

Instrumental Variable Identification of Dynamic Variance Decompositions*

Mikkel Plagborg-Møller
Princeton University

Christian K. Wolf
Princeton University

This version: June 13, 2018. 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 do not answer the key question of how important the shocks are in driving macro aggregates. We provide tools for doing inference on forecast variance decompositions in a general semiparametric moving average model, disciplined only by the availability of external instruments. The share of the forecast 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 much weaker assumptions than invertibility, a condition imposed in conventional Structural Vector Autoregressive (SVAR) analysis. In fact, external instruments make the invertibility assumption testable. To perform inference, we construct easily computable, partial identification robust confidence intervals. We illustrate our methods using (i) a structural macro model and (ii) an empirical study of the importance of monetary policy shocks.

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, Dario Caldara, Thorsten Drautzburg, Domenico Giannone, Yuriy Gorodnichenko, Ed Herbst, Marek Jarociński, Peter Karadi, Lutz Kilian, Michal Kolesár, Byoungchan Lee, Pepe Montiel Olea, Ulrich Müller, Emi Nakamura, Giorgio Primiceri, 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 without relying on dubious functional form assumptions or identifying restrictions. For example, local projections (LP) have become a popular direct regression-based alternative to Structural Vector Autoregression (SVAR) methods.¹ Additionally, instrumental variable (IV, also known as proxy variable) methods are now routinely used to conduct structural analysis under plausible identifying assumptions.² Several recent papers have combined these two ideas, yielding an appealingly semiparametric method, LP-IV, with an economically transparent framework for identification.³

However, researchers care not just about the effects of structural shocks on macro aggregates; they also want to know how *important* different shocks are in driving economic fluctuations. In the classical SVAR literature, this question was routinely addressed using forecast variance decompositions, which measure the fraction of the overall forecast variance for a variable that can be attributed to each of the driving shocks, and historical decompositions, which measure the contributions of each shock to observed fluctuations at certain points in time. Conveniently, in conventional SVAR analysis, identifying the impulse response functions also identifies the underlying shocks (and thus their importance). In contrast, in the semiparametric LP-IV setting, the extent to which external instruments are informative about forecast variance decompositions and historical decompositions has been an open question, to our knowledge. Applied researchers have thus faced an unfortunate dilemma between a need to quantify the importance of shocks and the desire to avoid imposing a restrictive SVAR structure (or assuming that shocks are directly observed).

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. Our model allows 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 forecast variance decompositions. *Point identification* of most parameters of interest, including historical decompositions, can be achieved under a further assumption that the shock of interest is recoverable from the

¹Jordà (2005); Angrist et al. (2017).

² Stock (2008); Stock & Watson (2012); Mertens & Ravn (2013); Gertler & Karadi (2015); Stock & Watson (2016); Caldara & Kamps (2017).

³Mertens (2015); Ramey (2016); Barnichon & Brownlees (2017); Stock & Watson (2017).

infinite past, present, and future values of the endogenous macro variables. This requirement is substantially weaker than the often questionable invertibility requirement of SVAR analysis – the ability to recover structural shocks only from past and present macro variables. In fact, we show that the availability of an external IV makes the invertibility assumption *testable*. Finally, to perform inference, we develop easily computable, partial identification robust confidence intervals for forecast variance decompositions and other objects of interest.

Following [Stock & Watson \(2017\)](#), the LP-IV model that we analyze, 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.⁴ 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. Unlike standard SVAR models, we do not restrict the shocks to be *invertible*, i.e., spanned by past and current (but not future) values of the observed endogenous variables.

In this baseline LP-IV model, we show that forecast 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, 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 forecast 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. We consider two definitions of the forecast variance decomposition that differ in the information set used for the forecast: either conditional on all past endogenous variables, or conditional on all past shocks. The identified sets for these two concepts differ, unless invertibility is imposed *a priori*.

As the LP-IV model does not assume that shocks are invertible *a priori*, we are able to sharply 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

⁴This Structural Vector Moving Average Model has been analyzed recently from a Bayesian viewpoint by [Barnichon & Matthes \(2017\)](#) and [Plagborg-Møller \(2017\)](#), although with little emphasis on IVs.

values of the endogenous variables. We show that this R^2 measure is partially identified, and we discuss which properties of the data would allow us to either confirm or reject invertibility with certainty. In particular, the data rule out invertibility (and thus the validity of SVAR-IV analysis) if and only if the IV Granger-causes the observed endogenous variables.⁵

Although the baseline model is partially identified, we additionally provide assumptions that guarantee point identification of certain variance decompositions and the degree of invertibility. 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.

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.⁶ 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 construct confidence intervals both for the unknown parameters and for the identified sets, including a Bayesian implementation of the latter. 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,

⁵We characterize the resulting population bias of SVAR-IV estimators in Online Appendix B.3.

⁶We 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.

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 invertibility-based (e.g., SVAR) 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 the short- or medium-run cycles of this shock.

Finally, we apply our method to U.S. data and show that monetary shocks account for at most a moderate fraction of the forecast variance of output growth and inflation. We use the specification of [Gertler & Karadi \(2015\)](#), whose external instrument is constructed from changes in interest rate futures during short time windows around Federal Open Market Committee announcements. Unlike [Gertler & Karadi \(2015\)](#), we do not assume invertibility (i.e., SVAR structure) since we find that the data reject this assumption. At all forecast horizons, our partial identification robust 90% confidence intervals rule out that the monetary shock contributes more than 31% of the forecast variance of output growth and more than 8% of the forecast variance of inflation. Moreover, monetary shocks do not contribute substantially to short-run fluctuations in a non-default-related corporate bond spread. Thus, despite our weak identifying assumptions, we obtain informative upper bounds on the importance of the monetary shock.

LITERATURE. A rapidly growing literature has provided inference tools for the LP-IV model, although variance decompositions and historical decompositions have been neglected. External IVs (also known as proxy variables) relax the assumption that the shock in a local projection is directly observed (or consistently estimable). The theoretical background for LP-IV estimation of impulse response functions was established by [Mertens \(2015\)](#), [Ramey \(2016\)](#), [Barnichon & Brownlees \(2017\)](#), [Jordà et al. \(2017\)](#), [Ramey & Zubairy \(2017\)](#), and most comprehensively [Stock & Watson \(2017\)](#). As we emphasize below, the usual IV exclusion restrictions only yield point identification of *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 implicit invertibility assumption ([Kilian & Lütkepohl,](#)

2017, Ch. 4). [Stock & Watson \(2017\)](#) 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 recoverability of the shock of interest is sufficient to yield 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.

The ability of the LP-IV framework to allow for noninvertible shocks is a key attraction. As is well known, the standard SVAR model imposes invertibility *a priori*. Because noninvertibility is a frequent feature of structural macro models, the issue has received a lot of attention in the SVAR literature (see references in [Plagborg-Møller, 2017](#), Sec. 2.3). [Stock & Watson \(2017\)](#) 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 data can conclusively reject invertibility. These conditions are related to Granger causality, as in the SVAR settings studied by [Giannone & Reichlin \(2006\)](#) and [Forni & Gambetti \(2014\)](#).

The concept of “recoverability” has precursors in the SVAR literature. Our definition of recoverability has independently been proposed by [Chahrour & Jurado \(2017\)](#), whose generic shock identification results do not specifically consider the LP-IV model and who do not consider our continuous measure of recoverability.⁷ 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 methods combine the literatures on partial identification and VAR-based spectral estimation. We employ the general confidence interval construction for interval-identified parameters developed by [Imbens & Manski \(2004\)](#) and [Stoye \(2009\)](#). Our partial identification methods are complementary to the literature on inference in sign-identified SVARs ([Moon et al., 2013](#); [Giacomini & Kitagawa, 2015](#); [Gafarov et al., 2017](#)).

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.⁸

⁷Moreover, we show that the recoverability assumption is testable in the LP-IV model.

⁸https://scholar.princeton.edu/mikkelpm/decomp_iv

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 does not assume shocks to be invertible, unlike standard SVAR analysis. We assume the availability of valid external IVs (proxy variables) – variables that correlate with the shock of interest, but not with the other shocks. In addition to impulse response functions, our main objects of interest are forecast variance decompositions, historical decompositions, and the degree of invertibility of the shock of interest.

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. 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,

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

where L is the lag operator. The (i, j) element $\Theta_{i,j,\ell}$ of the $n_y \times n_\varepsilon$ 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 $\Theta_{i,\bullet,\ell}$. To obtain a nonsingular stochastic process, we assume that $\Theta(x)$ has full row rank for all complex scalars x on the unit circle. This condition requires $n_\varepsilon \geq n_y$, but – crucially – we do not assume that the number of shocks n_ε is known. 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 principal appeal of LP-IV analysis is that it does not require an underlying SVAR structure or a fully-specified DSGE model.

Second, we assume the availability of external IVs for the shock of interest.⁹ We specify

⁹If instruments are not available, the model (1) is severely underidentified (Lippi & Reichlin, 1994).

the shock of interest to be the first one, $\varepsilon_{1,t}$. The n_z -dimensional vector $z_t = (z_{1,t}, \dots, z_{n_z,t})'$ of IVs is assumed to correlate with the first shock but not the other shocks,

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

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. Throughout the paper, $\|\cdot\|$ refers to the Euclidean norm. The key restriction on the IVs is that the shock of interest $\varepsilon_{1,t}$ is the only contemporaneous shock to enter into equation (2). In Online Appendix B.2 we discuss how our analysis changes if the exclusion restriction is relaxed. 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.

Finally, we assume that the structural shocks and IV disturbances are jointly i.i.d. standard Gaussian,

$$(\varepsilon'_t, v'_t)' \stackrel{i.i.d.}{\sim} N(0, I_{n_z+n_y}), \quad (3)$$

where I_n denotes the n -dimensional identity matrix. The mutual independence of the shocks is the standard assumption in empirical macroeconomics. We adopt the Gaussianity assumption for notational convenience. As usual, versions of all our results can be phrased using the language of linear projections.¹⁰ The sole meaningful restriction is that we only consider identification from the second-moment properties of the data, as is standard in the applied macro literature (and without loss of generality for Gaussian data). Also note that we have normalized the variances of all shocks to 1.¹¹

Throughout the paper we tacitly assume that the $(n_y + n_z)$ -dimensional data vector $(y'_t, z'_t)'$ is strictly stationary. This is achieved by assuming that the elements of Θ_{ℓ} , 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.

We allow for lagged values of z_t and y_t on the right-hand side of (2) because this is precisely enough to ensure that the LP-IV model is untestable (using second moments) in

¹⁰Replace conditional expectations by linear projections and replace conditional variances by variances of projection residuals.

¹¹Some authors prefer a different parametrization of the LP-IV model where the shocks have non-unit variances but instead some impulse responses $\Theta_{i,j,\ell}$ are normalized to 1 (Stock & Watson, 2016, Sec. 4.1.3). All our results can be translated one-for-one into this alternative parametrization by simple scaling.

the case of a single IV. In other words, when $n_z = 1$, there always exists a model of the form (1)–(3) that can match any given autocovariance structure of $(y'_t, z_t)'$, cf. [Proposition 1](#) below.¹² If the economic application allows the researcher to exclude certain lagged terms *a priori*, this provides testable restrictions that are straight-forward to impose in our VAR-based inference procedures in [Section 5](#). Except for the lag terms, our model for y_t and z_t is essentially identical to the LP-IV model studied by [Stock & Watson \(2017\)](#).

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. Furthermore, if the number of structural shocks n_ε strictly exceeds the number of endogenous variables 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

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

as an R^2 measure of invertibility of the shock of interest using data up to time $t + \ell$. If 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.3 \(Sims & Zha, 2006, pp. 243–245; Forni et al., 2016; Wolf, 2017\)](#). For noninvertible models, a plot of R_ℓ^2 for $\ell = 0, 1, 2, \dots$ 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

¹²The model with multiple instruments is testable, as further discussed in [Section 3.3](#).

Online Appendix B.4.

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 $R_\infty^2 = 1$.¹³ This property will become important when we consider assumptions that guarantee point identification.

VARIANCE DECOMPOSITIONS/RATIOS. In addition to the impulse responses Θ_ℓ , a primary object of interest is the forecast variance decomposition of the moving average model. Forecast variance decompositions capture the share of the forecast variance that can be attributed to particular shocks. In other words, they provide a quantitative measure of the importance of different shocks in generating macroeconomic fluctuations.

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

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

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 forecasting uncertainty. Writing out the denominator:

$$\begin{aligned} \text{Var}(y_{i,t+\ell} \mid \{y_\tau\}_{-\infty < \tau \leq t}) &= \text{Var} \left(\sum_{m=0}^{\infty} \Theta_{i,\bullet,m} \varepsilon_{t+\ell-m} \mid \{y_\tau\}_{-\infty < \tau \leq t} \right) \\ &= \sum_{m=0}^{\ell-1} \Theta_{i,\bullet,m} \Theta'_{i,\bullet,m} + \text{Var} \left(\sum_{m=\ell}^{\infty} \Theta_{i,\bullet,m} \varepsilon_{t+\ell-m} \mid \{y_\tau\}_{-\infty < \tau \leq t} \right). \end{aligned}$$

Thus, the denominator of the FVR captures both fundamental forecasting uncertainty (uncertainty related to future shock realizations) and noninvertibility-induced forecasting uncertainty (uncertainty related to imperfect knowledge about *past* shocks). In contrast, the numerator only reflects fundamental forecasting uncertainty. This means that, when the first

¹³Chahrour & Jurado (2017) 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).

shock is noninvertible, the FVR does not equal 1 even if the first shock 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 ,

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

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).

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.¹⁴

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

3 Identification

This section presents our main results on instrumental variable identification of variance decompositions and the degree of invertibility. For exposition, we start by deriving results for a static version of the LP-IV model. We then turn to the general dynamic model, which applies the static results to the frequency domain representation of the data. The dynamics involve additional nuances in characterizing the informativeness of the macro aggregates for the shock at all frequencies. Our main results assume availability of a single external IV for the shock of interest. We then show that identification analysis in a model with multiple

¹⁴Forni et al. (2016) point out the bias caused by noninvertibility when estimating the FVD using SVARs. They also propose a frequency-domain variance decomposition concept, to which our identification results can easily be applied, just as for the FVR. Details are available upon request.

IVs for the same shock can be reduced to the single-IV case.

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. The static model with a single IV assumes¹⁵

$$\begin{aligned} y_t &= \Theta_0 \varepsilon_t, \\ z_t &= \alpha \varepsilon_{1,t} + \sigma_v v_t, \\ (\varepsilon'_t, v_t)' &\overset{i.i.d.}{\sim} N(0, I_{n_\varepsilon+1}), \end{aligned}$$

where Θ_0 is $n_y \times n_\varepsilon$, and $\alpha, \sigma_v \geq 0$ are scalars. To avoid singularity, we assume that Θ_0 has full row rank, so in particular $n_\varepsilon \geq n_y$.

In the static case, the degree of invertibility of the model is fully summarized by the static projection R^2 , defined as

$$R_0^2 = 1 - \text{Var}(\varepsilon_{1,t} | y_t) = \Theta'_{\bullet,1,0} \text{Var}(y_t)^{-1} \Theta_{\bullet,1,0}.$$

Here the second equality follows from the usual linear projection formula.

As for variance decompositions, the static model does not distinguish between the FVR and FVD, and we can restrict attention to one-step prediction:

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

PARTIAL IDENTIFICATION. We now show that the impulse response functions, the degree of invertibility, and variance decompositions are all identified up to a scalar multiple. This factor of proportionality is interval-identified, with nontrivial and informative lower and upper

¹⁵While 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.

bounds. Since the data are i.i.d., identification in this model relies solely on contemporaneous covariance calculations.

It is immediate that impulse responses $\Theta_{i,1,0}$ to the shock of interest are identified up to the scale parameter α :

$$\text{Cov}(y_t, z_t) = \alpha \Theta_{\bullet,1,0}. \quad (5)$$

In particular, relative responses are identified, cf. [Stock & Watson \(2017\)](#). Since the vector $\Theta_{\bullet,1,0}$ is identified up to scale α , the degree of invertibility R_0^2 is identified up to the multiple $\frac{1}{\alpha^2}$, and the FVDs of different variables i are identified up to the same multiple $\frac{1}{\alpha^2}$:

$$R_0^2 = \frac{1}{\alpha^2} \text{Cov}(y_t, z_t)' \text{Var}(y_t)^{-1} \text{Cov}(y_t, z_t), \quad FVD_{i,1} = \frac{1}{\alpha^2} \frac{\text{Cov}(y_{i,t}, z_t)^2}{\text{Var}(y_{i,t})}.$$

Which values of the scale parameter α are consistent with the distribution of the data $w_t = (y_t', z_t)'$? First, the equation defining the IV z_t implies

$$\alpha^2 \leq \text{Var}(z_t) \equiv \alpha_{UB}^2. \quad (6)$$

The boundary case $\alpha = \alpha_{UB}$ corresponds to $\text{Var}(\varepsilon_{1,t} | z_t) = 0$, i.e., perfect instrument strength. Second, we find

$$\alpha_{LB}^2 \equiv \text{Var}(E(z_t | y_t)) = \alpha^2 \text{Var}(E(\varepsilon_{1,t} | y_t)) \leq \alpha^2 \text{Var}(\varepsilon_{1,t}) = \alpha^2, \quad (7)$$

where the inequality uses that “the total sum of squares exceeds the explained sum of squares”. The boundary case $\alpha = \alpha_{LB}$ corresponds to $\text{Var}(\varepsilon_{1,t} | y_t) = 0$, i.e., the observed macro aggregates y_t are perfectly informative about the hidden shock (invertibility).

It is not hard to show (and it follows from our general results below) that the bounds (6) and (7) on α^2 are *sharp*, in the following sense: Given any positive semidefinite variance-covariance matrix for $w_t = (y_t', z_t)'$, and given any value of α^2 in the interval between the bounds (6) and (7), we can construct a static model with the given value of α and which matches the given $\text{Var}(w_t)$ (we just have to choose Θ_0 and σ_v appropriately).

The width of the identified set $[\alpha_{LB}^2, \alpha_{UB}^2]$ for α^2 depends on the degree of invertibility and the strength of the instrument. The interval is never empty, and it collapses to a point only in the knife-edge case of a perfectly informative instrument *and* invertibility of the first shock. Generically, α – and so impulse responses, FVDs, and the degree of invertibility – are only partially identified, but with useful bounds that limit the range of admissible values.

Intuitively, the bounds arise because a large α requires a large variance of the IV, while a small α requires y_t and z_t to be nearly uncorrelated.

To interpret the identified set of $\frac{1}{\alpha^2}$, we can express it in terms of the underlying (unknown) model parameters:

$$\left[\underbrace{\frac{\alpha^2}{\alpha^2 + \sigma_v^2}}_{\text{instrument strength}} \times \frac{1}{\alpha^2}, \quad \underbrace{\frac{1}{R_0^2}}_{\text{recoverability}} \times \frac{1}{\alpha^2} \right].$$

The lower bound is more informative (i.e., 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. Conversely, the upper bound is more informative (i.e., smaller and closer to the true $\frac{1}{\alpha^2}$) when the model is closer to being invertible for the shock of interest.

Having partially identified α , we obtain identified sets for the FVD and degree of invertibility. By scaling the identified set for $\frac{1}{\alpha^2}$, we find the identified set for $FVD_{i,0}$:

$$\left[\underbrace{\frac{1}{\text{Var}(z_t)} \times \frac{\text{Cov}(y_{i,t}, z_t)^2}{\text{Var}(y_{i,t})}}_{\frac{\alpha^2}{\alpha^2 + \sigma_v^2} \times FVD_{i,0}}, \quad \underbrace{\frac{1}{\text{Var}(E(z_t | y_t))} \times \frac{\text{Cov}(y_{i,t}, z_t)^2}{\text{Var}(y_{i,t})}}_{\frac{1}{R_0^2} \times FVD_{i,0}} \right].$$

Instrument informativeness and invertibility thus map one-to-one into the width of the identified set for the FVD. The identified set for the degree of invertibility R_0^2 can similarly be obtained by scaling the identified set for $\frac{1}{\alpha^2}$. The identified set for R_0^2 always contains 1, i.e., the data can never reject invertibility in the static model.

Our bounds on variance decompositions are related to the literature on multivariate models with classical measurement error (Klepper & Leamer, 1984).¹⁶ First, the R^2 value from a regression of $y_{i,t}$ on $z_t = \alpha\varepsilon_{1,t} + \sigma_v v_t$ understates the importance of the shock $\varepsilon_{1,t}$, due to attenuation bias.¹⁷ Second, the R^2 from a regression of $y_{i,t}$ on $E(z_t | y_t) = \alpha E(\varepsilon_{1,t} | y_t)$ overstates the importance of the shock $\varepsilon_{1,t}$, due to bias caused by regressing on the “fitted value” $E(\varepsilon_{1,t} | y_t)$. Thus, we obtain lower and upper bounds on the R^2 in an infeasible regression of $y_{i,t}$ on $\varepsilon_{1,t}$.

¹⁶Although the static LP-IV model is just the multivariate classical measurement error model, our bounds are not formally implied by the Klepper & Leamer (1984) results, since we are interested in variance decompositions rather than regression coefficients.

¹⁷A related argument appears in Gorodnichenko & Lee (2017, Appendix D).

SUFFICIENT CONDITIONS FOR POINT IDENTIFICATION. Although the baseline model is partially identified, point identification obtains under a variety of auxiliary assumptions.

First, assume that the shock of interest is recoverable, which in the static model is the same as invertibility: $E(\varepsilon_{1,t} \mid y_t) = \varepsilon_{1,t}$, or equivalently $R_0^2 = 1$. Then α^2 equals the lower bound in (7); we can then identify the impulse responses $\Theta_{\bullet,1,0}$ from the covariance relationship (5), and σ_v from $\text{Var}(z_t)$. Hence, all objects of interest are point-identified under the recoverability assumption. A stronger condition than recoverability is that there are as many shocks as variables, $n_\varepsilon = n_y$. This condition implies that Θ_0 is square and invertible, so *all* shocks are recoverable, and point identification follows.

Second, restrictions on the instrument z_t can also be sufficient to ensure point identification. Point identification obtains if the instrument is perfect; that is, if $\sigma_v = 0$. In that case the shock $\varepsilon_{1,t}$ is effectively observed by the econometrician and all parameters can be identified directly from regressions of $y_{i,t}$ on z_t (local projections). Equivalently, the true α equals the upper bound in (6), and then all derivations follow as before. An alternative point-identifying assumption is the availability of multiple instruments with mutually independent “first-stage disturbances” v_t , as shown in Section 3.3.

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. The main building block result is that, exactly as in the static model, the identified set for the scale parameter α is an interval with informative bounds. From this result we derive identified sets for the main objects of interest: the degree of invertibility and variance decompositions. Relative to the static case, the dynamic case involves additional nuances in characterizing the informativeness of the data for the hidden shock at all frequencies.

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 (2). We write $\Sigma_v^{1/2} = \sigma_v \geq 0$, a scalar.

PRELIMINARIES. For the identification analysis, it will prove convenient to define and work with the IV projection residual

$$\tilde{z}_t \equiv z_t - E(z_t \mid \{y_\tau, z_\tau\}_{-\infty < \tau < t}) = \alpha \varepsilon_{1,t} + \sigma_v v_t. \quad (8)$$

We have thus removed any dependence on lagged observed variables, and \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)

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

This object is well-defined if the autocovariance function of $(a'_t, b'_t)'$ is absolutely summable. For any vector time series a_t , we denote its spectrum by $s_a(\omega) \equiv s_{aa}(\omega)$.

Just like identification in the static case proceeded through the variance-covariance matrix of the data, identification in the general dynamic model will rely heavily on the joint spectrum for $w_t = (y'_t, \tilde{z}'_t)'$ implied by the LP-IV model, i.e., equations (1), (3), and (8). This joint spectrum is given by

$$s_w(\omega) = \begin{pmatrix} s_y(\omega) & s_{y\tilde{z}}(\omega) \\ s_{y\tilde{z}}(\omega)^* & s_{\tilde{z}}(\omega) \end{pmatrix} = \frac{1}{2\pi} \begin{pmatrix} \Theta(e^{-i\omega})\Theta(e^{-i\omega})^* & \alpha\Theta(e^{-i\omega})e_1 \\ \alpha e'_1\Theta(e^{i\omega})' & \alpha^2 + \sigma_v^2 \end{pmatrix}, \quad \omega \in [0, 2\pi], \quad (9)$$

where e_1 is the unit vector with 1 as the first element and zeros elsewhere, an asterisk denotes complex conjugate transpose, and the matrix polynomial $\Theta(\cdot)$ was defined in equation (1).¹⁸ Note the similarity between the spectrum $s_w(\omega)$ and the covariance structure in the static model in Section 3.1. The main difference is that in the dynamic setting we have a matrix at each frequency $\omega \in [0, 2\pi]$.

IMPULSE RESPONSES. As in the static model, the impulse responses to the first shock are identified up to the scale parameter α :

$$\text{Cov}(y_t, \tilde{z}_{t-\ell}) = \alpha\Theta_{\bullet,1,\ell}. \quad (10)$$

This is also clear from expression (9) for the cross-spectrum of y_t and \tilde{z}_t . Thus, *relative* impulse responses $\Theta_{i,1,\ell}/\Theta_{11,0}$ are identified, as shown by Stock & Watson (2017) and others.

¹⁸See Brockwell & Davis (1991, Ch. 11) for references on spectral densities of moving average processes.

SCALE PARAMETER. We now show that, exactly as in the static case, the identified set for α is an interval with informative bounds. Although α itself is not a parameter of primary interest, the results in the next paragraphs are key ingredients to identification of variance decompositions and the degree of invertibility.

Just as in the static case, the variance of the instrument provides the upper bound:

$$\alpha^2 \leq \text{Var}(\tilde{z}_t) \equiv \alpha_{UB}^2. \quad (11)$$

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

To derive the lower bound, we apply 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 :¹⁹

$$\tilde{z}_t^\dagger \equiv E(\tilde{z}_t \mid \{y_\tau\}_{-\infty < \tau < \infty}), \quad (12)$$

$$\varepsilon_{1,t}^\dagger \equiv E(\varepsilon_{1,t} \mid \{y_\tau\}_{-\infty < \tau < \infty}).$$

Then, for every $\omega \in [0, 2\pi]$,

$$s_{\tilde{z}^\dagger}(\omega) = \alpha^2 s_{\varepsilon_1^\dagger}(\omega) \leq \alpha^2 s_{\varepsilon_1}(\omega) = \alpha^2 \times \frac{1}{2\pi}, \quad (13)$$

which is the frequency-domain analogue of the conditional variance inequality (7) in the static case.²⁰ Hence, we obtain the lower bound²¹

$$\alpha^2 \geq 2\pi \sup_{\omega \in [0, \pi]} s_{\tilde{z}^\dagger}(\omega) \equiv \alpha_{LB}^2. \quad (14)$$

This lower bound generalizes the lower bound (7) in the static model. Intuitively, in the static case a small value of α requires y_t and \tilde{z}_t to be nearly independent. In the dynamic case, 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 $\alpha = \alpha_{LB}$

¹⁹Brockwell & 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 the analogous formula applies for $s_{\varepsilon_1^\dagger}(\omega)$. They consider the bivariate case, but the multivariate generalization is straight-forward. Note that $s_{\tilde{z}^\dagger}(\omega)$ need not be constant in ω even though $s_{\tilde{z}}(\omega)$ is.

²⁰ $s_{\varepsilon_1}(\omega) \geq s_{\varepsilon_1^\dagger}(\omega) = s_{y\varepsilon_1}(\omega)^* s_y(\omega)^{-1} s_{y\varepsilon_1}(\omega)$ for all ω , since the joint spectrum is positive semidefinite.

²¹As always with univariate spectral densities, $s_{\tilde{z}^\dagger}(2\pi - \omega) = s_{\tilde{z}^\dagger}(\omega)$ for $\omega \in [0, \pi]$.

corresponds to the observed macro aggregates being perfectly informative about the hidden shock $\varepsilon_{1,t}$ at *some* frequency $\bar{\omega} \in [0, \pi]$, i.e., $s_{\varepsilon_1^\dagger}(\bar{\omega}) = s_{\varepsilon_1}(\cdot) = \frac{1}{2\pi}$.

The main building block 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 being unpredictable from $\{w_\tau\}_{-\infty < \tau < t}$. Choose any $\alpha \in (\alpha_{LB}, \alpha_{UB})$. Then there exists a model of the form (1), (3), and (8) with the given α such that the spectral density of w_t implied by the model matches the given spectral density.*

In words, the distribution of the data allows us to conclude that α^2 lies in the identified set $[\alpha_{LB}^2, \alpha_{UB}^2]$, but the data cannot rule out any values of α^2 in this interval. The proposition does not cover the knife-edge case $\alpha = \alpha_{LB}$, which presents some economically inessential technical difficulties.

The width of the identified set for α^2 depends on the application, although the set is never empty. To interpret the identified set, we can express it in terms of the underlying (unknown) model parameters. In almost perfect analogy to the static case, the identified set for $\frac{1}{\alpha^2}$ is then

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

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. Similar to the static case, the identified set for $\frac{1}{\alpha^2}$ does not collapse to a point unless the instrument is perfect *and* there exists a frequency $\bar{\omega}$ for which the data are perfectly informative about the frequency- $\bar{\omega}$ 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

$$\alpha_{LB}^2 = 2\pi \sup_{\omega \in [0, \pi]} s_{\tilde{z}^\dagger}(\omega) \geq \int_0^{2\pi} s_{\tilde{z}^\dagger}(\omega) d\omega = \text{Var}(\tilde{z}_t^\dagger). \quad (15)$$

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

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

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

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

Thus, if the shock is close to being invertible – or more generally, recoverable – α_{LB}^2 will be close to α^2 . As mentioned above, α_{LB}^2 will in fact be close to α^2 as long as the y_t process is highly informative about the $\varepsilon_{1,t}$ process at *some* frequency. For example, the observed macro variables y_t may not perfectly reveal the short-run fluctuations of an unobserved technology shock, so recoverability fails ($R_\infty^2 < 1$); yet a long-lag two-sided moving average of GDP growth may well approximate the low-frequency cycles of the technology shock. See [Section 4](#) for a concrete example.

DEGREE OF INVERTIBILITY. The identified set for the degree of invertibility at horizon ℓ follows directly from the identified set for α^2 , since

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

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

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

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_ℓ^2 at horizon ℓ equals

$$\left[\underbrace{\tilde{R}_\ell^2}_{\frac{\alpha^2}{\alpha^2 + \sigma_v^2} \times R_\ell^2}, \underbrace{\frac{\text{Var}(\tilde{z}_t)}{2\pi \sup_{\omega \in [0, \pi]} s_{\tilde{z}^\dagger}(\omega)} \times \tilde{R}_\ell^2}_{\frac{1}{2\pi \sup_{\omega \in [0, \pi]} s_{\varepsilon_1^\dagger}(\omega)} \times R_\ell^2} \right]. \quad (16)$$

This identified set implies conditions under which the data allow us to reject invertibility or recoverability with certainty.

Proposition 2. *Assume $\alpha_{LB}^2 > 0$. The identified set for R_0^2 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}_t^\dagger is serially uncorrelated.*

According to [Proposition 2](#), we know for sure that $\varepsilon_{1,t}$ is noninvertible if and only if \tilde{z}_t

Granger causes y_t .²² Moreover, we know for sure that $\varepsilon_{1,t}$ is non-recoverable if and only if \tilde{z}_t^\dagger , defined in (12), is serially correlated at some lag.

VARIANCE DECOMPOSITIONS/RATIOS. 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

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

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).

The identified set for the FVD requires more work. Intuitively, the (point-identified) full forecasting variance $\text{Var}(y_{i,t+\ell} \mid \{y_\tau\}_{-\infty < \tau \leq t})$ conflates pure forecasting uncertainty (which enters the denominator of the FVD) and invertibility-related forecasting uncertainty (which does not). We thus need to bound the contribution of pure forecasting uncertainty. The following proposition summarizes our results.

Proposition 3. *Let there be given a joint spectral density for $w_t = (y_t', \tilde{z}_t)'$ satisfying the assumptions in Proposition 1. Given knowledge of $\alpha \in (\alpha_{LB}, \alpha_{UB}]$, the largest possible value of the forecast variance decomposition $FVD_{i,\ell}$ is 1 (the trivial bound), while the smallest possible value is given by*

$$\frac{\sum_{m=0}^{\ell-1} \text{Cov}(y_{i,t}, \tilde{z}_{t-m})^2}{\sum_{m=0}^{\ell-1} \text{Cov}(y_{i,t}, \tilde{z}_{t-m})^2 + \alpha^2 \text{Var}(\tilde{y}_{i,t+\ell}^{(\alpha)} \mid \{\tilde{y}_\tau^{(\alpha)}\}_{-\infty < \tau \leq t})}. \quad (17)$$

Here $\tilde{y}_t^{(\alpha)} = (\tilde{y}_{1,t}^{(\alpha)}, \dots, \tilde{y}_{n_y,t}^{(\alpha)})'$ denotes a stationary Gaussian time series with spectral density $s_{\tilde{y}^{(\alpha)}}(\omega) = s_y(\omega) - \frac{2\pi}{\alpha^2} s_{y\tilde{z}}(\omega) s_{y\tilde{z}}(\omega)^*$, $\omega \in [0, 2\pi]$. Expression (17) is monotonically decreasing in α , so the overall lower bound on $FVD_{i,\ell}$ is attained by $\alpha = \alpha_{UB}$; in this boundary case we

²²As the upper bound in the identified set (16) for R_0^2 shows, consistency with invertibility in our model requires $R_0^2 = 2\pi \sup_{\omega \in [0, \pi]} s_{\varepsilon_1^\dagger}(\omega)$, which can only hold if future values of y_t do not help predict the current hidden shock $\varepsilon_{1,t}$. This is the case in the static model, so invertibility is untestable in the static model.

can represent $\tilde{y}_t^{(\alpha_{UB})} = y_t - E(y_t \mid \{\tilde{z}_\tau\}_{-\infty < \tau \leq t})$.

The upper bound on the ℓ -period-ahead FVD is always 1, for any $\ell \geq 1$. This is achieved by a model in which all shocks, except the first one, only affect y_t after an ℓ -period delay.

The expression for the lower bound (17) has a simple interpretation. Even if α is known, the denominator $\text{Var}(y_{i,t+\ell} \mid \{\varepsilon_\tau\}_{-\infty < \tau \leq t})$ of the FVD is not identified due to the lack of information about shocks other than the first. Although we can upper-bound this conditional variance by the denominator of the FVR, this upper bound is not sharp. Instead, to maximize the denominator, as much forecasting noise as possible should be of the pure forecasting variety, and not related to noninvertibility. For all shocks except for $\varepsilon_{1,t}$, this is achievable through a Wold decomposition construction (Hannan, 1970, Thm. 2'', p. 158). Given α , we know the contribution of the first shock to y_t ; the residual after removing this contribution has the distribution of $\tilde{y}_t^{(\alpha)}$, as defined in the proposition. If α is not known, the smallest possible value of the lower bound (17) is attained at the largest possible value of α , namely α_{UB} , for which $\varepsilon_{1,t}$ contributes the least to forecasts of y_t .

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 and degree of invertibility. We also discuss identification of historical decompositions. Proposition 3 showed that even point identification of α is insufficient to point-identify the FVD, although a sharp and informative lower bound can be computed.²³

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., there exists an $\bar{\omega} \in [0, \pi]$ such that $s_{\varepsilon_1^\dagger}(\bar{\omega}) = s_{\varepsilon_1}(\cdot) = \frac{1}{2\pi}$. Then $\alpha = \alpha_{LB}$, so the FVR and degree of invertibility are identified. This assumption is not testable. 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 for 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, $n_\varepsilon = n_y$.²⁴ Our proof of Proposition 1 shows how restrictive

²³Stock & Watson (2017) show that the FVD is point-identified if all shocks are assumed invertible.

²⁴If the joint spectrum of y_t and \tilde{z}_t is positive definite at every frequency, then $n_\varepsilon = n_y$ implies that $\Theta(L)^{-1}$

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 $n_\varepsilon = n_y$.

A second set of sufficient conditions for point identification relates to instrument informativeness. If the instrument is perfectly informative, so $\tilde{z}_t = \alpha\varepsilon_{1,t}$, then identification proceeds in accordance with the logic behind local projections (Jordà, 2005; Gorodnichenko & Lee, 2017). Alternatively, point identification obtains if multiple instruments with mutually independent “first-stage disturbances” are available, cf. Section 3.3.

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 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. If the distribution of the data is consistent with the model, the available IVs can be transformed to a single IV that captures all identifying power for the parameters of interest.

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

$$\tilde{z}_t \equiv z_t - E(z_t \mid \{y_\tau, z_\tau\}_{-\infty < \tau < t}) = \alpha\lambda\varepsilon_{1,t} + \Sigma_v^{1/2}v_t. \quad (18)$$

Online Appendix B.1 shows that the testable implication of the multiple-IV model is that the cross-spectrum $s_{yz}(\omega)$ 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.1 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

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.

instrument²⁵

$$\check{z}_t = \frac{1}{\lambda' \text{Var}(\tilde{z}_t)^{-1} \lambda} \lambda' \text{Var}(\tilde{z}_t)^{-1} \tilde{z}_t. \quad (19)$$

Because \check{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.2 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. We show how the width of the identified sets depends on the strength of the instrument and the informativeness of the macro variables for the unknown shock. Given our choice of observables, the model’s monetary policy shock is nearly invertible, so standard SVAR methods would deliver reasonable identification of this shock. In contrast, invertibility is a very poor approximation when identifying the effects of forward guidance (anticipated monetary) shocks or of technology shocks. Nevertheless, our sharp bounds on variance decompositions and the degree of invertibility are informative for all three shocks.

For clarity, we focus entirely on population bounds in this section, assuming the spectral density of the data is known. The econometrician uses our LP-IV techniques and does not exploit the underlying structure of the model.

MODEL. We employ the [Smets & Wouters \(2007\)](#) model. Throughout, we parametrize the model according to the posterior mode estimates of [Smets & Wouters \(2007\)](#).²⁶ 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

²⁵Intuitively, $\check{z}_t \propto E(\varepsilon_{1,t} | \tilde{z}_t)$.

²⁶Our 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>.

interest rate. These macro aggregates are all stationary in the model, so they should be viewed as deviations from trend. The model features seven unobserved shocks, so not all shocks can be invertible. In Online Appendix B.5, we discuss how the analysis would play out with alternative sets of observed macro aggregates.

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 consider three different shocks of interest: a monetary shock, a forward guidance shock, and a technology shock. 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 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 already shown by [Wolf \(2017\)](#). 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 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 ratio and the degree of invertibility are close to the truth, while the lower bounds depend on the informativeness of the IV. The identified set for R_ℓ^2 is given by

$$\left[\frac{1}{1 + \sigma_v^2} \times R_\ell^2, \quad \frac{1}{0.9465^2} \times R_\ell^2 \right].$$

MONETARY SHOCK: SPECTRAL DENSITY OF BEST 2-SIDED LINEAR PREDICTOR




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.

Since $R_0^2, R_\infty^2 < 0.9465^2$, the data reject both invertibility and recoverability. Moreover, the econometrician can rule out that the shock is highly noninvertible, provided the IV is not weak. [Figure 2](#) displays the identified set of the FVR at different forecast horizons.²⁷ 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, whereas the relative informativeness of the lower bounds depends entirely on the strength of the IV.

For FVDs, the lower bound of the identified set also depends on the informativeness of the IV, while the upper bound always equals the trivial value 1. [Figure 3](#) depicts the identified sets for FVDs, omitting the trivial upper bound. The lower bound is now not simply proportional to the true FVD, due to the intricacies of bounding the denominator of the FVD. In this application, the lower bound is nevertheless approximately equal to the true FVD scaled by instrument informativeness $\frac{1}{1+\sigma_v^2}$.

Due to the near-invertibility of the monetary shocks, incorrectly imposing invertibility (or recoverability) would cause the researcher to overstate the forecast variance ratios and historical decompositions by a modest 14 per cent each. As shown above, the data are in fact

²⁷Throughout 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 $\frac{1}{1+\sigma_v^2} = 0.5$.

MONETARY SHOCK: IDENTIFIED SET OF FVDs




Figure 3: Horizon-by-horizon identified sets for FVDs up to 10 quarters. The two lower bounds are for $\frac{1}{1+\sigma_v^2} = 0.25$ (lower dashed line) and $\frac{1}{1+\sigma_v^2} = 0.5$. Upper bound not shown.

sufficiently informative so as to reject both the invertibility and recoverability assumptions. If a researcher instead (also incorrectly) imposes the weaker and untestable assumption that the data are perfectly informative about the shock at *some* frequency, then all the aforementioned objects of interest would be overstated by around 11 per cent ($\frac{1}{\alpha_{LB}^2} \approx 1.11$).

FORWARD GUIDANCE SHOCK. We now augment the model to include forward guidance shocks, a type of news shock. A forward guidance shock is identical to a monetary shock, except it is anticipated two quarters in advance by economic agents.²⁸

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.²⁹

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 4](#).³⁰ 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 decompositions by a factor of $1/0.0792 \approx 13$ (!).³¹ Recoverability-based

²⁸Formally, 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 .

²⁹Another example is sluggish responses to news shocks. The decisions of forward-looking agents today reveal a nonzero but limited amount of information about the shock that is about to materialize in the future (unless highly responsive forward-looking observables are available). As the shock then hits, aggregates respond strongly, and the econometrician learns about the shock. In fact, in models with as many shocks as observables – e.g., the fiscal foresight models of [Leeper et al. \(2013\)](#) – all shocks are recoverable.

³⁰Note that the FVR of the interest rate is much lower than for the standard monetary shock ([Figure 2](#)). The interest rate now *drops* when news arrive, and overall it moves less because the endogenous rule-based response of the monetary authority counteracts the effects of the news prior to the policy impulse.

³¹To 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

FORWARD GUIDANCE SHOCK: IDENTIFIED SET OF FVRs




Figure 4: 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 $\frac{1}{1+\sigma_v^2} = 0.5$.

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; 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 technology shock is far from being recoverable, let alone invertible, with $R_0^2 = 0.2007$ 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 5](#). 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.

would not be proportional to true responses, due to the imposed VAR dynamics ([Stock & Watson, 2017](#)). We state formal results on the bias of SVAR-IV under noninvertibility in [Online Appendix B.3](#).

TECHNOLOGY SHOCK: SPECTRAL DENSITY OF BEST 2-SIDED LINEAR PREDICTOR




Figure 5: 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; a Bayesian implementation is available for the latter. We also discuss a test for 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 $T^{1/2}(\hat{\vartheta} - \vartheta) \xrightarrow{D} N(0, \Omega)$, 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 and of Bayesian inference procedures. 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.7 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 VAR(∞). In practice, we suggest estimating the lag length p by information criteria or likelihood ratio tests (Lütkepohl, 2005, Ch. 4). 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 e_t 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 data are 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 . Given the reduced-form VAR model, 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 (Lütkepohl, 2005, Ch. 2.3 and 3.6). This test has power against all Granger causal alternatives, so it has power against all falsifiable noninvertible alternatives by Proposition 2.³²

CONFIDENCE INTERVALS. We now construct partial identification robust confidence intervals for identified sets and for the true parameters. Here we rely heavily on the inference methods pioneered 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 $[\underline{h}(\vartheta), \bar{h}(\vartheta)]$, where $\underline{h}(\cdot)$ and $\bar{h}(\cdot)$ are continuous functions mapping the VAR parameter space into the real line, and such that $\underline{h}(\cdot) \leq \bar{h}(\cdot)$. A (pointwise) consistent estimator of the identified set $[\underline{h}(\vartheta), \bar{h}(\vartheta)]$ is then given by the plug-in interval

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

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

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

and their correlation,

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

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

We now construct a confidence interval for the entire identified set. The interval

$$\left[\underline{h}(\hat{\vartheta}) - \Phi^{-1}(1 - \beta/2) \hat{\sigma}, \bar{h}(\hat{\vartheta}) + \Phi^{-1}(1 - \beta/2) \hat{\bar{\sigma}} \right] \quad (20)$$

is a (pointwise) asymptotically valid level- $(1 - \beta)$ confidence interval for the identified set

³²Stock & Watson (2017) 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. Giannone & Reichlin (2006) and Forni & Gambetti (2014) propose Granger causality tests of invertibility in an SVAR context.

$[\underline{h}(\vartheta), \bar{h}(\vartheta)]$. That is, the above interval contains the *entire* identified set in at least $100(1 - \beta)\%$ of repeated experiments, asymptotically. This follows from the delta method and the arguments of [Imbens & Manski \(2004\)](#).

Next, we construct a confidence interval for the true parameter of interest. By definition of the identified set, the true parameter is contained in $[\underline{h}(\vartheta), \bar{h}(\vartheta)]$, but we know nothing else about the true parameter. Although the interval (20) 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. As in [Stoye \(2009, p. 1305\)](#), define the two scalars $\hat{\underline{c}}, \hat{\bar{c}}$ as the minimizers of the objective function

$$\hat{\underline{\sigma}} \times \hat{\underline{c}} + \hat{\bar{\sigma}} \times \hat{\bar{c}},$$

subject to the two constraints

$$\begin{aligned} \Pr \left(-\hat{\underline{c}} \leq U_1, \hat{\rho}U_1 \leq \hat{\bar{c}} + \frac{\hat{\Delta}}{\hat{\underline{\sigma}}} + \sqrt{1 - \hat{\rho}^2} \times U_2 \right) &\geq 1 - \beta, \\ \Pr \left(-\hat{\underline{c}} - \frac{\hat{\Delta}}{\hat{\underline{\sigma}}} - \sqrt{1 - \hat{\rho}^2} \times U_2 \leq \hat{\rho}U_1, U_1 \leq \hat{\bar{c}} \right) &\geq 1 - \beta. \end{aligned}$$

Here the probabilities are taken solely over the distribution of $(U_1, U_2)'$, which is bivariate standard normal. The above minimization problem is easy to solve numerically, cf. [Stoye \(2009, Appendix B\)](#). Given these definitions, the interval

$$\left[\underline{h}(\hat{\vartheta}) - \hat{\underline{c}} \times \hat{\underline{\sigma}}, \bar{h}(\hat{\vartheta}) + \hat{\bar{c}} \times \hat{\bar{\sigma}} \right]$$

is a (pointwise) asymptotically valid level- $(1 - \beta)$ confidence interval for the true parameter. Again, this result follows from the delta method and the results in [Stoye \(2009\)](#), who builds on [Imbens & Manski \(2004\)](#).

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 $\bar{h}(\hat{\vartheta})$, and the derivatives of the bounds $\dot{\underline{h}}(\hat{\vartheta})$ and $\dot{\bar{h}}(\hat{\vartheta})$. [Appendix A.1](#) provides formulas for the bounds and derivatives in terms of the VAR parameters. Simple bootstrap and Bayesian implementations are also available, see below.

We now discuss how to resolve the complication that the upper bound of the identified sets for $\frac{1}{\alpha^2}$, R_0^2 , and the FVR may not be continuously differentiable in the VAR parameters.

The issue arises because α_{LB}^2 is given by the maximum of a certain function, cf. (14). When this function has multiple maxima at the true VAR parameters (e.g., when the spectral density of \tilde{z}_t^\dagger is flat, as in the recoverable case), continuous differentiability of α_{LB}^2 in the VAR parameters ϑ may fail (Gafarov et al., 2017). In this case, delta method inference will be unreliable. As a 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 smaller average value $\text{Var}(\tilde{z}_t^\dagger) = \int_0^{2\pi} s_{\tilde{z}^\dagger}(\omega) d\omega$, cf. the inequality (15). The latter object is continuously differentiable in the VAR parameters, so inference using the above methods is unproblematic. Use of the non-sharp bound does lead to a power loss, but the loss is small if the shock $\varepsilon_{1,t}$ is close to being recoverable, or if the informational content of the data for the shock does not vary substantially across frequencies, as explained in Sections 3.2 and 4.³³ Note that continuous differentiability of the bounds for the FVD obtains without modifications.

Our 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.³⁴ 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.

BOOTSTRAP IMPLEMENTATION. The calculation of derivatives in the confidence interval formulas above is obviated by the bootstrap. Suppose we bootstrap the estimator $\hat{\vartheta}$ (Kilian & Lütkepohl, 2017, Ch. 12). Then we can compute $\hat{\sigma}$ as the bootstrap standard deviation of $\underline{h}(\hat{\vartheta})$, $\hat{\sigma}$ as the bootstrap standard deviation of $\bar{h}(\hat{\vartheta})$, and $\hat{\rho}$ as the bootstrap correlation of $\underline{h}(\hat{\vartheta})$ and $\bar{h}(\hat{\vartheta})$. 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.

³³More generally, we can lower-bound α_{LB}^2 by $\int_0^{2\pi} r(\omega) s_{\tilde{z}^\dagger}(\omega) d\omega$, where $r(\cdot)$ is a nonnegative function such that $\int_0^{2\pi} r(\omega) d\omega = 2\pi$. If the researcher has prior information about the frequencies ω at which y_t is particularly informative about $\varepsilon_{1,t}$, then $r(\omega)$ can be chosen to weight these frequencies more heavily. This yields a more informative bound than $\text{Var}(\tilde{z}_t^\dagger)$, while preserving continuous differentiability.

³⁴The Imbens & Manski (2004) and Stoye (2009) procedures are designed to control coverage uniformly over the width of the identified set. 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, in which case the width of the identified set (for all our objects of interest) would also be bounded away from zero. Hence, in this case, the uniform asymptotic validity of the confidence procedures is a trivial matter.

BAYESIAN IMPLEMENTATION. Finally, we discuss a Bayesian credible interval for the identified set. Bayesian inference on the identified set can be motivated by a desire for structural inference to be robust to the choice of prior on structural parameters, conditional on a fixed prior on reduced-form parameters (Giacomini & Kitagawa, 2015). Suppose we form a prior for the reduced-form VAR parameters ϑ , and we have a method for drawing from the posterior distribution of ϑ (e.g., Kilian & Lütkepohl, 2017, Ch. 5). Then we can compute the interval $[E(\underline{h}(\vartheta) \mid \text{data}), E(\overline{h}(\vartheta) \mid \text{data})]$, a Bayesian estimate of the identified set. Moreover, we can compute the shortest interval which covers at least $100(1 - \beta)\%$ of the posterior draws of the intervals $[\underline{h}(\vartheta), \overline{h}(\vartheta)]$. This smallest interval is then a probability- $(1 - \beta)$ credible interval for the identified set (Giacomini & Kitagawa, 2015).

SIMULATION STUDY. Online Appendix B.8 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 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 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 (2016) compute FVDs for a similar specification, but their analysis also assumes an SVAR model. We find that invertibility (i.e., SVAR structure) is rejected by the data. Nevertheless, our LP-IV confidence intervals still allow us to conclude that monetary shocks contribute at most moderately to fluctuations in output growth and inflation, as well as to financial stress in the short run.

Our specification follows Gertler & Karadi (2015) except 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 \(2012\)](#), 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.³⁵ Data are monthly from January 1990 to June 2012.³⁶ The Akaike Information Criterion selects $p = 6$ lags in the reduced-form VAR we use for inference. We employ a homoskedastic recursive residual VAR bootstrap with 10,000 draws.³⁷

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.³⁸ The data are consistent, however, with moderately large values of the degree of invertibility R_0^2 , as well as with very small values. The data are also consistent with small values of the degree of recoverability R_∞^2 .³⁹

The monetary shock accounts for a small to moderate fraction of the forecast variance of output growth and inflation, and of the Excess Bond Premium at short horizons. [Figure 6](#) 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 happen to be similar to the confidence intervals for the parameters themselves in this application. Because we compute Hall’s asymmetric bootstrap confidence interval, which has a built-in bias correction, some of the intervals and bias-adjusted point estimates

³⁵See [Gertler & Karadi \(2015\)](#) for details on the construction of the IV and a discussion of the exclusion restriction. [Nakamura & Steinsson \(2017\)](#) 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.2 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.

³⁶We could have estimated the y_t block of our reduced-form VAR for $(y_t', z_t)'$ on a longer sample, as [Gertler & Karadi \(2015\)](#) do, but we prefer the common sample for conceptual clarity.

³⁷All 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.

³⁸As discussed in [Section 5](#), this amounts to testing whether the IV Granger causes the other variables in the reduced-form VAR. Online Appendix B.6 provides p-values for such tests. [Stock & Watson \(2017\)](#) fail to reject invertibility in the [Gertler & Karadi \(2015\)](#) specification, apparently because they use a smaller lag length of 4 in their Granger causality tests.

³⁹As 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.

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

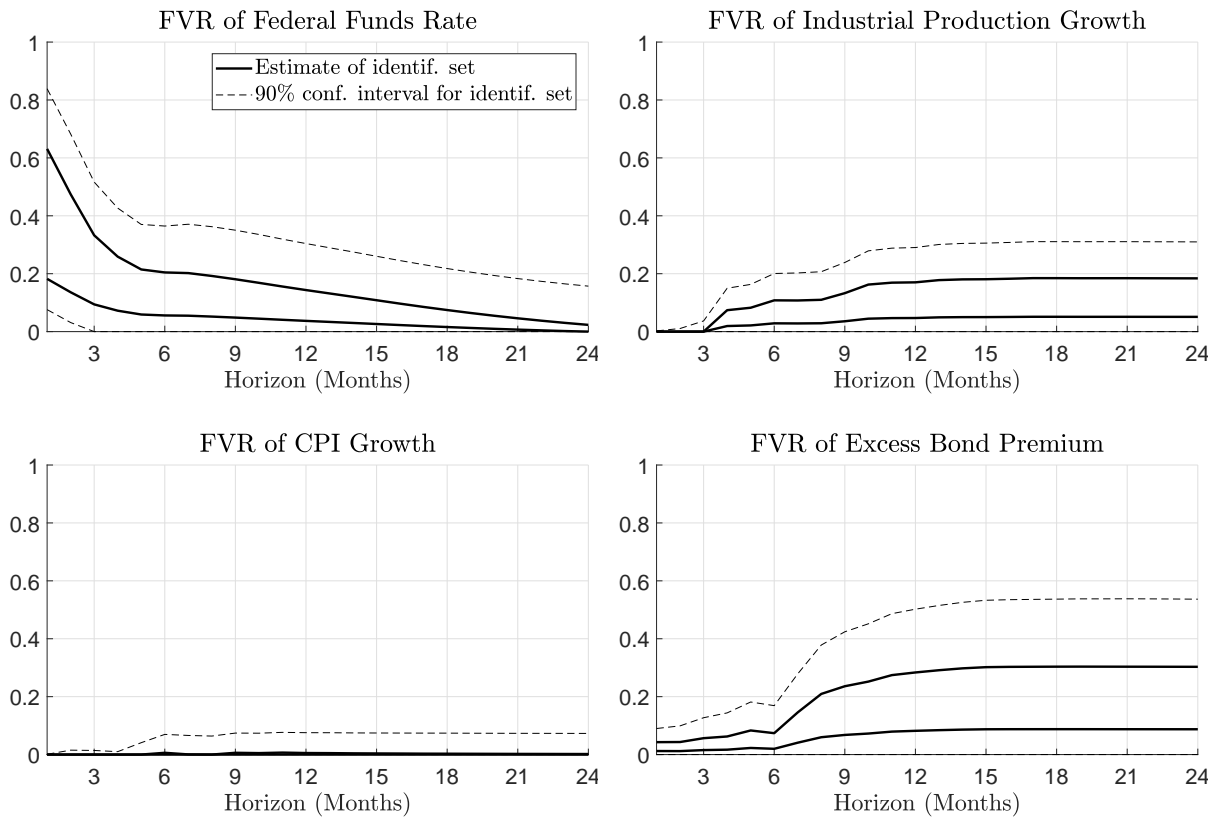


Figure 6: Point estimates and 90% confidence intervals for the identified sets of forecast variance ratios, across different variables and forecast horizons. For visual clarity, we force bias-corrected estimates/bounds to lie in $[0, 1]$.

go negative.⁴⁰ At all forecast horizons, the 90% confidence intervals rule out FVRs above 31% for output growth and 8% for inflation.⁴¹ 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, for essentially all variables and forecast horizons.

The application illustrates that LP-IV inference can be informative about important questions, in addition to being robust to violations of dubious identification assumptions such as invertibility (here rejected by the data) and the availability of a noise-less IV (an untestable assumption). We obtain tight upper bounds on forecast variance ratios, despite the finding by [Stock & Watson \(2017\)](#) that standard errors for LP-IV-estimated impulse response functions are large in this application. Since our analysis arguably weakens the identifying restrictions used in previous papers, it is reassuring that the finding of a modest role for monetary shocks in driving real output and inflation is in line with the existing literature ([Caldara & Herbst, 2016](#); [Ramey, 2016](#), Sec. 3).⁴²

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 forecast 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, bootstrap, or Bayesian inference in reduced-form VARs. The informativeness of our partial identification bounds does not depend on the degree of invertibility of the

⁴⁰This 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.

⁴¹Our upper bounds are even tighter if we run our analysis on the pre-crisis 1990–2006 sample.

⁴²For completeness, Online Appendix B.6 provides forecast variance decompositions implied by an estimated SVAR-IV model, although the latter is rejected by the data.

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., 2016). 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, to simplify the practical econometric procedure, our inference strategy relies on a slightly conservative identification bound that does not exploit the shape of the entire spectrum; it would be interesting to improve on this. Second, one could construct simultaneous (rather than pointwise) confidence bands for, say, forecast variance decompositions at multiple horizons. Third, future research should explore inference issues caused by parameters on the boundary, weak instruments, or near-unit roots. Lastly, our analysis imposed stationarity, but cointegration properties could be relevant for forecast variance decompositions of data in levels.

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 for example the standard VAR formulas in [Lütkepohl \(2005, Ch. 2.1.1\)](#). 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

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

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

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

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

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

BOUNDS FOR α . To compute the bounds for α , we need the quantities

$$\alpha_{UB}^2 = \text{Var}(\tilde{z}_t), \tag{21}$$

$$\alpha_{LB}^2 = 2\pi \max_{\omega \in [0, \pi]} s_{y\tilde{z}}(\omega)^* s_y(\omega)^{-1} s_{y\tilde{z}}(\omega), \quad (22)$$

$$\text{Var}(\tilde{z}_t^\dagger) = 2 \int_0^\pi s_{y\tilde{z}}(\omega)^* s_y(\omega)^{-1} s_{y\tilde{z}}(\omega) d\omega. \quad (23)$$

For the upper bound (21), we have

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

For the lower bound (22),

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

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 (22) by grid search.

Rather than explicitly computing the integral (23), 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 $M \rightarrow \infty$ 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 $(y'_{t+M}, \dots, y'_t, \dots, y'_{t-M})'$, and $\Sigma_{y, (M, M)}$ is the full variance-covariance matrix of $(y'_{t+M}, \dots, y'_t, \dots, y'_{t-M})'$. 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

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

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 $M \rightarrow \infty$ 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 $M \rightarrow \infty$ by

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

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 by⁴³

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

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

BOUNDS FOR FVD. To compute the overall lower bound for the FVD, we need $\text{Var}(\tilde{y}_{i,t+\ell}^{(\alpha_{UB})} \mid \{\tilde{y}_\tau^{(\alpha_{UB})}\}_{-\infty < \tau \leq t})$. As before, we approximate this by $\text{Var}(\tilde{y}_{i,t+\ell}^{(\alpha_{UB})} \mid \{\tilde{y}_\tau^{(\alpha_{UB})}\}_{t-M \leq \tau \leq t})$ 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

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

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.⁴⁴ Derivatives of all parameters with respect to ϑ can be computed by finite differences or automatic differentiation.

⁴³In 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.

⁴⁴Our 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 $x > (\geq) b^*B^{-1}b$.*

Proof. We focus on the semidefiniteness statement. Decompose $B = B^{1/2}B^{1/2*}$ and define $\tilde{b} = B^{-1/2}b$. The statement of the lemma is equivalent with the statement that $I_n - x^{-1}\tilde{b}\tilde{b}^*$ is positive semidefinite if and only if $x \geq \tilde{b}^*\tilde{b}$. Let ν be an arbitrary n -dimensional complex vector satisfying $\nu^*\nu = 1$. Then

$$\nu^* \left(I_n - x^{-1}\tilde{b}\tilde{b}^* \right) \nu = 1 - \frac{\tilde{b}^*\tilde{b}}{x} \cos^2 \left(\theta(\nu, \tilde{b}) \right),$$

where $\theta(\nu, \tilde{b})$ is the angle between ν and \tilde{b} . Evidently, $x^{-1}\tilde{b}^*\tilde{b} \leq 1$ is precisely the condition needed to ensure that the above display is nonnegative for every choice of ν . \square

A.2.2 Proof of **Proposition 1**

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

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

and the corresponding vector lag polynomial

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

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

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

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 that⁴⁵

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

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

$$\begin{aligned} y_t &= \bar{\Theta}_{\bullet,1}(L) \bar{\varepsilon}_{1,t} + \tilde{\Theta}(L) \tilde{\varepsilon}_t, \\ \tilde{z}_t &= \alpha \bar{\varepsilon}_{1,t} + \bar{\sigma}_v \bar{v}_t, \\ (\bar{\varepsilon}_{1,t}, \tilde{\varepsilon}_t, \bar{v}_t)' &\stackrel{i.i.d.}{\sim} N(0, I_{n_y+2}). \end{aligned}$$

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

A.2.3 Proof of Proposition 2

IDENTIFIED SET FOR R_0^2 . If the identified set contains 1, then there must exist an $\bar{\alpha} \in [\alpha_{LB}, \alpha_{UB}]$ and i.i.d., independent standard Gaussian processes $\bar{\varepsilon}_{1,t}$ and \bar{v}_t such that (i) $\tilde{z}_t = \bar{\alpha} \times \bar{\varepsilon}_{1,t} + \bar{v}_t$, (ii) \bar{v}_t is uncorrelated with y_t at all leads and lags, and (iii) $\bar{\varepsilon}_{1,t}$ 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 $\ell \geq 1$,

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

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 $\alpha_{LB}^2 = \text{Var}(E(\tilde{z}_t \mid \{y_{\tau}\}_{-\infty < \tau \leq t}))$, which further yields $\alpha_{LB}^2 = \text{Var}(\tilde{z}_t) \tilde{R}_0^2$. Hence, expression (16) 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 $2\pi \sup_{\omega \in [0, \pi]} s_{\tilde{z}^{\dagger}}(\omega) = \tilde{R}_{\infty}^2 \text{Var}(\tilde{z}_t)$, and the right-hand side equals $\text{Var}(\tilde{z}_t^{\dagger}) = \int_0^{2\pi} s_{\tilde{z}^{\dagger}}(\omega) d\omega$.

⁴⁵We 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).

But we have $\sup_{\omega \in [0, \pi]} s_{z^\dagger}(\omega) = \frac{1}{2\pi} \int_0^{2\pi} s_{z^\dagger}(\omega) d\omega$ if and only if $s_{z^\dagger}(\omega)$ is constant in ω almost everywhere, i.e., \tilde{z}_t^\dagger is white noise. \square

A.2.4 Proof of **Proposition 3**

Please see Online Appendix B.9.2.

References

- Angrist, J. D., Jordà, Ò., & Kuersteiner, G. M. (2017). Semiparametric Estimates of Monetary Policy Effects: String Theory Revisited. *Journal of Business & Economic Statistics*. Forthcoming.
- Barnichon, R. & Brownlees, C. (2017). Impulse Response Estimation by Smooth Local Projections. Manuscript, Universitat Pompeu Fabra.
- Barnichon, R. & Matthes, C. (2017). Functional Approximations of Impulse Responses (FAIR): New Insights into the Asymmetric Effects of Monetary Policy. Manuscript, Federal Reserve Bank of San Francisco.
- Brockwell, P. J. & Davis, R. A. (1991). *Time Series: Theory and Methods* (2nd ed.). Springer Series in Statistics. Springer.
- Caldara, D. & Herbst, E. (2016). Monetary Policy, Real Activity, and Credit Spreads: Evidence from Bayesian Proxy SVARs. Manuscript, Federal Reserve Board of Governors.
- Caldara, D. & Kamps, C. (2017). The Analytics of SVARs: A Unified Framework to Measure Fiscal Multipliers. *Review of Economic Studies*, 84(3), 1015–1040.
- Chahrour, R. & Jurado, K. (2017). Recoverability. Manuscript, Duke University.
- 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. (2016). Reassessing Structural VARs: Beyond the ABCs (and Ds). Manuscript, Universitat Autònoma de Barcelona.
- Gafarov, B., Meier, M., & Montiel Olea, J. L. (2017). Delta-Method Inference for a Class of Set-Identified SVARs. Manuscript, Columbia University.

- 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. (2015). Robust inference about partially identified SVARs. 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.
- Hannan, E. (1970). *Multiple Time Series*. Wiley Series in Probability and Statistics. John Wiley & Sons.
- Imbens, G. W. & Manski, C. F. (2004). Confidence Intervals for Partially Identified Parameters. *Econometrica*, 72(6), 1845–1857.
- 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. (2017). Large and state-dependent effects of quasi-random monetary experiments. CEPR Discussion Paper DP11801.
- 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.
- Klepper, S. & Leamer, E. E. (1984). Consistent Sets of Estimates for Regressions with Errors in All Variables. *Econometrica*, 52(1), 163–183.
- 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.

- Lütkepohl, H. (2005). *New Introduction to Multiple Time Series Analysis*. Springer.
- 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.
- Moon, H. R., Schorfheide, F., & Granziera, E. (2013). Inference for VARs Identified with Sign Restrictions. Manuscript, University of Pennsylvania.
- Nakamura, E. & Steinsson, J. (2017). High Frequency Identification of Monetary Non-Neutrality: The Information Effect. *Quarterly Journal of Economics*. Forthcoming.
- Plagborg-Møller, M. (2017). Bayesian Inference on Structural Impulse Response Functions. Manuscript, Princeton University.
- Ramey, V. & Zubairy, S. (2017). Government Spending Multipliers in Good Times and in Bad: Evidence from U.S. Historical Data. *Journal of Political Economy*. Forthcoming.
- 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.
- 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. (2016). Factor Models and Structural Vector Autoregressions in Macroeconomics. In J. B. Taylor & H. Uhlig (Eds.), *Handbook of Macroeconomics*, volume 2 chapter 8, (pp. 415–525). Elsevier.

- Stock, J. H. & Watson, M. W. (2017). Identification and Estimation of Dynamic Causal Effects in Macroeconomics. Manuscript based on 2017 Sargan Lecture, Royal Economic Society.
- Stoye, J. (2009). More on Confidence Intervals for Partially Identified Parameters. *Econometrica*, 77(4), 1299–1315.
- Wolf, C. K. (2017). Masquerading Shocks in Sign-Restricted VARs. Manuscript, Princeton University.