

IDENTIFICATION AND ESTIMATION OF DYNAMIC CAUSAL EFFECTS IN MACROECONOMICS USING EXTERNAL INSTRUMENTS*

James H. Stock and Mark W. Watson

External sources of as-if randomness — that is, external instruments — can be used to identify the dynamic causal effects of macroeconomic shocks. One method is a one-step instrumental variables regression (local projections – IV); a more efficient two-step method involves a vector autoregression. We show that, 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. If, however, lagged endogenous variables are needed as control variables in the one-step method, then the conditions for validity of the two methods are the same.

The identification and estimation of dynamic causal effects is a defining challenge of macroeconometrics. In the macroeconomic tradition dating to Frisch (1933) and Slutsky (1937), dynamic causal effects are conceived as the effect, over time, of an intervention that propagates through the economy, as modelled 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, 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 these models, structural shocks are the interventions of interest, and the goal is to estimate the dynamic causal effect of these shocks on macroeconomic outcomes.

In contrast, modern microeconomic identification strategies rely heavily on ‘external’ sources of variation that provide quasi-experiments to identify causal effects. Such external variation might be found, for example, in institutional idiosyncrasies that introduce as-if randomness in the variable of interest (the treatment). The use of

*Corresponding author: James H. Stock, Department of Economics, Harvard University, Cambridge MA 02138, USA. Email: james_stock@harvard.edu.

This work was presented by Stock as the Sargan Lecture to the Royal Economic Society on 11 April 2017. We thank Paul Beaudry, Mark Gertler, Oscar Jordà, Daniel Lewis, Karel Mertens, Mikkel Plagborg-Møller, Glenn Rudebusch, José Luis Montiel Olea, Valerie Ramey, Morten Ravn, Giovanni Ricco, Neil Shephard, Leif Anders Thorsrud, Christian Wolf and an anonymous referee for helpful comments and/or discussions.

such external instruments in microeconometrics has proven highly productive and has yielded compelling estimates of causal effects.

The subject of this article is the use of external instruments to estimate dynamic causal effects in macroeconomics. By an external instrument, we mean a variable that is correlated with a shock of interest, but not with other shocks, so that the instrument captures some exogenous variation in the shock of interest. These instruments are typically not a macro variable of ultimate interest, and as such they are external to the system. In referring to these instruments as external, we also connect with the original term for instruments, external factors (Wright, 1928).

External instruments can be used to estimate dynamic causal effects directly without an intervening VAR step. This method uses an instrumental variables (IV) version of what is called in the forecasting literature a direct multistep forecasting regression; in the impulse response literature, this method is called a local projection. Alternatively, the instruments can be used in conjunction with a VAR to identify structural 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 programme in macroeconometrics, 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. In many applications, the instrument is constructed as a partial measure of the shock of interest: the quantity of oil kept from market because of a political disruption, a change in fiscal policy not related to business cycle conditions, or the part of a monetary shock revealed during a monetary policy announcement window. Such constructed measures typically have measurement error, which in general leads to bias if the measure is treated as the true shock. However, that measurement error need not compromise the validity of the measure as an instrument. As in the microeconomic setting, finding such instruments is not easy. Still, in our view this research programme holds out the potential for more credible identification than is typically provided by SVARs identified using internal restrictions.

This article unifies and explicates a number of strands of recent work on external instruments in macroeconometrics. The idea that constructed shock series are best thought of as instruments is not new: Blanchard and Sims made this observation in the published general discussion of Romer and Romer (1989), but it seems not to have been followed up. To our knowledge, the earliest work to use constructed shocks as an instrument in a SVAR is Beaudry and Saito (1998), who use the Romer and Romer (1989) indicators to estimate impulse responses to monetary shocks. 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 Karadi (2015), Caldara and Kamps (2017) and a growing list of other researchers. Turning to single-equation methods, Hamilton (2003) developed a list of exogenous oil supply disruptions, which he used as an instrument for autoregressive-distributed lag estimation of the effect of oil supply shocks on GDP. The modern use of external instruments to estimate structural impulse response functions directly (that is, without estimating a VAR or iterating) dates to Jordà *et al.* (2015) and Ramey and Zubairy (2017), and is clearly

exposed in Ramey (2016). The condition for instrument validity in the direct regression without control variables, given in Section 1 below, appears in unpublished lecture notes by Mertens (2015). Those notes and Fieldhouse *et al.* (2017) discuss the extension of these conditions to control variables. Jordà *et al.* (2015) and Ramey (2016) call these direct IV regressions ‘local projections-IV’ (LP-IV) in reference to Jordà’s (2005) method of local projections (LP) on which it builds. We adopt this terminology while noting that these IV regressions emerge from the much older tradition of simultaneous equations estimation in macroeconomics pioneered by Sargan and his contemporaries. Although these methods increasingly are being used in applications, we are not aware of a unified presentation of the econometric theory of and connections between the SVAR-IV and LP-IV methods.

In addition to exposing the use of external instruments in macroeconomics, this article 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 recovered from current and lagged values of the observed data. Because of the dynamic nature of the macroeconomic problem, exogeneity of the instrument entails a strong ‘lead-lag exogeneity’ requirement that the instrument be uncorrelated with past and future 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 lead-lag exogeneity. But to be valid, this method requires invertibility. 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 of the structural impulse response function (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 the local asymptotic power of the test against this alternative. The focus of this test on the impulse response function – the estimand of interest – differs 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 lead-lag exogeneity condition because it is correlated with past shocks. A natural approach is to include additional regressors – lagged macro variables – that control for the lagged shocks. We show, however, that the condition for these control variables to produce valid inference in LP-IV is in general equivalent to assuming invertibility of the corresponding VAR, in which case SVAR-IV provides more efficient inference.

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

Following Ramey (2016), we illustrate these methods using Gertler and Karadi'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 substantial simplifications made throughout this article. First, we focus exclusively on linear models and identification through second moments, so that conditional expectations are typically replaced by projections. Second, we assume homogenous treatment effects so that valid instruments all have the same estimand (i.e. the local average treatment effect equals the average treatment effect). Both these simplifications are non-trivial. The assumption of non-linearity, in particular, rules out a frequent justification for using LP methods (either OLS or LP-IV), which is that LP methods can estimate non-linear effects without needing to model them as a system. That said, there is a tension between the assumption that the control variables and specification are correct in the single-equation specification, and what this must imply for the full system, and this tension is unresolved in the literature and merits further investigation. We return to this point in the conclusions.

Finally, we use two notational devices: the subscript '2:n' denotes the elements of a vector or matrix other than the first row or column, and {...} denotes a linear combination of the terms inside the braces.

1. 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 microeconomic setting.

1.1. Causal Effects and Instrumental Variables Regression

Our starting point is that the expected difference in outcomes between the treatment and control groups in a randomised controlled experiment with a binary treatment is the average treatment effect.¹ 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 linear model $Y = \gamma + X\beta + u$, where β is the treatment effect, random assignment implies that

¹ This starting point is actually a result, or conclusion, of a vast literature on defining causal effects for statistical analysis. See Imbens (2014) for a review, including discussion of both the potential outcomes framework and graphical models.

$E(u|X) = 0$ so that the population regression coefficient is the treatment effect. If randomisation 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 assumptions of linearity and homogeneous treatment effects, this treatment effect is estimated by ordinary least squares estimation of

$$Y = \beta X + \gamma'W + u, \quad (1)$$

where the intercept has been subsumed in $\gamma'W$.

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 ${}^\perp$ denote the residual from the population projection onto W , for example, $X^\perp = X - \text{Proj}(X|W)$. If the instrument satisfies the conditions:

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

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

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

1.2. Dynamic Causal Effects and the Structural Moving Average Model

In macroeconomics, we can imagine a counterpart of a randomised controlled experiment. For example, in the US, 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 US 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

$$E_t(Y_{2,t+h}|\varepsilon_{1,t} = 1) - E_t(Y_{2,t+h}|\varepsilon_{1,t} = 0).$$

We now assume linearity and stationarity, assumptions we maintain henceforth. With these assumptions, the h -lag treatment effect is the population coefficient in the regression:

$$Y_{2,t+h} = \Theta_{h,21}\varepsilon_{1,t} + u_{t+h}, \quad (4)$$

where throughout we omit constant terms for convenience. Because $\varepsilon_{1,t}$ is randomly assigned, $E(u_{t+h}|\varepsilon_{1,t}) = 0$, so $\Theta_{h,21} = E(Y_{2,t+h}|\varepsilon_{1,t} = 1) - E(Y_{2,t+h}|\varepsilon_{1,t} = 0)$. Thus, $\Theta_{h,21}$ is the causal effect of treatment 1 on variable 2, h periods after the treatment. Were $\varepsilon_{1,t}$ observed, this causal effect could be estimated by OLS estimation of (4).

The path of causal effects mapped out by $\Theta_{h,21}$ for $h = 0, 1, 2, \dots$ is the dynamic causal effect of treatment 1 on variable 2.²

The macroeconomic jargon for this random treatment $\varepsilon_{1,t}$ is a ‘structural shock’: a primitive, unanticipated economic force, or driving impulse, that is unforecastable and uncorrelated with other shocks.³ The macroeconomist’s shock is 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 Slutsky–Frisch paradigm represents the path of observed macroeconomic variables as arising from current and past shocks and measurement error. If we collect all such structural shocks and measurement error together in the $m \times 1$ vector ε_b , the $n \times 1$ vector of macroeconomic variables Y_t can be written in terms of current and past ε_t :

$$Y_t = \Theta(L)\varepsilon_t, \quad (5)$$

where L is the lag operator and $\Theta(L) = \Theta_0 + \Theta_1 L + \Theta_2 L^2 + \dots$, where Θ_h is an $n \times m$ matrix of coefficients. The shock variance matrix $\sum_{\varepsilon\varepsilon'} = E\varepsilon\varepsilon'$ is assumed to be positive definite to rule out trivial (non-varying) shocks. We assume that the shocks are mutually uncorrelated. 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.

The assumption that the structural shocks are mutually uncorrelated accords both with their interpretation as randomly assigned treatments and with their being primitive economic forces; see Ramey (2016) for a discussion. We assume that any measurement error included in ε_t is uncorrelated with the structural shocks, although measurement error could be correlated across variables. Because $\varepsilon_{1,t}$ is uncorrelated with the other shocks and with any measurement error, the causal effect can be written as $E(Y_{2,t+h}|\varepsilon_{1,t} = 1, \varepsilon_{2:n,t}, \varepsilon_s, s \neq t) - E(Y_{2,t+h}|\varepsilon_{1,t} = 0, \varepsilon_{2:n,t}, \varepsilon_s, s \neq t)$. Although conditioning on the other shocks is redundant by randomisation, this alternative expression connects with the definition, seen in the older macro literature, of the causal effect as the partial derivative $\partial Y_{2,t+h} / \partial \varepsilon_{1,t}$ holding all other shocks constant.

Representation (5) is the structural moving average representation of Y_t . The coefficients of $\Theta(L)$ are the structural impulse response functions, which are the dynamic causal effects of the shocks. In general, the number of shocks plus measurement errors, m , can exceed the number of observed variables, n .

The recognition that, if $\varepsilon_{1,t}$ were observed, $\Theta_{h,21}$ could be estimated by OLS estimation of (4) – or by OLS estimation of the distributed lag regression of Y_t on $\varepsilon_{1,t}, \varepsilon_{1,t-1}, \varepsilon_{1,t-2}, \dots$ – underpins a productive and insightful research programme. In this programme, which dates to Romer and Romer (1989), researchers aim to measure directly a specific macroeconomic shock. Influential examples include Rudebusch

² There is a literature that defines dynamic causal effects in terms of primitives and connects those to what can be identified in an experiment with data collected over time; see Lechner (2009), Angrist *et al.* (2018), Jordà *et al.* (2017) and especially Bojinov and Shephard (2017) for discussion and references. With the additional assumptions of linearity and stationarity, Bojinov and Shephard’s (2017) dynamic potential outcomes framework leads to (4).

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

(1998), who measured monetary shocks by Fed Funds surprises controlling for employment report announcements, and Kuttner (2001), Cochrane and Piazzesi (2002), Faust *et al.* (2003), Gürkaynak *et al.* (2005) and Bernanke and Kuttner (2005), all of whom used interest rate changes around Federal Reserve announcement dates to measure monetary policy shocks.

1.3. 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 Hamilton's (2003) and Kilian's (2008) lists of exogenous oil supply disruptions. In all of 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 normalisation 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:

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

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

The unit effect normalisation has advantages over the more common unit standard deviation normalisation, which sets $\text{var}(\varepsilon_{1,t}) = 1$. Most importantly, the unit effect normalisation allows for direct estimation of the dynamic causal effect in the native units relevant to policy analysis. While one can convert one normalisation to another, doing so entails rescaling by estimated values and care must be taken to conduct inference incorporating that normalisation (we elaborate on this below). As discussed in Stock and Watson (2016), the unit effect normalisation also allows for direct extension of SVAR methods to structural dynamic factor models.

The unit effect normalisation underpins the local projection approach because it allows the regression (4) to be rewritten in terms of an observable regressor, $Y_{1,t}$. Specifically, use the unit effect normalisation to write $Y_{1,t} = \varepsilon_{1,t} + \{\varepsilon_{2:n,t} \varepsilon_{t-1}, \varepsilon_{t-2}, \dots\}$ (recall the notational devices that $\varepsilon_{2:n,t} = (\varepsilon_{2,t}, \dots, \varepsilon_{n,t})'$ and that $\{\dots\}$ denotes a linear combination of the terms in braces). Rewriting this expression in terms of $\varepsilon_{1,t}$ and substituting it into (4) yields

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

where $u_{i,t+h}^h = \{\varepsilon_{t+h}, \dots, \varepsilon_{t+1}, \varepsilon_{2:n,t}, \varepsilon_{t-1}, \varepsilon_{t-2}, \dots\}$. Because $Y_{1,t}$ is endogenous, it is correlated with $u_{i,t+h}^h$, so 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 instrumental variables. These instruments can be used to estimate the dynamic causal effect using (7) if they satisfy:

CONDITION LP-IV

- (i) $E(\varepsilon_{1,t} Z_t') = \alpha' \neq 0$ (relevance);
- (ii) $E(\varepsilon_{2:n,t} Z_t') = 0$ (contemporaneous exogeneity);
- (iii) $E(\varepsilon_{t+j} Z_t') = 0$ for $j \neq 0$ (lead-lag exogeneity).

Conditions LP-IV (i) and (ii) are conventional IV relevance and exogeneity conditions, and are the counterparts of the microeconomic conditions equations (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 generally not restrictive: when Z_t contains only variables realised 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 a serially correlated 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 (5) imply that $E(Y_{i,t+h} Z_t') = \Theta_{h,i1} \alpha'$. Second, Condition LP-IV, the unit effect normalisation (6) and (5) imply that $E(Y_{1,t} Z_t') = \alpha'$. Thus, when Z_t is a scalar:

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

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

These moment expressions provide an intuitive interpretation of LP-IV. Suppose that $Y_{i,t}$ is GDP growth, $Y_{1,t}$ is the federal funds rate and Z_t is a monetary policy announcement instrument, constructed so that it satisfies Condition LP-IV. Then the causal effect of a monetary policy shock on GDP growth h periods hence is estimated by regressing $\Delta \ln \text{GDP}_{t+h}$ on FF_t using the announcement surprise Z_t as an instrument. In this two-stage least squares interpretation, the unit effect normalisation is imposed automatically.

Another interpretation of the moment condition (8) relates to the distributed lag representation of Y_i in terms of Z_i :

$$Y_t = \Pi(L)Z_t + v_t. \tag{9}$$

This is Theil and Boot’s (1962) final form of the dynamic model for (Y_t, Z_t) . It is also the time series counterpart to what is (somewhat confusingly) called the reduced form for non-dynamic simultaneous equations systems. In the non-dynamic setting with a single instrument, a familiar result is that the Wald IV estimator is the ratio of the reduced-form coefficients. Similarly, in the dynamic context, when Z_t is serially uncorrelated and a scalar, $\Theta_{h,i1}$ is the ratio of the h th distributed lag coefficient in the $Y_{i,t}$ equation, $\Pi_{h,i}$ to the impact effect on the first variable, $\Pi_{0,1}$; that is, $\Theta_{h,i1} = \Pi_{h,i}/\Pi_{0,1}$. In the monetary policy announcement example, $\Pi(L)$ is the impulse response function of Y_t with respect to the announcement surprise. The older literature treated this as the causal effect of interest, but as explained in Gertler and Karadi (2015), the surprise is better thought of as an instrument for the shock. Akin to the Wald estimator in the static setting, the IV estimator of the dynamic causal effect is the impulse response function of the effect of the shock on $\Delta \ln \text{GDP}$, divided by the impact effect of the announcement on the Federal Funds rate.

The lag exogeneity Condition LP-IV(iii) is testable: Z_t should be unforecastable in a regression of Z_t on lags of Y_t . If the lag exogeneity condition fails, then the LP-IV methods laid out in this subsection are not valid because Z_t will be correlated with the error u_{t+h} in (4). This problem can potentially be addressed by adding control variables to the LP-IV regression.

1.4. *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 4, this is the case in the Gertler and Karadi (2015) application.

Second, even if Condition LP-IV is satisfied, including control variables could reduce the sampling variance of the IV estimator by reducing the variance of the error term. The reasoning is standard: because the variance of the LP-IV estimator depends on the scale of the errors, including control variables that explain the error term can reduce the variance of the estimator. Here, the relevant variance is the long-run variance of the instrument-times-error, so the aim of including additional control variables is to reduce this long-run variance. Under Condition LP-IV, Y_{t-1}, Y_{t-2}, \dots and possibly future Z_{t+h}, \dots, Z_{t+1} are candidate control variables.

Adding control variables W_t to (7) yields

$$Y_{i,t+h} = \Theta_{h,i1} Y_{1,t} + \gamma_h' W_t + u_{i,t+h}^{\perp}, \tag{10}$$

where $x_t^{\perp} = x_t - \text{Proj}(x_t|W_t)$ for some variable x_t and $u_{i,t+h}^{\perp} = \{\varepsilon_{t+h}^{\perp}, \dots, \varepsilon_{t+1}^{\perp}, \varepsilon_{2:n,t}^{\perp}, \varepsilon_{t-1}^{\perp}, \varepsilon_{t-2}^{\perp}, \dots\}$.

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

CONDITION LP-IV[⊥]

- (i) $E(\varepsilon_{1,t}^\perp Z_t^{\perp'}) = \alpha' \neq 0$;
- (ii) $E(\varepsilon_{2:n,t}^\perp Z_t^{\perp'}) = 0$; and
- (iii) $E(\varepsilon_{t+j}^\perp Z_t^{\perp'}) = 0$ for $j \neq 0$.

By projecting on W_b , (10) can be written, $Y_{i,t+h}^\perp = \Theta_{h,i1} Y_{i,t}^\perp + u_{i,t+h}^{h\perp}$. For a scalar instrument, multiplying both sides of this expression by $Z_t^{\perp'}$ and using Condition LP-IV[⊥] and the unit effect normalisation (6) yields

$$\frac{E(Y_{i,t+h}^\perp Z_t^{\perp'})}{E(Y_{1,t}^\perp Z_t^{\perp'})} = \Theta_{h,i1}. \tag{11}$$

For a vector of instruments, $E(Y_{i,t+h}^\perp Z_t^{\perp'})\text{HE}(Z_t^\perp Y_{1,t}^\perp)/E(Y_{1,t}^\perp Z_t^{\perp'})\text{HE}(Z_t^\perp Y_{1,t}^\perp) = \Theta_{h,i1}$ for any positive definite matrix H . Equation (11) is the moment condition for IV estimation of (10) using instrument Z_t .

Equation (11) holds for all h , including the impact effect $h = 0$, with the proviso that for $h = 0$, the effect for the first variable is normalised to $\Theta_{0,11} = 1$. Under the unit effect normalisation, for $h = 0$ and $i = 1$, (10) become the identity $Y_{1,t} = Y_{1,t}$ (or $cY_{1,t}^\perp = Y_{1,t}^\perp$).

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 $Y_{i,t}$. Such control variables could include lagged values of Y_1 and Y_b , 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. Whether or not lagged Y s are used as controls, under Condition LP-IV(iii), leads and lags of Z_t can be included as controls to improve efficiency.

A 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[⊥] (iii) holds. It is useful to think of two cases.

In the first case, suppose Z_t is correlated with past values of $\varepsilon_{1,b}$ but not with past values of other shocks. As we discuss below, this situation arises in the Gertler and 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 is 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[⊥] (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_r . 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 subsection 2.2 that the requirement that Condition LP-IV[⊥] (iii) be satisfied by generic controls, when Condition LP-IV (iii) does not hold, is quite strong.

1.5. LP-IV: Econometric Odds and Ends

1.5.1. Levels, differences and cumulated impulse responses

In many applications, $Y_{i,t}$ will be specified in first differences, but interest is in 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:

$$\sum_{k=0}^h Y_{i,t+k} = \Theta_{h,i1}^{cum} Y_{1,t} + \gamma_h^{cum'} W_t + u_{i,t+h}^{h,cum\perp}, \tag{12}$$

where $\Theta_{h,i1}^{cum} = \sum_{k=0}^h \Theta_{k,i1}$. For example, if $Y_{i,t} = \Delta \ln \text{GDP}_t$, then the left-hand side of (12) is $\ln(\text{GDP}_{t+h}) - \ln(\text{GDP}_t)$, that is, the log-point change in GDP from t to $t+h$.

If Z_t satisfies LP-IV[⊥], it is a valid instrument for IV estimation of (12).

Another measure of a dynamic causal effect is the ratio of cumulative impulse responses. For example, a shock to government spending typically induces a flow over time of government outlays. As discussed by Ramey and Zubairy (2017, section 3.2.2), a useful measure of the effect on output of government spending is the cumulative GDP gain resulting from cumulative government spending over the same period. Fieldhouse *et al.* (2017) make a similar argument for considering ratios of cumulative multipliers in their study of the effect on residential investment of US housing agency purchases of mortgage-backed securities. As Ramey and Zubairy (2017) point out, this ratio of cumulative multipliers can be estimated in the LP-IV regression:

$$\sum_{k=0}^{h_i} Y_{i,t+k} = \rho_{i1}^{h_i,h_1} \sum_{k=0}^{h_1} Y_{1,t+k} + \gamma_{h_i,h_1}^{cum'} W_t + u_{i,t+h}^{h_i,h_1}, \tag{13}$$

where $\rho_{i1}^{h_i,h_1} = \sum_{k=0}^{h_i} \Theta_{k,i1} / \sum_{k=0}^{h_1} \Theta_{k,11}$ (in (13), we generalise Ramey and Zubairy, 2017 slightly to allow for different cumulative periods for Y_i and Y_1). When the instrument Z_t satisfies Condition LP-IV[⊥], $E \sum_{k=0}^{h_i} Y_{i,t+k} z_t^\perp = \sum_{k=0}^{h_i} \Theta_{k,i1} \alpha'$ and $E \sum_{k=0}^{h_1} Y_{1,t+k} z_t^\perp = \sum_{k=0}^{h_1} \Theta_{k,11} \alpha'$. Thus, when there is a single instrument, the IV moment condition is $E \sum_{k=0}^{h_i} Y_{i,t+k} z_t^\perp / E \sum_{k=0}^{h_1} Y_{1,t+k} z_t^\perp = \sum_{k=0}^{h_i} \Theta_{k,i1} / \sum_{k=0}^{h_1} \Theta_{k,11} = \rho_{i1}^{h_i,h_1}$. Thus, if Z_t satisfies LP-IV[⊥], it is a valid instrument for IV estimation of (13).

1.5.2. *HAC/HAR inference and long-horizon impulse responses*

When the instruments are strong, the validity of inference can be justified under standard assumptions of stationarity, weak dependence and existence of moments (Hayashi, 2000). However, the multistep nature of the direct regressions in general requires an adjustment for serial correlation of the instrument \times error process: the error terms in (7), (10) and (12) include future and lagged values of ε_t , and in general terms like $Z_t \varepsilon_{t+j}$ and $Z_{t+h} \varepsilon_t$ will be correlated. Inference based on standard heteroscedasticity and autocorrelation robust (HAR) covariance matrix estimators are valid at short to medium horizons.

One special case in which HAR inference is not needed is when the W s are lagged Y s, the VAR for Y is invertible and the Z s are serially uncorrelated conditional on the W s. In this case, $Z_t^\perp u_{t+h}^\perp$ is serially uncorrelated⁴ and standard heteroscedasticity-robust standard errors can be used. If in addition the errors are homoscedastic, homoscedasticity-only standard errors can be used.

1.5.3. *Historical and forecast error variance decompositions*

The historical decomposition decomposes the path of Y_t to the contributions of the individual shocks. The contribution of shock $\varepsilon_{1,t}$ to $Y_{i,t+h}$ can be read off the structural moving average representation (5):

$$\text{Historical contribution of } \varepsilon_{1,t} \text{ to } Y_{i,t+h} = \Theta_{h,i1} \varepsilon_{1,t}. \tag{14}$$

The forecast error variance decomposition (FEVD) decomposes the variance of the unforecasted change in a variable h periods hence to the variance contributions from the shocks that occurred between t and $t+h$. Because the shocks are uncorrelated over time and with each other, this decomposition, expressed in R^2 form is as follows:

$$FEVD_{h,i1} = \frac{\sum_{k=0}^{h-1} \Theta_{k,i1}^2 \sigma_{\varepsilon_1}^2}{\text{var}(Y_{i,t+h} | \varepsilon_t, \varepsilon_{t-1}, \dots)}. \tag{15}$$

If $\varepsilon_{1,t}$ can be recovered, then the historical decomposition can be computed using the LP-IV estimates of $\{\Theta_{h,j1}\}$, $h = 0, 1, 2, \dots$. Similarly, if $\sigma_{\varepsilon_1}^2$ and $\text{var}(Y_{i,t+h} | \varepsilon_t, \varepsilon_{t-1}, \dots)$ are identified, then the forecast error variance decomposition is identified and also can be computed using the LP-IV estimates of $\Theta_{h,j1}$, $h = 0, 1, 2, \dots$.

In general, even though Conditions LP-IV and LP-IV[⊥] serve to identify the impulse response function, they do not identify either $\varepsilon_{1,t}$ or $\sigma_{\varepsilon_1}^2$ without additional assumptions. A sufficient condition for identifying $\varepsilon_{1,t}$ and the FEVD is that the VAR for Y_t is invertible; a somewhat weaker condition for identifying $\varepsilon_{1,t}$ (but not the FEVD) is that Y_t is partially invertible. Weaker yet is the ‘recoverability’ condition discussed in Chahrouh and Jurado (2017) and Plagborg-Møller and Wolf (2017). Further discussion, including expressions for $\varepsilon_{1,t}$, $\sigma_{\varepsilon_1}^2$, and the FEVD, are deferred until the next Section.

⁴ This result follows by direct calculation using the invertibility results in subsection 2.2.

1.5.4. *Smoothness restrictions*

The IV estimator of (7), (10) and (12) 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 modelling 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 $\Theta(L)$ is invertible. A few recent papers develop methods for smoothing IRFs estimated by local projections using OLS. Barnichon and Brownlees (2016) and Plagborg-Møller (2016a) use smoothness priors to shrink the IRFs across horizons. Miranda-Agrippino and Ricco (2017) smooth LP IRFs by shrinking them towards SVAR IRVs. Although these papers develop these methods for OLS estimates of LP and SVARs, the extension to IV estimates seems straightforward.

1.5.5. *Weak instruments*

If the instruments are weak, then in general the distribution of the IV estimator in (7), (10) and (12) is not centred at $\Theta_{h,i1}$, and inference based on conventional IV standard errors is unreliable. However, a suite of heteroscedasticity and autocorrelation-robust methods now exists to detect weak instruments and to conduct inference robust to weak instruments in linear IV regression. For example, see Kleibergen (2005) for a HAR version of Moreira's (2003) conditional likelihood ratio statistics, and Montiel Olea and Pflueger (2013) and Andrews (2018) for HAR alternatives to first-stage F statistics for detecting weak identification.

As previously discussed, HAR inference is not needed in the special case that the W s are lagged Y s, the VAR for Y is invertible and the Z s are serially uncorrelated conditional on the W s. If in addition the errors are homoscedastic, then the suite of tools for weak identification in homoscedastic cross-section data can be applied, including the usual first-stage F statistic for assessing instrument strength.

1.5.6. *News shocks and the unit-effect normalisation*

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 legislative process but have direct effects on government spending and/or taxes 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 normalisation 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 normalisation ($\Theta_{0,11} = 1$) is inappropriate.

Instead, for a news shock, a k -period ahead unit-effect normalisation, $\Theta_{k,11} = 1$ 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 normalisation $\Theta_{12,11} = 1$ would be appropriate: this normalises the spending shock so that a 1 pp increase in the shock at time t corresponds to a 1 pp increase in observed government spending 12 months hence. With this k -period ahead normalisation, $Y_{1,t+k} = \varepsilon_{1,t} + \{\varepsilon_{t+k}, \dots, \varepsilon_{t+1}, \varepsilon_{2:n,t}, \varepsilon_{t+1}, \varepsilon_{t+2}, \dots\}$. Accordingly, $Y_{1,t+k}$ replaces $Y_{1,t}$ in the IV

regressions (7), (10) and (12). In practice, implementing this strategy requires a choice of the news lead-time k , and this choice would be informed by application-specific knowledge.

2. 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 modelling 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.

2.1. 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:

$$A(L)Y_t = v_t, \quad (16)$$

where $A(L) = I - A_1L - A_2L^2 - \dots$. 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 $v_t = Y_t - \text{Proj}(Y_t | Y_{t-1}, Y_{t-2}, \dots)$ 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:

$$v_t = \Theta_0 \varepsilon_t \text{ where } \Theta_0 \text{ is non-singular.} \quad (17)$$

A necessary condition for (17) to hold is that the number of variables in the VAR equal the number of shocks ($n = m$).

Because Y_t is second-order stationary, $A(L)$ is invertible. Thus, (16) and (17) yield a moving average representation in terms of the structural shocks:

$$Y_t = C(L)\Theta_0\varepsilon_t, \tag{18}$$

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

If (17) holds, then the SVAR impulse response function reveals the population dynamic causal effects; that is, $C(L)\Theta_0 = \Theta(L)$.⁵ Condition (17) is an implication of the assumption that the structural moving average is invertible. This ‘invertibility’ assumption, which underpins SVAR analysis, is non-trivial 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 summarise 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); and
- (ii) $E\varepsilon_{2:n,t}Z_t' = 0$ (exogeneity w.r.t. other current shocks).

Condition SVAR-IV and (17) imply that

$$E v_t Z_t = E(\Theta_0 \varepsilon_t Z_t) = \Theta_0 E \begin{pmatrix} \varepsilon_{1t} Z_t' \\ \varepsilon_{2:n,t} Z_t' \end{pmatrix} = \Theta_0 \begin{pmatrix} \alpha' \\ 0 \end{pmatrix} = \begin{pmatrix} \Theta_{0,11} \alpha' \\ \Theta_{0,2:n,1} \alpha' \end{pmatrix}. \tag{19}$$

With the help of the unit effect normalisation (6), it follows from (19) that, in the case of scalar Z_t :

$$\frac{E(v_{i,t} Z_t)}{E(v_{1,t} Z_t)} = \Theta_{0,i1}, \tag{20}$$

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

$$v_{i,t} = \Theta_{0,i1} v_{1,t} + \{\varepsilon_{2:n,t}\} \tag{21}$$

using the instrument Z_t .

Because the innovations v_i are not observed, the IV regression (21) is not feasible. One possibility is replacing the population innovations in (21) with their sample counterparts \hat{v}_i , 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 since $\hat{v}_{1,t}$ is a generated regressor.

Instead, $\Theta_{0,i1}$ can be estimated by an approach that directly yields the correct large-sample, strong-instrument standard errors. Because $v_{i,t} = Y_{i,t} - \text{Proj}(Y_{i,t} | Y_{t-1}, Y_{t-2}, \dots)$, (21) can be rewritten as

$$Y_{i,t} = \Theta_{0,i1} Y_{1,t} + \gamma_i(L) Y_{t-1} + \{\varepsilon_{2:n,t}\}, \tag{22}$$

⁵ Note that from (5) and (16), $v_t = A(L)\Theta(L)\varepsilon_t$. With the addition of condition (17), we have $\Theta_0\varepsilon_t = A(L)\Theta(L)\varepsilon_t$, so that $\Theta_0 = A(L)\Theta(L)$ and $\Theta(L) = A(L)^{-1}\Theta_0 = C(L)\Theta_0$.

where $\gamma_i(L)$ are the coefficients of $\text{Proj}(Y_{i,t} - \Theta_{0,i1} Y_{1,t} | Y_{t-1}, Y_{t-2}, \dots)$. The coefficients $\Theta_{0,i1}$ and $\gamma_i(L)$ 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 summarise, SVAR-IV proceeds in three steps:

- (i) Estimate (22) using instruments Z_t for the variables in Y_b using p lagged values of Y_t as controls. This, along with the unit effect normalisation $\Theta_{0,11} = 1$, yields the IV estimator of the first column of Θ_0 , $\hat{\Theta}_{0,1}^{\text{SVAR-IV}}$.
- (ii) Estimate a VAR(p) and invert the VAR to obtain $\hat{C}(L) = \hat{A}(L)^{-1}$.
- (iii) Estimate the dynamic causal effects of shock 1 on the vector of variables as

$$\hat{\Theta}_{h,1}^{\text{SVAR-IV}} = \hat{C}_h \hat{\Theta}_{0,1}^{\text{SVAR-IV}}. \tag{23}$$

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 $Y_{t-1}, Y_{t-2}, \dots, Y_{t-p}$. 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.

2.1.1. *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), $\sqrt{T}(\hat{\Gamma} - \Gamma) \xrightarrow{p} N(0, \Sigma_\Gamma)$. And, because estimator $\hat{\Theta}_{h,1}^{\text{SVAR-IV}}$ from Step (iii) is a smooth function of $\hat{\Gamma}$, $\sqrt{T}(\hat{\Theta}_{h,1}^{\text{SVAR-IV}} - \Theta_{h,1}) \xrightarrow{d} N(0, \Sigma_\Theta)$ where Σ_Θ can be calculated using the δ -method. Alternatively, and often more conveniently, confidence intervals can be computed using a parametric bootstrap. Doing so requires specifying an auxiliary process for Z_t . We provide some details in Appendix A in the context of our empirical illustration.

When instruments are weak, the asymptotic distribution of $\hat{\Theta}_{h,1}^{\text{SVAR-IV}}$ is not normal; Montiel Olea *et al.* (2017) discuss weak-instrument robust inference for SVARs identified by external instruments.

We stress that the normalisation of ultimate interest – typically the unit effect normalisation – needs to be incorporated into the computation of standard errors. In general, it is incorrect to use a different normalisation (such as the unit standard deviation normalisation), compute confidence bands, then rescale the bands and point estimates to obtain the unit effect normalisation. In practice, this means the unit effect normalisation must be ‘inside’ the bootstrap, not ‘outside’.

2.1.2. *Different data spans for Z and Y (‘unbalanced panels’)*

The SVAR-IV estimator of the impulse response function in (23) has two parts, \hat{C}_h and $\hat{\Theta}_{0,1}^{\text{SVAR-IV}}$. In general these can be estimated over different sample periods. For example, in Gertler and 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 at all horizons.

In contrast, there is less opportunity to improve efficiency by using the longer sample for Y using LP-IV. If Z satisfies Condition LP-IV, then the estimation must all be done on the shorter sample because the moments in (8) are only available over the period of overlap of the Y and Z samples. If control variables are included, the longer sample can be used to estimate Y_t^\perp and Y_{t+h}^\perp , but the moments in (11) must still be estimated over the period of overlap of the Y and Z samples.

A related limitation of LP-IV is that the number of observations available for estimation decreases with the horizon h . This is true regardless of whether the data samples for Z and Y are the same, but becomes more of an issue (compared to SVAR-IV) if the sample for Z is already short.

2.1.3. *News shocks and the unit-effect normalisation*

A structural moving average may be invertible even when it includes news shocks as long as Y_t contains forward-looking variables. But, as discussed in the previous Section, news variables require a change in the unit-effect normalisation from contemporaneous $\Theta_{0,11} = 1$ to k periods ahead $\Theta_{k,11} = 1$. To implement this normalisation in the SVAR, note that the effect of ε_t on Y_{t+k} is given by $\eta_t = \Theta_k \varepsilon_t = C_k \Theta_0 \varepsilon_t = C_k v_t$. The k -period ahead unit-effect normalisation is $\Theta_{k,11} = 1$, so $\eta_{1,t} = \varepsilon_{1,t} + \{\varepsilon_{2:n,t}\}$. Thus, letting $X_t = \hat{C}_k Y_t$, the normalisation is implemented by replacing $Y_{1,t}$ with $X_{1,t}$ in (22) and carrying out the three steps given above. Because $X_{1,t}$ is a generated regressor, standard errors differ from the model using $Y_{1,t}$ and are most easily calculated using simulation (parametric bootstrap) methods like those outlined in Appendix A.

2.1.4. *Historical and forecast error variance decompositions*

As discussed in subsection 1.4, if the shock is identified, then the historical decomposition can be computed using (14). The forecast error variance decomposition, given in (15), further requires identification of $\sigma_{\varepsilon_1}^2$ and the object in the denominator of that expression. The IRFs (Θ 's) appearing in (14) and (15) can be estimated using either LP-IV or SVAR-IV. By using the same estimator for the IRFs and the historical decompositions, the set of results will be internally consistent.

The shock $\varepsilon_{1,t}$, $\sigma_{\varepsilon_1}^2$ and the denominator of (15) are all identified from $\Theta_{0,1}$ if the VAR is invertible. Specifically, if (17) holds, then $\varepsilon_{1,t} = \lambda' v_t$ where $\lambda = \Theta_{0,1}' \Sigma_{vv}^{-1} / (\Theta_{0,1}' \Sigma_{vv}^{-1} \Theta_{0,1})$.⁶ It follows from this expression that $\sigma_{\varepsilon_1}^2 = \lambda' \Sigma_{vv} \lambda = (\Theta_{0,1}' \Sigma_{vv}^{-1} \Theta_{0,1})^{-1}$. Also, under invertibility the denominator of (15) is $\text{var}(Y_{i,t+h} | \varepsilon_t, \varepsilon_{t-1}, \dots) = \text{var}(Y_{i,t+h} | v_t, v_{t-1}, \dots) = \text{var}(Y_{i,t+h} | Y_t, Y_{t-1}, \dots)$, so the denominator is also identified. Thus, if $\Theta_{0,1}$ is identified and if the VAR is invertible, the historical decomposition and FEVD are also identified.

⁶ To show this result, first write $\Theta_{0,1}' \Sigma_{vv}^{-1} v_t = \Theta_{0,1}' (\Theta_0 \Sigma_{\varepsilon\varepsilon} \Theta_0')^{-1} v_t = \Theta_{0,1}' (\Theta_0')^{-1} \Sigma_{\varepsilon\varepsilon}^{-1} \Theta_0^{-1} v_t = e_1' \Sigma_{\varepsilon\varepsilon}^{-1} \varepsilon_t = \sigma_{\varepsilon_1}^2$, where the first line uses (17) to write $\Sigma_{vv} = \Theta_0 \Sigma_{\varepsilon\varepsilon} \Theta_0'$; the second line uses invertibility of Θ_0 ; the third line uses the fact that $B^{-1} B_1 = e_1$ (the first unit vector) where B_1 is the first column of the invertible matrix B and uses (17) plus invertibility to write $\varepsilon_t = \Theta_0^{-1} v_t$ and the final line uses the assumption that $\varepsilon_{1,t}$ is uncorrelated with $\varepsilon_{2:n,t}$. Similar algebra shows that $\Theta_{0,1}' \Sigma_{vv}^{-1} \Theta_{0,1} = 1/\sigma_{\varepsilon_1}^2$, and the result follows.

Recall that if LP-IV is implemented using the control variables $W_t = Y_{t-1}, Y_{t-2}, \dots$, then $\hat{\Theta}_{0,1}^{LP-IV} = \hat{\Theta}_{0,1}^{SVAR-IV}$. If so, the values of λ and $\sigma_{\varepsilon_1}^2$ computed using LP-IV and SVAR-IV are the same, as is the expression in the denominator of (15). Even if LP-IV is implemented using a reduced set of controls or, if Condition LP-IV holds, no controls, the full VAR must be used to obtain the innovations needed to compute λ and $\sigma_{\varepsilon_1}^2$.

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

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

$$\varepsilon_t = \text{Proj}(\varepsilon_t | Y_t, Y_{t-1}, \dots) \quad (\text{invertibility}). \tag{24}$$

In the linear models of this article, condition (24) is equivalent to saying that $\Theta(L)^{-1}$ exists.⁷ The reason we state the invertibility condition as (24) is that it is closer to the standard definition, $\varepsilon_t = E(\varepsilon_t | Y_t, Y_{t-1}, \dots)$, which applies to non-linear models as well.

In this subsection, we make four points. First, we show that (24), plus the assumption that the innovation covariance matrix is non-singular, implies (17). Second, we reframe (24) 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-IV¹ holds is equivalent to assuming that Condition SVAR-IV and invertibility (24) both hold.

2.2.1. *Demonstration that invertibility (24) implies (17)*

This result is well known but we show it here for completeness. Recall that by definition, $v_t = Y_t - \text{Proj}(Y_t | Y_{t-1}, Y_{t-2}, \dots) = \Theta(L)\varepsilon_t - \text{Proj}(\Theta(L)\varepsilon_t | Y_{t-1}, Y_{t-2}, \dots) = \Theta_0\varepsilon_t + \sum_{i=1}^{\infty} \Theta_i[\varepsilon_{t-i} - \text{Proj}(\varepsilon_{t-i} | Y_{t-1}, Y_{t-2}, \dots)]$, where the second equality uses (5), and the third equality uses the fact that $\text{Proj}(\varepsilon_t | Y_{t-1}, Y_{t-2}, \dots) = 0$ and collects terms. Equation (24) implies that $\text{Proj}(\varepsilon_{t-i} | Y_{t-1}, Y_{t-2}, \dots) = \varepsilon_{t-i}$, so the term in brackets in the final summation is zero for all i ; thus we have that $v_t = \Theta_0\varepsilon_t$ as in (17).

To see why (24) implies that Θ_0 is invertible, note that $\varepsilon_t = \text{Proj}(\varepsilon_t | Y_t, Y_{t-1}, \dots) = \text{Proj}(\varepsilon_t | v_t, v_{t-1}, \dots) = \text{Proj}(\varepsilon_t | \Theta_0\varepsilon_t, \Theta_0\varepsilon_{t-1}, \dots) = \text{Proj}(\varepsilon_t | \Theta_0\varepsilon_t) = \text{Proj}(\varepsilon_t | v_t)$, where the first equality is (24), the second follows because current and past innovations span the space of current and past Y 's, the third and fifth follow from $v_t = \Theta_0\varepsilon_t$, and the fourth follows from the serial independence of ε_t . Because $\varepsilon_t = \text{Proj}(\varepsilon_t | v_t)$, the

⁷ By $\Theta(L)^{-1}$ existing we mean that it is a square-summable limit of a sequence of matrix polynomials in positive powers of L .

equation $v_t = \Theta_0 \varepsilon_t$ must yield a unique solution for ε_t , so that Θ_0 has rank m . Moreover, because $\text{var}(v_t)$ is assumed to have full rank, $n \leq m$. Taken together these imply that $n = m$ and Θ_0 has rank n . Therefore, if (24) holds, then (17) holds.

2.2.2. *Invertibility as omitted variables*

One interpretation provided in the literature on invertibility is that invertibility implies that there are no omitted variables in the VAR (e.g. Fernández-Villaverde *et al.*, 2007): because invertibility implies that the spans of ε_t and v_t are the same, there is no forecasting gain from adding past shocks to the VAR. That is, the invertibility condition (24) implies that⁸

$$\text{Proj}(Y_t | Y_{t-1}, Y_{t-2}, \dots, \varepsilon_{t-1}, \varepsilon_{t-2}, \dots) = \text{Proj}(Y_t | Y_{t-1}, Y_{t-2}, \dots). \quad (25)$$

Condition (25) 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 past 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

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, a high-dimensional dynamic factor model or a factor-augmented vector autoregression (FAVAR). This is a potentially useful avenue to dealing with the invertibility problem; see, for example, Forni *et al.* (2009) and the survey in Stock and Watson (2016).¹⁰

It is important to note that expanding the number of variables will not necessarily result in (24) being satisfied, so that moving to large systems does not assure invertibility.

2.2.3. *Relation between assumptions SVAR-IV, LP-IV[⊥] and invertibility*

A major appeal of LP-IV is that the direct regression approach does not explicitly assume invertibility. If, however, the instrument depends on lagged shocks and lagged Y s are used as control variables, then in general the instrument is valid with these controls (i.e. Condition LP-IV[⊥] holds) if and only if Condition SVAR-IV holds and that the SVAR is invertible. Intuitively, if the instrument depends on lagged shocks, the

⁸ Equation (25) follows from (17) by writing, $\text{Proj}(Y_t | Y_{t-1}, Y_{t-2}, \dots, \varepsilon_{t-1}, \varepsilon_{t-2}, \dots) = \text{Proj}(Y_t | v_{t-1}, v_{t-2}, \dots, \varepsilon_{t-1}, \varepsilon_{t-2}, \dots) = \text{Proj}(Y_t | v_{t-1}, v_{t-2}, \dots) = \text{Proj}(Y_t | Y_{t-1}, Y_{t-2}, \dots)$, where the first and third equalities use the fact that the innovations are the Wold errors, and the second equality uses the implication of (17) that $\text{span}(\varepsilon_t) = \text{span}(v_t)$.

⁹ Condition (25) 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.

¹⁰ Aikman *et al.* (2016) use lagged macro factors as controls in local projection OLS regression, which they call factor-augmented local projections. This method is the local projection counterpart of FAVARs.

control variables must span the space of those shocks but the requirement that the Y s span the space of the shocks is simply the invertibility condition. This result is stated in the following theorem.

THEOREM 1. *Let \mathbf{Z} denote the set of scalar stochastic processes (instruments) such that for all $Z \in \mathbf{Z}$, Z satisfies LP-IV Conditions (i), (ii) and (iii for $j > 0$), but not (iii for $j < 0$). Let $W_t = \{Y_{t-1}, Y_{t-2}, \dots\}$. Then LP-IV $^\perp$ is satisfied for all $Z \in \mathbf{Z}$ if and only if (a) Z satisfies Condition SVAR-IV and (b) the invertibility condition (24) holds.*

Proof. We first show that condition SVAR-IV plus invertibility (24) implies Condition LP-IV $^\perp$. First note that for $j \geq 0$, $\text{Proj}(\varepsilon_{t+j}|Y_{t-1}, Y_{t-2}, \dots) = 0$ so $\varepsilon_{t+j}^\perp = \varepsilon_{t+j} - \text{Proj}(\varepsilon_{t+j}|Y_{t-1}, Y_{t-2}, \dots) = \varepsilon_{t+j}$. Thus, for $j \geq 0$, $E(\varepsilon_{t+j}^\perp Z_t^\perp) = E\{\varepsilon_{t+j}[Z_t - \text{Proj}(Z_t|Y_{t-1}, Y_{t-2}, \dots)]\} = E(\varepsilon_{t+j} Z_t)$. Setting $j = 0$, it follows that SVAR-IV (i) and (ii) are equivalent to LP-IV $^\perp$ (i) and (ii). In addition, Condition LP-IV (iii for $j > 0$) (which holds by definition of \mathbf{Z}) is equivalent to Condition LP-IV $^\perp$ (iii for $j > 0$). For $j < 0$, (24) directly implies that $\varepsilon_{t+j} = \text{Proj}(\varepsilon_{t+j}|Y_{t-1}, Y_{t-2}, \dots)$, so $\varepsilon_{t+j}^\perp = 0$ and thus $E(\varepsilon_{t+j}^\perp Z_t^\perp) = 0$ trivially; thus (24) implies LP-IV $^\perp$ (iii for $j < 0$). Thus condition SVAR-IV plus (24) implies Condition LP-IV $^\perp$ for all $Z \in \mathbf{Z}$.

We now show that, if Condition LP-IV $^\perp$ holds for all $Z \in \mathbf{Z}$, then Conditions SVAR-IV and (24) hold. First, as noted above, LP-IV $^\perp$ (i) and (ii) are equivalent to SVAR-IV (i) and (ii). It remains to show that, if LP-IV $^\perp$ (iii) holds for all $Z \in \mathbf{Z}$, then (24) holds. Consider $\tilde{Z} \in \mathbf{Z}$, and let $\tilde{Z}_t = \tilde{Z}_t + \varepsilon_{t-1}$; by construction, $\tilde{Z} \in \mathbf{Z}$. Because LP-IV $^\perp$ holds by assumption for all $Z \in \mathbf{Z}$, it holds in particular for \tilde{Z} and \tilde{Z} , so LP-IV $^\perp$ (iii, $j < 0$) implies that $E(\varepsilon_{t-1}^\perp \tilde{Z}_t^\perp) = E(\varepsilon_{t-1}^\perp \tilde{Z}_t^\perp) = 0$. But $E(\varepsilon_{t-1}^\perp \tilde{Z}_t^\perp) = E(\varepsilon_{t-1}^\perp \tilde{Z}_t^\perp) + E(\varepsilon_{t-1}^\perp)^2$, so it must be that $E(\varepsilon_{t-1}^\perp)^2 = 0$; but $E(\varepsilon_{t-1}^\perp)^2 = 0$ implies that (24) holds.

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 in general 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.¹¹

2.3. Observable Shocks, VAR Misspecification and Partial Invertibility

The external instrument approach to impulse response estimation treats shock measures, such as the Romer and Romer (1989) narrative shocks or a monetary announcement surprise as in Kuttner (2001), as instrumental variables. Originally, however, that literature treated those measures as the shocks directly. Given our focus on invertibility, we therefore briefly digress to consider issues of VAR specification

¹¹ 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 (2011) 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.

when the shock of interest is observed. We will refer to the situation in which $\varepsilon_{1,t}$ is observed, or at least is recoverable from the VAR innovations v_b as partial invertibility: we will say that the VAR is partially invertible if there is some λ such that $\varepsilon_{1,t} = \lambda'v_t$. The leading case is the observed shock case in which $\lambda = (1 \ 0 \ \dots \ 0)'$, with the observed shock ordered first in the VAR. Here, we first consider partial identification in the case that λ is identified without assuming full invertibility (the ‘observed shock’ case), so that the shock can be used directly as a regressor. We then contrast this with the case of identification by external instruments.

First, consider the case that $\varepsilon_{1,t}$ is observed, and let $Y_{1,t} = \varepsilon_{1,t}$, and as usual let $Y_{2:n,t}$ denote the remaining Y s. Write the structural moving average representation for $Y_{2:n,t}$ as $Y_{2:n,t} = \Theta_1(L)\varepsilon_{1,t} + \omega_t$, where ω_t is the distributed lag all the shocks other than $\varepsilon_{1,t}$. Because ω_t is stationary, it has a population VAR representation, $\omega_t = A_{22}(L)\omega_{t-1} + \zeta_t$. Premultiplying $Y_{2:n,t} = \Theta_1(L)\varepsilon_{1,t} + \omega_t$ by $I - LA_{22}(L)$ and rearranging yields, $Y_{2:n,t} = [I - LA_{22}(L)]\Theta_1(L)\varepsilon_{1,t} + A_{22}(L)Y_{2:n,t-1} + \zeta_t = A_{21}(L)\varepsilon_{1,t-1} + A_{22}(L)Y_{2:n,t-1} + \Theta_{0,1}\varepsilon_{1,t} + \zeta_t$, where $A_{21}(L) = L^{-1}\{[I - LA_{22}(L)]\Theta_1(L) - \Theta_{0,1}\}$ (note that the leading term of $[I - LA_{22}(L)]\Theta_1(L)$ is $\Theta_{0,1}$). The expressions for $Y_{1,t}$ and $Y_{2:n,t}$ combine to yield the VAR:

$$Y_t = \begin{pmatrix} Y_{1,t} \\ Y_{2:n,t} \end{pmatrix} = \begin{pmatrix} 0 & 0 \\ A_{21}(L) & A_{22}(L) \end{pmatrix} \begin{pmatrix} Y_{1,t-1} \\ Y_{2:n,t-1} \end{pmatrix} + \begin{pmatrix} v_{1,t} \\ v_{2:n,t} \end{pmatrix}, \tag{26}$$

$$\text{where } \begin{pmatrix} v_{1,t} \\ v_{2:n,t} \end{pmatrix} = \begin{pmatrix} 1 & 0 \\ \Theta_{0,1} & I \end{pmatrix} \begin{pmatrix} \varepsilon_{1,t} \\ \zeta_t \end{pmatrix}.$$

Assuming correct lag specification, the VAR coefficient estimator is consistent for the population lag matrix in (26). The lack of feedback in the population VAR coefficient matrix to the first variable, combined with the lower triangular error structure in (26), imply that the IRFs produced by a Cholesky factorisation of the VAR innovations, with the observed shock ordered first, produce an IRF that simply iterates on the second block of equations. That is, the IRF is computed from the difference equation $Y_{2:n,t} = [I - LA_{22}(L)]\Theta_1(L)\varepsilon_{1,t} + A_{22}(L)Y_{2:n,t-1}$, which yields the IRF $\Theta_1(L)$.

The conclusion that the VAR ‘ $\varepsilon_{1,t}$ first’ IRF is consistent for $\Theta_1(L)$ was reached without ever assuming that ζ_t spans the space of the remaining shocks: the VAR can have omitted variables in the sense that the shocks are not fully observable. The reason for this result is that $\varepsilon_{1,t}$ is strictly exogenous. Because of this strict exogeneity, $\Theta_1(L)$ can be consistently estimated by a distributed lag regression of $Y_{2:n,t}$ on $\varepsilon_{1,t}$, an autoregressive distributed lag regression, by GLS, or using a VAR with arbitrary choice of VAR variables, including a choice of VAR variables that differs from one variable of interest to the next.

These observations all extend to the case of partial invertibility, in which there is an identified λ such that $\varepsilon_{1,t} = \lambda'v_t$. Let $\tilde{\lambda}$ be a $n \times (n - 1)$ matrix such that $\tilde{\lambda}'\lambda = 0$ and $\tilde{\lambda}'\tilde{\lambda} = I$. Then the algebra of the preceding paragraph goes through using the transformed variables $\tilde{Y}_t = (\tilde{Y}_{1,t}, \tilde{Y}_{2:n,t}) = (\lambda'Y_t, \tilde{\lambda}'Y_t)$.

Returning to IV methods, an implication of these observations is that if the IV methods identify λ such that $\varepsilon_{1,t} = \lambda'v_b$, then the additional assumption of invertibility of the SVAR can be dispensed with for the validity of SVAR-IV. This said, as discussed in

subsection 2.2, identification of $\Theta_{0,1}$ is insufficient to identify λ , and the expression for λ given there (that $\lambda = \Theta_{0,1}'\Sigma_{vv}^{-1}/(\Theta_{0,1}'\Sigma_{vv}^{-1}\Theta_{0,1})$) was derived under the invertibility assumption (17). While the partial invertibility assumption that $\varepsilon_{1,t} = \lambda'v_t$ is weaker than invertibility assumption (17), it remains to be seen whether there are empirical applications in which this weaker condition would hold but invertibility does not.¹²

3. 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, at least under homoscedasticity. 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.

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

3.1. Null and Local Alternative

Under invertibility (24), the structural moving average can be written $Y_t = C(L)\Theta_0\varepsilon_t$ as in (18), where $C(L) = A(L)^{-1}$; that is, that $\Theta(L) = C(L)\Theta_0$. The null and alternative hypotheses thus are

$$H_0 : C_h\Theta_{0,1} = \Theta_{h,1}, \text{ all } h \text{ v. } H_1 : C_h\Theta_{0,1} \neq \Theta_{h,1}, \text{ some } h. \tag{27}$$

In addition to establishing the null distribution of the test, we wish to examine its distribution under an alternative to check that the test has power against non-invertibility. Beaudry *et al.* (2015) and Plagborg-Møller (2016*b*) provide numerical evidence that in many cases the non-invertible (non-fundamental) representation of a time series may be very close to its invertible representation. With this motivation, we focus on non-invertible IRFs that represent small departures from an invertible null.

Specifically, we consider the drifting sequence of alternatives:

$$C_{h,T}\Theta_{0,1} = \Theta_{h,1} + T^{-1/2}d_h + o(T^{-1/2}), \tag{28}$$

where under the null $d_h = 0$, while under the alternative d_h is a non-zero $n \times 1$ vector for at least some $h > 0$. In Appendix A.1, we construct a sequence of models that are non-invertible because of a small amount (specifically, $O_p(T^{-1/4})$) of measurement

¹² Evidently, without partial invertibility or recoverability, the historical and forecast error variance decompositions in (14) and (15) are not point-identified. Plagborg-Møller and Wolf (2017) derive set identification results for these decompositions using external instruments in the absence of invertibility or recoverability.

error contamination, and show that this sequence of models induces local non-invertibility of the form (28).

3.2. Test of Invertibility

We now turn to the test statistic. Let $\hat{\theta}^{\text{SVAR-IV}}$ denote an $m \times 1$ vector of SVAR-IV estimators (23), computed using a VAR(p), for different variables and/or horizons and let $\hat{\theta}^{\text{LP-IV}}$ 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 Appendix A that, with strong instruments and under standard moment/memory assumptions, under the null and local alternative:

$$\sqrt{T}(\hat{\theta}^{\text{LP-IV}} - \hat{\theta}^{\text{SVAR-IV}}) \xrightarrow{d} \text{N}(d, V), \quad (29)$$

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

The Hausman-type test statistic is

$$\xi = T(\hat{\theta}^{\text{LP-IV}} - \hat{\theta}^{\text{SVAR-IV}})' \hat{V}^{-1} (\hat{\theta}^{\text{LP-IV}} - \hat{\theta}^{\text{SVAR-IV}}), \quad (30)$$

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

We make four remarks about this test:

- (i) We suggest computation of the variance matrix \hat{V} using the parametric bootstrap, and we discuss some specifics in Appendix A.2.
- (ii) The LP-IV and SVAR-IV estimators for the impact effect ($h = 0$) are identical when lagged Y s 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.
- (iii) Under the local alternative (28), the test statistic has a non-central chi-squared distribution with m degrees of freedom and non-centrality parameter $\mu^2 = d' V^{-1} d$. The expressions in Appendix A show that, for a given local alternative d , the non-centrality parameter converges to zero as $\alpha \rightarrow 0$, and increases to a finite limit as $|\alpha|$ increases. Thus, the power of the test is increasing as the strength of the instrument increases, according to this local strong-instrument approximation.
- (iv) Existing tests for invertibility (Forni and Gambetti, 2014) test the implication of invertibility that Z_t does not Granger-cause Y_t . The test here differs because it focuses not on forecasting contribution, but on the object of interest in the

analysis, the impulse response function. In both approaches – directly testing Granger non-causality and the Hausman-type test approach here, the testable implications all stem from moments involving Z : second moments of Y alone cannot distinguish invertible from non-invertible processes.

4. Illustration: Gertler and 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 US 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 one-year US 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 the 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. However, the original literature treated such a measure as the shock and GK use it as an instrument, that is, $Z_t = FFF_t$.

Column (1) 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 normalisation. 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. 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 (5) to highlight the various components of $Y_{i,t+h}$:

$$Y_{i,t+h} = \Theta_{h,i1} \varepsilon_{1+h} + \{\varepsilon_{t+h}, \dots, \varepsilon_{t+1}\} + \{\varepsilon_{2:n,t}\} + \{\varepsilon_{t-1}, \dots\}, \quad (31)$$

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 (31), the error term in the first-stage regression is comprised of $\{\varepsilon_{2:n,t}\}$ and

Table 1

Estimated Causal Effect of Monetary Policy Shocks on Selected Economic Variables: Gertler and Karadi (2015) Variables, Instrument and Sample Period

	lag (<i>h</i>)	LP-IV			SVAR	SVAR – LP
		(1)	(2)	(3)	(4)	(5) ((4) – (2))
<i>R</i>	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)
<i>IP</i>	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)
<i>P</i>	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)
<i>EBP</i>	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 (<i>z</i> , <i>y</i>)	4 lags of (<i>z</i> , <i>y</i> , <i>f</i>)	12 lags of <i>y</i> 4 lags of <i>z</i>	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 to 2012m6; the other variables are available from 1979m1 to 2012m6. The LP-IV estimates in (1)–(3) use data from 1990m1 to 2012m6. The VAR for (4) 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 (5). 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 (1) and heteroscedasticity-robust (no lags) in (2), (3) and (4).

$\{\varepsilon_{t-1}, \dots\}$. 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 $\{\varepsilon_{t+h}, \dots, \varepsilon_{t+1}\} + \{\varepsilon_{2:n,t}\} + \{\varepsilon_{t-1}, \dots\}$.

Column (2) 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 lagged values of ε , they reduce the variance of the regression error term and, for example, the first-stage (partial) R^2 in (2) increases to $R^2 = 0.09$ with a first-stage F -statistic increasing to $F = 23.7$. Second, the controls adjust for a data processing issue that makes the FFF variable an invalid instrument in the LP-IV 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 $t - 1$.

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 a MA(1) structure, using lags of Z_t as controls eliminates the correlation with $\varepsilon_{1,t-1}$, so that Condition LP-IV[⊥] (iii) is satisfied. Despite the MA(1) structure, it is plausible that this instrument is uncorrelated with other shocks. Thus, to satisfy Condition LP-IV[⊥] (iii), it would suffice to include lags of Z s as a controls; including lagged Y s and additional lags of Z serves to improve precision (increase the first-stage F).¹³

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 Z s; 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 (3) adds lags of four factors computed from the FRED-MD data set (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 (2) and (3) in Table 1 improve on the model without controls, (1), by eliminating some of the variability associated with lagged ε and in particular by making Z satisfy LP-IV[⊥] (iii), whereas (1) does not satisfy LP-IV (iii). However, neither eliminates the variability associated with *future* ε 's, the $\{\varepsilon_{t+h}, \dots, \varepsilon_{t+1}\}$ component of the error term shown in (31). 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 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 (4) 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 subsection 2.1). Standard errors for the SVAR-IV estimate are computed by the parametric bootstrap described in Appendix A. Because the VAR uses 12 lags of Y instead of the four lags used as controls in the local projections, the first stage F -statistics differ slightly in columns (2) and (4). 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 (31)), and these terms are absent from the IV regression used in the SVAR, because

¹³ The construction of Z_t is described in footnote 6 in GK. The MA(1) structure invalidates the LP-IV regression reported in column (1), but it does not affect its validity in the SVAR-IV regression used by GK. An additional issue is that the weights used in GK's construction of Z_t are time varying because of floating FOMC meeting dates. In principle, this could yield a time-varying MA(1) structure but we approximate the MA coefficients as constant.

Table 2
Tests for VAR Invertibility (p-values)

	1 year 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 (30) 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.

only the impact effect, Θ_0 , is estimated by IV. Second, the VAR parameterisation 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 (2) and the SVAR-IV estimator in column (4). 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 (30) 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.

5. Conclusions

It is well known that, with Gaussian errors, every invertible model has multiple observationally equivalent non-invertible 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 non-Gaussian, and to exploit higher order moment restrictions to identify the causal structure (Lanne and Saikkonen, 2013; Gospodinov and Ng, 2015; Gouriéroux *et al.*, 2017). A second approach is to use *a priori* informative priors (Plagborg-Møller, 2016*b*). Here, we have shown that there is a third approach, which is to use an external instrument. Through an external instrument, additional information can be brought to bear to identify dynamic causal effects. Under a lead-lag exogeneity condition, the external instrument identifies the structural impulse response function without assuming invertibility.

A number of methodological issues concerning the use of external instruments merit further research. For example, this discussion assumes homogenous treatment effects. Although this assumption seems plausible in a macroeconomic setting (there is only one ‘subject’, although effects may be state-dependent), more work is warranted. Also, the usual weak-instrument toolkit does not cover all the methods used here, for example, one open question is how to robustify our test of invertibility to potentially weak instruments.

Additionally, an informal argument sometimes made in favour of the local projections method is that it is robust to VAR misspecification concerning lag length, non-linearities and state dependence. In this article, we have put these arguments to one side by assuming a linear, constant-coefficient structural moving average representation. To us, the robustness of LP-IV to non-linearities is not obvious, particularly when the instrument depends in part on lagged shocks: if so, the control variables would need to span the space of those shocks, and it seems that there would be a non-linear counterpart to our no free lunch theorem (Theorem 1). In any event, it would be of interest to see these arguments made precise.

In our view, the most exciting work to be done in this area is empirical. We look forward to the development of new external instruments that provide plausibly exogenous variation to provide more credible identification of dynamic causal effects.

Appendix A. Asymptotics and Bootstrap Implementation

A.1. Asymptotic Distribution of the Hausman Test Statistic for Invertibility

This Appendix derives the asymptotic distribution (29) under the null of invertibility and under a sequence of local alternatives. For simplicity, we consider the case that the test is based on all impulse responses for a single horizon h and that the instrument is a scalar; extensions to multiple horizons and a vector of instruments is straightforward. Accordingly, we show that $T^{1/2}(\hat{\Theta}_{h,1}^{SVAR-IV} - \hat{\Theta}_{h,1}^{LP-IV}) \xrightarrow{d} N(d_h, V_h)$. This result implies that the test statistic ζ given in (30) has an asymptotic chi-squared distribution with n degrees of freedom under the null, and a non-central chi-squared distribution with non-centrality parameter $\mu^2 = d_h' V_h^{-1} d_h$ under the local alternative.

We begin with the analysis under the null of invertibility. The SVAR is

$$A(L)Y_t = \Theta(L)\varepsilon_t, \tag{A.1}$$

where $A(L)$ is a polynomial of order p . The Wold moving average polynomial is $C(L) = A(L)^{-1} = I + C_1L + \dots$. Under the null hypothesis of invertibility (17), with the maintained hypothesis that Y_t has the linear structural MA representation (5), the structural IRF satisfies H_0 in (27), that is, $\Theta_{h,1} = C_h\Theta_{0,1}$ for all h .

For future reference, we note that the SVAR can be written in state-space form as

$$\begin{aligned} Y_t &= SX_t, \\ X_t &= AX_{t-1} + G\varepsilon_t, \end{aligned} \tag{A.2}$$

where $X_t = (Y_t' \ Y_{t-1}' \ \dots \ Y_{t-p+1}')'$, A is the VAR companion matrix, the upper block of G is Θ_0 and all other elements of G are zero and $S = (I_n \ 0 \ \dots \ 0)$ is a selection matrix.

The local projection equation, written for the vector Y , is

$$Y_{t+h} = \Theta_{h,1} Y_{1,t} + \Gamma_h W_t + u_{t+h}^\perp, \tag{A.3}$$

where the control variables are $W_t = X_{t-1}$ and, from (A.2), $\Gamma_h = SA^{h+1}$. Consistent with Assumption LP-IV $^\perp$, we represent Z_t as

$$Z_t = \beta \varepsilon_{1,t} + B' W_t + e_t, \tag{A.4}$$

where e_t is uncorrelated with ε_s for all t and s . All variables are assumed to be second-order stationary with sample moments that satisfy

$$T^{-1/2} \sum \text{vec}(a_t b_t' - E(a_t b_t')) \xrightarrow{d} N(0, \Sigma_{ab}), \tag{A.5}$$

for any variables (a_b, b_t) .

The LP-IV estimator is

$$\hat{\Theta}_{h,1}^{\text{LP-IV}} = (Z' M_W Y_{1,0})^{-1} (Y_h' M_W Z), \tag{A.6}$$

where Z denotes the $T \times 1$ vector of instruments, $Y_{1,0}$ denotes the $T \times 1$ vector $(Y_{1,1} \dots Y_{1,T})'$, Y_h denotes the $T \times n$ matrix with t 'th row Y_{t+h}' and $M_W = I - W(W'W)^{-1}W'$, where W is a $T \times (np)$ matrix with t th row W_t' . The SVAR-IV estimator is as follows:

$$\hat{\Theta}_{h,1}^{\text{SVAR-IV}} = \hat{C}_h \hat{\Theta}_{0,1}^{\text{LP-IV}}, \tag{A.7}$$

where $\hat{C}(L) = \hat{A}(L)^{-1}$ where $\hat{A}(L)$ is the OLS estimator of $A(L)$.

Under H_0 in (27) and the assumption that Z_t is a strong instrument, a straightforward calculation then yields

$$\begin{aligned} T^{1/2} (\hat{\Theta}_{h,1}^{\text{SVAR-IV}} - \hat{\Theta}_{h,1}^{\text{LP-IV}}) &= T^{1/2} (\hat{C}_h - C_h) \Theta_{0,1}^{\text{LP-IV}} \\ &+ \alpha^{-1} \left(C_h T^{-1/2} \sum Z_t^\perp u_t^{0\perp} - T^{-1/2} \sum Z_t^\perp u_{t+h}^{\perp} \right) + o_p(1) \\ &\xrightarrow{d} N(0, V_h), \end{aligned} \tag{A.8}$$

where the result uses $Z' M_W Y_{1,0} = Z^\perp M_W Y_{1,0}^\perp = Z^\perp Y_{1,0}^\perp + O_p(1)$ and similarly for $Y_h' M_W Z$, $T^{-1} Z^\perp Y_{1,0}^\perp \xrightarrow{p} \alpha = E(Z_t^\perp \varepsilon_{1,t})$, and the delta-method.

We now consider a sequence of stochastic processes that are local to the invertible model and the resulting estimators. Specifically, maintain the definitions of all of the variables and parameters given above (so that Y_t is generated by the invertible model, etc.), but now consider the sequence of stochastic processes $Y_{t,T}$:

$$Y_{t,T} = S X_t + T^{-1/4} \eta_t, \tag{A.9}$$

where η_t is white noise and uncorrelated with ε_τ for all t and τ . Notice that $Y_{t,T} = Y_t + T^{-1/4} \eta_t$, so that $\text{var}(Y_{t,T}) = \text{var}(Y_t) + T^{-1/2} \text{var}(\eta_t)$, and the autocovariances of $Y_{t,T}$ and Y_t coincide for all non-zero lags. The measurement error $T^{-1/4} \eta_t$ in (A.9) means that X_t cannot be perfectly recovered from current and lagged values of $Y_{t,T}$ and $\varepsilon_t \neq \text{Proj}(\varepsilon_t | Y_{t,T}, Y_{t-1,T}, \dots)$, so the model is not invertible.

The implied p -th order VAR for $Y_{t,T}$ is local to the VAR for Y_t ; that is

$$A_T(L) = A(L) + T^{-1/2} a(L) + o(T^{-1/2}), \tag{A.10}$$

where $A_T(L)$ denotes the projection of $Y_{t,T}$ onto $(Y_{t-1,T}, \dots, Y_{t-p,T})$. Similarly, the implied moving coefficients, $A_T(L)^{-1} = I + C_{1,T}L + \dots$ satisfy $C_{h,T} = C_h + T^{-1/2} c_h + o(T^{-1/2})$. Because $C_h \Theta_0 = \Theta_h$ (the invertible null), we have that $C_{h,T} \Theta_0 = C_h \Theta_0 + T^{-1/2} c_h \Theta_0 + o(T^{-1/2}) = \Theta_h +$

$T^{-1/2}d_h + o(T^{-1/2})$, where $d_h = c_h\Theta_0$. Thus, the local contamination in (A.9) implies the nearly invertible moving average sequence (28).

Let $\hat{A}_T(L)$ denote the OLS estimator of $A_T(L)$ using $Y_{i,T}$. A calculation shows that $T^{1/2}(\hat{A}_T(L) - \hat{A}(L)) = a(L) + o_p(1)$ and:

$$T^{1/2}(\hat{C}_{h,T} - \hat{C}_h) = c_h + o_p(1). \tag{A.11}$$

Although the VAR and MA models for Y_i and $Y_{i,T}$ differ by a $T^{-1/2}$ component, the LP-IV estimators using $Y_{i,T}$ and Y_i are equivalent to order $T^{-1/2}$. To see this, write the LP equation as:

$$Y_{t+h,T} = \Theta_{h,1}Y_{1,t,T} + \Gamma_{h,T}W_{i,T} + u_{t+h,T}^{\perp}, \tag{A.12}$$

where $W_{i,T} = (Y_{t-1,T} \dots Y_{t-p,T})$. From (A.2) and (A.12), $\Gamma_{h,T}W_{i,T} = SA^{h+1} \times \text{Proj}(X_{t-1}|W_{i,T})$ and $u_{t+h,T}^{\perp} = u_{t+h}^{\perp} + T^{-1/4}\eta_{t+h} + g_{t+h}^h$, where $g_{t+h}^h = SA^{h+1} \times [X_{t-1} - \text{Proj}(X_{t-1}|W_{i,T})] = O_p(T^{-1/4})$. Similarly, let the instruments satisfy:

$$Z_{i,T} = \beta\varepsilon_{1,t} + B'W_{i,T} + e_t, \tag{A.13}$$

where now e_t is assumed to be uncorrelated with ε_τ and η_τ for all t and τ . Using instruments that satisfy (A.13) ensures that Condition LP-IV^L holds under both the null and local alternative. Let $\hat{\Theta}_{h,1}^{LP-IV}(\{Y_{i,T}, Z_{i,T}\})$ denote the LP-IV estimators using $\{Y_{i,T}, Z_{i,T}\}$. Using (A.12) and (A.13), it follows that:

$$T^{1/2} \left[\hat{\Theta}_{h,1}^{LP-IV}(\{Y_{i,T}, Z_{i,T}\}) - \hat{\Theta}_{h,1}^{LP-IV}(\{Y_t, Z_t\}) \right] = o_p(1). \tag{A.14}$$

Finally, the SVAR estimator constructed from $\{Y_{i,T}, Z_{i,T}\}$ is

$$\hat{\Theta}_h^{SVAR-IV}(\{Y_{i,T}, Z_{i,T}\}) = \hat{C}_{h,T} \hat{\Theta}_0^{LP-IV}(\{Y_{i,T}, Z_{i,T}\}). \tag{A.15}$$

Equations (A.8), (A.11), (A.14) and (A.15) imply:

$$\begin{aligned} & T^{1/2} \left[\hat{\Theta}_{h,1}^{SVAR-IV}(\{Y_{i,T}, Z_{i,T}\}) - \hat{\Theta}_{h,1}^{LP-IV}(\{Y_{i,T}, Z_{i,T}\}) \right] \\ &= T^{1/2} \left[\hat{\Theta}_{h,1}^{SVAR-IV}(\{Y_t, Z_t\}) - \hat{\Theta}_{h,1}^{LP-IV}(\{Y_t, Z_t\}) \right] + c_h\Theta_{0,1} + o_p(1) \\ &\xrightarrow{d} N(d_h, V_h), \end{aligned} \tag{A.16}$$

where $d_h = c_h\Theta_{0,1}$.

A.2. Parametric Bootstrap Estimation of V_h

The standard errors of the estimators in Tables 1 and 2 were computed using the sample variances computed from 1,000 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:

$$\begin{bmatrix} \hat{A}(L) & 0 \\ 0 & \hat{\rho}(L) \end{bmatrix} \begin{bmatrix} \tilde{Y}_t \\ \tilde{Z}_t \end{bmatrix} = \begin{bmatrix} \tilde{v}_t \\ \tilde{e}_t \end{bmatrix}, \text{ where } \begin{bmatrix} \tilde{v}_t \\ \tilde{e}_t \end{bmatrix} \sim i.i.d. N \left(\begin{bmatrix} 0 \\ 0 \end{bmatrix}, \begin{bmatrix} S_{\tilde{v}\tilde{v}} & S_{\tilde{v}\tilde{e}} \\ S_{\tilde{e}\tilde{v}} & S_{\tilde{e}\tilde{e}} \end{bmatrix} \right), \tag{A.17}$$

where $\hat{A}(L)$ is estimated from a VAR(12), $\hat{\rho}(L)$ is estimated from an AR(4) and $S_{\tilde{v}\tilde{v}}$, $S_{\tilde{v}\tilde{e}}$ and $S_{\tilde{e}\tilde{e}}$ are sample covariances for the VAR/AR residuals. These samples are used to compute the SVAR-IV and LP-IV estimates of $\Theta_{h,1}$.

*Harvard University and the National Bureau of Economic Research
Woodrow Wilson School, Princeton University and the National Bureau of Economic Research*

Additional Supporting Information may be found in the online version of this article:

Data S1.

References

- Aikman, D., Bush, O. and Taylor, A.M. (2016), 'Monetary versus macroprudential policies: causal impacts of interest rates and credit controls in the era of the UK Radcliffe report', NBER Working Paper No. 22380.
- Andrews, I. (2018), 'Valid two-step identification-robust confidence sets for GMM', *Review of Economics and Statistics*. https://doi.org/10.1162/REST_a_00682
- Angrist, J.D., Jordà, O. and Kuersteiner, G.M. (2018), 'Semiparametric estimates of monetary policy effects: string theory revisited', *Journal of Business and Economic Statistics*. <https://doi.org/10.1080/07350015.2016.1204919>
- Barnichon, R. and Brownlees, C. (2016), 'Impulse response estimation by smooth local projections', CEPR Discussion Paper No. DP11726.
- Beaudry, P., Fève, P., Guay, A. and Portier, F. (2015), 'When is non-fundamentalness in VARs a real problem? An application to news shocks', manuscript, University of British Columbia.
- Beaudry, P. and Saito, M. (1998), 'Estimating the effects of monetary policy shocks: an evaluation of different approaches', *Journal of Monetary Economics*, vol. 42(2), pp. 241–60.
- Bernanke, B.S. and Kuttner, K.N. (2005), 'What explains the stock market's reaction to federal reserve policy?', *Journal of Finance*, vol. 60(3), pp. 1221–57.
- Bojinov, I. and Shephard, N. (2017), 'Time series experiments, causal estimands, exact p-values, and trading', manuscript, Harvard University.
- Caldara, D. and Kamps, C. (2017), 'The analytics of SVARs: a unified framework to measure fiscal multipliers', *Review of Economic Studies*, vol. 73 (1), pp. 195–218.
- Chahrour, R. and Jurado, K. (2017), 'Recoverability', manuscript, Duke University.
- Cochrane, J.H. and Piazzesi, M. (2002), 'The fed and interest rates: a high-frequency identification', *American Economic Review*, vol. 92(May), pp. 90–5.
- Faust, J., Rogers, J.H., Swanson, E. and Wright, J.H. (2003), 'Identifying the effects of monetary policy shocks on exchange rates using high frequency data', *Journal of the European Economic Association*, vol. 1(5), pp. 1031–57.
- Fernández-Villaverde, J., Rubio-Ramírez, J.F., Sargent, T.J. and Watson, M.W. (2007), 'The ABCs (and Ds) of understanding VARs', *American Economic Review*, vol. 97(3), pp. 1021–6.
- Fieldhouse, A., Mertens, K. and Ravn, M.O. (2017), 'The macroeconomic effects of government asset purchases: evidence from postwar U.S. housing credit policy', NBER Working Paper No. 23154.
- Forni, M. and Gambetti, L. (2014), 'Sufficient information in structural VARs', *Journal of Monetary Economics*, vol. 66(C), pp. 124–36.
- Forni, M., Giannone, D., Lippi, M. and Reichlin, L. (2009), 'Opening the black box: structural factor models with large cross sections', *Econometric Theory*, vol. 25(5), pp. 1319–47.
- Frisch, R. (1933), 'Propagation problems and impulse problems in dynamic economics', in (R. Krammer, ed.), *Economic Essays in Honor of Gustav Cassel*, pp. 171–205, London: Allen & Unwin.
- Gertler, M. and Karadi, P. (2015), 'Monetary policy surprises, credit costs, and economic activity', *American Economic Journal: Macroeconomics*, vol. 7(1), pp. 44–76.
- Gilchrist, S. and Zakrajšek, E. (2012), 'Credit spreads and business cycle fluctuations', *American Economic Review*, vol. 102(4), pp. 1692–720.
- Gospodinov, N. and Ng, S. (2015), 'Minimum distance estimation of possibly noninvertible moving average models', *Journal of Business & Economic Statistics*, vol. 33(3), pp. 403–17.
- Gouriéroux, C., Monfort, A. and Renne, J.-P. (2017), 'Statistical inference for independent component analysis: application to structural VAR models', *Journal of Econometrics*, vol. 196(1), pp. 111–26.
- Gürkaynak, R.S., Sack, B. and Swanson, E. (2005), 'The sensitivity of long-term interest rates to economic news: evidence and implications for macroeconomic models', *American Economic Review*, vol. 95(1), pp. 425–36.
- Hamilton, J.D. (2003), 'What is an oil shock?', *Journal of Econometrics*, vol. 113(2), pp. 363–98.
- Hausman, J.A. (1978), 'Specification tests in econometrics', *Econometrica*, vol. 46(6), pp. 1251–71.
- Hayashi, F. (2000), *Econometrics*, Princeton, NJ: Princeton University Press.
- Imbens, G. (2014), 'Instrumental variables: an econometrician's perspective', *Statistical Science*, vol. 29(3), pp. 323–58.
- Jordà, O. (2005), 'Estimation and inference of impulse responses by local projections', *American Economic Review*, vol. 95(1), pp. 161–82.
- Jordà, O., Schularick, M. and Taylor, A.M. (2015), 'Betting the house', *Journal of International Economics*, vol. 96(Supplement 1), pp. S2–18.
- Jordà, O., Schularick, M. and Taylor, A.M. (2017), 'Large and state-dependent effects of quasi-random monetary experiments', NBER Working Paper No. 23074.

- 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*, vol. 90(2), pp. 216–40.
- Kim, Y.J. and Kilian, L. (2011). 'How reliable are local projection estimators of impulse responses?', *Review of Economics and Statistics*, vol. 93(4), pp. 1460–6.
- Kleibergen, F. (2005). 'Testing parameters in GMM without assuming they are identified', *Econometrica*, vol. 73(4), pp. 1103–23.
- Kuttner, K.N. (2001). 'Monetary policy surprises and interest rates: evidence from the Fed funds futures market', *Journal of Monetary Economics*, vol. 47(3), pp. 523–44.
- Lanne, M. and Saikkonen, P. (2013). 'Noncausal vector autoregression', *Econometric Theory*, vol. 29(3), pp. 447–81.
- Lechner, M. (2009). 'Sequential causal models for the evaluation of labor market programs', *Journal of the American Statistical Association*, vol. 27(1), pp. 71–83.
- McCracken, M. and Ng, S. (2016). 'FRED-MD: a monthly database for macroeconomic research', *Journal of Business and Economic Statistics*, vol. 34(4), pp. 574–89.
- Mertens, K. (2015). 'Bonn summer school: advances in empirical macroeconomics, lecture 2' (slide deck). <https://karelmertens.com.files.wordpress.com/2017/09/lecture2.pdf> (last accessed 22 March 2018)
- Mertens, K. and Ravn, M.O. (2013). 'The dynamic effects of personal and corporate income tax changes in the United States', *American Economic Review*, vol. 103(4), pp. 1212–47.
- Miranda-Agrippino, S. and Ricco, G. (2017). 'The transmission of monetary policy shocks', manuscript, Department of Economics, University of Warwick.
- Montiel Olea, J.L. and Pflueger, C. (2013). 'A robust test for weak instruments', *Journal of Business and Economic Statistics*, vol. 31(3), pp. 358–69.
- Montiel Olea, J., Stock, J.H. and Watson, M.W. (2017). 'Inference in structural VARs with external instruments', manuscript.
- Moreira, M. (2003). 'A conditional likelihood ratio test for structural models', *Econometrica*, vol. 71(4), pp. 1027–48.
- 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.
- Plagborg-Møller, M. and Wolf, C. (2017), 'On Structural Inference with External Instruments', manuscript, Princeton University.
- Ramey, V. (2011). 'Identifying government spending shocks: it's all in the timing', *Quarterly Journal of Economics*, vol. 126(1), pp. 1–50.
- Ramey, V. (2016). 'Macroeconomic shocks and their propagation', in (J. B. Taylor and H. Uhlig, eds.), *Handbook of Macroeconomics*, vol. 2A, pp. 71–162, Amsterdam: Elsevier.
- Ramey, V. and Zubairy, S. (2017). 'Government spending multipliers in good times and in bad: evidence from U.S. historical data', *Journal of Political Economy*, vol. 119(1), pp. 78–121.
- Romer, C.D. and Romer, D.H. (1989). "Does monetary policy matter? A new test in the spirit of Friedman and Schwartz" (with discussion), in (O.J. Blanchard and S. Fischer, eds.), *NBER Macroeconomics Annual 1989*, pp. 121–70, Cambridge, MA: MIT Press.
- Romer, C.D. and Romer, D.H. (2010). 'The macroeconomic effects of tax changes: estimates based on a new measure of fiscal shocks', *American Economic Review*, vol. 100(3), pp. 763–801.
- Rothenberg, T.J. and Leenders, C.T. (1964). 'Efficient estimation of simultaneous equation systems', *Econometrica*, vol. 32(4), pp. 57–76.
- Rudebusch, G.D. (1998). 'Do measures of monetary policy in a VAR make sense?', *International Economic Review*, vol. 39(4), pp. 907–48.
- Sargan, D. (1964). 'Three-stage least-squares and full information maximum likelihood estimates', *Econometrica*, vol. 32(1/2), pp. 77–81.
- Sims, C.A. (1980). 'Macroeconomics and reality', *Econometrica*, vol. 48(1), pp. 1–48.
- Slutsky, E. (1937). 'The summation of random causes as the source of cyclic processes', *Econometrica*, vol. 5(2), pp. 105–46.
- Stock, J.H. (2008). 'What's new in econometrics: time series', lecture 7. Short course lectures, NBER Summer Institute. http://www.nber.org/minicourse_2008.html (last accessed 22 March 2018).
- Stock, J.H. and Watson, M.W. (2012). 'Disentangling the channels of the 2007–2009 recession', *Brookings Papers on Economic Activity*, vol. 1(1), pp. 81–135.
- Stock, J.H. and Watson, M.W. (2016). 'Dynamic factor models, factor-augmented vector autoregressions, and structural vector autoregressions in macroeconomics', in (J.B. Taylor and H. Uhlig, eds.), *Handbook of Macroeconomics*, vol. 2A, pp. 415–525, Amsterdam: Elsevier.
- Theil, H. and Boot, J.C.G. (1962). 'Revue de l'Institut International de Statistique', *Review of the International Statistical Institute*, vol. 30(2), pp. 136–52.
- Wright, P.G. (1928). *The Tariff on Animal and Vegetable Oils*, New York: McMillan.
- Zellner, A. (1962). 'An efficient method of estimating seemingly unrelated regressions and tests for aggregation bias', *Journal of the American Statistical Association*, vol. 57(298), pp. 348–68.
- Zellner, A. and Theil, H. (1962). 'Three-stage least squares: simultaneous estimation of simultaneous equations', *Econometrica*, vol. 30(1), pp. 54–78.