Identification and Estimation of Dynamic Causal Effects in Macroeconomics

June 8, 2017

James H. Stock
Department of Economics, Harvard University
and the National Bureau of Economic Research

and

Mark W. Watson*
Department of Economics and the Woodrow Wilson School, Princeton University
and the National Bureau of Economic Research

Abstract

An exciting development in empirical macroeconometrics is the increasing use of external sources of as-if randomness to identify the dynamic causal effects of macroeconomic shocks. This approach – the use of external instruments – is the dynamic, macroeconometric counterpart of the highly successful strategy in microeconometrics of using external as-if randomness to provide instruments that identify causal effects. This lecture provides conditions on instruments and control variables under which external instrument methods produce valid inference on dynamic causal effects, that is, structural impulse response function; these conditions can help guide the search for valid instruments in applications. We consider two methods, a one-step instrumental variables regression and a two-step method that entails estimation of a vector autoregression, using the same instrument. Under a restrictive instrument validity condition, the one-step method is valid even if the vector autoregression is not invertible, so comparing the two estimates provides a test of invertibility. Under a less restrictive condition, in which multiple lagged endogenous variables are needed as control variables in the one-step method, the conditions for validity of the two methods are the same.

*This work was presented by James Stock as the Sargan Lecture to the Royal Economic Society on April 11, 2017. We thank Daniel Lewis, Mikkel Plagborg-Møller, José Luis Montiel Olea, Valerie Ramey, and Christian Wolf for helpful comments and discussion.

1. Introduction

The identification and estimation of dynamic causal effects is a defining challenge of macroeconometrics. In the macroeconomic tradition dating to Slutzky (1927) and Frisch (1933), dynamic causal effects are conceived as the effect, over time, of an intervention that propagates through the economy, as modeled by a system of simultaneous equations. Restrictions on that system can be used to identify its parameters.

In a classic result by the namesake of this lecture, Denis Sargan (1964) (along with Rothenberg and Leenders (1964)) showed that full information maximum likelihood estimation, subject to identifying restrictions, is asymptotically equivalent to instrumental variables (IV) estimation by three stage least squares. The three stage least squares instruments are obtained from restrictions on the system, typically that some variables and/or their lags enter some equations but not others, and thus are *internal* instruments – they are internal to the system. The massive modern literature since Sims (1980) on point-identified structural vector autoregressions (SVARs) descends from this tradition, and nearly all the papers in that literature can be interpreted as achieving identification through internal instruments.

In contrast, modern microeconometric identification strategies exploit *external* sources of variation that provide quasi-experiments to identify causal effects. Such external variation might be found, for example, in institutional idiosyncracies that introduce as-if randomness in the variable of interest (the treatment). The use of such external instruments has proven highly successful and has yielded compelling estimates of causal effects.

The subject of this lecture is the use of external instruments – instruments not in the VAR – to estimate dynamic causal effects in macroeconomics. The instruments can be used to estimate dynamic causal effects directly without an intervening VAR step; this methods uses an IV version of what is called in the forecasting literature a direct multistep forecasting regression. Alternatively, the instruments can be used in conjunction with a VAR to identify SVAR impulse response functions; this is the IV version of an iterated multistep forecast.

The use of external instruments has opened a new and rapidly growing research program in macroeconomics, in which credible identification is obtained using as-if random variation in the shock of interest that is distinct from – external to – the macroeconomic shocks hitting the economy. As in the microeconometric setting, finding such instruments is not easy. Still, in our view this research program holds out the potential for more credible identification than is typically provided by SVARs identified using internal restrictions.

This lecture unifies and explicates a number of strands of recent work on external instruments in macroeconometrics, and also makes several new contributions. The method of external instruments for SVAR identification (SVAR-IV) was introduced by Stock (2008), and has been used by Stock and Watson (2012), Mertens and Ravn (2013), Gertler and Koradi (2015), Caldara and Kamps (2017), and a growing list of other researchers. The modern use of external instruments to estimate structural impulse response functions directly (that is, without estimating a VAR step) dates to independent contributions by Jordà,Schularick, and Taylor (2015) and Ramey and Zubairy (2017), and is clearly exposited in Ramey (2016). The condition for instrument validity in the direct regression without control variables, given in Section 2, appears in unpublished lecture notes by Mertens (2015). Ramey (2016) calls these direct IV regressions "local projections-IV" (LP-IV) in reference to Jordà's (2005) method of local projections (LP) on which it builds, and we adopt Ramey's (2016) terminology while noting that these IV regressions emerge from the much older tradition of simultaneous equations estimation in macroeconomics. Although these methods are increasingly being used in applications, we are not aware of a unified presentation of the econometric theory and theoretical connections between the SVAR-IV and LP-IV methods.

In addition to expositing the use of external instruments in macroeconomics, this lecture makes five contributions to this literature.

First, we provide conditions for instrument validity for LP-IV, and show that under those conditions LP-IV can estimate dynamic causal effects without assuming invertibility, that is, without assuming that the structural shocks can be accurately recovered from current and lagged values of the observed data. Because of the dynamic nature of the macroeconometric problem, exogeneity of the instrument entails a strong "lag exogeneity" requirement that the instrument be uncorrelated with past shocks, at least after including control variables. This condition provides concrete guidance for the construction of instruments, and choice of control variables, when undertaking LP-IV.

Second, we recapitulate how IV estimation can be undertaken in a SVAR (the SVAR-IV method). This method is more efficient asymptotically than LP-IV under strong-instrument asymptotics, and it does not require lag exogeneity. But to be valid, this method requires invertibility, that is, the space of the VAR innovations (errors) spans the space of the structural shocks. Invertibility is a very strong, albeit commonly made, assumption: under invertibility, a forecaster using a VAR would find no value in augmenting her system with data on the true macroeconomic shocks, were they magically to become available.

Third, having a more efficient estimator (SVAR-IV) that requires invertibility for consistency, and a less efficient estimator (LP-IV) that does not, gives rise to a Hausman (1978) -type test for whether the SVAR is invertible. We provide this test statistic, obtain its large-sample null distribution, introduce the concept of local non-invertibility, and derive its local asymptotic power against the alternative of local non-invertibility. This test differs conceptually from existing tests for invertibility, which examine the no-omitted-variables implication by adding variables, see for example Forni and Gambetti (2014).

Fourth, lest one think that LP-IV is too good to be true, we provide a "no free lunch" result. Suppose an instrument satisfies a contemporaneous exogeneity condition, but not the no-lag exogeneity condition because it is correlated with lagged shocks. A natural approach would be to include variables that capture the lagged shocks as control variables. We show, however, that the condition for these control variables to produce valid inference in LP-IV is equivalent to assuming invertibility, in which case SVAR-IV provides more efficient inference.

Fifth, we discuss some econometric odds and ends, such as heteroskedasticity- and autocorrelation-robust (HAR) standard errors, what to do if the external instruments are weak, estimation of cumulative dynamic effects, and the pros and cons of using generic controls including factors (factor-augmented LP-IV).

Following Ramey (2016), we illustrate these methods using Gertler and Koradi's (2015) application, in which they estimate the dynamic causal effect of a monetary policy shock using SVAR-IV, with an instrument that captures the news revealed in regularly scheduled monetary policy announcements by the Federal Open Market Committee.

Before proceeding, we note two simplifications made throughout this lecture. First, we focus exclusively on linear models, so that conditional expectations are replaced by projections. Extension to nonlinearities and conditional expectations is straightforward and secondary for the points being made here. Second, we assume homogenous treatment effects. Doing so has the nontrivial implication that valid instruments all have the same estimand. We return to the issue of heterogeneous treatment effects in the conclusion. In addition, we use two notational devices: a

"•" subscript denotes all elements of a vector or matrix other than the first row or column, and {…} denotes linear combinations of the terms inside the braces.

2. Identifying Dynamic Causal Effects using External Instruments and Local Projections

The LP-IV method emerges naturally from the modern microeconometrics use of instrumental variables. Making this connection requires some translation between two sets of jargon, however, so we start with a brief review of causal effects and instrumental variables regression in the microeconometric setting.

2.1 Causal effects and instrumental variables regression

The theoretical foundations of causal inference are a staple of modern microeconometrics; see for example Imbens (2014) for a review. In brief, if a binary treatment X is randomly assigned, then all other determinants of Y are independent of X, which implies that the (average) treatment effect is $E(Y|X=1) - E(Y|X=0)$. In the regression model $Y = \gamma + X\beta + u$, random assignment implies that $E(u|X) = 0$ and the regression coefficient β is the treatment effect. If randomization is conditional on covariates W, then the treatment effect for an individual with covariates $W=w$ is estimated by the outcome of a random experiment on a group of subjects with the same value of W, that is, it is $E(Y|X=1,W=w) - E(Y|X=0,W=w)$. With the additional assumption of linearity, this treatment effect is estimated by ordinary least squares estimation of

$$Y = \beta X + \gamma' W + u, \tag{1}$$

where the intercept has been subsumed in $\gamma' W$. This basic idea that a randomized controlled experiment estimates the average treatment effect lies at the core of statistical definitions of causality.

In observational data, the treatment level X is often endogenous. This is generally the case when the subject has some control over receiving the treatment in an experiment. But if there is some source of variation Z that is correlated with treatment, such as random assignment to the treatment or control group, conditional on observed covariates W, then the causal effect can be estimated by instrumental variables. Let "$\perp$" denotes the residual from the population projection onto W, for example $X^\perp = X - \text{Proj}(X \mid W)$. If the instrument satisfies the conditions

$$\text{(i) } E(X^\perp Z^\perp) \neq 0 \text{ (relevance)} \tag{2}$$

$$\text{(ii) } E(u^\perp Z^\perp) = 0 \text{ (exogeneity),} \tag{3}$$

and if the instruments are strong, then instrumental variables estimation of (1) consistently estimates the causal effect β.

2.2 Dynamic causal effects and the structural moving average model[1]

In macroeconomics, we can imagine a counterpart of randomized controlled experiment. For example, in the United States, the Federal Open Market Committee (FOMC) could set the Federal Funds rate according to a rule, such as the Taylor rule, perturbed by a randomly chosen amount. Although we have only one subject (the U.S. macroeconomy), by repeating this experiment through time, the FOMC could generate data on the effect of these random interventions.

More generally, let $\varepsilon_{1,t}$ denote the mean-zero random treatment at date t. Then the causal effect on the value of a variable Y_2, h periods hence, of a unit intervention in ε_1 is $E\left(Y_{2,t+h} \mid \varepsilon_{1,t} = 1\right) - E\left(Y_{2,t+h} \mid \varepsilon_{1,t} = 0\right)$. We denote this treatment effect $\Theta_{h,21}$, the effect of treatment 1 on variable 2, h periods after the treatment.

Assuming linearity (as we do throughout), this treatment effect is the coefficient in the regression,

$$Y_{2,t+h} = \Theta_{h,21}\varepsilon_{1,t} + u_{t+h}. \tag{4}$$

Because $\varepsilon_{1,t}$ is randomly assigned, it is independent of other factors comprising u_{t+h}. Thus, were $\varepsilon_{1,t}$ observed, this causal effect could be estimated by OLS estimation of (4).

[1] See Lechner (2009), Angrist, Jordà, and Kuersteiner (2017) and Bojinov and Shephard (2017) for discussion (and references) of identification of dynamic causal effects from primitive assumptions in the potential outcomes context.

The macroeconometric jargon for this random treatment $\varepsilon_{1,t}$ is a *structural shock*: a primitive, unanticipated economic force, or driving impulse, that is uncorrelated with other shocks.[3] The macroeconomist's shock can be thought of as the microeconomists' random treatment, and impulse response functions are the causal effects of those treatments on variables of interest over time, that is, dynamic causal effects.

The path of observed macroeconomic variables can be thought of as arising solely from current and past shocks and measurement error. If we collect all such structural shocks and measurement error together in the vector ε_t, and if we assume linearity and stationarity, the vector of macroeconomic variables Y_t can be written in terms of current and past ε_t:

$$Y_t = \Theta(\mathrm{L})\varepsilon_t \qquad (5)$$

where L is the lag operator and $\Theta(\mathrm{L}) = \Theta_0 + \Theta_1\mathrm{L} + \Theta_2\mathrm{L}^2 + \ldots$, where Θ_h is an $n_Y \times n_\varepsilon$ matrix of coefficients. The shock variance matrix $\Sigma_\varepsilon = E\varepsilon_t\varepsilon_t'$ is assumed to be positive definite to rule out trivial (non-varying) shocks.[4] Throughout, we treat Y_t as having been transformed so that it is second order stationary, for example real activity variables would appear in growth rates, and for notational convenience we suppress the contant term in (5).[5]

Representation (5) is the structural moving average representation of Y_t. The coefficients of $\Theta(\mathrm{L})$ are the structural impulse response functions, which are the dynamic causal effects of the shocks. In general, the number of shocks plus measurement error terms, n_ε, can exceed the number of observed variables, n_Y.

The recognition that, if $\varepsilon_{1,t}$ were observed, $\Theta_{h,21}$ could be estimated by OLS estimation of (4), has led to a productive and insightful research program. In this program, which dates to

[3] For an extensive discussion, see Ramey (2016).

[4] We have defined ε_t to include both structural shocks and measurement error, in that order. Then Σ_ε is block diagonal, with the block corresponding to the shocks being diagonal and the block corresponding to the measurement error being positive definite. In principle, measurement error can be correlated across observable variables if, for example, they are computed in part from the same survey.

[5] Because $\varepsilon_{1,t}$ is randomly assigned, it is independent of all other shocks in the economy, the causal effect can be written as $E\left(Y_{2,t+h} \mid \varepsilon_{1,t} = 1, \varepsilon_{\bullet,t}, \varepsilon_s, s \neq t\right) - E\left(Y_{2,t+h} \mid \varepsilon_{1,t} = 0, \varepsilon_{\bullet,t}, \varepsilon_s, s \neq t\right)$. Although conditioning on the other shocks is redundant by randomization, this alternative expression connects with the definition of a causal effect with the partial derivative $\partial Y_{2,t+h}/\partial\varepsilon_{1,t}$ from (5).

Romer and Romer (1989), researchers aim to measure directly a specific macroeconomic shock. Influential examples include Kuttner (2001), Cochrane and Piazzesi (2002), and Faust, Rogers, Swanson, and Wright (2003), Gürkaynak, Sack, and Swanson (2005), and Bernanke and Kuttner (2005), all of whom used interest rate changes around Federal Reserve announcement dates to measure monetary policy shocks.

2.3 Direct estimation of structural IRFs using external instruments (LP-IV)

One difficulty with directly measured shocks is that they capture only part of the shock, or are measured with error. For example, Kuttner (2001)-type variables measure that part of a shock revealed in a monetary policy announcement but not the part revealed, for example, in speeches by FOMC members. This concern applies to other examples, including Romer and Romer's (1989) binary indicators, Romer and Romer's (2010) measure of exogenous changes in fiscal policy, and Kilian's (2008) list of exogenous oil supply disruptions. In all these cases, the constructed variable is correlated with the true (unobserved) shock and, if the author's argument for exogeneity is correct, the constructed variable is uncorrelated with other shocks. That is, the constructed variable is not the shock, but is an instrument for the shock. This instrument is not obtained from restrictions internal to a VAR (or some other dynamic simultaneous equations model); rather, it is an external instrument.

This reasoning suggests using instrumental variables methods to estimate the dynamic causal effects of the shock. To do so, however, requires resolving a difficulty not normally encountered in microeconometrics, which is that the shock/treatment $\varepsilon_{1,t}$ is unobserved. As a result, the scale of $\varepsilon_{1,t}$ is indeterminate, that is, (4) holds for all h if $\varepsilon_{1,t}$ is replaced by $c\varepsilon_{1,t}$ and $\Theta_{h,21}$ is replaced by $c^{-1}\Theta_{h,21}$. This scale ambiguity is resolved by adopting, without loss of generality, a normalization for the scale of $\varepsilon_{1,t}$. Specifically, we assume that $\varepsilon_{1,t}$ is such that a unit increase in $\varepsilon_{1,t}$ increases $Y_{1,t}$ by one unit:

$$\Theta_{0,11} = 1 \text{ (unit effect normalization).} \tag{6}$$

For example, if $\varepsilon_{1,t}$ is the monetary policy shock and $Y_{1,t}$ is the federal funds rate, (6) sets the scale of $\varepsilon_{1,t}$ so that a 1 percentage point monetary policy shock increases the federal funds rate by 1 percentage point.[6]

With the unit effect normalization in hand, the regression (4) can be rewritten with an observable regressor, $Y_{1,t}$. Specifically, use the unit effect normalization to write $Y_{1,t} = \varepsilon_{1,t} + \{\varepsilon_{\bullet,t}, \varepsilon_{t-1}, \varepsilon_{t-2}, \ldots\}$, where we use the "•" notational devices, $\varepsilon_{\bullet,t} = (\varepsilon_{2,t}, \ldots, \varepsilon_{n_\varepsilon,t})'$, and $\{\ldots\}$ to denote linear combinations of the terms in braces. Rewriting this expression in terms of $\varepsilon_{1,t}$ and substituting it into (4) yields,

$$Y_{i,t+h} = \Theta_{h,i1} Y_{1,t} + u^h_{i,t+h} \tag{7}$$

where $u^h_{i,t+h} = \{\varepsilon_{t+h}, \ldots, \varepsilon_{t+1}, \varepsilon_{\bullet,t}, \varepsilon_{t-1}, \varepsilon_{t-2}, \ldots\}$. Because $Y_{1,t}$ is endogenous, OLS estimation of (7) is not valid. But, with a suitable instrument, (7) can be estimated by IV.

Let Z_t be a vector of n_Z instrumental variables. These instruments can be used to estimate the dynamic causal effect using (7) if they satisfy:

Condition LP-IV

(i) $E\left(\varepsilon_{1,t} Z_t'\right) = \alpha' \neq 0$ (relevance)

(ii) $E\left(\varepsilon_{\bullet,t} Z_t'\right) = 0$ (contemporaneous exogeneity)

(iii) $E\left(\varepsilon_{t+j} Z_t'\right) = 0$ for $j \neq 0$ (lead/lag exogeneity).

[6] The unit effect normalization has several advantages over the more familiar unit standard deviation normalization, which sets $\mathrm{var}(\varepsilon_{1,t}) = 1$. Most importantly, the unit effect normalization allows for direct estimation of the dynamic causal effect in the native units, which are the relevant units for policy analysis. While one can convert one scale normalization to another, doing so entails rescaling by estimated values and care must be taken to conduct inference incorporating that normalization (e.g., bootstrapping then rescaling differs from bootstrapping the rescaled estimate). The unit effect normalization is the natural one for the IV methods here, and as discussed in Stock and Watson (2016), it allows for direct extension of SVAR methods do structural dynamic factor models.

Conditions LP-IV (i) and (ii) are conventional IV relevance and exogeneity conditions, and are the counterparts of the microeconometric conditions (2) and (3) in the absence of control variables.

Condition LP-IV (iii) arises because of the dynamics. The key idea of this condition is that $Y_{2,t+h}$ generally depends on the entire history of the shocks, so if Z_t is to identify the effect of shock $\varepsilon_{1,t}$ alone, it must be uncorrelated with all shocks at all leads and lags. The requirement that Z_t be uncorrelated with future ε's is not restrictive: when Z_t contains only variables realized at date t or earlier, it follows from the definition of shocks as unanticipated structural disturbances. In contrast, the requirement that Z_t be uncorrelated with past ε's is restrictive and strong.

We will refer to Condition LP-IV (iii) as requiring that Z_t be unpredictable given past ε's, although strictly the requirement is that it not be linearly predictable given past ε's. Note that Z_t could be serially correlated yet satisfy this condition. For example, suppose $Z_t = \delta\varepsilon_{1,t} + \zeta_t$, where ζ_t is serially correlated measurement error that is independent of $\{\varepsilon_t\}$; then Z_t satisfies Condition LP-IV.

The IV estimator of $\Theta_{h,i1}$ obtains by noting two implications of the assumptions. First, condition LP-IV and equation (5) imply that $\mathrm{E}(Y_{i,t+h}Z_t') = \Theta_{h,i1}\alpha'$. Second, condition LP-IV, the unit effect normalization (6), and equation (5) imply that $\mathrm{E}(Y_{1,t}Z_t') = \alpha'$. Thus when Z_t is a scalar,

$$\frac{\mathrm{E}(Y_{i,t+h}Z_t)}{\mathrm{E}(Y_{1,t}Z_t)} = \Theta_{h,i1}\,. \tag{8}$$

For a vector of instruments, $\mathrm{E}(Y_{2,t+h}Z_t')\Lambda\,\mathrm{E}(Z_tY_{1,t})\big/\mathrm{E}(Y_{1,t}Z_t')\Lambda\,\mathrm{E}(Z_tY_{1,t}) = \Theta_{h,i1}$ for any positive definite matrix Λ. These are the moment expressions for IV estimation of (7) using the instrument Z_t.

2.4. Extension of LP-IV to Control Variables

There are two reasons to consider adding control variables to the IV regression (7).

First, although an instrument might not satisfy Condition LP-IV, it might do so after including suitable control variables; that is, the instruments might satisfy the exogeneity conditions only after controlling for some observable factors. As discussed in Section 5, this is the case in the Gertler-Karadi (2015) application.

Second, including control variables could reduce the sampling variance of the IV estimator by reducing the variance of the error term, even if Condition LP-IV is satisfied.

Recalling the notation $x_t^{\perp} = x_t - \text{Proj}(x_t \mid W_t)$ for some variable x_t, adding control variables W_t to (7) yields,

$$Y_{i,t+h} = \Theta_{h,i1} Y_{1,t} + \gamma_h{}' W_t + u_{i,t+h}^{h\perp} . \tag{9}$$

With control variables W, the conditions for instrument validity are,

<u>Condition LP-IV$^{\perp}$</u>

(i) $E\left(\varepsilon_{1,t}^{\perp} Z_t^{\perp\prime}\right) = \alpha' \neq 0$

(ii) $E\left(\varepsilon_{\bullet,t}^{\perp} Z_t^{\perp\prime}\right) = 0$

(iii) $E\left(\varepsilon_{t+j}^{\perp} Z_t^{\perp\prime}\right) = 0$ for $j \neq 0$.

Under Condition LP-IV$^{\perp}$ and the unit effect normalization (6), when Z_t is a scalar,

$$\frac{\mathrm{E}(Y_{i,t+h}^{\perp} Z_t^{\perp})}{\mathrm{E}(Y_{1,t}^{\perp} Z_t^{\perp})} = \Theta_{h,i1} , \tag{10}$$

with the extension following (8) (modified for "$^{\perp}$" variables) when there are multiple instruments. Equation (10) is the moment condition for IV estimation of (9) using instrument Z_t.

The question of what control variables to include, if any, is a critical one that depends on the application.

Even if condition LP-IV (iii) holds, including control variables could reduce the variance of the regression error and thus improve estimator efficiency. This suggests using control variables aimed at capturing some of the dynamics of $Y_{1,t}$ and $Y_{2,t}$. Such control variables could include lagged values of Y_1 and Y_2, or additionally lagged values of other macro variables. Such control variables could also include generic controls, such as lagged factors from a dynamic factor model.

The more difficult problem arises if Conditions LP-IV (i) and (ii) hold, but Condition LP-IV (iii) fails because Z_t is correlated with one or more lagged shocks. Then instrument validity hinges upon including in W variables that control for those lagged shocks, so that Condition LP-IV$^{\perp}$ (iii) holds. It is useful to think of two cases.

In the first case, suppose Z_t is correlated with past values of $\varepsilon_{1,t}$, but not with past values of other shocks. As we discuss below, this situation arises in the Gertler-Karadi (2015) application, where the construction of Z_t induces a first-order moving average structure. In this case, including lagged values of Z as controls would be appropriate. Another example might be oil supply disruptions arising from political disturbances as in Hamilton (2003) and Kilian (2008), where the onset of the disruption might plausibly be unpredictable using lagged ε's, but the disruption indicator could exhibit time series correlation because any given disruption could last more than one period. If so, it could be appropriate to include lagged values of Z as controls, or otherwise to modify the instrument so that it satisfies condition LP-IV$^{\perp}$ (iii).

A second case arises when Z_t is correlated with past shocks including those other than $\varepsilon_{1,t}$. If so, instrument validity given the controls requires that the controls span the space of those shocks. If it were known which past shocks were correlated with Z, then application-specific reasoning could guide the choice of controls, akin to the first case. But without such information, the controls would need to span the space of all past shocks. This reasoning suggests using generic controls. One such set of generic controls would be a vector of macro variables, say Y_t. Another such set could be factors estimated from a dynamic factor model; using such factors would provide a factor-augmented IV estimate of the structural impulse response function. We show in Section 3.2 that the requirement that these generic controls lead to Condition LP-IV$^{\perp}$ (iii) is quite strong.

2.4 LP-IV: Econometric odds and ends

Levels, differences, and cumulated impulse responses. In many applications, $Y_{i,t}$ will be specified in first differences, but interest is on impulse responses for its levels. Impulse responses for levels are cumulated impulse responses for first differences. The cumulated impulse responses can be computed from the IV regression,

$$(Y_{i,t+h} + Y_{i,t+h-1} + \ldots + Y_{i,t}) = \Theta^{cum}_{h,i1} Y_{1,t} + \gamma^{cum\prime}_h W_t + u^{h,cum\perp}_{i,t+h} \tag{11}$$

where $\Theta^{cum}_{h,i1} = \sum_{s=0}^{h} \Theta_{s,i1}$. If Z_t satisfies LP-IV$^{\perp}$, it is a valid instrument for IV estimation of (11).

HAC/HAR inference and long-horizon impulse reponses. When the instruments are strong, the validity of inference can be justified under standard assumptions of stationarity, weak dependence, and existence of moments (see for example Hayashi (2000)). However, the multistep nature of the direct regressions requires an adjustment for serial correlation of the instrument×error process: The error terms in (7), (9), and (11) include future and lagged values of ε_t, and in general terms like $Z_t\varepsilon_{t+j}$ and $Z_{t+j}\varepsilon_t$ will be correlated. Inference based on standard heteroskedasticity- and autocorrelation robust (HAR) covariance matrix estimators are valid at short to medium horizons.

Smoothness restrictions. The IV estimator of (7), (9), and (11) impose no restrictions across the values of the dynamic causal effects for different horizons. In many applications, smoothness across horizons is sensible. The VAR methods discussed in the next section impose smoothness by modeling the structural moving average (5) as the inverse of a low-order VAR, however as is discussed in that section those methods require the additional assumption that Θ(L) is invertible. Alternatively, it is possible to impose smoothness restrictions directly on the coefficients in the LP-IV regressions, see Plagborg-Møller (2016a) and Barnichon and Brownless (2017).

Weak instruments. If the instruments are weak, then in general distribution of the IV estimator in (7), (9), and (11) is not centered at $\Theta_{h,i1}$, and inference based on conventional IV standard errors is unreliable. However, a suite of heteroskedasticity- and autocorrelation-robust methods now exists to detect weak instruments and to conduct inference robust to weak instruments in linear IV regression; see for example Kleibergen (2005) for a HAR version of

Moreira's (2003) conditional likelihood ratio statistics, and Andrews (2017) and Montiel Olea and Pflueger (2013) for HAR alternatives to first-stage F statistics for detecting weak identification.

News shocks and the unit-effect normalization. In some applications interest focuses on a "news shock," which is defined to be a shock that is revealed at time t, but has a delayed effect on its natural indicator. For example, Ramey (2011) argues that many fiscal shocks are news shocks because they are revealed during the legislature process but have direct effects on government spending only with a lag. Despite this lag, forward looking variables, like consumption, investment, prices, and interest rates may respond immediately to the shock. This differential timing changes the scale normalization for the shock because $\Theta_{0,11}$ may equal zero; that is, the news shock $\varepsilon_{1,t}$ affects its indicator $Y_{1,t}$ only with a lag. Thus, the contemporaneous unit-effect normalization ($\Theta_{0,11} = 1$) is inappropriate.

Instead, for a news shock, a k-period ahead unit-effect normalization, $\Theta_{k,11} = 1$ for pre-specified k, should be used. For example, if government spending reacts to news about spending with a 12-month lag, then the 12-month-ahead unit-effect normalization $\Theta_{12,11} = 1$ would be appropriate. With this k-period ahead normalization, $Y_{1,t+k} = \varepsilon_{1,t} + \{\varepsilon_{t+k}, \ldots, \varepsilon_{t+1}, \varepsilon_{\cdot t}, \varepsilon_{t-1}, \ldots\}$. Accordingly, $Y_{1,t+k}$ replaces $Y_{1,t}$ in the IV regressions (7), (9), and (11). In practice, implementing this strategy requires a choice of the news lead-time k, and this choice would be informed by application-specific knowledge.

3. Identifying Dynamic Causal Effects using External Instruments and VARs

Since Sims (1980), the standard approach in macroeconomics to estimation of the structural moving average representation (5) has been to estimate a structural vector autoregression (SVAR), then to invert the SVAR to estimate $\Theta(L)$. This approach has several virtues. Macroeconomists are in general interested in responses to multiple shocks, and the SVAR approach provides estimates of the full system of responses. It emerges from the long tradition, dating from the Cowles Commission, of simultaneous equation modeling of time series variables. It imposes parametric restrictions on the high-dimensional moving average representation that, if correct, can improve estimation efficiency. And, importantly, it replaces the computationally difficult problem of estimating a multivariate moving average with the straightforward task of single-equation estimation by OLS.

These many advantages come with two requirements. The first is that the researcher has some scheme to identify the relation between the VAR innovations and the structural shocks, assuming that the two span the same space; this is generally known as the SVAR identification problem. The second is that, in fact, this spanning condition holds, a condition that is generally referred to as invertibility. Here, we begin by discussing how IV methods can be used to solve the thorny SVAR identification problem. We then turn to a discussion of invertibility, which we interpret as an omitted variable problem.

3.1. SVAR-IV

A vector autoregression expresses Y_t as its projection on its past values, plus an innovation v_t that is linearly unpredictable from its past:

$$\mathrm{A(L)}Y_t = v_t, \tag{12}$$

where $\mathrm{A(L)} = \mathrm{I} - \mathrm{A}_1\mathrm{L} - \mathrm{A}_2\mathrm{L}^2 - \dots$. We assume that the VAR innovations have a non-singular covariance matrix (otherwise a linear combination of Y could be perfectly predicted). Because the construction of $v_t = Y_t - \mathrm{Proj}(Y_t|Y_{t-1}, Y_{t-2},\ldots)$ is the first step in the proof of the Wold decomposition, the innovations are also called the Wold errors.

In a structural VAR, the innovations are assumed to be linear combinations of the shocks and, moreover, the spaces spanned by the innovations and the structural shocks are assumed to coincide:

$$v_t = \Theta_0\varepsilon_t \quad \text{where } \Theta_0 \text{ is nonsingular .} \tag{13}$$

Because Y_t is by assumption second order stationary, A(L) is invertible. Thus (12) and (13) yield a moving average representation in terms of the structural shocks,

$$Y_t = \mathrm{C}(L)\Theta_0\varepsilon_t, \tag{14}$$

where C(L)=A(L)$^{-1}$.

If (13) holds, then the SVAR impulse response function reveals the population dynamic causal effects; that is, $A(L)^{-1}\Theta_o = \Theta(L)$.[7] Condition (13) is an implication of the assumption that the structural moving average is invertible. This "invertibility" assumption, which underpins SVAR analysis, is nontrivial and we discuss it in more detail in the next subsection.

Under the assumption of invertibility, the SVAR identification problem is to identify Θ_0. Here, we summarize SVAR identification using external instruments.

Suppose there is an instrument Z_t that satisfies the first two conditions of condition LP-IV, which we relabel as Condition SVAR-IV:

Condition SVAR-IV

(i) $E\varepsilon_{1t}Z_t' = \alpha' \neq 0$ (relevance)

(ii) $E\varepsilon_{\bullet t}Z_t' = 0$ (exogeneity w.r.t. other current shocks)

Condition SVAR-IV and (13) imply that,

$$Ev_tZ_t = E\left(\Theta_0\varepsilon_tZ_t\right) = \Theta_0 E\begin{pmatrix}\varepsilon_{1t}Z_t' \\ \varepsilon_{\bullet t}Z_t'\end{pmatrix} = \Theta_0\begin{pmatrix}\alpha' \\ 0\end{pmatrix} = \begin{pmatrix}\Theta_{0,11}\alpha' \\ \Theta_{0,\bullet 1}\alpha'\end{pmatrix}. \tag{15}$$

With the help of the unit effect normalization (6), it follows from (15) that, in the case of scalar Z_t,

$$\frac{E(v_{i,t}Z_t)}{E(v_{1,t}Z_t)} = \Theta_{0,i1}, \tag{16}$$

with the extension to multiple instruments as follows (8). Thus $\Theta_{0,i1}$ is the population estimand of the IV regression,

[7] To show this, note that from (5) and (12), $v_t = A(L)\Theta(L)\varepsilon_t$. With the addition of condition (13), we have $\Theta_0\varepsilon_t = A(L)\Theta(L)\varepsilon_t$, so that $\Theta_0 = A(L)\Theta(L)$, that is, $A(L)^{-1}\Theta_0 = \Theta(L)$.

$$\nu_{i,t} = \Theta_{0,i1}\nu_{1,t} + \{\varepsilon_{\bullet,t}\} \qquad (17)$$

using the instrument Z_t.

Because the innovations ν_t are not observed, the IV regression (17) is not feasible. One possibility is replacing the population innovations in (17) with their sample counterparts $\hat{\nu}_t$, which are the VAR residuals. However, while doing so would provide a consistent estimator with strong instruments, the resulting standard errors would need to be adjusted because of potential correlation between Z_t and lagged values of Y_t ($\hat{\nu}_{1,t}$ is a generated regressor).

Instead, $\Theta_{0,i1}$ can be estimated by an approach that yields the correct large-sample, strong-instrument standard errors. Because $\nu_{i,t} = Y_{i,t} - \text{Proj}\left(Y_{i,t} \mid Y_{t-1}, Y_{t-2}, ...\right)$, equation (17) can be rewritten as

$$Y_{i,t} = \Theta_{0,i1}Y_{1,t} + \gamma_i(\text{L})Y_{t-1} + \{\varepsilon_{\bullet t}\}, \qquad (18)$$

where $\gamma_i(\text{L})$ are the coefficients of $\text{Proj}\left(Y_{i,t} - \Theta_{0,i1}Y_{1,t} \mid Y_{t-1}, Y_{t-2}, ...\right)$. The coefficients $\Theta_{0,i1}$ and $\gamma_i(\text{L})$ can be estimated by two-stage least squares equation-by-equation using the instrument Z_t. By classic results of Zellner and Theil (1962) and Zellner (1962), this equation-by-equation estimation by two stage least squares entails no efficiency loss – is in fact equivalent to – system estimation by three stage least squares.

To summarize, SVAR-IV proceeds in three steps:

1. Estimate (18) using instruments Z_t for the variables in Y_t, using p lagged values of Y_t as controls. This, along with the unit effect normalization $\Theta_{0,11} = 1$, yields the IV estimator of the first column of Θ_0, $\hat{\Theta}_{0,1}^{SVAR-IV}$.
2. Estimate a VAR(p) and invert the VAR to obtain $\hat{\text{C}}(\text{L}) = \hat{\text{A}}(\text{L})^{-1}$.
3. Estimate the dynamic causal effects of shock 1 on the vector of variables as

$$\hat{\Theta}_{h,1}^{SVAR-IV} = \hat{\text{C}}_h\hat{\Theta}_{0,1}^{SVAR-IV} \qquad (19)$$

It is useful to compare the SVAR-IV and LP-IV estimators. For $h = 0$, the SVAR-IV and LP-IV estimators of $\Theta_{0,i1}$ are the same when the control variables W_t are Y_{t-1}, Y_{t-2},…, Y_{t-p}. For $h > 0$, however, the SVAR-IV and LP-IV estimators differ. In the SVAR-IV estimator, the impulse response functions are generated from the VAR dynamics. In contrast, the LP-IV estimator does not use the VAR parametric restriction: the dynamic causal effect is estimated by h distinct IV regressions, with no parametric restrictions tying together the estimates across horizons.

Inference. Let Γ denote the unknown parameters in A(L) and $\Theta_{0,1}$ (the first column of Θ_0). Under standard regression and strong instrument assumptions (e.g., Hayashi (2000)), $\sqrt{T}\left(\hat{\Gamma}-\Gamma\right)\xrightarrow{p}$ N(0,Σ_Γ). And, because estimator $\hat{\Theta}_{h,1}^{SVAR-IV}$ from Step 3 is a smooth function of $\hat{\Gamma}$, $\sqrt{T}\left(\hat{\Theta}_{h,1}^{SVAR-IV}-\Theta_{h,1}\right)\xrightarrow{d}$ N(0,Σ_Θ) where Σ_Θ is readily computed from the δ-method. When instruments are weak, the asymptotic distribution of $\hat{\Theta}_{h,1}^{SVAR-IV}$ is not normal; Montiel-Olea, Stock and Watson (2017) discuss weak-instrument robust inference for SVARs identified by external instruments.

Different data spans for Z and Y. The SVAR-IV estimator of the impulse response function in (19) has two parts, $\hat{\mathrm{C}}_h$ and $\hat{\Theta}_{0,1}^{SVAR-IV}$. In general these can be estimated over different sample periods. For example, in Gertler-Karadi (2015), the data on the macro variables Y_t are available for a longer period than are data on the instruments, and they estimate the VAR coefficients A(L) over the longer sample and $\hat{\Theta}_{0,1}^{SVAR-IV}$ over the shorter sample when Z_t is available. Using the longer sample for the VAR improves efficiency. Note that the flexibility of using different samples for the dynamics and the IV step is possible in SVAR-IV, but not in LP-IV, because LP-IV directly estimates $\Theta_{h,1}$ in a single step.

News shocks and the unit-effect normalization. A structural moving average may be invertible even when it includes news shocks as long as Y_t contains forward-looking variables. But, as in analysis in the previous section, news variables require a change in the unit-effect normalization from contemporaneous $\Theta_{0,11} = 1$ to k periods ahead $\Theta_{k,11} = 1$. To implement this normalization in the SVAR, note that the effect of ε_t on Y_{t+k} is given by $\eta_t = \Theta_k\varepsilon_t = C_k\Theta_0\varepsilon_t = C_k\nu_t$. The k-period ahead unit-effect normalization says $\eta_{1,t} = \varepsilon_{1,t} + \{\varepsilon_{\bullet t}\}$. Thus, letting $X_t = \hat{C}_k Y_t$, the normalization is implemented by replacing $Y_{1,t}$ with $X_{1,t}$ in (18) and carrying out the three steps given above.

3.2. Invertibility, Omitted Variable Bias, and the Relation between Assumptions SVAR-IV and LP-IV

The structural moving average Θ(L) in (5) is said to be invertible if ε_t can be linearly determined from current and lagged values of Y_t:

$$\varepsilon_t = \text{Proj}(\varepsilon_t \mid Y_t, Y_{t-1}, \ldots). \qquad \text{(invertibility)} \tag{20}$$

In the linear models of this lecture, condition (20) is equivalent to saying that $\Theta(L)^{-1}$ exists. The reason we state the invertibility condition as (20) is that it is closer to the standard definition, $\varepsilon_t = E(\varepsilon_t \mid Y_t, Y_{t-1}, \ldots)$, which applies to nonlinear models as well.

In this subsection, we make four points. First, we show that (20), plus the assumption that the innovation covariance matrix is nonsingular, implies (13). Second, we reframe (20) to show how very strong this condition is: under invertibility, a forecaster using a VAR who magically stumbled upon the history of true shocks would have no interest in adding those shocks to her forecasting equations. Third, this reframing provides a natural reinterpretation of invertibility as a problem of omitted variables; thus LP-IV can be seen as a solution to omitted variables bias, akin to a standard motivation for IV regression in microeconometrics. Fourth, we show that there is, at a formal level, a close connection between the choice of control variables in LP-IV and invertibility. Specifically, we show that, for a generic instrument Z_t, using lagged Y_t as control variables to ensure that Condition LP-IV$^{\perp}$ holds is equivalent to assuming that Condition SVAR-IV and invertibility (20) both hold.

***Demonstration that invertibility* (20) *implies* (13).** This result is well known but we show it here for completeness. Recall that by definition, $\nu_t = Y_t - \text{Proj}\left(Y_t \mid Y_{t-1}, Y_{t-2}, \ldots\right) = \Theta(\text{L})\varepsilon_t - \text{Proj}\left[\Theta(\text{L})\varepsilon_t \mid Y_{t-1}, Y_{t-2}, \ldots\right] = \Theta_0\varepsilon_t + \sum_{i=1}^{\infty}\Theta_i\left[\varepsilon_{t-i} - \text{Proj}\left(\varepsilon_{t-i} \mid Y_{t-1}, Y_{t-2}, \ldots\right)\right]$, where the second equality uses (5), and the third equality uses the fact that $\text{Proj}\left(\varepsilon_t \mid Y_{t-1}, Y_{t-2}, \ldots\right) = 0$ and collects terms. Equation (20) implies that $\text{Proj}\left(\varepsilon_{t-i} \mid Y_{t-1}, Y_{t-2}, \ldots\right) = \varepsilon_{t-i}$, so the term in brackets in the final summation is zero for all i; thus we have that $\nu_t = \Theta_0\varepsilon_t$ as in (13).

To see why (20) implies that Θ_0 is invertible, note that $\varepsilon_t = \text{Proj}\left(\varepsilon_t \mid Y_t, Y_{t-1}, \ldots\right) = \text{Proj}\left(\varepsilon_t \mid \nu_t, \nu_{t-1}, \ldots\right) = \text{Proj}\left(\varepsilon_t \mid \Theta_0\varepsilon_t, \Theta_0\varepsilon_{t-1}, \ldots\right) = \text{Proj}\left(\varepsilon_t \mid \Theta_0\varepsilon_t\right) = \text{Proj}\left(\varepsilon_t \mid \nu_t\right)$, where the first equality is (20), the second follows because current and past innovations span the space of current and past Y's, the third and fifth follows from $\nu_t = \Theta_0\varepsilon_t$, and the fourth follows from the serial independence of ε_t. Because $\varepsilon_t = \text{Proj}\left(\varepsilon_t \mid \nu_t\right)$, the equation $\nu_t = \Theta_0\varepsilon_t$ must yield a unique solution for ε_t, so that Θ_0 has rank n_ε. Moreover, because var(ν_t) is assumed to have full rank, $n_y \le n_\varepsilon$. Taken together these imply that $n_y = n_\varepsilon$ and Θ_0 has rank n_ε. Therefore, if (20) holds, then (13) holds.

Invertibility as omitted variables. The invertibility condition can be interpreted as a condition that there are no omitted variables in the VAR: because invertibility implies that the spans of ε_t and ν_t are the same, there is no forecasting gain from adding past shocks to the VAR. That is, the invertibility condition (20) implies that,[8]

$$\text{Proj}(Y_t \mid Y_{t-1}, Y_{t-2}, \ldots, \varepsilon_{t-1}, \varepsilon_{t-2}, \ldots) = \text{Proj}(Y_t \mid Y_{t-1}, Y_{t-2}, \ldots). \tag{21}$$

Condition (21) both shows how strong the assumption of invertibility is, and provides an interpretation of invertibility as a problem of omitted variables. If invertibility holds, then knowledge of the history true shocks would not improve the VAR forecast. If instead those forecasts were improved by adding the shocks to the regression – infeasible, of course, but a thought experiment – then the VAR has omitted some variables, and that omission is an indication of the failure of the invertibility assumption.[9]

---

[8] Formally, equation (21) follows by writing, $\text{Proj}\left(Y_t \mid Y_{t-1}, Y_{t-2}, \ldots, \varepsilon_{t-1}, \varepsilon_{t-2}, \ldots\right) = \text{Proj}\left(Y_t \mid \nu_{t-1}, \nu_{t-2}, \ldots, \varepsilon_{t-1}, \varepsilon_{t-2}, \ldots\right) = \text{Proj}\left(Y_t \mid \nu_{t-1}, \nu_{t-2}, \ldots\right) = \text{Proj}\left(Y_t \mid Y_{t-1}, Y_{t-2}, \ldots\right)$, where the first and third equalities uses the fact that the innovations are the Wold errors, and the second equality uses the implication of (13) that span(ε_t) = span(ν_t) .

[9] Condition (21) is closely related to Proposition 3 in Forni and Gambetti (2014), which states (with some refinements) that the structural moving average is invertible if no added state variable in a VAR have predictive content for Y_t. That observation leads to their test for invertibility, which involves estimating factors using a dynamic factor model and including them in the VAR.

In general, one solution to omitted variable problems is to include the omitted variables in the regression. In the case at hand, that is challenging, because the omitted variables are the unobserved structural shocks. Pursuing this line of reasoning suggests using a large number of variables in the VAR, or using a high dimensional dynamic factor model. This is a potentially useful avenue to dealing with the invertibility problem, see for example Forni, Giannone, Lippi, and Reichlin (2009) and the survey in Stock and Watson (2016).

It is important to note that expanding the number of variables will not necessarily result in (20) being satisfied, so that moving to large systems does not assure invertibility.

Relation between assumptions SVAR-IV, LP-IV$^{\perp}$, and invertibility. As stressed in the introduction, the appeal of LP-IV is that the direct regression approach does not explicitly assume invertibility. There is, however, a close connection between the LP-IV and SVAR-IV methods, if the control variables in LP-IV are lagged Ys.

Specifically, suppose Z is an instrument that satisfies parts (i) and (ii) of condition LP-IV, but is predictable from past values of ε, and so does not satisfy LP-IV (iii). Then assuming that this problem is resolved by using lagged Ys as control variables in LP-IV is equivalent to assuming that condition SVAR-IV holds and that the SVAR is invertible. This result is stated formally in the following theorem.

> ***Theorem 1.*** Let Z_t be an instrument that is correlated with ε_{t-j} for some $j \geq 1$. Let the control variables W_t in Condition LP-IV$^{\perp}$ be Y_{t-1}, Y_{t-2}, … . Then in general Z_t satisfies Condition LP-IV$^{\perp}$ if and only if both Z_t satisfies Condition SVAR-IV and the invertibility condition (20) holds.

Proof. The equivalence of parts (i) and (ii) of Condition LP-IV$^{\perp}$ and SVAR-IV is immediate: because $\text{Proj}(\varepsilon_t \mid Y_{t-1}, Y_{t-2}, \ldots) = 0$, $\varepsilon_t^{\perp} = \varepsilon_t$ and $E\left(\varepsilon_t^{\perp} Z_t^{\perp}\right) =$ $E\left[\varepsilon_t\left(Z_t - \text{Proj}(Z_t \mid Y_{t-1}, Y_{t-2}, \ldots)\right)\right] = E\left(\varepsilon_t Z_t\right)$.

We now show that Z_t satisfies Condition LP-IV$^{\perp}$ (iii) if and only if the invertibility condition (20) holds. For convenience suppose that Z_t is a scalar; the extension to the vector case is direct. First we show that condition (20) implies condition LP-IV$^{\perp}$ (iii). For $j > 0$, $\text{Proj}\left(\varepsilon_{t+j} \mid W_t\right) = 0$, so that $E\left(Z_t^{\perp} \varepsilon_{t+j}^{\perp}\right) = E\left(Z_t^{\perp} \varepsilon_{t+j}\right) = 0$. Thus LP-IV$^{\perp}$ (i), (ii) and $E(Z_t^{\perp} \varepsilon_{t+j}^{\perp}) = 0$

for $j > 0$ are satisfied. Because condition (20) implies (13), $\text{Proj}\left(\varepsilon_{t-j} \mid Y_{t-1}, Y_{t-2}, \ldots\right) =$ $\text{Proj}\left(\varepsilon_{t-j} \mid \varepsilon_{t-1}, \varepsilon_{t-2}, \ldots\right) = \varepsilon_{t-j}$, where the final equality holds for $j \geq 1$.Thus for $j \geq 1$, $\varepsilon_{t-j}^{\perp} = 0$ so $E(Z_t^{\perp}\varepsilon_{t-j}^{\perp}) = 0$. Thus LP-IV$^{\perp}$ is satisfied if condition (20) holds.

Next we show by contradiction that condition LP-IV$^{\perp}$ (iii) implies condition (20). Assume LP-IV$^{\perp}$ is not satisfied because $E\left(Z_t^{\perp}\varepsilon_{t-j}^{\perp}\right) \neq 0$ for some $j \geq 1$; this implies $\text{var}(\varepsilon_{t-j}^{\perp}) >$ 0 (by Cauchy-Schwartz). But $\text{var}(\varepsilon_{t-j}^{\perp}) = \text{var}\left[\varepsilon_{t-j}\text{-Proj}\left(\varepsilon_{t-j} \mid Y_{t-1}, Y_{t-2}, \ldots\right)\right] \leq$ $\text{var}\left[\varepsilon_{t-j}\text{-Proj}\left(\varepsilon_{t-j} \mid Y_{t-j}, Y_{t-j-1}, \ldots\right)\right]$, so $\text{var}\left[\varepsilon_{t-j}\text{-Proj}\left(\varepsilon_{t-j} \mid Y_{t-j}, Y_{t-j-1}, \ldots\right)\right] > 0$ and (20) does not hold. □

We interpret this theorem as a "no free lunch" result. Although LP-IV can estimate the impulse response function without assuming invertibility, to do so requires an instrument that either satisfies LP-IV (iii) or that can be made to do so by adding control variables that are specific to the application. Simply including past Y's out of concern that Z_t is correlated with past shocks is valid if and only if the VAR with those past Y's is invertible; but if so, it is more efficient to use SVAR-IV.[10]

4. A Test of Invertibility

Suppose one has an instrument that satisfies condition LP-IV. Under invertibility, SVAR-IV and LP-IV are both consistent, but SVAR-IV is more efficient. If, however, invertibility fails, LP-IV is consistent but SVAR-IV is not. This observation suggests that comparing the SVAR-IV and LP-IV estimators provides a Hausman (1978)-type test of the null hypothesis of invertibility. Throughout, we maintain the assumption that Y_t has the linear structural moving average (5). We additionally assume the VAR lag length p is finite and known.

[10] It is well known that in VARs, distributions of estimators of impulse response functions are generally not well approximated by their asymptotic distributions in sample sizes typically found in practice. A more relevant comparison would be of the efficiency of the estimators in a simulation calibrated to empirical data. Kim and Kilian did such an exercise comparing LP and SVAR estimators, with identification by a Cholesky decomposition (what we would call internal instruments). Their results are consistent with improvements in efficiency, and tighter confidence intervals, for SVARs than LP.

Before introducing the test, we make precise the null and alternative hypothesis. We also provide a novel nesting of local departures from the null, which we refer to as local non-invertibility.

Null and local alternative. Under invertibility (20), the structural moving average can be written $Y_t = \mathrm{C(L)}\Theta_0\varepsilon_t$ as in (14), where $\mathrm{C(L)} = \mathrm{A(L)}^{-1}$; that is, that $\Theta(\mathrm{L}) = \mathrm{C(L)}\Theta_0$. The null and alternative hypotheses thus are,

$$H_0\text{: } \Theta_{h,1} = \mathrm{C}_h\Theta_{0,1}\text{, all } h \text{ v. } H_1\text{: } \Theta_{h,1} \neq \mathrm{C}_h\Theta_{0,1}\text{, some } h. \tag{22}$$

In addition to establishing the null distribution of the test, we wish to examine its distribution under an alternative to demonstrate that the test has power against non-invertibility. To do so, we follow standard practice and consider a drifting sequence of alternatives. Specifically, we introduce the concept of a nearly invertible process, that is, a process that is not invertible but is within a $T^{-1/2}$ neighborhood of an invertible process.

To motivate the specification of the nearly invertible process, use the structural moving average representation $Y_t = \Theta(\mathrm{L})\varepsilon_t$ and the Wold representation $Y_t = \mathrm{C(L)}\nu_t$ to write $\nu_t =$ $\mathrm{A(L)}\Theta(\mathrm{L})\varepsilon_t = (\mathrm{I} - \mathrm{A}_1\mathrm{L} - \mathrm{A}_2\mathrm{L}^2 - \ldots)(\Theta_0 + \Theta_1\mathrm{L} + \ldots) = \Theta_0\varepsilon_t + \varphi(\mathrm{L})\varepsilon_{t-1}$, where $\varphi_j = -\sum_{i=0}^{j+1} A_{j+1-i}\Theta_i$. Under invertibility, $\varphi(\mathrm{L}) = 0$ (from (13)), so that past ε's have no predictive power for current Wold errors, given current ε's. Our local model allows a predictive contribution, but requires it to be small, specifically, we set $\varphi(\mathrm{L}) = T^{-1/2} f(\mathrm{L})$, where $f(\mathrm{L})$ is a non-drifting lag polynomial:

$$\nu_t = \Theta_0\varepsilon_t + T^{-1/2} f(\mathrm{L})\varepsilon_{t-1}, \tag{23}$$

that is, $\mathrm{A(L)}\Theta(\mathrm{L}) = \Theta_0 + T^{-1/2} f(\mathrm{L})\mathrm{L}$. Premultiplying this latter expression by C(L) yields $\Theta(\mathrm{L})$ $= \mathrm{C(L)}\Theta_0 + T^{-1/2}\mathrm{C(L)} f(\mathrm{L})$. Letting $d(\mathrm{L}) = \mathrm{C(L)}f(\mathrm{L})\mathrm{L}$, this yields the local alternative,

$$\Theta_h = C_h\Theta_0 + T^{-1/2} d_h . \tag{24}$$

where under the null $d_h = 0$, while under the alternative d_h is nonzero for at least some $h > 0$.

Although (24) is a standard expression for a local alternative that could have been reached by purely statistical reasoning, its derivation here via (23) connects with the discussion in Section 3.2: under the nearly invertible alternative, forecasts of Y_t would be improved by only a small amount, were past values of ε_t actually available for forecasting.

Test of invertibility. We now turn to the test statistic. Let $\hat{\theta}^{SVAR-IV}$ denote an $m\times 1$ vector of SVAR -IV estimators (19), computed using a VAR(p), for different variables and/or horizons, and let $\hat{\theta}^{LP-IV}$ denote the corresponding LP-IV estimators. Compute the LP-IV estimator using as control variables the p lags of Y that appear in the VAR; because Z_t satisfies condition LP-IV, including these lags as controls is not necessary for consistency but makes the two statistics comparable for use in the same test statistic.

It is shown in the appendix that, with strong instruments and under standard moment/memory assumptions, under the null and local alternative,

$$\sqrt{T}\left(\hat{\theta}^{LP-IV}-\hat{\theta}^{SVAR-IV}\right)\xrightarrow{d}N(d,V), \tag{25}$$

where d consists of the elements of $\{d_h\}$ corresponding to the variable-horizon combinations that comprise $\hat{\theta}^{LP-IV}$ and $\hat{\theta}^{SVAR-IV}$.

The Hausman-type test statistic is,

$$\xi = T(\hat{\theta}^{LP-IV}-\hat{\theta}^{SVAR-IV})'\hat{V}^{-1}(\hat{\theta}^{LP-IV}-\hat{\theta}^{SVAR-IV}), \tag{26}$$

where $\hat{V}$ is a consistent estimator of V. Under the null of invertibility, $\xi \xrightarrow{d} \chi^2_m$.

We make four remarks about this test.

1. We suggest computation of the variance matrix $\hat{V}$ using the parametric bootstrap, and we discuss some specifics in the appendix.
2. The LP-IV and SVAR-IV estimators for the impact effect (h = 0) are identical when lagged Ys are used as controls. Thus this test compares the LP-IV and SVAR-IV estimates of the impulse responses for $h \geq 1$. This test therefore assesses the validity of the parametric restrictions imposed by inverting the SVAR, compared to direct estimation

of the impulse response function by LP-IV. Here, we have maintained the assumption that the structural moving average is linear and the VAR lag length is finite and known. Under these maintained assumptions, any divergence between the SVAR impulse responses and the direct estimates, in population, is attributable to non-invertibility.

3. Under the local alternative (24), the test statistic has a noncentral chi-squared distribution with m degrees of freedom and noncentrality parameter $\mu^2 = d'V^{-1}d$. In the appendix, it is shown that V has the form, $V = \Omega/\alpha^2$, where Ω does not depend on α. Thus the noncentrality parameter can be written, $\mu^2 = \alpha^2 d'\Omega^{-1}d$. Thus the power of the test increases with the norm of α. The stronger the instrument, the greater the power of the test for non-invertibility.
4. The forecasting expression for the local alternative (23) suggests an alternative test for invertibility. Because the instrument is correlated with ε_t, past Z_t should have forecasting power for v_t. Thus invertibility implies that Z_t does not Granger-cause Y_t. In contrast to the Granger-causality test, the Hausman-type-test focuses on departures from the null in the direction of economic interest, which here is estimation of the dynamic causal effect, not forecasting.

5. Illustration: Gertler-Karadi (2015) identification of the dynamic causal effect of monetary policy

Gertler and Karadi (2015) use the SVAR-IV method to estimate the effect of a monetary policy shock on real output, prices, and various credit variables, and Ramey (2016) applies LP-IV to their data to illustrate the differences between the two methods. Here, we extend Ramey's comparison and formally test invertibility. We use this application to discuss several implementation details.

Gertler and Karadi's (2015) benchmark analysis uses U.S. monthly data to estimate the effect of Federal Reserve policy shocks on four variables: the index of industrial production and the consumer price index (both in logarithms, denoted here as *IP* and *P*), the interest rate on 1-year U.S. Treasury bonds (R_t), and a financial stress indicator, the Gilchrist and Zakrajšek (2012) excess bond premium (*EBP*). We first-difference *IP* and *P*, so the vector of variables is $Y_t = (R_t, 100\Delta IP, 100\Delta P, EBP)$, where *R* and *EBP* are measured in percentage points at annual rate and

ΔIP and ΔP are multiplied by 100 so these variables are measured in percentage point growth rates.

Gertler and Karadi (GK) identify the monetary policy shock using changes in Federal Funds futures rates (*FFF*) around FOMC announcement dates. In doing so, they draw on insights from Kuttner (2001) and others who argued that this measure is plausibly uncorrelated with other shocks because they are changes across a short announcement window. Whereas the original literature used such a measure as the shock, Gertler and Karadi (2015) use it as an instrument; that is, $Z_t = FFF_t$.

Column (a) of Table 1 shows results for the LP-IV regression (7), the equation without controls, using the GK data that span 1990m1 – 2012m6. Standard errors in Table 1 for LP-IV impulse responses are Newey-West with h+1 lags. We highlight three results. First, the table shows that the estimated contemporaneous ($h = 0$) effect of monetary policy shocks on interest rates (R) is $\Theta_{0,11} = 1.0$; this is the unit-effect normalization. Second, the first-stage F-statistic – that is the (standard) F-statistic from the regression of R_t onto FFF_t – is small, only 1.7, raising weak instrument concerns. And third, the estimated standard errors for the estimated causal effects are large, particularly for large values of h.

These final two results are related. To see why, rewrite equation (5) to highlight the various components of $Y_{i,t+h}$:

$$Y_{i,t+h} = \Theta_{h,i1}\varepsilon_{1,t} + \{\varepsilon_{t+h},...,\varepsilon_{t+1}\} + \{\varepsilon_{\bullet,t}\} + \{\varepsilon_{t-1},...\} \tag{27}$$

where, again, the notation $\{\cdot\}$ denotes a linear function of the variables included in the braces. The first-stage F-statistic is from the regression of $Y_{1,t}$ (= R_t) onto Z_t (= FFF_t). From (27), the error term in the first-stage regression is comprised of $\{\varepsilon_{\bullet,t}\}$ and $\{\varepsilon_{t-1},...\}$. Because interest rates are very persistent, only a small fraction of the variance is attributable to contemporaneous shocks, ε_t; a fraction of this contemporaneous effect is associated with the monetary policy shock $\varepsilon_{1,t}$, and only a fraction of $\varepsilon_{1,t}$ can be explained by the instrument Z_t. Taken together, these effects yield a first-stage regression with $R^2 = 0.006$ and a correspondingly small F-statistic. Similar logic explains the large standard errors for the estimated causal effects because these are associated with IV regressions with error terms comprised of $\{\varepsilon_{t+h},...,\varepsilon_{t+1}\} + \{\varepsilon_{\bullet,t}\} + \{\varepsilon_{t-1},...\}$.

Column (b) of Table 1 repeats the estimation, but now using four lags of Y_t and Z_t as controls. The controls serve two purposes. First, because these controls are correlated with $\{\varepsilon_{t-1}, \ldots\}$, they reduce the variance of the regression error term and, for example, the first-stage (partial) R^2 in (b) increases to $R^2 = 0.09$ with a first-stage F-statistic increases to $F = 23.7$. Second, the controls adjust for a data processing issue that makes the *FFF* variable an invalid instrument in the regression without controls. Specifically, as pointed out by Ramey (2016), Gertler and Karadi (2015) form their *FFF* instrument as a moving average of returns from month t and month $t-1$.[11] Thus, FFF_t will be correlated with both $\varepsilon_{1,t}$ *and* $\varepsilon_{1,t-1}$, violating Assumption LP-IV (iii). Because Z_t has an MA(1) structure, using lags of Z_t as controls eliminates the correlation with $\varepsilon_{1,t-1}$, so that Condition LP-IV$^{\perp}$ (iii) is satisfied. Note that it is plausible that this instrument is uncorrelated with other shocks. Thus, to satisfy Condition LP-IV$^{\perp}$ (iii), it would suffice to include Z_{t-1} as a control; including lagged Ys and additional lags of Z serves to improve precision (increase the first-stage F).

If there are more than four shocks that affect Y_t, or if some elements of Y_t are measured with error (as *IP* and *P* surely are), then the innovations to the four variables making up Y_t will not span the space of the shocks. This is not a problem for the validity of LP-IV with lagged Zs, however it does suggest that including additional variables that are correlated with the shocks could further reduce the regression standard error and thus result in smaller standard errors. One plausible set of such variables are principal components (factors) computed from a large set of macro variables. With this motivation, column (c) adds lags of four factors computed from the FRED-MD dataset (McCracken and Ng (2016)). In this illustration, these additional controls yield results that are largely consistent with the results using lags of Z and Y.

Both specification (b) and (c) in Table 1 improve on the model without controls, (a), by eliminating some of the variability associated with lagged ε and in particular by making Z satisfy LP-IV$^{\perp}$ (iii), whereas (a) does not satisfy LP-IV (iii). However, neither eliminates the variability associated with of *future* ε's, the $\{\varepsilon_{t+h}, \ldots, \varepsilon_{t+1}\}$ component of the error term shown in (27). The variability of this component increases with the horizon h, and this is evident in the large standard errors in estimates associated with long-horizons. When the structural moving average

[11] This is described in their footnote 6. We stress that, while this data processing procedure invalidates their instrument in the LP-IV regression reported in column (1), it does not affect its validity in the SVAR-IV regression used by GK.

model is invertible, it is in effect possible to control for both lagged and future values of ε in the IV regression using VAR methods.

Column (d) of Table 1 shows results from a SVAR with 12 lags, with monetary policy identified by the *FFF* instrument. Because the data on the *Y*s are available for a longer span than the data on the instrument, we follow Gertler and Karadi (2015) and estimate the VAR over the sample 1980m7-2012m6, while $\Theta_{0,1}$ is estimated over the sample 1990m1-2012m6 (see the discussion of data spans towards the end of Section 3.1). Standard errors for the SVAR-IV estimate are computed by the parametric bootstrap described in the Appendix. Because the VAR uses 12 lags of *Y* instead of the 4 lags used as controls in the local projections, the first stage *F*-statistics differ slightly in columns (b) and (d). As expected, the standard errors for the estimated dynamic causal effects are smaller for the SVAR than for the local projections, particularly for large values of *h*, for two reasons. First, the local projections are estimated using regressions with error terms that include leads and lags of ε (see (27)), and these terms are absent from the IV regression used in the SVAR, because only the impact effect, Θ_0, is estimated by IV. Second, the VAR parameterization imposes smoothness and damping on the moving average coefficients in C_h, which further reduces the standard errors. Still, in this empirical application, the standard errors in the SVAR remain large.

The final column of Table 1 shows the difference in estimates of dynamic causal effects from the LP-IV estimator in column (b) and the SVAR-IV estimator in column (d). These differences form the basis for the invertibility test developed in the last section, and the standard errors shown in final column are computed from the parametric bootstrap, which imposes invertibility. Some of the differences between the SVAR and LP estimates are large, but so are their estimated errors, and none of the differences are statistically significant. Relative to the sampling uncertainty, the differences in the LP and SVAR estimates shown in Table 1 are not large enough to conclude that the SVAR suffers from misspecification associated with a lack of invertibility.

Table 2 shows results for two additional tests for invertibility. The first row shows results for the test ξ in (26) for the differences of the LP-IV and SVAR-IV estimates jointly across the lags shown in Table 1. The second row shows results from Granger-causality tests that include four lags of *Z* in each of the VAR equation. Despite the large differences, in economic terms, between the two estimates of the impulse responses, the table indicates that there is no statistically significant evidence against the null of hypothesis of invertibility.

6. Conclusions

It is well known that, with Gaussian errors, every invertible model has are a multiplicity of observationally equivalent noninvertible representations, so if one is to distinguish among them, some external information must be brought to bear. One approach is to assume that the shocks are independent and nonGaussian, and to exploit higher moment restrictions to identify the causal structure (cf. Lanne and Saikkonen (2013), Gospodinov and Ng (2015) and Gouriéroux, Monfort, and Renne (2017)). A second approach is to use *a-priori* informative priors (Plagborg-Møller (2016b)). Here, we have shown that there is a third approach, which is to use an external instrument.

There remain a number of methodological issues concerning the use of external instruments. For example, this discussion assumes homogenous treatment effects, which on the surface is plausible in a macroeconomic setting (there is only one "subject," although effects may vary over time), and more work remains concerning heterogeneous treatment effects in this setting. Also, the usual weak-instrument devices do not cover all the methods used here, for example one question is how to robustify the test of invertibility to the case that instruments are potentially weak. But in our view, the most exciting work to be done in this area is empirical: the new external instruments, yet to be developed, that provide sufficient plausibly exogenous variation to provide more credible identification of dynamic causal effects.

Appendix

The Test for Invertibility: Distribution and Bootstrap Variance Matrix

Let $W_t = \{Y_{t-1},\ldots, Y_{t-p}\}$ be the set of lags in the VAR, and also the conditioning variables in the LP-IV regression. Let $Y_{t+h}^{\perp}$ denote the residuals from a regression of Y_{t+h} onto W_t, etc; this notational shift departs from the text where "⊥" means population projection residual, here is denotes sample residual. Note that $Y_t^{\perp} = \hat{v}_t$. This appendix considers the case of scalar instrument denoted by z_t.

The SVAR-IV estimator is:

$$\hat{\Theta}_{h,1}^{SVAR-IV} = \hat{C}_h \frac{\sum Y_t^{\perp} z_t^{\perp}}{\sum \hat{v}_{1t} z_t^{\perp}} \tag{28}$$

The LP-IV estimator is:

$$\hat{\Theta}_{h,1}^{LP-IV} = \frac{\sum Y_{t+h}^{\perp} z_t^{\perp}}{\sum Y_{1t}^{\perp} z_t^{\perp}} = \frac{\sum (Y_{t+h} - \hat{Y}_{t+h|t-1}^{d}) z_t^{\perp}}{\sum \hat{v}_{1t} z_t^{\perp}} \tag{29}$$

The difference between the two estimators is,

$$\sqrt{T}\left(\hat{\Theta}_{h,1}^{LP-IV} - \hat{\Theta}_{h,1}^{SVAR-IV}\right) = \frac{\frac{1}{\sqrt{T}} \sum \left[Y_{t+h}^{\perp} - \hat{C}_h Y_t^{\perp} \right] z_t^{\perp}}{\frac{1}{T} \sum \hat{v}_{1t} z_t^{\perp}}. \tag{30}$$

First consider the numerator of (30), which we denote Ψ_T:

$$
\begin{aligned}
\Psi_T &= \frac{1}{\sqrt{T}} \sum \left[Y_{t+h}^{\perp} - \hat{C}_h Y_t^{\perp} \right] z_t^{\perp} \\
&= T^{-1/2} \sum \left[(Y_{t+h} - \hat{Y}_{t+h|t-1}^{d}) - \hat{C}_h (Y_t - \hat{Y}_{t|t-1}) \right] z_t^{\perp} \\
&= T^{-1/2} \sum \left[(Y_{t+h} - Y_{t+h|t-1}^{d}) - (Y_{t+h|t-1}^{d} - \hat{Y}_{t+h|t-1}^{d}) - \hat{C}_h (Y_t - Y_{t|t-1}) + \hat{C}_h (Y_{t|t-1} - \hat{Y}_{t|t-1}) \right] z_t^{\perp} \\
&= T^{-1/2} \sum \left(Y_{t+h} - Y_{t+h|t-1}^{d} \right) z_t^{\perp} - \hat{C}_h T^{-1/2} \sum \left(Y_t - Y_{t|t-1} \right) z_t^{\perp} + o_p(1)
\end{aligned}
\tag{31}
$$

where the second line inserts the definitions of $Y_{t+h}^{\perp}$ and $Y_t^{\perp}$, where $\hat{Y}_{t+h|t-1}^{d}$ is the direct forecast of Y_{t+h} computed using $W_t = Y_{t-1}$ as a regressor; the third line adds and subtracts population projections onto W_t; and the fourth line uses the fact that $z_t^{\perp}$ is orthogonal to W_t to eliminate the second and fourth terms in the third line.

Consider the first term in the final line of (31):

$$
\begin{aligned}
&T^{-1/2} \sum \left(Y_{t+h} - Y_{t+h|t-1}^{d} \right) z_t^{\perp} \\
&= T^{-1/2} \sum \left(Y_{t+h} - E(Y_{t+h} \mid \varepsilon_{t-1}, \varepsilon_{t-2}, \ldots) \right) z_t^{\perp} + T^{-1/2} \sum \left(E(Y_{t+h} \mid \varepsilon_{t-1}, \varepsilon_{t-2}, \ldots) - E(Y_{t+h} \mid \nu_{t-1}, \nu_{t-2}, \ldots) \right) z_t^{\perp} \\
&= T^{-1/2} \sum \left(\Theta_0 \varepsilon_{t+h} + \ldots + \Theta_{h-1} \varepsilon_{t+1} + \Theta_h \varepsilon_t \right) z_t^{\perp} + T^{-1/2} \sum \left((\Theta_{h+1} \varepsilon_{t-1} + \Theta_{h+2} \varepsilon_{t-2} + \ldots) - (C_{h+1} \nu_{t-1} + C_{h+2} \nu_{t-2} + \ldots) \right) z_t^{\perp} \\
&= T^{-1/2} \sum \zeta_{t+h}^{(h)} z_t^{\perp} + \Theta_h T^{-1/2} \sum \varepsilon_t z_t^{\perp} + T^{-1/2} \sum_t \sum_{j \geq 1} \left(\Theta_{h+j} \varepsilon_{t-j} - C_{h+j} \nu_{t-j} \right) z_t^{\perp}
\end{aligned}
$$

where $\zeta_{t+h}^{(h)} = \Theta_0 \varepsilon_{t+h} + \ldots + \Theta_{h-1} \varepsilon_{t+1}$. Consider the second term in the final line of (31):

$$
\hat{C}_h T^{-1/2} \sum \left(Y_t - Y_{t|t-1} \right) z_t^{\perp} = T^{1/2} \left(\hat{C}_h - C_h \right) T^{-1} \sum \nu_t z_t^{\perp} + T^{1/2} C_h T^{-1} \sum \nu_t z_t^{\perp} .
$$

Substitution of the final lines of the previous two expressions into the final line of (31) yields,

$$
\begin{aligned}
\Psi_T &= T^{-1/2}\sum \zeta_{t+h}^{(h)} z_t^{\perp} + T^{1/2}\Theta_h T^{-1}\sum \varepsilon_t z_t^{\perp} + T^{-1/2}\sum_t \sum_{j\geq 1}\left(\Theta_{h+1}\varepsilon_{t-1} - C_{h+1}\nu_{t-1}\right) z_t^{\perp} \\
&\quad - T^{1/2}\left(\hat{C}_h - C_h\right) T^{-1}\sum \nu_t z_t^{\perp} - T^{1/2} C_h T^{-1}\sum \nu_t z_t^{\perp} \\
&= T^{-1/2}\sum \zeta_{t+h}^{(h)} z_t^{\perp} + \left(T^{1/2}\Theta_h T^{-1}\sum \varepsilon_t z_t^{\perp} - T^{1/2} C_h T^{-1}\sum \nu_t z_t^{\perp}\right) - T^{1/2}\left(\hat{C}_h - C_h\right) T^{-1}\sum \nu_t z_t^{\perp} \\
&\quad + T^{-1/2}\sum_t \sum_{j\geq 1}\left(\Theta_{h+1}\varepsilon_{t-1} - C_{h+1}\nu_{t-1}\right) z_t^{\perp} \\
&= T^{-1/2}\sum \zeta_{t+h}^{(h)} z_t{}_t^{\perp} - T^{1/2}\left(\hat{C}_h - C_h\right)\Theta_{0,1} + T^{1/2}\left(\Theta_h - C_h\Theta_{0,1}\right)\alpha + o_p(1)
\end{aligned}
\tag{32}
$$

where the final line follows from manipulations using the near-invertibility nesting (23) and (24). For example, consider the term in the second equality,

$$
\begin{aligned}
T^{1/2}\Theta_h T^{-1}\sum \varepsilon_t z_t^{\perp} - T^{1/2} C_h T^{-1}\sum \nu_t z_t^{\perp} &= T^{1/2}\Theta_h T^{-1}\sum \varepsilon_t z_t^{\perp} - T^{1/2} C_h T^{-1}\sum\left(\Theta_0\varepsilon_t + T^{-1/2} f(L)\varepsilon_{t-1}\right) z_t^{\perp} \\
&= T^{1/2}\left(\Theta_h - C_h\Theta_0\right) T^{-1}\sum \varepsilon_t z_t^{\perp} - C_h T^{-1}\sum\left(f(L)\varepsilon_{t-1}\right) z_t^{\perp} \\
&= T^{1/2}\left(\Theta_{h,1} - C_h\Theta_{0,1}\right)\alpha + o_p(1) \\
&= d_h\alpha + o_p(1)
\end{aligned}
$$

where the final line follows from $T^{-1}\sum \varepsilon_t z_t^{\perp} \xrightarrow{p} e_1\alpha$ under Condition LP-IV (i) and (ii), and $C_h T^{-1}\sum\left(f(L)\varepsilon_{t-1}\right) z_t^{\perp} \xrightarrow{p} 0$ under LP-IV (iii).

The first two terms in the final line of (32) are jointly normally distributed with a variance matrix V_h that is the same under the null and local alternative. Thus,

$$
\Psi_T \xrightarrow{d} N\left(d_h\alpha, \Omega_h\right), \tag{33}
$$

where Ω_h is the asymptotic covariance matrix of $T^{-1/2}\sum \zeta_{t+h}^{(h)} z_t{}_t^{\perp} - T^{1/2}\left(\hat{C}_h - C_h\right)\Theta_{0,1}$. Note that neither of these terms involve terms in $\varepsilon_t Z_t$ so that the variance Ω_h does not depend on α.

Now turn to the denominator of (30). Under strong instruments,

$$
\frac{1}{T}\sum \hat{\nu}_{1t} z_t^{\perp} = \frac{1}{T}\sum \nu_{1t} z_t^{\perp} + o_p(1) = \frac{1}{T}\sum e_1'\left(\Theta_0\varepsilon_t + T^{-1/2} f(L)\varepsilon_{t-1}\right) z_t^{\perp} + o_p(1) = \Theta_{0,11}\alpha = \alpha,
$$

where the second, third, and fourth equality use the calculations above and the final equality uses the unit effect normalization. Thus,

$$\sqrt{T}\left(\hat{\Theta}_{h,1}^{LP-IV}-\hat{\Theta}_{h,1}^{SVAR-IV}\right)\xrightarrow{d}N\left(d_h,V_h\right)\text{, where } V_h = \Omega_h/\alpha^2. \tag{34}$$

We do not provide an expression for Ω_h, which in its general form is quite complex. Recall that Ω_h is the asymptotic variance of $T^{-1/2}\sum\zeta_{t+h}^{(h)}z_t{}_t^{\perp}-T^{1/2}\left(\hat{C}_h-C_h\right)\Theta_{0,1}$, where $\zeta_{t+h}^{(h)}=\Theta_0\varepsilon_{t+h}+...+\Theta_{h-1}\varepsilon_{t+1}$. The first term is the contribution to the uncertainty from the multistep forecast error arising from the LP-IV regression. The variance of this term is straightforward to derive under condition LP-IV. The second term is the contribution to the sampling uncertainty from estimating, then inverting, the VAR. The formula for this variance, which is the variance of the VAR estimator of the impulse response function times $\Theta_{0,1}$, is complex, see Lütkepohl (2005), and entails linearization of the inverse of the VAR lag polynomial. This linearized expression, along with (23), allows writing this second term in terms of the shocks, from which one can obtain an expression for the covariance between the two terms. This expression is unlikely to be used in practice and we do not provide it here. Instead, we suggest computing $\hat{V}_h$ by the bootstrap.

Parametric bootstrap evaluation of $\hat{V}_h$. The variances (and hence standard errors) of the estimators in Tables 1 and 2 were computed using the sample variances computed from 1000 draws from a parametric bootstrap. For each draw, we generated samples of size T for ($\tilde{Y}_t \;\; \tilde{Z}_t$) from the stationary VAR:

$$\begin{bmatrix}\hat{A}(\mathrm{L}) & 0\\ 0 & \hat{\rho}(\mathrm{L})\end{bmatrix}\begin{bmatrix}\tilde{Y}_t\\ \tilde{Z}_t\end{bmatrix}=\begin{bmatrix}\tilde{v}_t\\ \tilde{e}_t\end{bmatrix},\text{ where }\begin{bmatrix}\tilde{v}_t\\ \tilde{e}_t\end{bmatrix}\sim i.i.d.N\left(\begin{bmatrix}0\\0\end{bmatrix},\begin{bmatrix}S_{\hat{v}\hat{v}} & S_{\hat{v}\hat{e}}\\ S_{\hat{e}\hat{v}} & S_{\hat{e}\hat{e}}\end{bmatrix}\right)$$

where $\hat{A}(\mathrm{L})$ is estimated from a VAR(12), $\hat{\rho}(\mathrm{L})$ is estimated from an AR(4), and $S_{\hat{v}\hat{v}}$, $S_{\hat{v}\hat{v}}$, and $S_{\hat{e}\hat{e}}$ are sample covariances for the VAR/AR residuals. These samples are used to compute the SVAR-IV and LP-IV estimates of $\Theta_{h,i1}$.

References

Andrews, I. (2017), “Valid Two-Step Identification-Robust Confidence Sets for GMM.” Forthcoming, *The Review of Economics and Statistics*.

Angrist, J.D., Ò. Jordà, and G.M. Kuersteiner (2017). “Semiparametric Estimates of Monetary Policy Effects: String Theory Revisited,” *Journal of Business and Economic Statistics*, forthcoming.

Barnichon, R. and C. Brownlees (2016). “Impulse Response Estimation by Smooth Local Projections,” CEPR Discussion paper DP11726.

Bernanke, B.S. and K.N. Kuttner (2005). “What Explains the Stock Market’s Reaction to Federal Reserve Policy?”, *The Journal of Finance* 40, 1221-1257.

Bojinov, I. and N. Shephard (2017). “Time Series Experiments, Causal Estimands, Exact *p*-values, and Trading.” manuscript, Harvard University.

Caldara, D. and C. Kamps (2107). “The Analytics of SVARs: A Unified Framework to Measure Fiscal Multipliers,” forthcoming, *The Review of Economic Studies*.

Cochrane, J.H., and M. Piazzesi (2002). “The Fed and Interest Rates: A High-Frequency Identification,” *American Economic Review* 92 (May), 90-95.

Faust, J., Rogers, J.H., Swanson, E., and J.H. Wright (2003). “Identifying the Effects of Monetary Policy Shocks on Exchange Rates Using High Frequency Data,” *Journal of the European Economic Association* 1(5), 1031-57.

Forni, M. and L. Gambetti (2014). “Sufficient Information in Structural VARs,” *Journal of Monetary Economics* 66C), 124-136.

Forni, M., D. Giannone, M. Lippi, and L. Reichlin (2009). “Opening the Black Box: Structural Factor Models with Large Cross Sections,” *Econometric Theory*, 25, 1319-1347.

Frisch, R. (1933). “Propagation problems and impulse problems in dynamic economics,” in: *Economic Essays in Honor of Gustav Cassel*. Allen & Unwin, London, 171–205.

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

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

Gospodinov, N. and S. Ng (2015). “Minimum Distance Estimation of Possibly Noninvertible Moving Average Models.” *Journal of Business & Economic Statistics* 33, 403–417.

Gouriéroux, C., A. Monfort and J-P. Renne (2017), "Statistical inference for independent component analysis: Application to structural VAR models," *Journal of Econometrics*, 196, pp 111-126.

Gürkaynak, R.S., Sack, B., and E. Swanson (2005). “The sensitivity of long-term interest rates to economic news: Evidence and implications for macroeconomic models,” *American Economic Review* 95, 425–436.

Hamilton, J. D. (2003) “What Is an Oil Shock?” *Journal of Econometrics* 113: 363–98.

Hausman, J.A. (1978), "Specification Tests in Econometrics," *Econometrica*, 46(6), 1251-1271.

Hayashi, F., *Econometrics*. Princeton: Princeton University Press, 2000.

Imbens, G. (2014). “Instrumental Variables: An Econometrician’s Perspective.” *Statistical Science* 29, 323-58.

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

Jordà, Ò., M. Schularick, and A.M. Taylor (2015). “Betting the House,” *Journal of International Economics* 96, S2-S18.

Kilian, L. (2008). “Exogenous Oil Supply Shocks: How Big Are They and How Much Do They Matter for the U.S. Economy?” *Review of Economics and Statistics* 90, 216–40.

Kim, Y.J. and L. Kilian (2011). “How Reliable are Local Projection Estimators of Impulse Responses?” *The Review of Economics and Statistics* 93, 1460-1466.

Kleibergen, F. (2005). “Testing Parameters in GMM without Assuming They Are Identified,” *Econometrica* 73, 1103-1123.

Kuttner, K.N. (2001). “Monetary policy surprises and interest rates: Evidence fropm the Fed funds futures market.” *Journal of Monetary Economics* 47, 523-544.

Lanne, M. and P. Saikkonen (2013). “Noncausal Vector Autoregression.” *Econometric Theory* 29, 447-481.

Lechner, M. (2009). “Sequential Causal Models for the Evaluation of Labor Market Programs,” *Journal of the American Statistical Association* 27, 71-83.

Lütkepohl, H. (2005). *New Introduction to Multiple Time Series Analysis*. Berlin: Spring-Verlag.

Mertens, K. (2015). “Bonn Summer School – Advances in Empirical Macroeconomics, Lecture 2” (slide deck).

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

McCracken, M. and S. Ng (2016). “FRED-MD: A Monthly Database for Macroeconomic Research,” *Journal of Business and Economic Statistics* 34, 574-589.

Montiel Olea, J.L. and C. Pflueger (2013). “A Robust Test for Weak Instruments,” *Journal of Business and Economic Statistics* 31, 358-369.

Montiel Olea, J., J.H. Stock, and M.W. Watson (2017). “Inference in Structural VARs with External Instruments,” in preparation.

Moreira, M. (2003). “A Conditional Likelihood Ratio Test for Structural Models,” *Econometrica* 71, 1027-1048.

Plagborg-Møller, M. (2016a), “Estimation of Smooth Impulse Response Functions,” manuscript, Harvard University.

Plagborg-Møller, M. (2016b), “Bayesian Inference on Structural Impulse Response Functions,” manuscript, Harvard University.

Ramey, V. (2011). “Identifying government spending shocks: it’s all in the timing,” *Quarterly Journal of Economics* 126, 1–50

Ramey, V. (2016). “Macroeconomic Shocks and their Propagation,” *Handbook of Macroeconomics*, *Vol. 2A.* Amsterdam: Elsevier, pp. 71-162.

Ramey, V. and S. Zubairy, (2017). “Government Spending Multipliers in Good Times and in Bad: Evidence from U.S. Historical Data,” *Journal of Political Economy*, forthcoming.

Romer, C.D. and D.H. Romer (1989). “Does Monetary Policy Matter? A New Test in the Spirit of Friedman and Schwartz," in O.J. Blanchard and S. Fischer (eds.), *NBER Macroeconomics Annual* 1989. Cambridge, MA: MIT Press, 121-70.

Romer, C.D. and D.H. Romer (2010). “The Macroeconomic Effects of Tax Changes: Estimates Based on a New Measure of Fiscal Shocks,” *American Economic Review* 100, 763-801.

Rothenberg, T.J. and C.T. Leenders (1964). “Efficient Estimation of Simultaneous Equation Systems,” *Econometrica* 32, 57-76.

Sargan, D. (1964). “Three-Stage Least-Squares and Full Information Maximum Likelihoood Estimates,” *Econometrica* 32, 77-81.

Sims, C.A. (1980). “Macroeconomics and Reality,” *Econometrica* 48, 1-48.

Slutsky, E. (1937). The summation of random causes as the source of cyclic processes," *Econometrica* 5, 105-146.

Stock, J.H. (2008). *What's New in Econometrics: Time Series, Lecture 7*. Short course lectures, NBER Summer Institute, at http://www.nber.org/minicourse_2008.html.

Stock, J.H. and M.W. Watson (2012). "Disentangling the Channels of the 2007-09 Recession." *Brookings Papers on Economic Activity*, no. 1: 81-135.

Stock, J.H. and M.W. Watson (2016). "Dynamic Factor Models, Factor-Augmented Vector Autoregressions, and Structural Vector Autoregressions in Macroeconomics," *Handbook of Macroeconomics*, *Vol. 2A.* Amsterdam: Elsevier, pp. 415-525.

Zellner, A. (1962). "An Efficient Method of Estimating Seemingly Unrleated Regressions and Tests for Aggregation Bias," *Journal of the American Statistical Association* 63, 502-511.

Zellner, A. and H. Theil (1962). "Three-Stage Least Squares: Simultaneous Estimation of Simultaneous Equations." *Econometrica* 30, 54-78.

Table 1: Estimated causal effect of monetary policy shocks on selected economic variables: Gertler-Karadi (2015) variables, instrument and sample period

		LP-IV			SVAR		SVAR – LP
	lag (h)	**(a)**	**(b)**	**(c)**	**(d)**		**(d)-(b)**
R	0	1.00 (0.00)	1.00 (0.00)	1.00 (0.00)	1.00 (0.00)		0.00 (0.00)
	6	-0.07 (1.34)	1.12 (0.52)	0.67 (0.57)	0.89 (0.31)		-0.23 (1.19)
	12	-1.05 (2.51)	0.78 (1.02)	-0.12 (1.07)	0.78 (0.46)		0.00 (1.79)
	24	-2.09 (5.66)	-0.80 (1.53)	-1.57 (1.48)	0.40 (0.49)		1.19 (2.57)
IP	0	-0.59 (0.71)	0.21 (0.40)	0.03 (0.55)	0.16 (0.59)		-0.06 (0.35)
	6	-2.15 (3.42)	-3.80 (3.14)	-4.05 (3.65)	-0.81 (1.19)		3.00 (2.32)
	12	-3.60 (6.23)	-6.70 (4.70)	-6.86 (5.49)	-1.87 (1.54)		4.83 (4.00)
	24	-2.99 (10.21)	-9.51 (7.70)	-8.13 (7.62)	-2.16 (1.65)		7.35 (6.40)
P	0	0.02 (0.07)	-0.08 (0.25)	-0.04 (0.25)	0.02 (0.23)		0.10 (0.13)
	6	0.16 (0.42)	-0.39 (0.52)	-0.79 (0.83)	0.31 (0.41)		0.71 (0.98)
	12	-0.26 (0.88)	-1.35 (1.03)	-1.37 (1.23)	0.45 (0.54)		1.80 (1.53)
	24	-0.88 (3.08)	-2.26 (1.31)	-2.58 (1.69)	0.50 (0.65)		2.76 (2.60)
EBP	0	0.51 (0.61)	0.67 (0.40)	0.82 (0.49)	0.77 (0.29)		0.09 (0.24)
	6	0.22 (0.30)	1.33 (0.81)	1.66 (1.04)	0.48 (0.20)		-0.85 (0.51)
	12	0.56 (0.91)	0.84 (0.65)	0.91 (0.80)	0.18 (0.13)		-0.66 (0.55)
	24	-0.44 (1.29)	0.94 (0.66)	0.85 (0.76)	0.06 (0.07)		-0.88 (0.62)
Controls		none	4 lags of (z,y)	4 lags of (z,y,f)	12 lags of y 4 lags of z		na
First-stage F^{Hom}		1.7	23.7	18.6	20.5		na
First-stage F^{HAC}		1.1	15.5	12.7	19.2		na

Notes: The instrument, Z_t, is available from 1990m1-2012m6; the other variables are available from 1979m1-2012m6. The LP-IV estimates in (a)-(c) use data from 1990m1-2012m6. The VAR for (d) is computed over 1980m7-2012m6; and the IV-regression computed over 1990m5-2012m6. The numbers in parentheses are standard errors computed by Newey-West HAC with h+1 lags for the local projections, and using a parametric Gaussian bootstrap for the SVAR and the SVAR – LP differences shown in (e). In the final two rows F^{Hom} is the standard (conditional homoscedasticity, no serial correlation) first-stage F-statistic, while F^{HAC} is the Newey-West version using 12 lags in (a) and heteroskedasticity-robust (no lags) in (b), (c), and (d).

Table 2: Tests for VAR Invertibility (*p*-values)

	1Year Rate	**ln(IP)**	**ln(CPI)**	**GZ EBP**
VAR-LP difference (lags 0,6,12,24)	0.95	0.55	0.75	0.26
VAR Z-GC test	0.16	0.09	0.38	0.97

Notes: The first row is the bootstrap p-value for the test ξ in (26) of the null hypothesis that IV-LP and IV-SVAR causal effects are same for h = 0, 6, 12, and 24. The second row shows p-values for the F-statistic testing the null hypothesis that the coefficients on four lags of Z are jointly equal to zero in each of the VAR equations.