_\tau\}_{-\infty<\tau\leq t}$. This explains why the SVAR literature has not made the distinction between the two concepts.[11]

For completeness, we also consider the frequency-specific *unconditional* variance decomposition (VD) of Forni et al. (2018, Sec. 3.4). The VD for variable i over the frequency band $[\omega_1, \omega_2]$ is given by

$$VD_i(\omega_1, \omega_2) \equiv \frac{\int_{\omega_1}^{\omega_2} |\Theta_{i,1}(e^{-i\omega})|^2 \, d\omega}{\sum_{j=1}^{n_\varepsilon} \int_{\omega_1}^{\omega_2} |\Theta_{i,j}(e^{-i\omega})|^2 \, d\omega}, \quad 0 \leq \omega_1 < \omega_2 \leq \pi, \tag{5}$$

where $\Theta_{i,j}(L)$ is the (i,j) element of the lag polynomial $\Theta(L)$. $VD_i(\omega_1, \omega_2)$ is the percentage reduction in the variance of $y_{i,t}$ – after passing the data through a bandpass filter that retains only cyclical frequencies $[\omega_1, \omega_2]$ – caused by entirely "shutting off" the shock of interest $\varepsilon_{1,t}$. The software package Dynare automatically reports $VD_i(0, \pi)$ after solving a DSGE model.

Historical decomposition. The *historical decomposition* of variable $y_{i,t}$ at time t attributable to the shock of interest is defined as $E(y_{i,t} \mid \{\varepsilon_{1,\tau}\}_{-\infty<\tau\leq t}) = \sum_{\ell=0}^{\infty} \Theta_{i,1,\ell}\varepsilon_{1,t-\ell}$.

[11] Forni et al. (2018) point out the bias caused by noninvertibility when estimating the FVD using SVARs.

3 Identification

This section studies the identification of our parameters of interest in the LP-IV model. For exposition, we start by deriving results for a static version of the model. We then turn to the general dynamic model, which applies the static results to the frequency domain representation of the data.

3.1 Static model

We use an illustrative static model to motivate why variance decompositions are partially identified in the general case but can be point-identified under additional assumptions. Although the static model does not capture all the nuances of the dynamic LP-IV model, it provides useful intuition for the general case.

Model. Consider a static version of our model with a single IV:[12]

$$
\begin{aligned}
y_t &= \Theta_{\bullet,1,0}\varepsilon_{1,t} + \xi_t, \\
z_t &= \alpha\varepsilon_{1,t} + \sigma_v v_t, \\
(\varepsilon_{1,t}, v_t, \xi_t)' &\overset{i.i.d.}{\sim} N\left(0, \begin{pmatrix} I_2 & 0_{2\times n_y} \\ 0_{n_y\times 2} & \Sigma_\xi \end{pmatrix}\right).
\end{aligned}
$$

Here $\alpha, \sigma_v \geq 0$ are scalars, $\xi_t \equiv \sum_{j=2}^{n_\varepsilon} \Theta_{\bullet,j,0}\varepsilon_{j,t}$ is an n_y-dimensional random vector that captures all the structural shocks other than the one of interest, and $\Sigma_\xi \equiv \text{Var}(\xi_t)$.

Note that the static model is nothing but a multivariate classical measurement error model: To gauge the importance of $\varepsilon_{1,t}$ in driving y_t, we would like to regress y_t on $\varepsilon_{1,t}$; however, instead of the shock, we only observe the noisy proxy z_t. By the usual measurement error logic, the importance of $\varepsilon_{1,t}$ for y_t is not point-identified because the signal-to-noise ratio α^2/σ_v^2 for the proxy z_t is not known *a priori*: A high (low) observed correlation $\text{Corr}(y_{i,t}, z_t)$ may be due to $\varepsilon_{1,t}$ being an important (unimportant) driver of $y_{i,t}$ or due to the proxy z_t being strong (weak). However, as in the measurement error literature, we now show that we

[12]While the static model is primarily intended for gaining intuition, the results in this subsection are directly relevant for SVAR analysis with an external IV. In that framework, y_t would be the reduced-form VAR residuals, which are a linear function of the vector ε_t of contemporaneous structural shocks. Textbook SVAR analysis further assumes that $n_\varepsilon = n_y$, so the model is identified up to an orthogonal rotation matrix.

can obtain *bounds* on the importance of $\varepsilon_{1,t}$.[13]

In the static model there is no distinction between the FVR, FVD, and VD, so our parameter of interest for now is

$$FVD_{i,1} = 1 - \frac{\operatorname{Var}(y_{i,t} \mid \varepsilon_{1,t})}{\operatorname{Var}(y_{i,t})} = \frac{\Theta_{i,1,0}^2}{\operatorname{Var}(y_{i,t})}.$$

This is the usual population R^2 value in the (infeasible) regression of $y_{i,t}$ on $\varepsilon_{1,t}$.

IDENTIFICATION. We can easily construct lower and upper bounds for $FVD_{i,1}$. For a lower bound, consider a simple regression of $y_{i,t}$ on the proxy $z_t = \alpha\varepsilon_{1,t} + \sigma_v v_t$. The R^2 from this regression is

$$\frac{\operatorname{Var}(E[y_{i,t} \mid z_t])}{\operatorname{Var}(y_{i,t})} = \frac{\operatorname{Cov}(y_{i,t}, z_t)^2}{\operatorname{Var}(z_t)\operatorname{Var}(y_{i,t})} = \frac{\alpha^2}{\alpha^2 + \sigma_v^2} \times FVD_{i,1}.$$

This is the familiar attenuation bias: We are regressing on a noisy measure of the shock, thus understating its importance in driving fluctuations in $y_{i,t}$.[14] For an upper bound, consider a regression of $y_{i,t}$ on the "index" $E(z_t \mid y_t)$. The R^2 from this regression is

$$\frac{\operatorname{Var}(E[y_{i,t} \mid E(z_t \mid y_t)])}{\operatorname{Var}(y_{i,t})} = \frac{\operatorname{Cov}(y_{i,t}, E(z_t \mid y_t))^2}{\operatorname{Var}(E(z_t \mid y_t))\operatorname{Var}(y_{i,t})} = \frac{\operatorname{Var}(\varepsilon_{1,t})}{\operatorname{Var}(E(\varepsilon_{1,t} \mid y_t))} \times FVD_{i,1}.$$

We are now overstating the importance of the shock. Intuitively, by constructing the first-step fitted value $E(z_t \mid y_t) = \alpha E(\varepsilon_{1,t} \mid y_t)$, we remove the measurement error v_t that causes attenuation bias. However, since $1 = \operatorname{Var}(\varepsilon_{1,t}) \geq \operatorname{Var}(E(\varepsilon_{1,t} \mid y_t))$, we are effectively removing too much variation, thus overstating the importance of $\varepsilon_{1,t}$.

Putting the pieces together, we have thus obtained the bounds

$$1 - \frac{\operatorname{Var}(y_{i,t} \mid z_t)}{\operatorname{Var}(y_{i,t})} \leq FVD_{i,1} \leq 1 - \frac{\operatorname{Var}(y_{i,t} \mid E(z_t \mid y_t))}{\operatorname{Var}(y_{i,t})}. \tag{6}$$

It is not hard to show (and it follows from our general results below) that these bounds are *sharp*, i.e., exploit all information about $FVD_{i,1}$ afforded by the distribution of the data.[15]

[13] Because our parameters of interest are not regression coefficients as in Klepper & Leamer (1984), these bounds do not follow immediately from the existing literature, to our knowledge.

[14] A related argument appears in Gorodnichenko & Lee (2017, Appendix D).

[15] Given any positive semidefinite variance-covariance matrix for $w_t = (y_t', z_t)'$, and given any value of $FVD_{i,1}$ in the interval between the bounds, we can construct a static model which matches the given $\operatorname{Var}(w_t)$ and $FVD_{i,1}$ (we just have to choose α, Θ_0, and σ_v appropriately).

The width of the identified set for $FVD_{i,1}$ depends on the degree of invertibility of $\varepsilon_{1,t}$ and on the strength of the proxy/IV z_t. Intuitively, attenuation bias decreases (so the lower bound in (6) is larger) when the IV strength α^2/σ_v is larger. And if $\varepsilon_{1,t}$ is invertible, meaning $E(\varepsilon_{1,t} \mid y_{i,t}) = \varepsilon_{1,t}$, then $E(z_t \mid y_t) \propto \varepsilon_{1,t}$, so the upper bound in (6) equals $FVD_{i,1}$.

It is immediate from the previous discussion that point identification obtains under a variety of auxiliary assumptions. For example, if we assume that the shock of interest is invertible (which in the static model is the same as being recoverable), then we have argued that the upper bound for $FVD_{i,1}$ binds. Alternatively, point identification obtains if the instrument is assumed to be perfect, i.e., $\sigma_v = 0$; then the lower bound for $FVD_{i,1}$ binds.

3.2 General dynamic model

We now present our main identification results for the general dynamic model, applying the logic of the static model frequency-by-frequency to the frequency domain representation of the data. Relative to the static case, the dynamic case involves additional challenges in characterizing the informativeness of the data for the hidden shock at all frequencies.

We maintain Assumptions 1 to 3 throughout, but for the moment, we carry out the analysis for the case of a single IV ($n_z = 1$), leaving the generalization to Section 3.3. That is, z_t is a scalar and $\lambda = 1$ in equation (3). We write $\Sigma_v^{1/2} = \sigma_v \geq 0$, a scalar.

Preliminaries. For the identification analysis, it will prove convenient to define the IV projection residual that removes any dependence on lagged observed variables:

$$\tilde{z}_t \equiv z_t - E(z_t \mid \{y_\tau, z_\tau\}_{-\infty<\tau<t}) = \alpha\varepsilon_{1,t} + \sigma_v v_t. \tag{7}$$

Note that $\tilde{z}_t$ is serially uncorrelated by construction.

Next, we need to define our notation for spectral density matrices. For any two jointly stationary vector time series a_t and b_t of dimensions n_a and n_b, respectively, define the $n_a \times n_b$ cross-spectral density matrix function (Brockwell & Davis, 1991, Ch. 4 and 11)

$$s_{ab}(\omega) \equiv \frac{1}{2\pi} \sum_{\ell=-\infty}^{\infty} e^{-i\omega\ell} \operatorname{Cov}(a_t, b_{t-\ell}), \quad \omega \in [0, 2\pi].$$

For any vector time series a_t, we denote its spectrum by $s_a(\omega) \equiv s_{aa}(\omega)$.

Relative impulse responses. *Absolute* impulse responses to the first shock are identified up to the scale parameter α:

$$\operatorname{Cov}(y_t, \tilde{z}_{t-\ell}) = \alpha \Theta_{\bullet,1,\ell}. \tag{8}$$

Thus, *relative* impulse responses $\Theta_{i,1,\ell}/\Theta_{11,0}$ are point-identified, as shown by Stock & Watson (2018) and others.[16] Relative impulse responses answer a useful question: How much does $y_{i,t+\ell}$ increase following a shock $\varepsilon_{1,t}$ of a magnitude that increases $y_{1,t}$ by one unit on impact?

However, variance decompositions and the degree of invertibility depend on the absolute impulse responses, not just the relative ones. Hence, as is clear from the above display, although the scale parameter α is not economically interesting in itself, it is key to identification of our parameters of interest.

Scale parameter. We now show that the identified set for the scale parameter α is an interval with informative bounds, proceeding as we did in the simple static model.

First, the variance of the instrument provides an upper bound for the scale parameter:

$$\alpha^2 \leq \operatorname{Var}(\tilde{z}_t) \equiv \alpha_{UB}^2. \tag{9}$$

As in the static model, the boundary case $\alpha = \alpha_{UB}$ corresponds to perfect IV informativeness.

Second, for the lower bound, we apply a version of the argument from the static case to the joint spectrum of the data at every frequency. Define first the projections of $\tilde{z}_t$ and $\varepsilon_{1,t}$, respectively, onto all lags and leads of the endogenous variables y_t:

$$\begin{aligned} \tilde{z}_t^\dagger &\equiv E(\tilde{z}_t \mid \{y_\tau\}_{-\infty<\tau<\infty}), \\ \varepsilon_{1,t}^\dagger &\equiv E(\varepsilon_{1,t} \mid \{y_\tau\}_{-\infty<\tau<\infty}). \end{aligned} \tag{10}$$

Note that $\tilde{z}_t^\dagger$ is the analogue of $E(z_t \mid y_t)$ in the static identification analysis. Then, for every frequency $\omega \in [0, 2\pi]$, we have

$$s_{\tilde{z}^\dagger}(\omega) = \alpha^2 s_{\varepsilon_1^\dagger}(\omega) \leq \alpha^2 s_{\varepsilon_1}(\omega) = \alpha^2 \times \tfrac{1}{2\pi}, \tag{11}$$

where the inequality is the frequency-domain analogue of "the explained sum of squares is

[16] Some authors employ the alternative but equivalent unit effect normalization where $\Theta_{11,0} = 1$, but shocks have non-unit variances.

smaller than the total sum of squares".[17] Hence, we obtain the lower bound

$$\alpha^2 \geq 2\pi \sup_{\omega\in[0,\pi]} s_{\tilde{z}^\dagger}(\omega) \equiv \alpha_{LB}^2. \tag{12}$$

Intuitively, a small value of α requires $\tilde{z}_t$ to be nearly unpredictable by y_t at *every* frequency ω, e.g., both in the long run and at business cycle frequencies. The boundary case $\alpha = \alpha_{LB}$ corresponds to the observed macro aggregates being perfectly informative about the hidden shock $\varepsilon_{1,t}$ at *some* frequency $\overline{\omega} \in [0, \pi]$, i.e., the projection residual $\varepsilon_{1,t} - \varepsilon_{1,t}^\dagger$ has a spectral density $s_{\varepsilon_1 - \varepsilon_1^\dagger}(\omega) = s_{\varepsilon_1}(\omega) - s_{\varepsilon_1^\dagger}(\omega)$ that vanishes at frequency $\overline{\omega}$.

The main theoretical result of this paper is that the above bounds $\alpha_{LB}^2, \alpha_{UB}^2$ are *sharp*.

Proposition 1. *Let there be given a joint spectral density for $w_t = (y_t', \tilde{z}_t)'$, continuous and positive definite at every frequency, with $\tilde{z}_t$ unpredictable from $\{w_\tau\}_{-\infty<\tau<t}$. Choose any $\alpha \in (\alpha_{LB}, \alpha_{UB}]$. Then there exists an LP-IV model as in Assumptions 1 and 2 with the given α such that the model-implied spectral density of w_t matches the given spectral density.*

Recall that the previous discussion has already shown that any value of $\alpha^2 \notin [\alpha_{LB}^2, \alpha_{UB}^2]$ is impossible. Thus, the distribution of the data tells us that $\alpha^2 \in [\alpha_{LB}^2, \alpha_{UB}^2]$, but the data cannot rule out any values of α^2 in this interval. The proposition does not cover the knife-edge case $\alpha = \alpha_{LB}$ due to economically inessential technicalities.

The width of the identified set for the scale parameter α depends on the application, although the set is never empty. To interpret the identified set, we express it in terms of the (unknown) model parameters. We focus on the identified set for $\frac{1}{\alpha^2}$, as this transformation is most relevant for identifying $FVD_{i,\ell}$ and R_ℓ^2, as shown below. This identified set equals

$$\left[\underbrace{\frac{\alpha^2}{\alpha^2+\sigma_v^2}}_{\text{instrument strength}} \times \frac{1}{\alpha^2},\ \underbrace{\frac{1}{1-2\pi \inf_{\omega\in[0,\pi]} s_{\varepsilon_1-\varepsilon_1^\dagger}(\omega)}}_{\text{informativeness of data for shock}} \times \frac{1}{\alpha^2}\right].$$

Analogously to the static case, the lower bound of the identified set for $\frac{1}{\alpha^2}$ is larger (and closer to the true $\frac{1}{\alpha^2}$) when the instrument is stronger in the sense of a higher signal-to-noise ratio. The upper bound of the identified set for $\frac{1}{\alpha^2}$ is smaller (and closer to the true $\frac{1}{\alpha^2}$) when the data are more informative about the shock of interest, at least at *some* frequency. Compare with the static case, where only $\text{Var}(\varepsilon_{1,t} \mid y_t)$ mattered for the tightness of the bound. The

[17] Brockwell & Davis (1991, Remark 3, p. 439) show that $s_{\tilde{z}^\dagger}(\omega) = s_{y\tilde{z}}(\omega)^* s_y(\omega)^{-1} s_{y\tilde{z}}(\omega)$ and $s_{\varepsilon_1^\dagger}(\omega) = s_{y\varepsilon_1}(\omega)^* s_y(\omega)^{-1} s_{y\varepsilon_1}(\omega)$. Since the joint spectrum is positive semidefinite, $s_{\varepsilon_1}(\omega) \geq s_{\varepsilon_1^\dagger}(\omega)$ for all ω.

identified set for $\frac{1}{\alpha^2}$ does not collapse to a point unless the instrument is perfect *and* there exists a frequency $\overline{\omega}$ for which the data are perfectly informative about the frequency-$\overline{\omega}$ cyclical component of the shock.

To further interpret α_{LB}^2, we derive a lower bound to this object that is explicitly tied to the degree of recoverability/invertibility. First, we have

$$\alpha_{LB}^2 = 2\pi \sup_{\omega \in [0,\pi]} s_{\tilde{z}^\dagger}(\omega) \geq \int_0^{2\pi} s_{\tilde{z}^\dagger}(\omega)\, d\omega = \mathrm{Var}(\tilde{z}_t^\dagger). \tag{13}$$

The far right-hand side above depends on the degree of recoverability of the shock:

$$\mathrm{Var}(\tilde{z}_t^\dagger) = \mathrm{Var}(E(\tilde{z}_t \mid \{y_\tau\}_{-\infty<\tau<\infty})) = \alpha^2(1 - \mathrm{Var}(\varepsilon_{1t} \mid \{y_\tau\}_{-\infty<\tau<\infty})) = \alpha^2 \times R_\infty^2.$$

An even lower bound on α_{LB}^2 is given by

$$\mathrm{Var}(E(\tilde{z}_t \mid \{y_\tau\}_{-\infty<\tau\leq t})) = \alpha^2(1 - \mathrm{Var}(\varepsilon_{1t} \mid \{y_\tau\}_{-\infty<\tau\leq t})) = \alpha^2 \times R_0^2.$$

Thus, if the shock is close to being invertible – or more generally, recoverable – α_{LB}^2 will be close to α^2. In fact, as discussed above, the bound will be tight provided the full time series $\{y_t\}$ of the macro variables is informative about $\varepsilon_{1,t}$ at least at *some* frequency (e.g., low frequencies or high frequencies). These comparisons illustrate how our lower bound on α is a natural generalization of the familiar invertibility requirement. In Section 4 we provide several concrete examples where invertibility fails, but nevertheless α_{LB}^2 is close to α^2.

Degree of invertibility. The identified set for the degree of invertibility at horizon ℓ follows directly from the identified set for α^2, since (recalling $\mathrm{Var}(\varepsilon_{1,t}) = 1$)

$$R_\ell^2 = 1 - \mathrm{Var}(\varepsilon_{1,t} \mid \{y_\tau\}_{-\infty<\tau\leq t+\ell}) = \frac{1}{\alpha^2} \times \mathrm{Var}(E(\tilde{z}_t \mid \{y_\tau\}_{-\infty<\tau\leq t+\ell})),$$

and the variance on the right-hand side above is point-identified. Now similarly define

$$\tilde{R}_\ell^2 \equiv 1 - \frac{\mathrm{Var}(\tilde{z}_t \mid \{y_\tau\}_{-\infty<\tau\leq t+\ell})}{\mathrm{Var}(\tilde{z}_t)} = \frac{\mathrm{Var}(E(\tilde{z}_t \mid \{y_\tau\}_{-\infty<\tau\leq t+\ell}))}{\mathrm{Var}(\tilde{z}_t)}$$

as the (point-identified) R^2 in a population regression of $\tilde{z}_t$ on lags and leads of y_τ up to time $\tau = t + \ell$. Then the identified set for the degree of invertibility R^2_ℓ at horizon ℓ equals

$$\left[\tilde{R}^2_\ell,\ \frac{\operatorname{Var}(\tilde{z}_t)}{2\pi \sup_{\omega\in[0,\pi]} s_{\tilde{z}^\dagger}(\omega)} \times\ \tilde{R}^2_\ell\right]. \tag{14}$$

This identified set implies conditions under which the distribution of the observable data is consistent with invertibility or recoverability.

Proposition 2. *Assume $\alpha^2_{LB} > 0$. The identified set for R^2_0 contains 1 if and only if the instrument residual $\tilde{z}_t$ does not Granger cause the macro observables y_t. The identified set for R^2_∞ contains 1 if and only if the projection $\tilde{z}^\dagger_t$ is serially uncorrelated.*

According to Proposition 2, $\varepsilon_{1,t}$ is certain to be noninvertible if and only if $\tilde{z}_t$ Granger causes y_t. Moreover, $\varepsilon_{1,t}$ is certain to be non-recoverable if and only if $\tilde{z}^\dagger_t$, defined in (10), is serially correlated at some lag. Thus, the IV makes the extraneous invertibility and recoverability assumptions testable.

Variance decompositions. We now turn to the identification of variance decompositions, the main parameters of interest. The identified sets for the FVR and FVD defined in Section 2 are different. For the FVR, simply observe that

$$FVR_{i,\ell} = \frac{\sum_{m=0}^{\ell-1}\Theta^2_{i,1,m}}{\operatorname{Var}(y_{i,t+\ell}\mid\{y_\tau\}_{-\infty<\tau\le t})} = \frac{1}{\alpha^2}\times\frac{\sum_{m=0}^{\ell-1}\operatorname{Cov}(y_{i,t},\tilde{z}_{t-m})^2}{\operatorname{Var}(y_{i,t+\ell}\mid\{y_\tau\}_{-\infty<\tau\le t})}.$$

Hence, as in the static case, the identified set for $FVR_{i,\ell}$ equals the identified set for $\frac{1}{\alpha^2}$, scaled by the (point-identified) second fraction on the far right-hand side above. In particular, the upper bound on the FVR depends on the informativeness of the macro variables for the shock of interest. Hence, adding more variables to the vector y_t of endogenous observables leads to a weakly narrower identified set (in relative terms, since the estimand itself changes if y_t is changed).

Bounding the FVD requires more work. Online Appendix B.1 states the identified set for the FVD, but we here summarize the result. The sharp upper bound on the ℓ-period-ahead FVD equals the trivial bound of 1, for any $\ell \ge 1$. This is achieved by a model in which all shocks, except the first one, only affect y_t after an ℓ-period delay. The sharp lower bound on the FVD is nontrivial and informative; see Online Appendix B.1 for the precise expression. The reason that identification of the FVD is more challenging than for the FVR is that, even if we knew α, the IV z_t provides no information about the other structural shocks $\varepsilon_{j,t}$, $j \neq 1$. This matters because the definition (4) of the FVD, unlike that of the FVR, conditions on knowing all past shocks.

Finally, and analogously to the FVR, the frequency-specific unconditional variance decomposition $VD_i(\omega_1, \omega_2)$ is interval-identified with informative lower and upper bounds, since

$$VD_i(\omega_1, \omega_2) = \frac{1}{\alpha^2} \times \frac{\int_{\omega_1}^{\omega_2} |s_{y_i \tilde{z}}(\omega)|^2 \, d\omega}{\int_{\omega_1}^{\omega_2} s_{y_i}(\omega) \, d\omega},$$

where $s_{y_i \tilde{z}}(\omega) = \alpha \Theta_{i,1}(e^{-i\omega})$ is the i-th element of $s_{y\tilde{z}}(\omega)$, cf. equation (5).

Absolute impulse responses. The identified set for the *absolute* impulse response $\Theta_{i,1,\ell}$ is obtained by scaling the identified set for $\frac{1}{\alpha}$, cf. equation (8). This generalizes existing results on *relative* impulse responses, as discussed above (Stock & Watson, 2018).

Sufficient conditions for point identification. Although we have shown that partial identification analysis is informative in the general model, we now give a variety of sufficient conditions that ensure point identification of α and thus the FVR, VD, and degree of invertibility.[18] We also discuss identification of historical decompositions.

The first set of sufficient conditions relates to the informativeness of the macro aggregates y_t for the hidden shock $\varepsilon_{1,t}$. In this category, our weakest condition for point identification is that the data y_t is perfectly informative about $\varepsilon_{1,t}$ at *some* frequency, i.e., the spectral density of the projection residual $\varepsilon_{1,t} - \varepsilon_{1,t}^\dagger$ vanishes at some frequency $\overline{\omega}$. Then $\alpha = \alpha_{LB}$, so the FVR, VD, and degree of invertibility are identified. This assumption is not testable; as emphasized above, it is the weakest generalization of the familiar invertibility assumption of SVAR analysis. A stronger but more easily interpretable identifying assumption is recoverability, i.e., $\varepsilon_{1,t}^\dagger \equiv E(\varepsilon_{1,t} \mid \{y_\tau\}_{-\infty<\tau<\infty}) = \varepsilon_{1,t}$. This assumption is testable, cf. Proposition 2. Under recoverability, we have both $\alpha = \alpha_{LB}$ and $\tilde{z}_t^\dagger = \alpha \varepsilon_{1,t}$. Recoverability is a restrictive assumption, but at least it is a meaningfully weaker requirement than invertibility in many economic applications, as discussed further in Section 4. In fact, recoverability is implied by the usual SVAR assumption that there are as many shocks as variables, $n_\varepsilon = n_y$.[19] Our proof

[18] Online Appendix B.1 shows that even point identification of α is insufficient to point-identify the FVD, although a sharp and informative lower bound can be computed.

[19] Since we have ruled out singularities, $n_\varepsilon = n_y$ implies that $\Theta(L)^{-1}$ is a well-defined two-sided lag polynomial (Brockwell & Davis, 1991, Thm. 3.1.3), so that $\varepsilon_t = \Theta(L)^{-1} y_t$ and *all* shocks are recoverable.

of Proposition 1 shows how restrictive this assumption really is: α is partially identified with the same sharp bounds as above even if we know that the number of shocks n_ε can be at most $n_y + 1$. Thus, no identifying power is gained from the knowledge that the number of shocks is "small", unless that means $n_\varepsilon = n_y$.

Point identification can also be achieved by assuming that the instrument is perfect, i.e., $\sigma_v = 0$. Then $\tilde{z}_t = \alpha\varepsilon_{1,t}$ and identification proceeds in accordance with the logic behind local projections (Jordà, 2005; Gorodnichenko & Lee, 2017). This assumption is not testable.

Under either recoverability or perfect instrument informativeness, we can point-identify the historical decomposition corresponding to the identified shock, cf. the definition in Section 2. This object is identified because both the absolute impulse responses and the time series of the shock itself are identified, as argued above.

3.3 Extension: multiple instruments

We now argue that identification analysis in the model with multiple IVs for the shock of interest ($n_z \geq 2$) can be reduced to the single-IV setting without loss of generality.

The multiple-IV model is testable, unlike the single-IV model. As in the single-IV case, define the projection residual

$$\tilde{z}_t \equiv z_t - E(z_t \mid \{y_\tau, z_\tau\}_{-\infty<\tau<t}) = \alpha\lambda\varepsilon_{1,t} + \Sigma_v^{1/2} v_t. \tag{15}$$

Online Appendix B.2 shows that the testable implication of the multiple-IV model is that the cross-spectrum $s_{y\tilde{z}}(\omega)$ has a rank-1 factor structure. The validity of the multiple-IV model can be rejected if and only if this factor structure fails.

When the multiple-IV model is consistent with the distribution of the data, identification analysis can be reduced to the single-IV case in Section 3.2. Specifically, Online Appendix B.2 shows that (i) λ is point-identified, and (ii) the identified sets for α, variance decompositions, and the degree of invertibility are the same as the identified sets that exploit only the scalar instrument

$$\breve{z}_t \equiv \frac{1}{\lambda' \operatorname{Var}(\tilde{z}_t)^{-1}\lambda} \lambda' \operatorname{Var}(\tilde{z}_t)^{-1} \tilde{z}_t. \tag{16}$$

Intuitively, $\breve{z}_t \propto E(\varepsilon_{1,t} \mid \tilde{z}_t)$. Because $\breve{z}_t$ is a linear combination of all n_z instruments, the identified sets are narrower than if we had used any one instrument $z_{k,t}$ in isolation.

In Online Appendix B.3 we consider the more general case of multiple instruments being correlated with *multiple* structural shocks. In particular, we allow the instrument set to be correlated with a pre-specified number of structural shocks and then bound the forecast variance contribution of this combination of shocks to the macro aggregates of interest. The derived bounds would for example be informative in the application of Mertens & Ravn (2013), who use two external IVs plausibly correlated with two latent tax shocks.

4 Illustration using a structural macro model

We use the workhorse business cycle model of Smets & Wouters (2007) to illustrate the informativeness of our partial identification bounds on the degree of invertibility and variance decompositions. The nature of our exercise is as follows: We consider an econometrician observing a small set of macroeconomic aggregates generated from the Smets-Wouters model as well as noisy measures of some of the model's true underlying structural shocks (i.e., valid external instruments). For clarity, we abstract from any sampling uncertainty and assume that the econometrician observes an infinite amount of data, so the joint spectral density of observed macro aggregates and external IVs is perfectly known to her. Given this spectral density, she uses our LP-IV bounds to draw conclusions about variance decompositions and the degree of invertibility, without exploiting the underlying structure of the model.

We stress that the purpose of this section is merely to illustrate the workings of our identification bounds in an economically interpretable setting. Hence, we deliberately consider a small number of observable variables. Our results below are necessarily sensitive to the set of observables, as shown through robustness checks in Online Appendix B.6. However, we caution against the belief that the use of a large number of observable variables will automatically guarantee that shocks are recoverable (or invertible) in realistic applications.[20]

Model. We employ the Smets & Wouters (2007) model. Throughout, we parametrize the model according to the posterior mode estimates of Smets & Wouters (2007).[21] Following the canonical trivariate VAR in the empirical literature on monetary policy shock transmission, we assume the econometrician observes aggregate output, inflation, and the short-term policy interest rate. These macro aggregates are all stationary in the model, so they should be

[20] Although most DSGE models in the literature feature a small number of shocks for simplicity, in reality the addition of new observables will likely contaminate the analysis with additional nuisance shocks (including, but not limited to, measurement error). While the recoverability assumption may in some settings be justified from a theoretical standpoint, it should not be taken for granted.

[21] Our implementation of the Smets-Wouters model is based on Dynare replication code kindly provided by Johannes Pfeifer. The code is available at https://sites.google.com/site/pfeiferecon/dynare.

viewed as deviations from trend. The model features seven unobserved shocks, so not all shocks can be invertible.

The econometrician observes a single external instrument z_t for the shock of interest $\varepsilon_{1,t}$:

$$z_t = \alpha\varepsilon_{1,t} + \sigma_v v_t.$$

We normalize $\alpha = 1$ throughout and compute identified sets for two different degrees of informativeness of the external instrument, $\frac{1}{1+\sigma_v^2} \in \{0.25, 0.5\}$. We do not attach any specific economic interpretation to the IV in the context of the Smets & Wouters (2007) model.

We separately consider three different shocks of interest: a monetary shock, a forward guidance shock, and a technology shock. These shocks have been chosen to illustrate how the informational requirements of our procedure are substantively less demanding than those of familiar SVAR-IV analysis. The conventional monetary policy shock as well as the technology shock are already included in the original Smets & Wouters (2007) model. For the forward guidance shock, we depart from the original model by assuming that monetary policy shocks are known two quarters in advance.[22] With our set of observables, the monetary shock is nearly invertible, but the others are not. The forward guidance shock is instead nearly recoverable, whereas only the *long-run* cycles of the technology shock can be accurately recovered from the data. Nevertheless, we show that our partial identification analysis is informative about the effects of all three shocks.

Monetary shock. We first consider identification of monetary policy shocks. These are defined as shocks to the serially correlated disturbance in the model's Taylor rule.

The monetary shock is nearly invertible in our parametrization. Specifically, the collection of all past and current values of the observable macro variables explain a fraction $R_0^2 = 0.8705$ of the variance of the shock, as shown by Wolf (2018).[23] The infinite past, present, and *future* of the observables yield only slightly sharper identification, with $R_\infty^2 = 0.8767$. Figure 1 shows the spectral density $s_{\varepsilon_1^\dagger}(\cdot)$ of the two-sided best linear predictor of the monetary shock based on all macro variables. The data are essentially equally informative about medium and high frequencies of the monetary shock, whereas the long-run cycles of the shock cannot

[22]Formally, we implement forward guidance by changing the baseline Smets & Wouters (2007) model so that the monetary shock has time subscript $t-2$ instead of t. This is the notion of forward guidance discussed, for example, in Del Negro et al. (2012).

[23]Wolf (2018) argues that the R_0^2 of monetary policy shocks is robustly high because such shocks uniquely move nominal interest rates and inflation in opposite directions.

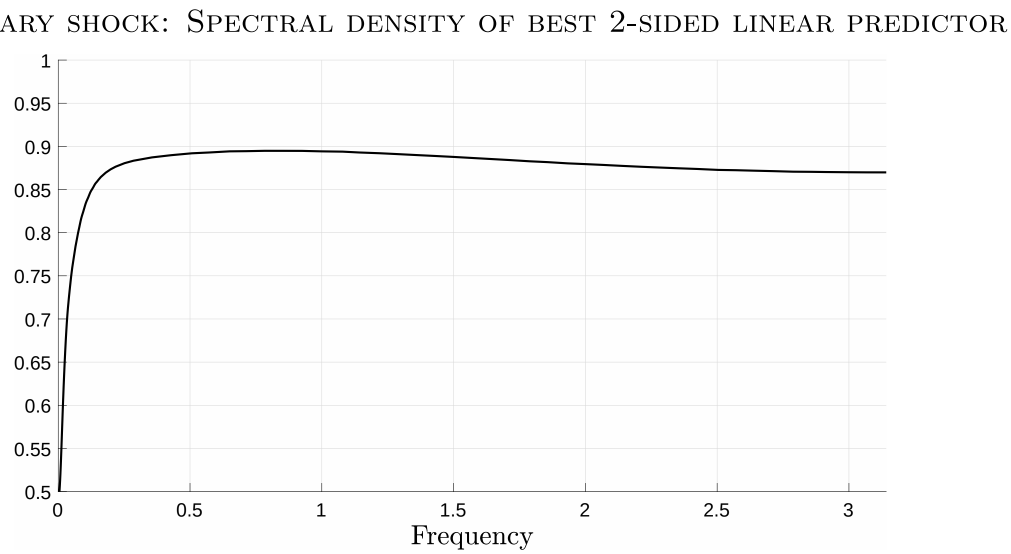

Figure 1: Scaled spectral density $2\pi s_{\varepsilon_1^\dagger}(\cdot)$ of the best two-sided linear predictor of the monetary shock. A frequency ω corresponds to a cycle of length $\frac{2\pi}{\omega}$ quarters.

be accurately recovered from the data. At the peak of the spectral density, the observables explain a fraction 0.8958 of the variance of that particular cyclical component of the monetary shock; hence, $\alpha_{LB} = \sqrt{0.8958} = 0.9465$, which is close to the truth of 1.

Because the shock is nearly invertible, the upper bounds of the identified sets for the forecast variance ratios are close to the truth, while the lower bounds depend on the informativeness of the IV. Figure 2 displays the identified set of the FVR at different forecast horizons.[24][25] The upper and lower bounds are proportional to the true FVRs. The lower bound scales one-for-one with instrument informativeness, while the upper bound scales one-for-one with the maximal informativeness of the data for the shock across frequencies. The upper bounds are thus close to the true FVRs in this application with a near-invertible shock. The informativeness of the lower bounds depends entirely on the strength of the IV.

Due to the near-invertibility of the shock, SVAR-IV identification of the monetary shock would only be slightly biased (Forni et al., 2018; Wolf, 2018). This, however, is not the case for the next two shocks we consider.

[24] We omit identified sets for FVDs as the upper bound is trivial, cf. Online Appendix B.1.

[25] Throughout this paper, the identified sets for FVRs are constructed horizon by horizon. However, the *joint* uncertainty about FVRs at different horizons is caused by uncertainty about the single parameter α.

Monetary shock: Identified set of FVRs

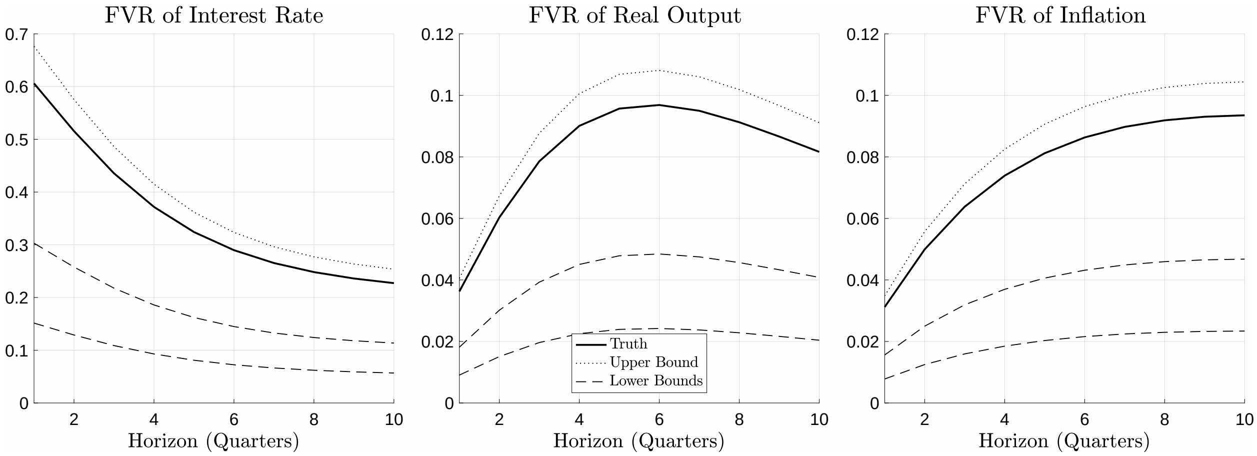

Figure 2: Horizon-by-horizon identified sets for FVRs up to 10 quarters. The two lower bounds are for $\frac{1}{1+\sigma_v^2} = 0.25$ (lower dashed line) and $\frac{1}{1+\sigma_v^2} = 0.5$.

Forward guidance shock. We now modify the model to include forward guidance shocks, a type of news shock. As discussed above, a forward guidance shock is identical to a monetary shock, except it is anticipated two quarters in advance by economic agents.

As is common with news shocks, the forward guidance shock is highly noninvertible but approximately recoverable. The wedge between information contained in the infinite past and information contained in the entire time series of observables is sizable: Contemporaneous informativeness is limited, with $R_0^2 = 0.0792$, but looking two quarters ahead basically returns us to the level of informativeness for the standard monetary shock, with $R_2^2 = 0.8731$ and $R_\infty^2 = 0.8813$. Intuitively, on impact, all macro aggregates move in the same direction, suggesting to the econometrician that the economy was probably buffeted by a demand shock. But two quarters from now, when the anticipated innovation finally hits, the interest rate response suddenly switches sign, sending a strong signal that in fact a monetary policy shock – and not some other kind of demand shock – had occurred. This is one example of why, with news shocks, the incremental bite of two-sided analysis can be substantial.

Despite the high degree of noninvertibility, the identified sets for the FVRs of the forward guidance shock are as informative as those for the monetary shock, as shown in Figure 3. This demonstrates that our partial identification analysis is not only robust to noninvertibility – its quantitative usefulness does not depend on the degree of invertibility *per se*. In stark contrast, identification that incorrectly imposes invertibility (e.g., SVARs) would overstate variance

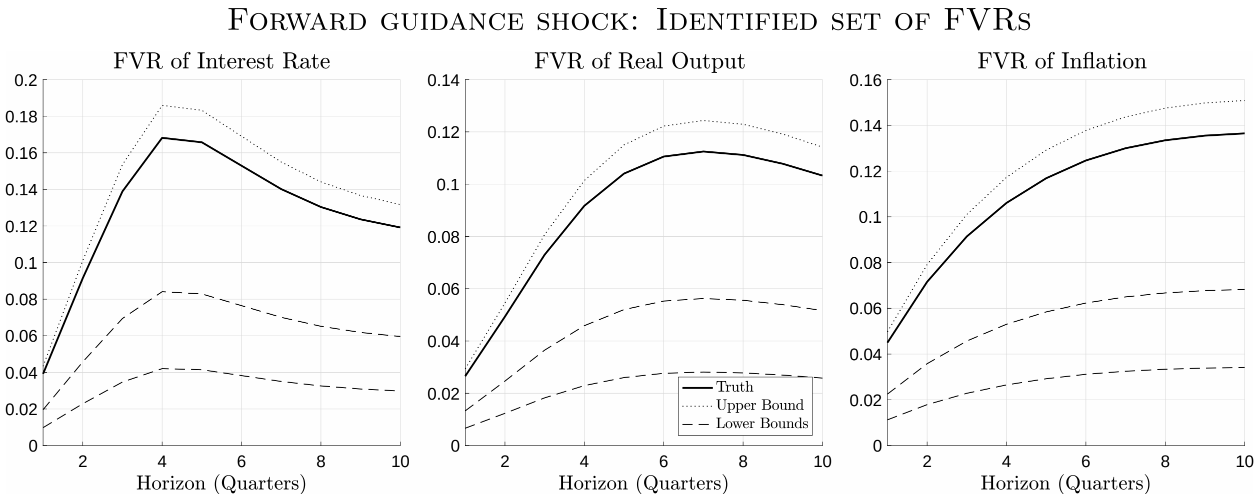

Figure 3: Horizon-by-horizon identified sets for FVRs up to 10 quarters. The two lower bounds are for $\frac{1}{1+\sigma_v^2} = 0.25$ (lower dashed line) and $\frac{1}{1+\sigma_v^2} = 0.5$.

decompositions by a factor of $1/0.0792 \approx 13$ (!).[26] Recoverability-based identification would err by a more modest factor of $1/0.8813 \approx 1.13$.

Technology shock. Finally, we consider identification of technology shocks, defined as an innovation to the autoregressive process of total factor productivity.

Unlike the monetary and forward guidance shocks, the technology shock is far from recoverable, using our baseline set of observables; nevertheless, our bounds remain informative. In the model, the technology shock is much more important in accounting for low-frequency cycles of the data than it is for high-frequency cycles. The degrees of invertibility and recoverability are low: $R_0^2 = 0.2007$ and $R_\infty^2 = 0.2209$. However, the data are very informative about the lowest-frequency cycles of the technology shock, as shown in Figure 4. As a result, $\alpha_{LB}^2 = 0.9092$ is close to the true value of 1, and the upper bounds of our identified sets for FVRs and the degree of invertibility (not shown) yield tight identification. In contrast, identification that incorrectly imposes either invertibility or recoverability of the shock overstates the FVR by a factor of about 5.

[26] To be exact, standard SVAR-IV methods would overstate *impact* impulse responses by a factor of $1/\sqrt{0.0792} \approx 3.6$ and so impact variance decompositions by a factor of 13. Subsequent impulse responses would not be proportional to true responses, due to the imposed VAR dynamics (Stock & Watson, 2018). We state formal results on the bias of SVAR-IV under noninvertibility in Online Appendix B.4.

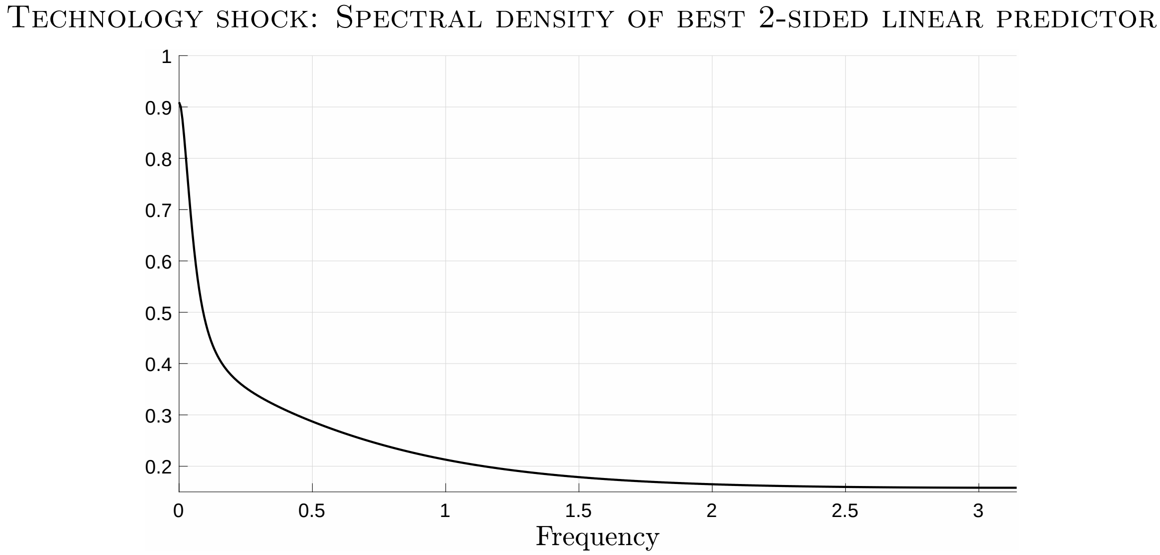

Figure 4: Scaled spectral density $2\pi s_{\varepsilon_1^\dagger}(\cdot)$ of the best two-sided linear predictor of the technology shock. A frequency ω corresponds to a cycle of length $\frac{2\pi}{\omega}$ quarters.

5 Inference

To make the identification analysis practically useful, we develop partial identification robust confidence intervals and tests. In a first step, the researcher estimates a reduced-form VAR model, which is then used in a second step to derive sample analogues of our population bounds. Using the general partial identification confidence procedures of Imbens & Manski (2004) and Stoye (2009), we construct confidence intervals for both the parameters and for the identified sets. We also discuss a test of invertibility. The confidence intervals are shown to be asymptotically valid under nonparametric regularity conditions. Finally, we show through simulations that the LP-IV confidence intervals perform well in finite samples.

We assume the availability of a single instrument z_t for notational simplicity. The generalization to multiple instruments is straight-forward, as discussed in Section 3.3.

Reduced-form VAR. We assume that the second-moment properties of the data are captured by a reduced-form VAR in $(y_t', z_t)'$. The lag length p is initially assumed to be finite and known, but this is relaxed below. Thus, assume that there exist $(n_y+1)\times(n_y+1)$ matrices A_ℓ, $\ell = 1, 2, \dots, p$, and a symmetric positive definite $(n_y+1)\times(n_y+1)$ matrix Σ

such that the spectral density of $W_t \equiv (y_t', z_t)'$ is given by

$$s_W(\omega) = \left(I_{n_y+1} - \sum_{\ell=1}^{p} A_\ell e^{-i\omega\ell}\right)^{-1} \Sigma \left(I_{n_y+1} - \sum_{\ell=1}^{p} A_\ell e^{-i\omega\ell}\right)^{-1*}, \quad \omega \in [0, 2\pi],$$

and such that all roots of the polynomial $x \mapsto \det(I_{n_y+1} - \sum_{\ell=1}^{p} A_\ell x^\ell)$ are outside the unit circle. Let $\vartheta \equiv (\text{vec}(A_1)', \dots, \text{vec}(A_p)', \text{vech}(\Sigma)')'$ denote the collection of true reduced-form VAR parameters, and let $\hat{\vartheta} \equiv (\text{vec}(\hat{A}_1)', \dots, \text{vec}(\hat{A}_p)', \text{vech}(\hat{\Sigma})')'$ denote the least-squares estimators. Under standard conditions, we have $T^{1/2}(\hat{\vartheta} - \vartheta) \xrightarrow{p} N(0, \Omega)$, where T is the sample size, and the asymptotic variance Ω can be estimated consistently by $\hat{\Omega}$, say (Kilian & Lütkepohl, 2017, Ch. 2.3).

Although not essential to our approach, we assume a reduced-form VAR structure for three reasons. First, VARs are known to be able to approximate any spectral density function arbitrarily well as the VAR lag length tends to infinity. Second, the VAR structure facilitates the development of a test of invertibility. Third, VAR-based inference amounts to applying our population calculations from Section 3 to a spectrum of a particular functional form (namely a VAR spectrum with the particular estimated parameters $\hat{\vartheta}$). All inequalities satisfied in the population must then also hold in any finite sample, thus guaranteeing nonempty identified sets, for example (up to numerical error, but not statistical error). The advantages of the VAR approach notwithstanding, we remark that one could in principle use any well-behaved estimator of the spectrum of $(y_t', z_t)'$.

While we here assume a finite VAR lag length for expositional simplicity, Online Appendix B.9 proves that our VAR-based inference strategy is asymptotically valid under *nonparametric* regularity conditions, provided that the VAR lag length $p = p_T$ used for estimation diverges with the sample size T at an appropriate rate. That is, the inference strategy is valid even if the true data generating process (DGP) is a possibly non-Gaussian VAR(∞). In practice, we suggest estimating the lag length p by information criteria or likelihood ratio tests. We emphasize that assuming a reduced-form VAR is less restrictive than doing SVAR-IV inference: We do not assume that the reduced-form VAR residuals span the true structural shocks ε_t. For example, we continue to allow the number of structural shocks to possibly exceed the number of variables in the VAR.

Invertibility test. It is straight-forward to test for invertibility of the shock of interest using the estimated reduced-form VAR. We showed in Proposition 2 that the distribution of the data is consistent with invertibility of $\varepsilon_{1,t}$ if and only if $\tilde{z}_t$ does not Granger cause y_t.

Granger non-casuality of $\tilde{z}_t$ for y_t is equivalent with Granger non-causality of z_t for y_t. A test of the Granger non-causality null hypothesis amounts to a test of the exclusion restrictions that lags of z_t do not enter the reduced-form VAR equations for y_t. This test has power against all Granger causal alternatives, so it has power against all falsifiable noninvertible alternatives by Proposition 2.[27]

Confidence intervals. We now construct partial identification robust confidence intervals for identified sets and for the true parameters. Here we simply apply the general inference methods developed by Imbens & Manski (2004) and refined by Stoye (2009).

We start by defining notation. Given the reduced-form VAR model, all identified sets derived in Section 3.2 are of the form $[\underline{h}(\vartheta), \overline{h}(\vartheta)]$, where $\underline{h}(\cdot)$ and $\overline{h}(\cdot)$ are continuous functions mapping the VAR parameter space into the real line, and such that $\underline{h}(\cdot) \leq \overline{h}(\cdot)$. A (pointwise) consistent estimator of the identified set $[\underline{h}(\vartheta), \overline{h}(\vartheta)]$ is then given by the plug-in interval

$$[\underline{h}(\hat{\vartheta})\, ,\ \overline{h}(\hat{\vartheta})].$$

Let $\hat{\Delta} \equiv \overline{h}(\hat{\vartheta}) - \underline{h}(\hat{\vartheta})$ denote the width of the estimate of the identified set. Assume $\underline{h}(\cdot)$ and $\overline{h}(\cdot)$ are continuously differentiable at the true VAR parameters ϑ with $1 \times \dim(\vartheta)$ dimensional Jacobian functions $\dot{\underline{h}}(\cdot)$ and $\dot{\overline{h}}(\cdot)$. Define the standard errors of $\underline{h}(\hat{\vartheta})$ and $\overline{h}(\hat{\vartheta})$,

$$\hat{\underline{\sigma}} \equiv \sqrt{T^{-1}\dot{\underline{h}}(\hat{\vartheta})\hat{\Omega}\dot{\underline{h}}(\hat{\vartheta})'}\,, \quad \hat{\overline{\sigma}} \equiv \sqrt{T^{-1}\dot{\overline{h}}(\hat{\vartheta})\hat{\Omega}\dot{\overline{h}}(\hat{\vartheta})'},$$

and their correlation,

$$\hat{\rho} \equiv \frac{T^{-1}\dot{\underline{h}}(\hat{\vartheta})\hat{\Omega}\dot{\overline{h}}(\hat{\vartheta})'}{\hat{\underline{\sigma}} \times \hat{\overline{\sigma}}}.$$

Finally, let $\Phi(\cdot)$ denote the standard normal cumulative distribution function.

The interval

$$\left[\underline{h}(\hat{\vartheta}) - \Phi^{-1}(1-\beta/2)\hat{\underline{\sigma}}\,,\ \overline{h}(\hat{\vartheta}) + \Phi^{-1}(1-\beta/2)\hat{\overline{\sigma}}\right] \tag{17}$$

is a (pointwise) asymptotically valid level-$(1-\beta)$ confidence interval for the identified set $[\underline{h}(\vartheta), \overline{h}(\vartheta)]$. That is, the above interval contains the *entire* identified set in at least $100(1-\beta)\%$ of repeated experiments, asymptotically. This follows immediately from the

[27] Stock & Watson (2018) develop an LP-IV invertibility test which directs power against alternatives with impulse response functions that differ substantially from the invertible null. They do not discuss whether their test has power against all falsifiable noninvertible alternatives.

delta method and the Bonferroni argument of Imbens & Manski (2004).

It is also possible to construct a confidence interval for the true parameter of interest. By definition of the identified set, the true parameter is contained in $[\underline{h}(\vartheta), \overline{h}(\vartheta)]$, but we know nothing else about the true parameter. Although the interval (17) trivially has asymptotic coverage of at least $1-\beta$ for the true parameter, Imbens & Manski (2004) showed that it is possible to develop a narrower interval with the same property. Online Appendix B.8 provides details on the simple computation, using Stoye's (2009) refinement of the Imbens & Manski procedure. In practical macroeconomic applications, however, we have found that the Stoye (2009) confidence interval for the parameter itself is often only a few percent narrower than the confidence interval (17) for the entire identified set. This is because the estimated width $\hat{\Delta}$ of the identified set is usually large relative to the sampling uncertainty.

To implement the above confidence interval procedures, the researcher needs to compute the VAR estimator $\hat{\vartheta}$, the asymptotic variance matrix estimate $\hat{\Omega}$, the bound estimates $\underline{h}(\hat{\vartheta})$ and $\overline{h}(\hat{\vartheta})$, and the derivatives of the bounds $\dot{\underline{h}}(\hat{\vartheta})$ and $\dot{\overline{h}}(\hat{\vartheta})$. Appendix A.1 provides formulas for the bounds and derivatives in terms of the VAR parameters. A simple bootstrap implementation is also available, as we discuss below.

In practice, we recommend that inference be based on an alternative, conservative lower bound on α^2. Recall that α_{LB}^2 is given by the maximum of a certain function, cf. (12). When this function has multiple maxima, continuous differentiability of α_{LB}^2 in the VAR parameters ϑ may fail, rendering the delta method inapplicable (Gafarov et al., 2018). As a simple remedy, we suggest replacing the maximum $\alpha_{LB}^2 = 2\pi \sup_{\omega \in [0,\pi]} s_{\tilde{z}^\dagger}(\omega)$ in all our bounds with the weakly smaller *average* $\tilde{\alpha}_{LB}^2 \equiv \int_0^{2\pi} s_{\tilde{z}^\dagger}(\omega)\, d\omega = \text{Var}(\tilde{z}_t^\dagger)$. The latter object is continuously differentiable in the VAR parameters, so inference using the above methods is unproblematic. If the true shock of interest is near-recoverable, then $\tilde{z}_t^\dagger$ is nearly white noise (cf. Proposition 2), so the average $\tilde{\alpha}_{LB}^2$ will be close to the supremum α_{LB}^2, and little power will be lost by using the conservative bound.[28] We recommend the use of the conservative bound $\tilde{\alpha}_{LB}^2$ because it yields valid confidence intervals (regardless of recoverability) that are easy to compute using the Imbens-Manski-Stoye procedures, and which are likely to perform well when the researcher has taken care to employ a set of macro variables y_t that are very

[28] An alternative, but more involved approach would be to apply the intersection bound confidence intervals of Chernozhukov et al. (2013) or Andrews & Shi (2013, 2017). This would lead to better power properties against a wider range of non-recoverable alternatives, although likely at the cost of lower power in the near-recoverable case. If the researcher has prior information about the frequencies at which y_t is particularly informative about $\varepsilon_{1,t}$, then a simpler approach would be to lower-bound α_{LB}^2 by $2\pi \exp(\frac{1}{2\pi}\int_0^{2\pi} r(\omega) \log s_{\tilde{z}^\dagger}(\omega)\, d\omega)$, where $r(\cdot)$ is a nonnegative weight function such that $\int_0^{2\pi} r(\omega)\, d\omega = 2\pi$.

informative about the shock $\varepsilon_{1,t}$ (in the sense of near-recoverability).

The resulting confidence intervals are pointwise valid in both senses of the word. First, we focus on constructing a confidence interval for each parameter of interest separately, as opposed to capturing the joint uncertainty of several parameters at once. Second, our asymptotics are pointwise in the true parameters; we do not derive the coverage under the worst-case data generating process.[29] In particular, we ignore finite-sample issues caused by weak instruments, i.e., $\alpha_{LB} \approx 0$. We also ignore the familiar parameter-on-the-boundary issue that may arise if one of the population bounds is at the boundary of its parameter space (this issue can also arise in standard SVAR inference on variance decompositions).

Bootstrap implementation. The calculation of derivatives in the confidence interval formulas above is obviated by the bootstrap. Suppose we bootstrap the VAR estimator $\hat{\vartheta}$ (Kilian & Lütkepohl, 2017, Ch. 12). Then we can compute $\underline{\hat{\sigma}}$ as the bootstrap standard deviation of $\underline{h}(\hat{\vartheta})$, $\hat{\overline{\sigma}}$ as the bootstrap standard deviation of $\overline{h}(\hat{\vartheta})$, and $\hat{\rho}$ as the bootstrap correlation of $\underline{h}(\hat{\vartheta})$ and $\overline{h}(\hat{\vartheta})$. By plugging into the same confidence interval formulas as above, we achieve the same (pointwise) asymptotic coverage probability as the delta method confidence intervals, provided an appropriate bootstrap consistency condition holds.

Simulation study. Online Appendix B.10 presents a simulation study of the LP-IV bootstrap confidence intervals for parameters and identified sets. We consider a variety of structural VARMA DGPs, including a non-invertible one. We find that the finite-sample coverage rates of our confidence intervals are close to the nominal level throughout, except when parameter-on-the-boundary issues cause over-coverage. In particular, the LP-IV confidence intervals have at least as accurate coverage as corresponding SVAR-IV confidence intervals for invertible DGPs, and of course perform much better in the non-invertible case.

6 Empirical application

To illustrate our inference procedure, we study the importance of monetary policy shocks in U.S. data. We use the empirical setting of Gertler & Karadi (2015), whose external instrument for the monetary shock is obtained from high-frequency financial data. Using

[29] We do not discuss uniform asymptotics here because this seems to require bounding the magnitude of the largest eigenvalue of the VAR polynomial away from 1 to ensure stationarity, in which case the width of the identified set (for all our objects of interest) would also be bounded away from zero. This effectively assumes away the uniformity issue highlighted by Imbens & Manski (2004).

an SVAR-IV approach, Gertler & Karadi (2015) estimated impulse responses to a 20 basis point shock in the short-term interest rate, but they did not consider the *importance* of the monetary shock using variance decompositions. Caldara & Herbst (2019) compute FVDs for a similar specification, but their analysis also assumes an SVAR model. Ramey (2016), citing likely non-invertibility due to the increasing prevalence of forward guidance in the conduct of U.S. monetary policy, cautions against such standard SVAR-IV analysis. Consistent with her concerns, we reject invertibility (i.e., SVAR structure) in our specification. Despite our weak identifying assumptions, we in fact further tighten the upper bounds on the contribution of monetary shocks to inflation dynamics obtained by Caldara & Herbst (2019).

Model. Our specification largely follows Gertler & Karadi (2015), except of course in that we do not impose a SVAR structure. We consider four endogenous macro variables y_t: output growth (log growth rate of industrial production), inflation (log growth rate of CPI inflation), the Federal Funds Rate, and the Excess Bond Premium of Gilchrist & Zakrajšek (2012) as a measure of the non-default-related corporate bond spread. 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.[30] Data are monthly from January 1990 to June 2012. The Akaike Information Criterion selects $p = 6$ lags in the reduced-form VAR we use for inference. To construct confidence sets we employ a homoskedastic recursive residual VAR bootstrap with 10,000 draws.[31]

Results. The data reject invertibility. Table 1 shows point estimates and 90% confidence intervals for the identified sets of the degree of invertibility and the degree of recoverability. We also show partial identification robust 90% confidence intervals for these parameters themselves, cf. Section 5. Since the confidence sets for the degree of invertibility exclude 1, we can reject invertibility at the 10% level, though the data are consistent with moderately

[30]See Gertler & Karadi (2015) for details on the construction of the IV and a discussion of the exclusion restriction. Nakamura & Steinsson (2018) argue that the monetary shock identified using this IV partially captures revelation of the Federal Reserve's superior information about economic fundamentals. Online Appendix B.3 shows that our FVR bounds can generally be interpreted as bounding the importance of the particular linear combination of shocks that tend to hit during FOMC announcements.

[31]All empirical results reported in this section, including unreported confidence intervals for the FVR parameters themselves, can be produced in about 6 minutes per 1,000 bootstrap draws, using Matlab R2017b without parallelization on a personal laptop with 1.60 GHz processor and 8 GB RAM.

large values of the degree of invertibility R_0^2.[32] We remark that the rejection of invertibility is sensitive to the choice of VAR lag length and to the choice of data series; still, the estimated *lower* bound on the degree of invertibility is always close to 0. The data are consistent with a wide range of values for the degree of recoverability R_∞^2.[33]

Figure 5 shows partial identification robust confidence intervals for the forecast variance ratio of the four endogenous macro variables with respect to the monetary shock. We report point estimates and confidence intervals for the identified sets (at each horizon separately); these intervals are similar to the confidence intervals for the parameters themselves in this application.[34] At all forecast horizons, the 90% confidence intervals rule out FVRs above 31% for output growth and 8% for inflation.[35] At forecast horizons up to 6 months, we can rule out that the monetary shock accounts for more than 18% of the forecast variance of the Excess Bond Premium. However, we cannot rule out that the monetary shock is an important contributor to medium- or long-run forecasts of the bond premium. On the other hand, we cannot rule out that the monetary shock is completely unimportant either.[36]

Hence, we have shown that the weak assumptions of the LP-IV model suffice to obtain tight upper bounds on the forecast variance contribution of monetary shocks for several variables, especially inflation. This is despite the finding by Stock & Watson (2018) that standard errors for LP-IV-estimated impulse response functions are large in this application. Many commentators have documented a recent divorce between inflation and output dynamics (Hall, 2011); our results document a similar divorce in dynamics *conditional* on monetary policy shocks. Although this finding is not novel in itself (Christiano et al., 1999; Ramey, 2016), our identifying assumptions are much weaker than the existing literature. Intuitively, we are able to tightly bound the importance of monetary shocks for inflation because the IV z_t correlates more with measures of real activity than measures of prices. We conclude that,

[32]As discussed in Section 5, testing invertibility amounts to testing whether the IV Granger causes the other variables in the reduced-form VAR. Online Appendix B.7 provides p-values for such tests. Note that Stock & Watson (2018) fail to reject invertibility in a somewhat different specification.

[33]As discussed in Section 5, we base inference on a slightly conservative lower bound for α, so our confidence intervals for R_∞^2 include 1 by construction.

[34]Because we compute Hall's asymmetric percentile bootstrap confidence interval, which has a built-in bias correction, some of the intervals and bias-adjusted point estimates go negative. This could be avoided by using Efron's percentile interval, but bias correction is desirable in VAR contexts (Kilian & Lütkepohl, 2017, Ch. 12). Our qualitative conclusions are not sensitive to the bias correction.

[35]If we run our analysis on the pre-crisis 1990–2006 sample only, our upper bounds rule out FVRs above 13% for inflation, at all horizons.

[36]For completeness, Online Appendix B.7 provides forecast variance decompositions implied by an estimated SVAR-IV model, although the latter is rejected by the data.

Empirical application: Degree of invertibility/recoverability

R_0^2	Estimate of IS	$[0.197, 0.687]$
	Conf. int. for IS	$[0.090, 0.877]$
	Conf. int. for param.	$[0.117, 0.828]$
R_∞^2	Estimate of IS	$[0.283, 1.000]$
	Conf. int. for IS	$[0.186, 1.000]$
	Conf. int. for param.	$[0.207, 1.000]$

Table 1: 90% confidence intervals for the degree of invertibility R_0^2 and the degree of recoverability R_∞^2, along with point estimates and 90% confidence intervals for the identified sets of these parameters. IS = identified set. All numbers are bootstrap bias corrected. Upper bound of IS for R_∞^2 equals 1 by construction.

Empirical application: Forecast variance ratios

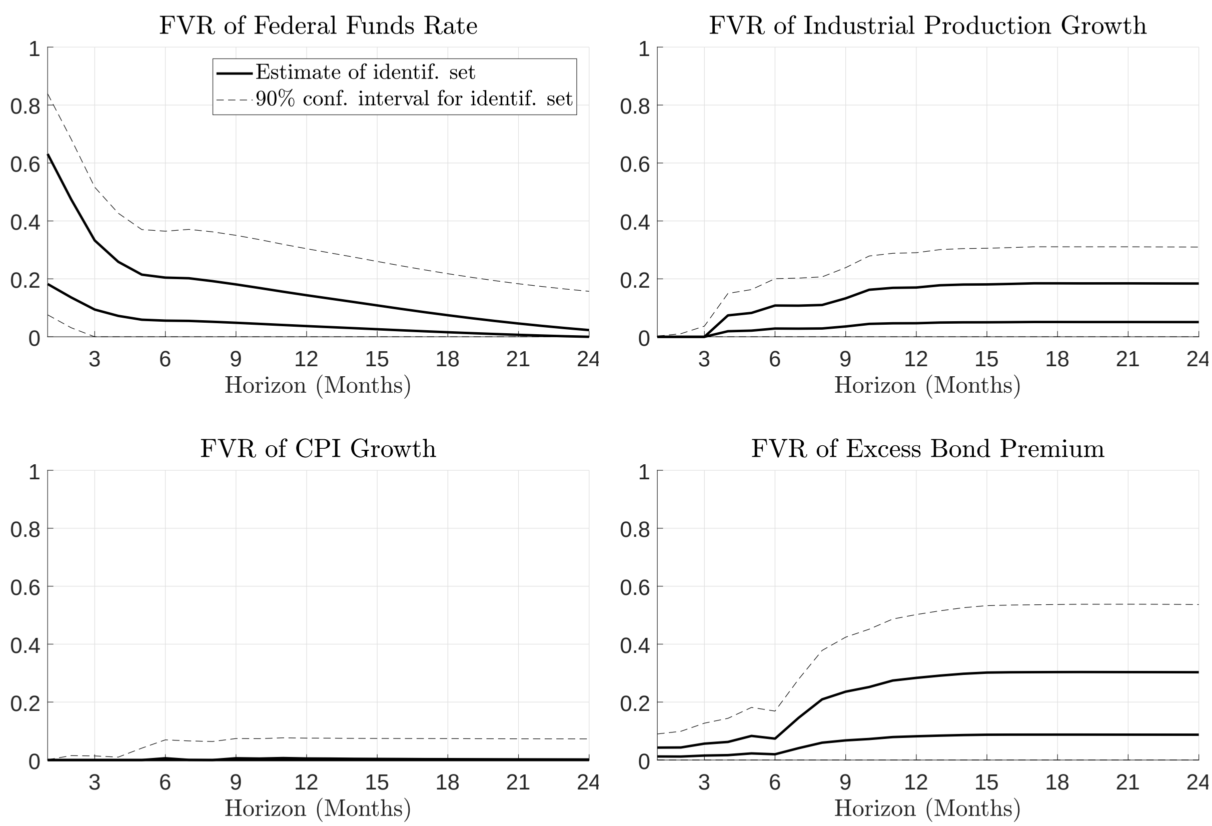

Figure 5: Point estimates and 90% confidence intervals for the identified sets of forecast variance ratios, across different variables and forecast horizons. For visual clarity, we force bias-corrected estimates/bounds to lie in $[0, 1]$.

to the extent that inflation is a monetary phenomenon, it is so because of the systematic component of monetary policy, not because of erratic policy shocks.

7 Conclusion

We expand the toolkit of the LP-IV approach to causal inference in macroeconometrics. LP-IV has recently become a popular method for estimating impulse response functions by exploiting interpretable exclusion restrictions, without imposing invertibility or functional form assumptions on shock transmission. However, existing methods did not allow researchers to quantify the importance of individual shocks. We fill this gap by providing identification results and inference techniques for variance decompositions, historical decompositions, and the degree of invertibility. Our partial identification robust confidence interval procedure is computationally straight-forward and relies on familiar methods for delta method or bootstrap inference in reduced-form VARs. The informativeness of our partial identification bounds does not depend on the degree of invertibility of the shocks *per se*, but rather on the strength of the instrument and the informativeness of the macro variables for *some* short-, medium-, or long-run cycles of the shock of interest. In contrast, the validity of SVAR-IV analysis relies on the testable assumption that the shock of interest is nearly invertible (Forni et al., 2018). Finally, we show that if researchers are willing to assume that the shock of interest is recoverable – a substantively weaker assumption than invertibility – most objects of interest are point-identified.

Our work points to several potential future research directions. First, one could construct simultaneous (rather than pointwise) confidence bands for, say, forecast variance decompositions at multiple horizons. Second, methods from the moment inequality literature could be applied to develop confidence intervals with better power properties in applications where the shock of interest is likely to be highly non-recoverable (Andrews & Shi, 2013, 2017; Chernozhukov et al., 2013). Third, future research should explore inference issues caused by parameters on the boundary, weak instruments, or near-unit roots. Fourth, our analysis imposed stationarity, but cointegration properties could be relevant for forecast variance decompositions of data in levels. Finally, one could perform Bayesian inference on the identified sets of the structural parameters.

A Appendix

A.1 Formulas for implementing the confidence intervals

Here we provide formulas needed to construct the partial identification robust confidence intervals in Section 5. Assume the spectrum of $W_t = (y_t', z_t)'$ has VAR structure as in Section 5. We now show how to compute the interval bounds from the reduced-form VAR parameters $\vartheta = (\text{vec}(A_1)', \dots, \text{vec}(A_p)', \text{vech}(\Sigma))'$.

Preparations. We first map the reduced-form VAR parameters ϑ into the various variances and covariances needed to compute our objects of interest. In principle, it is possible to directly compute these mappings from ϑ using standard VAR formulas. However, we prefer working with the vector moving average representation, as outlined below. Our first objective is then to map the VAR representation into a VMA representation

$$W_t = B(L)e_t,$$

where

$$B(L) = \textstyle\sum_{\ell=0}^{\infty} B_\ell L^\ell, \quad e_t \overset{i.i.d.}{\sim} N(0, I_{n_W}),$$

and $n_W \equiv n_y + 1$. This is achieved by setting

$$B_0 = \Sigma^{\frac{1}{2}}, \quad B_h = \textstyle\sum_{\ell=1}^{h} A_\ell B_{h-\ell}, \;\; h \geq 1,$$

where $A_\ell = 0_{n_W \times n_W}$ for $\ell > p$. In practice, we truncate this recursion at some large $\hat{p}$, and set $B_h = 0_{n_W \times n_W}$ for $h > \hat{p}$. For each h, denote the top $n_y \times n_W$ block of the $n_W \times n_W$ matrix B_h by $B_{y,h}$, and write $B_y(L) = \sum_{\ell=0}^{\infty} B_{y,\ell} L^\ell$ for the entire lag polynomial. Then

$$y_t = B_y(L)e_t.$$

Let $B_{\tilde{z}}$ be the bottom row of B_0, so that

$$\tilde{z}_t = B_{\tilde{z}} e_t.$$

Bounds for α. To compute the bounds for α, we need the quantities

$$\alpha_{UB}^2 = \text{Var}(\tilde{z}_t), \tag{18}$$

$$\alpha_{LB}^2 = 2\pi \max_{\omega \in [0,\pi]} s_{y\tilde{z}}(\omega)^* s_y(\omega)^{-1} s_{y\tilde{z}}(\omega), \tag{19}$$

$$\text{Var}(\tilde{z}_t^\dagger) = 2 \int_0^\pi s_{y\tilde{z}}(\omega)^* s_y(\omega)^{-1} s_{y\tilde{z}}(\omega)\, d\omega. \tag{20}$$

For the upper bound (18), we have

$$\text{Var}(\tilde{z}_t) = B_{\tilde{z}} B_{\tilde{z}}' \equiv \Sigma_{\tilde{z}}.$$

For the lower bound (19),

$$\begin{aligned} s_y(\omega) &= B_y(e^{-i\omega}) B_y(e^{-i\omega})^*, \\ s_{y\tilde{z}}(\omega) &= \textstyle\sum_{\ell=0}^{\hat{p}} \Sigma_{y,\tilde{z},\ell} e^{-i\omega\ell}, \end{aligned}$$

where

$$\Sigma_{y,\tilde{z},\ell} \equiv \text{Cov}(y_t, \tilde{z}_{t-\ell}) = B_{y,\ell} B_{\tilde{z}}'.$$

In practice, we compute the maximum in (19) by grid search.

Rather than explicitly computing the integral (20), we note that we can approximate $\text{Var}(\tilde{z}_t^\dagger) = \text{Var}(E(\tilde{z}_t \mid \{y_\tau\}_{t-\infty<\tau<t+\infty}))$ arbitrarily well as $M \to \infty$ by

$$\text{Var}(E(\tilde{z}_t \mid \{y_\tau\}_{t-M\leq\tau\leq t+M})) = \Sigma_{\tilde{z},y,(M,M)} \Sigma_{y,(M,M)}^{-1} \Sigma_{\tilde{z},y,(M,M)}',$$

where $\Sigma_{\tilde{z},y,(M,M)}$ is the covariance vector of $\tilde{z}_t$ and $(y_{t+M}', \ldots, y_t', \ldots, y_{t-M}')'$, and $\Sigma_{y,(M,M)}$ is the full variance-covariance matrix of $(y_{t+M}', \ldots, y_t', \ldots, y_{t-M}')'$. For any given M, we can construct these matrices from

$$\text{Cov}(\tilde{z}_t, y_{t+h}) = \begin{cases} B_{\tilde{z}} B_{y,h}' & \text{if } h \geq 0, \\ 0 & \text{otherwise,} \end{cases}$$

and

$$\text{Cov}(y_t, y_{t-h}) = \textstyle\sum_{\ell=0}^{\hat{p}} B_{y,\ell} B_{y,\ell+h}'.$$

BOUNDS FOR R_0^2. The only missing ingredient to computing the identified set for the degree of invertibility is $\text{Var}(\tilde{z}_t \mid \{y_\tau\}_{-\infty<\tau\leq t})$. We can approximate this quantity arbitrarily well as $M \to \infty$ by

$$\text{Var}(\tilde{z}_t \mid \{y_\tau\}_{t-M\leq\tau\leq t}) = \Sigma_{\tilde{z}} - (\Sigma_{y,\tilde{z},0}', 0_{1\times n_y M}) \Sigma_{y,(M)}^{-1} (\Sigma_{y,\tilde{z},0}', 0_{1\times n_y M})',$$

where $\Sigma_{y,(M)}$ is the full variance-covariance matrix of $(y_t', y_{t-1}', \dots, y_{t-M}')'$.

Bounds for R^2_∞. The only missing ingredient to computing the identified set for the degree of recoverability is $\text{Var}(\tilde{z}_t \mid \{y_\tau\}_{-\infty<\tau<\infty})$. We can approximate this quantity arbitrarily well as $M \to \infty$ by

$$\text{Var}(\tilde{z}_t \mid \{y_\tau\}_{t-M\leq\tau\leq t+M}) = \Sigma_{\tilde{z}} - \Sigma_{\tilde{z},y,(M,M)}\Sigma^{-1}_{y,(M,M)}\Sigma'_{\tilde{z},y,(M,M)},$$

where all objects were already defined above.

Bounds for FVR. To compute the identified set for $FVR_{i,\ell}$, we need $\text{Cov}(y_t, \tilde{z}_{t-h})$ as well as $\text{Var}(y_{i,t+\ell} \mid \{y_\tau\}_{-\infty<\tau\leq t})$. The first object was discussed above, and the second object is well approximated for large M by[37]

$$\begin{aligned}\text{Var}(y_{i,t+\ell} \mid \{y_\tau\}_{t-M\leq\tau\leq t}) = \text{Var}(y_{i,t}) - &(\text{Cov}(y_{i,t+\ell}, y_t), \dots, \text{Cov}(y_{i,t+\ell}, y_{t-M}))\Sigma^{-1}_{y,(M)} \\ &\times (\text{Cov}(y_{i,t+\ell}, y_t), \dots, \text{Cov}(y_{i,t+\ell}, y_{t-M}))',\end{aligned}$$

where $\Sigma_{y,(M)}$ was defined above.

Bounds for FVD. To compute the overall lower bound for the FVD, we need $\text{Var}(\tilde{y}^{(\alpha_{UB})}_{i,t+\ell} \mid \{\tilde{y}^{(\alpha_{UB})}_\tau\}_{-\infty<\tau\leq t})$. As before, we approximate this by $\text{Var}(\tilde{y}^{(\alpha_{UB})}_{i,t+\ell} \mid \{\tilde{y}^{(\alpha_{UB})}_\tau\}_{t-M\leq\tau\leq t})$ for large M. The same formula used above for $\text{Var}(y_{i,t+\ell} \mid \{y_\tau\}_{t-M\leq\tau\leq t})$ applies, where covariances are obtained from

$$\text{Cov}(\tilde{y}^{(\alpha_{UB})}_{t+\ell}, \tilde{y}^{(\alpha_{UB})}_t) = \text{Cov}(y_{t+\ell}, y_t) - \frac{1}{\alpha^2_{UB}} \textstyle\sum_{m=0}^\infty \text{Cov}(y_t, \tilde{z}_{t-m-\ell}) \, \text{Cov}(y_t, \tilde{z}_{t-m})'.$$

The sum can be truncated when the contribution of additional terms is small.

Remarks. Our computations require two truncation choices: the maximal VMA horizon $\hat{p}$, and the maximal prediction horizons M. In all codes, we set these truncation parameters large enough to leave results unaffected by further increases.[38] Derivatives of all parameters with respect to ϑ can be computed by finite differences or automatic differentiation.

[37] In our computations for this paper we use the Kalman filter to compute the conditional variance, but there is little difference in numerical accuracy or speed relative to the formula stated here.

[38] Our choices of truncation parameters therefore differs according to the application, depending on the persistence of the studied processes.

A.2 Proofs of main results

A.2.1 Auxiliary lemma

Lemma 1. *Let B be an $n \times n$ Hermitian positive definite complex-valued matrix and b an n-dimensional complex-valued column vector. Let x be a nonnegative real scalar. Then $B - x^{-1}bb^*$ is positive (semi)definite if and only if $x > (\geq)\, b^* B^{-1} b$.*

Please find the proof in Online Appendix B.11.1.

A.2.2 Proof of Proposition 1

Let α and the spectrum $s_w(\omega)$ be given. Define the n_y-dimensional vectors

$$\overline{\Theta}_{\bullet,1,\ell} = \alpha^{-1} \operatorname{Cov}(y_t, \tilde{z}_{t-\ell}), \quad \ell \geq 0,$$

and the corresponding vector lag polynomial

$$\overline{\Theta}_{\bullet,1}(L) = \sum_{\ell=0}^{\infty} \overline{\Theta}_{\bullet,1,\ell} L^{\ell}.$$

Since $\alpha^2 \leq \alpha_{UB}^2$, we may define $\overline{\sigma}_v = \sqrt{\operatorname{Var}(\tilde{z}_t) - \alpha^2}$. Since $\alpha^2 > \alpha_{LB}^2$, Lemma 1 implies that

$$s_y(\omega) - \frac{2\pi}{\alpha^2} s_{y\tilde{z}}(\omega) s_{y\tilde{z}}(\omega)^* = s_y(\omega) - \frac{1}{2\pi} \overline{\Theta}_{\bullet,1}(e^{-i\omega}) \overline{\Theta}_{\bullet,1}(e^{-i\omega})^*$$

is positive definite for every $\omega \in [0, 2\pi]$. Hence, the Wold decomposition theorem (Hannan, 1970, Thm. 2″, p. 158) implies that there exists an $n_y \times n_y$ matrix lag polynomial $\tilde{\Theta}(L) = \sum_{\ell=0}^{\infty} \tilde{\Theta}_\ell L^\ell$ such that[39]

$$s_y(\omega) - \frac{1}{2\pi} \overline{\Theta}_{\bullet,1}(e^{-i\omega}) \overline{\Theta}_{\bullet,1}(e^{-i\omega})^* = \frac{1}{2\pi} \tilde{\Theta}(e^{-i\omega}) \tilde{\Theta}(e^{-i\omega})^*, \quad \omega \in [0, 2\pi].$$

Thus, the following model for $w_t = (y_t', \tilde{z}_t)'$ generates the desired spectrum $s_w(\omega)$:

$$\begin{aligned} y_t &= \overline{\Theta}_{\bullet,1}(L) \overline{\varepsilon}_{1,t} + \tilde{\Theta}(L) \tilde{\varepsilon}_t, \\ \tilde{z}_t &= \alpha \overline{\varepsilon}_{1,t} + \overline{\sigma}_v \overline{v}_t, \end{aligned}$$

[39] We can rule out a deterministic term in the Wold decomposition because a continuous and positive definite spectral density satisfies the full-rank condition of Hannan (1970, p. 162).

$$(\overline{\varepsilon}_{1,t}, \tilde{\varepsilon}_t', \overline{v}_t)' \overset{i.i.d.}{\sim} N(0, I_{n_y+2}).$$

Note that the construction requires only $n_\varepsilon = n_y + 1$ shocks, $\overline{\varepsilon}_{1,t} \in \mathbb{R}$ and $\tilde{\varepsilon}_t \in \mathbb{R}^{n_y}$. □

A.2.3 Proof of Proposition 2

IDENTIFIED SET FOR R_0^2. If the identified set contains 1, then there must exist an $\overline{\alpha} \in [\alpha_{LB}, \alpha_{UB}]$ and i.i.d., independent standard Gaussian processes $\overline{\varepsilon}_{1,t}$ and $\overline{v}_t$ such that (i) $\tilde{z}_t = \overline{\alpha} \times \overline{\varepsilon}_{1,t} + \overline{v}_t$, (ii) $\overline{v}_t$ is uncorrelated with y_t at all leads and lags, and (iii) $\overline{\varepsilon}_{1,t}$ lies in the closed linear span of $\{y_\tau\}_{-\infty<\tau\leq t}$. This immediately implies the "only if" statement.

For the "if" part, assume $\tilde{z}_t$ does not Granger cause y_t. By the equivalence of Sims and Granger causality, $\tilde{z}_t^\dagger = E(\tilde{z}_t \mid \{y_\tau\}_{-\infty<\tau<\infty}) = E(\tilde{z}_t \mid \{y_\tau\}_{-\infty<\tau\leq t})$. Note that the latter best linear predictor is white noise since, for any $\ell \geq 1$,

$$\begin{aligned}\operatorname{Cov}\left(E(\tilde{z}_t \mid \{y_\tau\}_{-\infty<\tau\leq t}), y_{t-\ell}\right) &= \operatorname{Cov}(\tilde{z}_t, y_{t-\ell}) - \operatorname{Cov}\left(\tilde{z}_t - E(\tilde{z}_t \mid \{y_\tau\}_{-\infty<\tau\leq t}), y_{t-\ell}\right) \\ &= 0 - 0,\end{aligned}$$

using the fact that $\tilde{z}_t$ is a projection residual. In conclusion, the best linear predictor $\tilde{z}_t^\dagger$ of $\tilde{z}_t$ given $\{y_\tau\}_{-\infty<\tau<\infty}$ depends only on $\{y_\tau\}_{-\infty<\tau\leq t}$ and it has a constant spectrum. From the expression for α_{LB}^2, we get that $\alpha_{LB}^2 = \operatorname{Var}(E(\tilde{z}_t \mid \{y_\tau\}_{-\infty<\tau\leq t}))$, which further yields $\alpha_{LB}^2 = \operatorname{Var}(\tilde{z}_t)\tilde{R}_0^2$. Hence, expression (14) implies that the upper bound of the identified set for R_0^2 equals 1.

IDENTIFIED SET FOR R_∞^2. The upper bound of the identified set for R_∞^2 equals 1 if and only if $2\pi \sup_{\omega\in[0,\pi]} s_{\tilde{z}^\dagger}(\omega) = \tilde{R}_\infty^2 \operatorname{Var}(\tilde{z}_t)$, and the right-hand side equals $\operatorname{Var}(\tilde{z}_t^\dagger) = \int_0^{2\pi} s_{\tilde{z}^\dagger}(\omega)\, d\omega$. But we have $\sup_{\omega\in[0,\pi]} s_{\tilde{z}^\dagger}(\omega) = \frac{1}{2\pi}\int_0^{2\pi} s_{\tilde{z}^\dagger}(\omega)\, d\omega$ if and only if $s_{\tilde{z}^\dagger}(\omega)$ is constant in ω almost everywhere, i.e., $\tilde{z}_t^\dagger$ is white noise. □

References

Andrews, D. W. & Shi, X. (2013). Inference Based on Conditional Moment Inequalities. *Econometrica*, *81*(2), 609–666.

Andrews, D. W. & Shi, X. (2017). Inference based on many conditional moment inequalities. *Journal of Econometrics*, *196*(2), 275–287.

Angeletos, G.-M., Collard, F., & Dellas, H. (2018). Business Cycle Anatomy. Manuscript, MIT.

Angrist, J. D., Jordà, Ò., & Kuersteiner, G. M. (2018). Semiparametric Estimates of Monetary Policy Effects: String Theory Revisited. *Journal of Business & Economic Statistics*, *36*(3), 371–387.

Barnichon, R. & Brownlees, C. (2018). Impulse Response Estimation by Smooth Local Projections. *Review of Economics and Statistics*. Forthcoming.

Barnichon, R. & Matthes, C. (2018). Functional Approximation of Impulse Responses. *Journal of Monetary Economics*, *99*, 41–55.

Blanchard, O. J., L'Huillier, J. P., & Lorenzoni, G. (2013). News, Noise, and Fluctuations: An Empirical Exploration. *American Economic Review*, *103*(7), 3045–3070.

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

Caldara, D. & Herbst, E. (2019). Monetary policy, real activity, and credit spreads: Evidence from bayesian proxy svars. *American Economic Journal: Macroeconomics*, *11*(1), 157–92.

Chahrour, R. & Jurado, K. (2018). Recoverability. Manuscript, Duke University.

Chernozhukov, V., Lee, S., & Rosen, A. M. (2013). Intersection Bounds: Estimation and Inference. *Econometrica*, *81*(2), 667–737.

Christiano, L., Eichenbaum, M., & Evans, C. (1999). Monetary Policy Shocks: What Have We Learned and to What End? In J. B. Taylor & M. Woodford (Eds.), *Handbook of Macroeconomics, Volume 1A* chapter 2, (pp. 65–148). Elsevier.

Christiano, L. J., Motto, R., & Rostagno, M. (2014). Risk Shocks. *American Economic Review*, *104*(1), 27–65.

Cochrane, J. H. (1994). Shocks. In *Carnegie-Rochester Conference Series on Public Policy*, volume 41 (pp. 295–364). Elsevier.

Del Negro, M., Giannoni, M. P., & Patterson, C. (2012). The Forward Guidance Puzzle. Federal Reserve Bank of New York Staff Report No. 574.

Forni, M. & Gambetti, L. (2014). Sufficient information in structural VARs. *Journal of Monetary Economics*, *66*(Supplement C), 124–136.

Forni, M., Gambetti, L., Lippi, M., & Sala, L. (2017a). Noise Bubbles. *Economic Journal*, *127*(604), 1940–1976.

Forni, M., Gambetti, L., Lippi, M., & Sala, L. (2017b). Noisy News in Business Cycles. *American Economic Journal: Macroeconomics*, *9*(4), 122–152.

Forni, M., Gambetti, L., & Sala, L. (2018). Structural VARs and Non-invertible Macroeconomic Models. *Journal of Applied Econometrics*. Forthcoming.

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. & Reichlin, L. (2006). Does Information Help Recovering Structural Shocks from Past Observations? *Journal of the European Economic Association*, *4*(2/3), 455–465.

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

Gorodnichenko, Y. & Lee, B. (2017). A Note on Variance Decomposition with Local Projections. Manuscript, University of California Berkeley.

Granziera, E., Moon, H. R., & Schorfheide, F. (2018). Inference for VARs identified with sign restrictions. *Quantitative Economics*, *9*(3), 1087–1121.

Hall, R. E. (2011). The Long Slump. *American Economic Review*, *101*(2), 431–469.

Hannan, E. (1970). *Multiple Time Series*. Wiley Series in Probability and Statistics. John Wiley & Sons.

Hansen, L. P. & Sargent, T. J. (1991). Two Difficulties in Interpreting Vector Autoregressions. In L. P. Hansen & T. J. Sargent (Eds.), *Rational Expectations Econometrics*, Underground Classics in Economics chapter 4, (pp. 77–119). Westview Press.

Imbens, G. W. & Manski, C. F. (2004). Confidence Intervals for Partially Identified Parameters. *Econometrica*, *72*(6), 1845–1857.

Jermann, U. & Quadrini, V. (2012). Macroeconomic Effects of Financial Shocks. *American Economic Review*, *102*(1), 238–271.

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. (2018). The effects of quasi-random monetary experiments. Manuscript, Federal Reserve Bank of San Francisco.

Justiniano, A., Primiceri, G. E., & Tambalotti, A. (2010). Investment shocks and business cycles. *Journal of Monetary Economics*, *57*(2), 132–145.

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.

King, R. G., Plosser, C. I., Stock, J. H., & Watson, M. W. (1991). Stochastic Trends and Economic Fluctuations. *American Economic Review*, *81*(4), 819–840.

Klepper, S. & Leamer, E. E. (1984). Consistent Sets of Estimates for Regressions with Errors in All Variables. *Econometrica*, *52*(1), 163–183.

Kydland, F. E. & Prescott, E. C. (1982). Time to Build and Aggregate Fluctuations. *Econometrica*, *50*(6), 1345–1370.

Leeper, E. M., Walker, T. B., & Yang, S.-C. S. (2013). Fiscal Foresight and Information Flows. *Econometrica*, *81*(3), 1115–1145.

Lippi, M. & Reichlin, L. (1994). VAR analysis, nonfundamental representations, Blaschke matrices. *Journal of Econometrics*, *63*(1), 307–325.

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

Mertens, K. & Ravn, M. O. (2010). Measuring the Impact of Fiscal Policy in the Face of Anticipation: A Structural VAR Approach. *Economic Journal*, *120*(544), 393–413.

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.

Nakamura, E. & Steinsson, J. (2018). High Frequency Identification of Monetary Non-Neutrality: The Information Effect. *Quarterly Journal of Economics*, *133*(3), 1283–1330.

Plagborg-Møller, M. (2019). Bayesian inference on structural impulse response functions. *Quantitative Economics*, *10*(1), 145–184.

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. (1989). Does Monetary Policy Matter? A New Test in the Spirit of Friedman and Schwartz. In O. J. Blanchard & S. Fischer (Eds.), *NBER Macroeconomics Annual*, volume 4 (pp. 121–170). University of Chicago Press.

Schmitt-Grohé, S. & Uribe, M. (2012). What's News in Business Cycles. *Econometrica*, *80*(6), 2733–2764.

Sims, C. A. & Zha, T. (2006). Does Monetary Policy Generate Recessions? *Macroeconomic Dynamics*, *10*(02), 231–272.

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*, *28*(610), 917–948.

Stoye, J. (2009). More on Confidence Intervals for Partially Identified Parameters. *Econometrica*, *77*(4), 1299–1315.

Wolf, C. K. (2018). SVAR (Mis)Identification and the Real Effects of Monetary Policy. Manuscript, Princeton University.